

Alcatel-Lucent Security Management Server (SMS)

Release 9.4

Policy Guide

260-100-016R9.4
Issue 1
June 2009

Alcatel-Lucent - Proprietary

This document contains proprietary information of Alcatel-Lucent and is not to be disclosed or used except in accordance with applicable agreements.

Copyright © 2009 Alcatel-Lucent
Unpublished and Not for Publication
All Rights Reserved

Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners.

The information presented is subject to change without notice. Alcatel-Lucent assumes no responsibility for inaccuracies contained herein.

Copyright © 2009 Alcatel-Lucent. All Rights Reserved.

Contents

About this information product

Purpose	xxiii
Reason for reissue	xxiii
Who Should Read This Book?	xxiii
What Is in this Book	xxiii
What is Not in this Book	xxv
Supported Brick devices	xxvi
Where to Find Technical Support	xxvi
How to comment	xxvi

1 Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets

Overview	1-1
What is a Brick Zone Ruleset?	1-2
Preventing Spoofing Attacks	1-12
To Create a Brick Zone Ruleset	1-14
To Create a Security Rule	1-16
To Add Time and Day Restrictions to a Rule	1-26
To Add Advanced Features to a Rule	1-28
To Add Bandwidth Management to a Rule	1-34
To Add TOS/Alarm Capability to a Rule	1-38
To Configure IP Address(es) for Rules-Based Routing	1-40

	To Enable SCTP Processing on a Rule	1-42
	To Assign a Brick Zone Ruleset to a Physical Port	1-46
	Maintaining Brick Zone Rulesets	1-51
	Maintaining Security Rules	1-59
	Brick Zone Ruleset Templates Provided with the SMS Application for VoIP/NOE Traffic	1-66
2	Host Groups	
	Overview	2-1
	What is a Host Group?	2-2
	To Set Up a Host Group	2-4
	Global Host Groups	2-8
	Dynamic Host Groups	2-10
	Nested Host Groups	2-12
	Maintaining Host Groups	2-15
	Host Groups Provided with the SMS Application	2-22
	Host Group Templates Provided with the SMS Application for VoIP/NOE Traffic	2-25
3	Domain Name Groups	
	Overview	3-1
	Domain Name Groups	3-2
	To Set Up a Domain Name Group	3-3
	Global Domain Name Groups	3-7
	Nested Domain Name Groups	3-10
	Maintaining Domain Name Groups	3-12
4	Service Groups	
	Overview	4-1
	What is a Service Group?	4-2
	To Set Up a Service Group	4-4

Global Service Groups 4-9

Nested Service Groups 4-11

Maintaining Service Groups 4-13

Service Groups Provided with the SMS Application 4-23

Service Group Templates Provided with the SMS Application for VoIP/NOE Traffic 4-27

5 Application Filters

Overview 5-1

Process Flow 5-4

DHCP Relay Application Filter 5-5

DNS Application Filters 5-8

ESP Application Filter 5-18

FTP Application Filter 5-21

GTP Application Filter 5-31

H.323 Application Filters 5-47

HTTP Application Filter 5-53

NOE Application Filter 5-62

NOE Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic 5-68

RPC Application Filters 5-74

RTSP Application Filter 5-79

SIP Application Filters 5-87

SMTP Application Filters 5-105

SQL*Net Application Filter 5-114

TFTP Application Filter 5-117

TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic 5-123

Global Application Filters 5-127

To Copy an Application Filter 5-129

6 Network Address Translation

Overview 6-1

What is Network Address Translation? 6-2

To Set Up Source Address Mapping 6-4

To Set Up Destination Address Mapping 6-8

To Set Up Destination Port Mapping 6-12

Dynamic NAT 6-13

To Perform Source Address Mapping with a Router 6-16

To Perform Source Address Mapping without a Router 6-18

To Perform Destination Address Mapping without a Router 6-20

Other Examples of a Brick Responding to ARPs 6-22

7 Dependency Masks

Overview 7-1

What is a Dependency Mask? 7-2

To Set Up a Dependency Mask 7-3

Maintaining Dependency Masks 7-8

Example: RealAudio Session 7-15

8 User Authentication

Overview 8-1

What is User Authentication? 8-2

How Authentication Works 8-4

Before Setting Up User Authentication 8-8

To Set Up Local Password Authentication 8-11

To Set Up RADIUS and SecurID Authentication 8-21

To Set Up VPN Certificate Authentication 8-32

9 Digital Certificates

Overview 9-1

What is a Digital Certificate? 9-2

What is the Certificate Manager? 9-3

To Start the Certificate Manager 9-5

To Obtain and Install a Server Certificate 9-7

To Obtain and Import a VPN Certificate 9-13

Group Assignment for VPN Certificates 9-21

Pending Certificate Signing Requests 9-25

Server Versus VPN Certificate Signing Request 9-26

10 LAN-LAN Tunnels

Overview 10-1

What is a LAN-LAN Tunnel? 10-2

To Set LAN-LAN Tunnel Defaults 10-6

To Set Up a LAN-LAN Tunnel 10-13

To Set Up a LAN-LAN Tunnel with UDP Encapsulation 10-25

Redundant LAN-LAN Tunnels 10-27

Maintaining LAN-LAN Tunnels 10-28

To Set Up Service Level Agreements 10-32

11 Client Tunnel Endpoints

Overview 11-1

What is a Client Tunnel? 11-2

To Set the Client Tunnel Defaults 11-8

To Set Up a Client Tunnel Endpoint 11-24

To Set Up a Client Tunnel with UDP Encapsulation 11-33

What to Do Next 11-35

	Maintaining Client Tunnel Endpoints	11-40
	To Create a Message for Alcatel-Lucent IPSec Client Users	11-46
A	Local Presence	
	Overview	A-1
	How Local Presence Works	A-2
	Why Set Up Local Presence?	A-3
	How Local Presence Works	A-4
	Local Presence - Scenario 1	A-5
	Local Presence - Scenario 2	A-7
B	Pre-Configured Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets	
	Overview	B-1
	administrativezone	B-3
	dhcpzone_on_inside_if	B-11
	dhcpzone_on_internet_if	B-17
	firewall	B-23
	nocgwzone	B-30
C	Denial of Service Attacks	
	Overview	C-1
	Syn Flood Protection	C-2
	Intelligent Cache Management	C-4
	Robust Fragment Reassembly	C-6
D	RADIUS Attributes	
	Overview	D-1
	User Group Information	D-2
	Admin Key	D-4
	Client Program Information	D-5

IKEv1 Preshared Key D-7

Timeout Information D-8

Local Presence IP Information D-10

Primary WINS/Secondary DNS D-13

Sample Configuration D-15

E **Configuring the Brick for VoIP/NOE Traffic Using Pre-Defined SMS Templates**

Overview E-1

Configuring the Bricks and SMS Locations E-3

VPN Hub and Spoke Configurations E-7

Brick VPN Configurations E-19

Multiple Zones Per Site E-29

NATed Configurations E-33

Index

List of tables

8 User Authentication

8-1 [SMS-defined user groups](#) 8-20

List of figures

1	Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets	
1-1	Packet Filtering Process	1-3
1-2	Brick Zone Ruleset Editor	1-14
1-3	Brick Zone Rule Editor (Basic Tab)	1-17
1-4	Browse: Select a Brick Window	1-19
1-5	Brick Zone Rule Editor (Time and Day Tab)	1-27
1-6	Brick Zone Rule Editor (Advanced Tab)	1-28
1-7	Brick Zone Rule Editor (Bandwidth Tab)	1-35
1-8	Bandwidth Tab (Asymmetric Bandwidth Configuration)	1-36
1-9	Brick Zone Rule Editor (TOS/Alarms Tab)	1-38
1-10	Brick Zone Rule Editor (Rules Based Routing Tab)	1-40
1-11	Brick Zone Rule Editor (SCTP Tab)	1-42
1-12	Brick Zone Rule Editor (SCTP Feature Enabled With Processing Options)	1-43
1-13	Brick Policy Assignment Editor	1-47
1-14	Apply Brick Window	1-49
1-15	View Brick Zone Rulesets	1-52
1-16	Brick Zone Ruleset Editor (with Traffic Matcher Tool Enabled)	1-53
1-17	Brick Zone Ruleset Editor (Traffic Match Search Performed)	1-54
1-18	Copy Window	1-55
1-19	Confirmation Window (Brick Zone Rulesets)	1-57

1-20	Apply Policy Window	1-58
1-21	Brick Zone Ruleset Editor (Viewing Rules in a Zone Ruleset)	1-59
1-22	Brick Zone Ruleset Editor (Copying a Rule)	1-62
1-23	Brick Zone Ruleset Editor (Pasted Rule)	1-63
1-24	Confirmation Window (Rules)	1-64
1-25	Deployment Sites for Brick Serving as Application Layer Gateway (ALG) in VoIP Network	1-67
1-26	Pre-Defined NOE Brick Zone Rulesets (NOE Subfolder in SMS Navigator)	1-69

2 Host Groups

2-1	Host Group Editor	2-4
2-2	Host Group Entry Window	2-5
2-3	IP Addresses in a Host Group	2-7
2-4	Locate Original Window	2-9
2-5	Host Group Entry Window (Nested Host Groups)	2-13
2-6	Nested Host Groups	2-14
2-7	Navigator Window (View Host Groups)	2-15
2-8	Entities Found Using this Host Group Window	2-17
2-9	Copy Window	2-19
2-10	Confirmation Window (Host Groups)	2-21
2-11	Pre-Defined NOE Host Group Templates (NOE Subfolder in Navigator)	2-26

3 Domain Name Groups

3-1	Domain Name Group Editor	3-3
3-2	Domain Name Group Entry Window	3-4
3-3	Domain names in a Domain Name Group	3-5
3-4	Locate Original Window	3-8
3-5	Domain Name Group Entry Window (Nested Domain Name Groups)	3-11
3-6	Navigator Window (View Domain Name Groups)	3-12

3-7	Entities Found Window	3-13
3-8	Copy Window	3-15
3-9	Confirmation Window (Domain Name Groups)	3-17
4	Service Groups	
4-1	Service Group Editor	4-4
4-2	Service Editor	4-6
4-3	Service Group with Two Services	4-8
4-4	Locate Original Window (Service Groups)	4-10
4-5	Service Editor (Nested Service Groups)	4-12
4-6	Nested Service Groups	4-12
4-7	Navigator Window (View Service Groups)	4-14
4-8	Service Group Editor (Service Group Details)	4-15
4-9	Service Editor (View Protocol/Service Details)	4-16
4-10	Entities Found Using this Service Group Window	4-17
4-11	Service Group Editor (Edit Mode)	4-18
4-12	Copy Window	4-19
4-13	Confirmation Window (Service Groups)	4-21
4-14	Pre-Defined NOE Service Group Templates (NOE subfolder in Navigator)	4-28
5	Application Filters	
5-1	DHCP Application Filter Editor	5-6
5-2	DNS Application Filter Editor (Protected Domain Names Tab)	5-9
5-3	DNS Application Filter Editor Screen (Names and Addresses Tab)	5-11
5-4	DNS Application Filter Editor Screen (RR Type Tab)	5-12
5-5	DNS Application Filter Editor (RR Class Tab)	5-13
5-6	DNS Application Filter Editor (Advanced Tab)	5-14
5-7	ESP Application Filter	5-18

5-8	FTP Application Filter Editor (Commands Filter Tab)	5-23
5-9	FTP Application Editor (Protocol Anomaly Check Tab)	5-25
5-10	FTP Application Filter Editor (Data Port & User Blocking Tab)	5-27
5-11	FTP Application Editor (Miscellaneous Options Tab)	5-29
5-12	GTP Application Filter -Stateful (GTP PDP Context Deletion Monitoring)	5-32
5-13	GTP Application Filter Editor (GTP Version 0 Parameters)	5-34
5-14	GTP Application Filter Editor (GP Interface Message List Tab)	5-36
5-15	GTP Application Filter Editor (Non-Gp Interface Message Tab)	5-37
5-16	GTP Application Filter Editor (IMSI) Prefix list Tab)	5-38
5-17	GTP Application Filter (MSISDN Prefix List Tab)	5-39
5-18	GTP Application Filter Editor(APN List Tab)	5-40
5-19	GTP Application Filter Editor (GTP Nesting Tab)	5-41
5-20	GTP Application Filter Editor (Stateful GTP Tab)	5-42
5-21	GTP Application Filter Editor (Removable R6 Information Elements Tab) (Only Available on GTP Version 1)	5-43
5-22	Removal R6 Information Elements (GTP Messages Sub-Tab)	5-44
5-23	Removable R6 Information Elements (Information Elements Sub-Tab)	5-44
5-24	H.323_VOIP Application Filter (Address Translation Tab)	5-49
5-25	H.323_RAS Application Filter	5-50
5-26	Network Topology - Example 3	5-52
5-27	Application Filter Editor	5-54
5-28	URI Path Pattern Editor	5-56
5-29	Application Filter Editor (Keyword Tab)	5-58
5-30	Pre-Defined NOE_HTTP application filter for VoIP/NOE Traffic	5-61
5-31	NOE Application Filter Editor	5-64
5-32	Pre-Defined NOE Application Filter Templates (NOE Subfolder in Navigator)	5-69
5-33	Application Filter Editor (RPC Tab)	5-75

5-34	Application Filter Editor (Portmapper Tab)	5-76
5-35	RTSP Application Filter (Options Tab)	5-81
5-36	RTSP Application Filter Editor (Methods Tab)	5-85
5-37	SIP Proxy in the private network	5-91
5-38	SIP Proxy on public side of the Brick	5-92
5-39	SIP Proxy in the DMZ	5-92
5-40	SIP Application Filter Editor (Options Default Tab)	5-93
5-41	SIP Application Filter Editor (Options Tab)- both zones	5-95
5-42	SIP Application Filter Editor (Names and Other Addresses Tab)	5-97
5-43	SIP Editor (Names Entry Screen)	5-98
5-44	SIP Application Filter Editor (DNS Default Tab)	5-100
5-45	SIP Application Filter Editor (DNS Tab)-both zones	5-101
5-46	SIP Application Filter Editor (Methods Default Tab)	5-102
5-47	SIP Application Filter Editor (Methods Tab) - both zones	5-103
5-48	SMTP Application Filter Editor (Commands Tab)	5-106
5-49	SMTP Application Filter Editor (MIME Types Tab)	5-108
5-50	SMTP Application Editor (Attachment Tab)	5-109
5-51	SMTP Application Filter Editor (Anomalies Tab)	5-110
5-52	SMTP Application Filter Editor (Other Tab)	5-112
5-53	SQL*Net Application Filter Editor	5-115
5-54	TFTP Application Filter Editor	5-119
5-55	Select a Application Filter Window	5-120
5-56	Pre-Defined TFTP Application Filter Templates (noe Subfolder in Navigator)	5-124
5-57	Locate Original Window (Application Filters)	5-128
5-58	Copy Window	5-129

6 Network Address Translation

6-1	Web Access Rule	6-4
-----	-----------------------	-----

6-2	Source Address Mapping Box	6-5
6-3	Local Access Rule	6-9
6-4	Destination Address Mapping Box	6-9
7	Dependency Masks	
7-1	Dependency Masks Editor	7-3
7-2	Navigator Window (View Dependency Masks)	7-9
7-3	Entities Found Using this Host Group Window	7-10
7-4	Dependency Masks Editor (Edit Mode)	7-11
7-5	Copy Window	7-12
7-6	Confirmation Window (Dependency Masks)	7-14
7-7	RealAudio Sessions	7-15
8	User Authentication	
8-1	User Editor	8-12
8-2	Password Panel	8-13
8-3	Copy Window	8-15
8-4	User Group Editor	8-17
8-5	Copy Window	8-18
8-6	Authentication Service Editor	8-22
8-7	Authentication Service Editor (Primary RADIUS Tab)	8-23
8-8	Authentication Service Editor (Attribute Codes Tab)	8-25
8-9	RADIUS Attributes Window	8-26
8-10	Attribute Codes tab (Configured SMS Parameters With RADIUS Attributes)	8-27
8-11	Authentication Service Editor	8-28
8-12	SecurID Options	8-29
8-13	Authentication Service Editor	8-33
8-14	VPN Certificate Options	8-34

8-15	Attributes	8-34
8-16	Copy Window	8-35
9	Digital Certificates	
9-1	Certificate Manager	9-3
9-2	Certificate Manager Window	9-5
9-3	Certificate Signing Request Window	9-9
9-4	Certificate Signing Request Window	9-15
10	LAN-LAN Tunnels	
10-1	Types of LAN-LAN Tunnel Endpoints	10-2
10-2	LAN-LAN Defaults Editor (Parameters Tab)	10-7
10-3	LAN-LAN Defaults Editor (Policy Tab)	10-9
10-4	LAN-LAN Tunnel Editor (Tunnel Tab)	10-15
10-5	All Group Tunnels Viewer	10-29
10-6	Default SLA Parameters	10-33
10-7	LAN-LAN Tunnel Editor (SLA Probes tab)	10-34
10-8	Tunnel Round Trip Delay Statistics	10-36
10-9	Group Round Trip Delay Statistics	10-37
11	Client Tunnel Endpoints	
11-1	Client Tunnels	11-3
11-2	Client Defaults Editor (Parameters Tab)	11-9
11-3	Client Defaults Editor (Remote Client ID Tab)	11-16
11-4	Client Defaults Editor (Policy Tab)	11-18
11-5	Apply Group Window	11-22
11-6	Client Tunnel Endpoint Editor (Endpoint Tab)	11-25
11-7	Client Tunnel Endpoint Editor (PDG Accounting Tab)	11-30
11-8	Client License Allocation Window	11-36

11-9	Group License Editor Window	11-37
11-10	Tunnel Endpoint Editor Window	11-38
11-11	Client Tunnel Endpoints Window	11-40
11-12	Search Client Tunnels Window	11-42
11-13	Confirmation Window (Client Tunnel Endpoint)	11-45
11-14	Client Message	11-47
C	Denial of Service Attacks	
C-1	Syn Flood Protection	C-3
C-2	Intelligent Cache Management	C-5
C-3	Robust Fragment Reassembly	C-6
E	Configuring the Brick for VoIP/NOE Traffic Using Pre-Defined SMS Templates	
E-1	Deployment of Brick devices and SMSs in a Customer's Enterprise VoIP Network	E-3
E-2	Brick Editor (Assigning Public IP Address for Brick <-> SMS Communications)	E-4
E-3	Brick Editor (Physical Ports Tab, Showing Private IP Address Assigned to Ether1)	E-5
E-4	Brick Editor (Policy Assignment Tab, Showing administrativezone Ruleset and LSMS Host Group Assigned)	E-6
E-5	Brick Editor (Policy Assignment Tab, Showing Pre-Defined Regional Office Zone Rulesets Applied)	E-11
E-6	Brick Editor (Policy Assignment Tab, Showing Pre-Defined Branch Office Zone Rulesets Applied)	E-12
E-7	Brick Editor (Policy Assignment Tab, Showing Pre-Defined Remote Office Zone Rulesets Applied)	E-16
E-8	LAN-LAN Tunnel Editor (Showing Parameter Settings for HQ_to_Regional Tunnel)	E-20
E-9	LAN-LAN Tunnel Editor (Endpoint1 Tab, Showing Headquarters Brick and Endpoint 1 Assignment)	E-21
E-10	LAN-LAN Tunnel Editor (Endpoint2 Tab, Showing Regional Brick and Endpoint 2 Assignment)	E-22
E-11	LAN-LAN Tunnel Editor (Showing Parameter Settings for HQ_to_Branch_Office Tunnel)	E-23

List of figures

E-12	LAN-LAN Tunnel Editor (Endpoint 1 Tab, Showing Headquarters Brick and Endpoint 1 Assignment)	E-24
E-13	LAN-LAN Tunnel Editor (Endpoint 2 Tab, Showing Branch Office Brick and Endpoint 2 Assignment)	E-25
E-14	LAN-LAN Tunnel Editor (Showing Parameter Settings for HQ_to_Remote Tunnel)	E-26
E-15	LAN-LAN Tunnel Editor (Endpoint 1 Tab, Showing Headquarters Brick and Endpoint 1 Assignment)	E-27
E-16	LAN-LAN Tunnel Editor (Endpoint 2 Tab, Showing Branch Office Brick and Endpoint 2 Assignment)	E-28
E-17	Copy Window (Renaming Duplicate Zone Ruleset)	E-29
E-18	Brick Policy Assignment Editor (Showing Range of Phone IP Addresses for Duplicate Zone Ruleset in Hosts Behind Tunnel Field)	E-31
E-19	Hosted PCX with Local PSTN Access	E-33
E-20	Hosted PCX With Full VoIP Deployment	E-33
E-21	Brick Zone Ruleset Editor (Showing Pre-Defined Host Groups Assigned within noe_branch Ruleset)	E-35
E-22	Brick Editor (Policy Assignment Tab, Showing noe_branch Ruleset Assigned to Brick Interface)	E-36
E-23	Brick Zone Ruleset Editor (Showing Pre-Defined Host Groups Assigned within noe_remote Ruleset)	E-37
E-24	Brick Editor (Policy Assignment Tab, Showing noe_remote Ruleset Assigned to Brick Interface)	E-38

About this information product

Purpose

This document explains how to use Release 9.4 of the Alcatel-Lucent Security Management Server (SMS) application to create security policies and set up VPN tunnels on Alcatel-Lucent *VPN Firewall Brick*[™] Security Appliances.

Reason for reissue

Updated for Release 9.4 features.

Who Should Read This Book?

The *SMS Policy Guide* is intended to be read by network administrators who will be using the SMS application to:

- Set up and manage security policies on one or more Brick devices and
- Configure one or more Brick devices to serve as tunnel endpoints of a virtual private network (VPN).

In the terminology used by the SMS, these administrators are referred to as SMS Administrators and Group Administrators, depending on the privileges they have been given when their profiles were created.

What Is in this Book

The *SMS Policy Guide* explains how to configure a Brick device to act as a firewall, screening incoming and outgoing traffic, and passing some of that traffic on to one or more proxy hosts.

It also explains how to configure a Brick device to act as an endpoint in both LAN-LAN and client tunnels.

The *SMS Policy Guide* covers the following topics:

Chapter	Purpose
“About this information product”	This chapter provides an introduction to this document.
Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”	This chapter explains how to set up and manage Brick zone rulesets. Brick zone rulesets determine which traffic is permitted to pass through a Brick device, and which traffic is dropped or proxied by the Brick device.
Chapter 2, “Host Groups”	This chapter explains how to set up, use, and maintain host groups. A host group is a collection of IP addresses that can be used when creating rules.
Chapter 3, “Domain Name Groups”	This chapter explains how to set up, use, and maintain Domain Name Groups. A Domain Name Group is a collection of domain names.
Chapter 4, “Service Groups”	This chapter explains how to set up, use, and maintain service groups. A service group is a collection of services that can be used when creating rules.
Chapter 5, “Application Filters”	This chapter explains how to use application filters. Application filters allow additional application layer validation, inspection and access control directly on the Brick device.
Chapter 6, “Network Address Translation”	This chapter explains how to set up Network Address Translation (NAT). NAT enables the Brick device to map the source or destination addresses of inbound and outbound sessions to other addresses.
Chapter 7, “Dependency Masks”	This chapter explains how to set up, use and maintain dependency masks. Dependency masks allow Administrators to set up a dependency between a Brick rule and a specific session in the session cache.
Chapter 8, “User Authentication”	This chapter explains how to set up user authentication, utilizing either a database residing on the SMS, a RADIUS or SecurID server, or X.509 digital certificates.
Chapter 9, “Digital Certificates”	This chapter explains how to obtain and install X.509 digital certificates. This requires use of the Certificate Manager, a separate application installed with the SMS.
Chapter 10, “LAN-LAN Tunnels”	This chapter explains how to set up a LAN-LAN tunnel between two Brick devices. This allows hosts connected to either device to communicate securely over the Internet.

Chapter	Purpose
Chapter 11, “Client Tunnel Endpoints”	This chapter explains how to configure a Brick device to serve as the endpoint of a client tunnel. A client tunnel is a tunnel between a host running the Alcatel-Lucent IPSec Client application and a Brick device.
Appendix A, “Local Presence”	This appendix provides information on how to create a pool of addresses that can be used by Alcatel-Lucent IPSec Client users to gain access to a local LAN, known as <i>local presence</i> .
Appendix B, “Pre-Configured Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”	This appendix contains a description of pre-configured Brick Zone rulesets.
Appendix C, “Denial of Service Attacks”	This appendix describes how to enable a Brick device to fend off denial-of-service (DoS) attacks.
Appendix D, “RADIUS Attributes”	This appendix provides step-by-step instructions and illustrations on how to set up the SMS application and Brick device to serve as an Application Layer Gateway (ALG) in various VoIP deployment scenarios using the pre-defined Brick zone ruleset, host group, service group, and application filter templates that are provided with the SMS application.

What is Not in this Book

If you are looking for information on any of the following topics, you should refer to the *SMS Administration Guide*:

- How to log on and off the SMS
- How to use the Navigator window and other system conventions
- How to connect a Brick device to your network and configure the Brick device so that it is communicating with the SMS
- How to create groups and set up additional Administrator accounts

These and other topics are covered in the *SMS Administration Guide*. Since these topics pertain primarily to the set up and administration of the hardware, we recommend that you read the *SMS Administration Guide*, and perform all required tasks, before you approach the *SMS Policy Guide*.

Supported Brick devices

The following available Brick models are supported by the current SMS release:

- Alcatel-Lucent *VPN Firewall Brick*® Model 50 Security Appliance
- Alcatel-Lucent *VPN Firewall Brick*® Model 150 Security Appliance
- Alcatel-Lucent *VPN Firewall Brick*® Model 350 Security Appliance
- Alcatel-Lucent *VPN Firewall Brick*® Model 1100/1100A Security Appliance
- Alcatel-Lucent *VPN Firewall Brick*® Model 700 Security Appliance
- Alcatel-Lucent *VPN Firewall Brick*® Model 1200 Standard and HS VPN Security Appliances

Important! *Note: only later Model 20 Brick devices are supported with this release. Model 20 Bricks that have 6-8 MB of RAM and 8 MB of flash are not supported with this release.*

Where to Find Technical Support

Technical assistance and additional information can be acquired by telephone or e-mail. If you require technical assistance, first collect information that technical support staff can use to diagnose the problem. This includes:

- Software version of the SMS.
- Model number and serial number of the Brick device.
- The SMS server platform (Microsoft *Windows*® or Sun *Solaris*® operating systems).
- Description of problem.
- Layout of your network. For example, is the Brick device connected to a device such as a hub or router? Is the Brick device operating as a bridge or is it using static routes? What is connected to the Brick device ports? What is the IP address range and VBA for each zone? What is the security policy for each port?

After gathering the information, contact Alcatel-Lucent Security Customer Care at 1-866-582-3688.

How to comment

To comment on this information product, go to the [Online Comment Form](http://www.lucent-info.com/comments/enus/) (<http://www.lucent-info.com/comments/enus/>) or e-mail your comments to the Comments Hotline (comments@alcatel-lucent.com).

1 Alcatel-Lucent *VPN Firewall* *Brick*[™] Security Appliance Zone Rulesets

Overview

Purpose

This chapter explains how to set up and manage Brick zone rulesets. Brick zone rulesets determine whether traffic is passed, dropped, proxied, or tunneled by the Brick device.

Contents

What is a Brick Zone Ruleset?	1-2
Preventing Spoofing Attacks	1-12
To Create a Brick Zone Ruleset	1-14
To Create a Security Rule	1-16
To Add Time and Day Restrictions to a Rule	1-26
To Add Advanced Features to a Rule	1-28
To Add Bandwidth Management to a Rule	1-34
To Add TOS/Alarm Capability to a Rule	1-38
To Configure IP Address(es) for Rules-Based Routing	1-40
To Enable SCTP Processing on a Rule	1-42
To Assign a Brick Zone Ruleset to a Physical Port	1-46
Maintaining Brick Zone Rulesets	1-51
Maintaining Security Rules	1-59
Brick Zone Ruleset Templates Provided with the SMS Application for VoIP/NOE Traffic	1-66



What is a Brick Zone Ruleset?

Definition

A Brick zone ruleset is a set of security rules that is assigned to one or more ports on one or more Brick devices. The purpose of these security rules is to specify the conditions under which inbound and outbound sessions will be permitted to pass through those ports.

When you assign a Brick zone ruleset to a port, you may specify the IP addresses of the hosts connected to that port. These are the hosts that will be protected by the ruleset. This combination of rules and IP addresses is referred to as a "zone."

Important! Internet Protocol Version 6 (IPv6) addresses are not supported in the current SMS release.

Packet Filtering

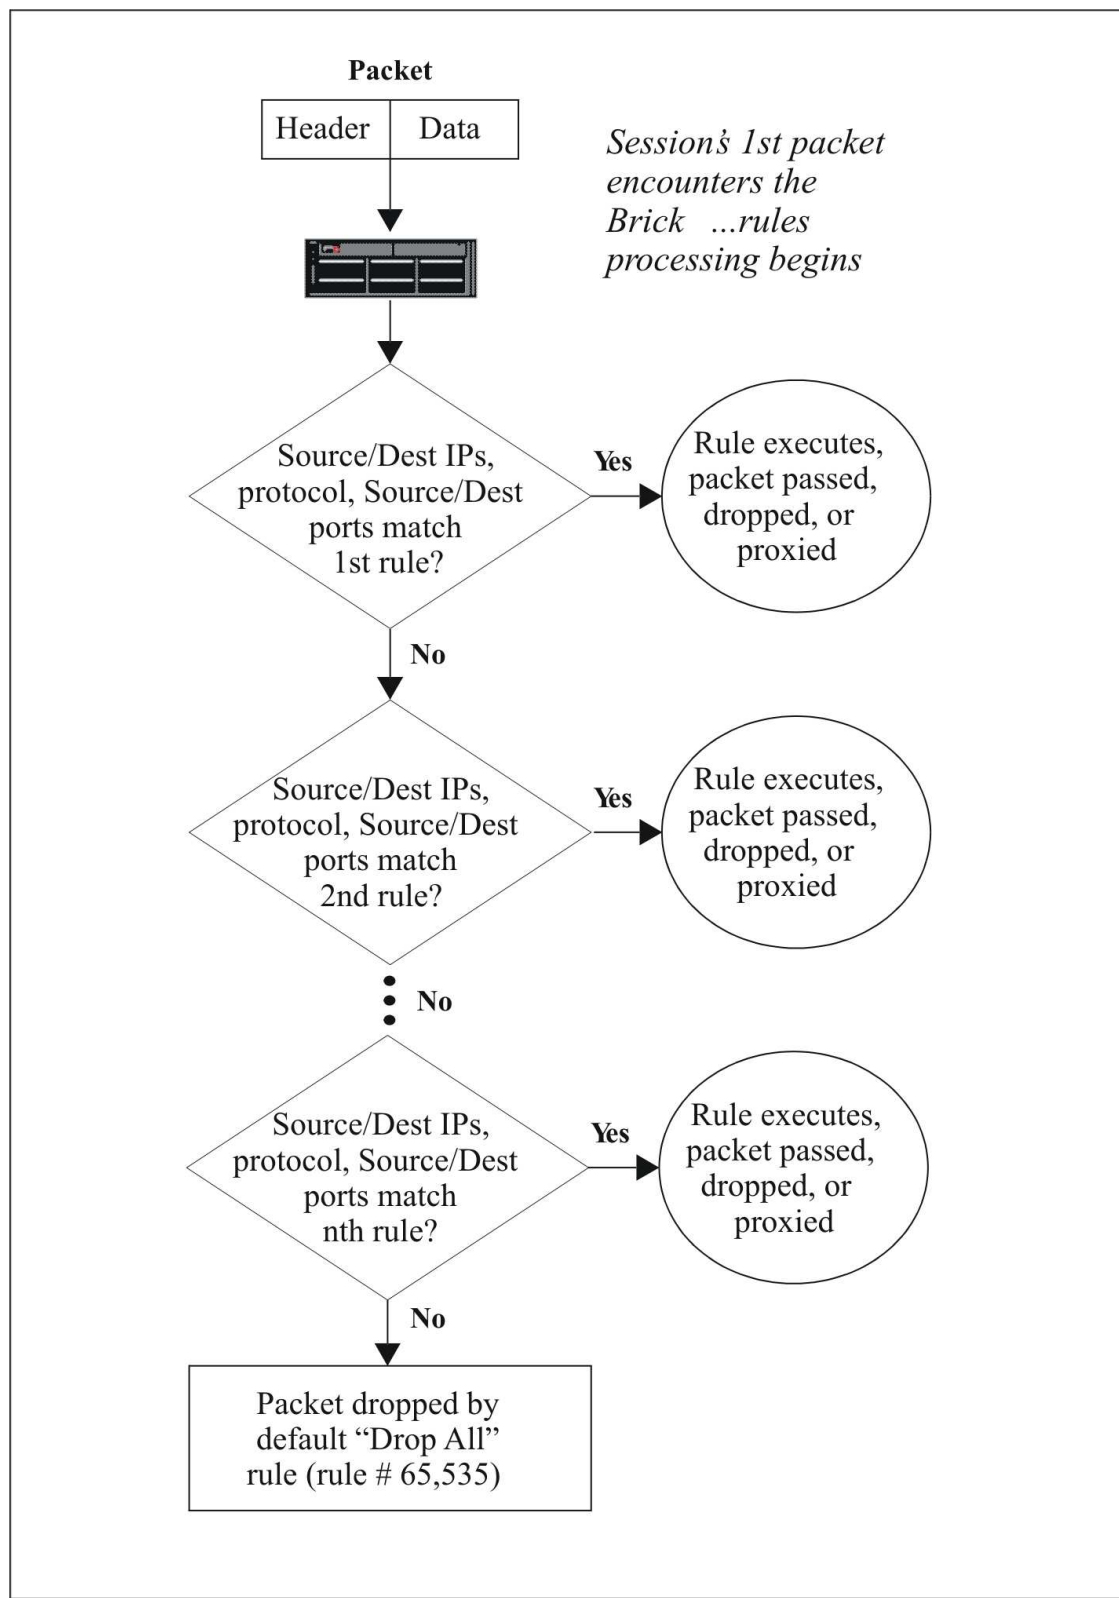
The Brick device intercepts all inbound and outbound traffic traveling through each port and examines each packet header. The Brick device then compares the information that it finds in the headers of the packets with the security rules in the zones that have been assigned to the ports.

If the Brick finds a rule that matches the information in the packet headers, it takes the action associated with the rule: it drops, passes, proxies or tunnels the session. For some protocols, the Brick can inspect the content of the packets further for application layer filtering before determining whether to pass or drop the packet.

If the Brick cannot find a rule that matches the packet headers, the packets are dropped by the last rule, a "drop all packets" rule that is automatically inserted into each ruleset. If the Brick cannot find a zone ruleset that processes the packet, the Brick drops it.

This process is illustrated in [Figure 1-1, "Packet Filtering Process"](#) (p. 1-3) on the next page.

Figure 1-1 Packet Filtering Process



Customized Zones with Brick-Specific Rules

Once a Brick zone ruleset is defined, the same zone ruleset, with its set of security rules and associated host port IP addresses, can be applied to one or more Bricks in your network. In some cases, however, the required security constraints for various types of traffic, even within the same zone, may be different for each Brick and the related portion of the network being protected. SMS offers flexibility in defining security rules for traffic at each point of the network, by providing the optional capability to add new rules in a zone ruleset for a specific Brick, or to modify the rule parameters of an existing rule in a zone for a specific Brick. In this way, a Brick zone ruleset can be defined with generic rules that apply to all Bricks in a zone and Brick-specific rules that only apply to a particular Brick in a zone.

As an example, consider three Bricks, named *Northeast*, *Southeast*, and *Northwest*, that have a zone ruleset with generic rules for FTP service that blocks FTP access to certain sites and allows it to all other sites, as well as Brick-specific rules, such as a longer or shorter session idle time for FTP access. If, for example, the users connected to the *Northwest* Brick need a longer session idle time, a new rule can be inserted into the zone ruleset for that specific Brick by selecting the generic FTP rule for that Brick in the Zone Ruleset Editor, selecting the Duplicate function to copy and create a new rule, and making the required modification(s) to the new rule. Conversely, if the users connected to the *Northeast* Brick need a shorter session idle time for FTP access, while FTP access needs to be blocked for all users connected to the *Southeast* Brick, new rules can be created and applied to that specific Brick.

When a zone ruleset is applied to the Brick(s), each Brick uses the generic rules that apply to all Bricks within that zone, as well as any Brick-specific rules.

Rules-based routing

The SMS provides the option to configure a rule for any type of protocol traffic to route all packets that match the rule to a proxy server, router, or other device, utilizing third party software, to perform content filtering functions such as command blocking, URL filtering, and virus scanning. The exception is for dynamic pinholes that are opened for RTSP, H.323 and SIP protocol traffic; this traffic will not be routed by the rules-based routing feature.

A rule can be set up to route all packets of a certain traffic type that are coming into the zone, going out of the zone, or both, to a content filter.

To set up rules-based routing, two IP address fields are provided on the Brick Zone Rule Editor to configure two IP addresses for a rule—one used for forward routing of packets, the other used for reverse routing of packets. In either address field, the IP address entered is, typically, the address of a proxy server that is performing content filtering or router.

A host group can be selected in place of an individual IP address for forward routing of packets, to load balance between multiple content filters (proxy servers or routers). If a host group is selected for the forward route address, the Brick device forwards packets to each host in the host group in round-robin fashion. Options can be enabled to selectively skip hosts in the host group that are pinged by the Brick device and do not respond for three consecutive pings (in 10 second intervals) or do not respond to a SYN message with a SYN-ACK message.

In order to selectively forward packets to hosts in a host group, one of the following must be configured:

- The zone ruleset using rules-based routing must have a VBA assigned to the interface (port).
- The zone ruleset must be on the port facing the content filters (proxy servers), and that zone ruleset (or the entire Brick device) must be configured to route the return packet(s) back to the cached source media/MAC.
- The address of the port facing the content filter (proxy server) must be included in the Zone IP addresses of the zone with the rules-based routing feature enabled.

When a packet that matches the rule arrives at the zone, the Brick device routes the packet as if the destination was the Forward Route Address specified in the Rules Based Routing tab. This may be, for example, the IP address of a proxy server or router performing content filtering, or the VBA of a firewall in another partition. If the Forward Route Address field is blank, the Brick device routes the packet normally, to the packet destination address.

When the response packet arrives at the zone, the Brick device routes the packet as if the destination was the Reverse Route Address specified on the Rules Based Routing tab. If the Reverse Route Address field is blank, the Brick device routes the packet normally, to the packet destination address.

Rules-based routing can be applied in networks where:

- The proxy server (content filter) device is on one leg between two firewalls
- The firewall is situated between the client and the proxy server (content filter)
- The proxy server (content filter) is in a loopback configuration, to handle traffic coming into and out of the Brick

When configuring the rules-based routing feature on the Brick Zone Rule Editor, the following restrictions apply with respect to creating rules:

- A rule cannot have a forward routing IP address and an **Action** of **Proxy**
- A rule cannot be configured with both an IP address for routing packets for content filtering (in either direction) and an **Action** of **VPN**
- A rule cannot have a forward routing IP address, an **Action** of **VPN**, a **Virtual Private Network** setting on the Advanced tab of **Internal**, and a **Direction** of **In to Zone**

- A rule cannot have a forward routing IP address, an **Action** of **VPN**, a **Virtual Private Network** setting on the Advanced tab of **External**, and a **Direction** of **Out of Zone**
- A rule cannot have a reverse routing IP address, an **Action** of **VPN**, a **Virtual Private Network** setting on the Advanced tab of **Internal**, and a **Direction** of **Out of Zone**
- A rule cannot have a reverse routing IP address, an **Action** of **VPN**, **Virtual Private Network** setting on the Advanced tab of **External**, and a **Direction** of **In to Zone**

Important! To prevent a customer in one group from using rules-based routing to launch a source routing attack on a customer in another group on the same Brick device, make sure that the zones for each customer are assigned to separate partitions.

If a data packet is tunneled, UDP encapsulated, or crosses a partition boundary, any routing IP address information used for filtering content is removed.

SCTP port filtering

Stream Control Transport Protocol (SCTP) is a protocol designed to transport Public Switched Telephony Network (PSTN) signaling messages over IP networks, with broader application. SCTP provides a reliable transport protocol operating on top of IP.

SCTP offers the following services (in accordance with RFC 2960):

- Acknowledged error-free non-duplicated transfer of user data
- Data fragmentation to conform to discovered path MTU size
- Sequenced delivery of user messages within multiple streams
- Bundling of multiple user messages into a single SCTP packet
- Network-level fault tolerance through multi-homing support at one or both ends of an association
- Appropriate congestion avoidance behavior and resistance to flooding and masquerading attacks

The SCTP-aware feature, when enabled for a rule, enables a Brick device to decode SCTP protocol packets exchanged between communication endpoints, allow/deny pre-configured SIGTRAN applications, and dynamically open the exact pinhole needed for a specific communication.

Refer to [“To Enable SCTP Processing on a Rule”](#) (p. 1-42) for instructions on how to enable the SCTP-aware feature on a rule.

Brick zone rulesets and application layer gateway (ALG)/NOE feature

When the NOE traffic filtering option is enabled within a TFTP application filter, the Brick creates dynamic rules in the zone ruleset for inspecting RTP and NOE signalling traffic between the IP phones and other network elements (call servers, MGW) in a

VoIP network. The Brick also creates a dynamic rule for any telnet debug sessions to IP phones. An SMS administrator should create an additional rule within the assigned Brick zone ruleset to allow HTTP traffic from the IP phone address(es) to the presentation server. Other rules should be created within the assigned zone ruleset to allow/drop other traffic that the Brick is expected to handle, according to the customer-specific policy.

Traffic matcher tool

The Traffic Matcher tool allows administrators to test Brick zone rulesets (policies) by entering a specification for simulated traffic, and then finding which rules would be triggered by that traffic pattern. For example, the tool can provide an answer for such a question as, 'what rules in this policy, if any, will match telnet traffic'.

The Traffic Matcher tool is integrated into the viewer portion of the Brick Zone Ruleset Editor. To use the tool, click the checkbox **Show Traffic Matcher** to place a check in it. A Traffic Matcher box is inserted over the table of rules. Inside the Traffic Match Criteria box are traffic-related fields such as Direction, Source Host, Destination Host, Service, and other fields, which allow you to specify the pattern of "virtual traffic" you wish to "run through" the zone ruleset. The traffic data that you specify is compared to existing rules in that zone ruleset, and the rules that match the specified traffic criteria are displayed.

When you click the **Match** button, the "virtual traffic" specifications that you entered are "run through" the existing rules, and the system finds the rules that would be "triggered by" or match the specified traffic criteria. Rules are matched and displayed based on the general traffic characteristics specified, such as traffic type and direction, *not* on other specific details, such as the ACTION taken on the traffic, or traffic originating from a certain NAT address.

The tool allows you to display only the rules that match the "virtual traffic" details you supplied, or to display all rules with those that match the "virtual traffic" details highlighted in yellow.

Some of the traffic fields in the Traffic Matcher tool can accept multiple or aggregate values. For example, the **VLAN ID** field can accept multiple IDs, separated by commas. If, for example, you enter "1,3" as traffic criteria in the **VLAN ID** field, which represents some traffic being sent with a VLAN ID of 1 and other traffic being sent with a VLAN ID of 3, the Traffic Matcher tool only retrieves the rule(s) that would apply to traffic being sent with a VLAN ID of 1 and 3. If an aggregate value such as a host group, a service group, or '*' is specified in a search field, only rules that match all possible values within the group are retrieved. Therefore, only rules with '*' in their specification would be retrieved, in the case of aggregate entities such as host groups or service groups.

The Source and Destination fields may contain either hosts or users. Consequently, the tool offers a way to specify traffic match criteria for both. The user field is found immediately below the host field and the two do not interfere with each other in the matching algorithm. For example, if you specify a Source Host and leave Source User blank, at a minimum, all rules with Source fields that are designated as user fields are matched.

Due to limitations in the matching algorithm, some matching rules may be displayed as non-matching rules. For example, a rule might have a service group that is defined as **6/70/0-65535**, but if the traffic Service field is set to **7/70/***, the system may not determine it to be a match, because the tool does not detect that **0-65535** is in the range of **'*'**. In this case, in order to match the rule, the Service field in the traffic matcher would have to specify a range that is equal to or is a subset of the range that is specified in the policy rule.

Brick Session Cache

The Brick maintains a session cache that contains a record of every session that was passed, dropped, or proxied.

When the Brick examines the first packet of a session and, based on its security rules, decides to pass the packet, it places an entry in the session cache.

For each subsequent packet *from that session*, the Brick looks in the cache first and, if it finds a matching packet, it passes all subsequent packets.

Using the session cache instead of the rules database allows the Brick to speed up the packet filtering process.

Rule Types and Numbering

A Brick zone ruleset can contain up to 65,535 rules. This includes two kinds of rules — rules created by an Administrator, and rules automatically created by the SMS.

Every rule, regardless of its type, is assigned a number. The table below explains the numbering conventions:

Rule #	Purpose
1 - 199	Reserved for future features.
200 - 299	Firewall, administration, and proxy rules. These rules are automatically added to the pre-configured Brick zone rulesets when the SMS application is installed.
300 - 399	User authentication rules. These rules are automatically added to a Brick zone ruleset when an Administrator creates a rule with a user group as the source or destination.

Rule #	Purpose
400 - 499	VPN rules. These rules are automatically added to a Brick zone ruleset when an Administrator creates a rule establishing a LAN-LAN or client tunnel.
500 - 999	Reserved for future features.
1000 - 64999	Administrator created rules. These are the rules that SMS or Group Administrators create. The numbers are assigned automatically by the SMS when the rule is created.
65000 - 65534	Reserved for future features.
65535	Drop-all rule. This rule is automatically added to every zone when it is created. Its purpose is to ensure that no traffic is permitted through the Brick unless it is specifically authorized by a previous rule.

The rules automatically created by the SMS (rules #1 - 999 and 65000-65534) are referred to as *system* rules. The Administrator has the option of displaying these rules in the Brick Zone Ruleset Editor or hiding them. By default, they are displayed. To hide them, click the **Hide System Rules** checkbox at the bottom right of the Brick Zone Ruleset Editor

The numbers of the system rules appear in light blue. System rules can be activated or deactivated, and can be edited. They cannot be deleted, nor can they be re-numbered or re-ordered.

Examples

The table below lists several sample Administrator-created security rules that might appear in a typical zone security policy. The sections that follow explain each rule in detail.

Rule #	Direction	Source	Destination	Service	Action
1000	Out of zone	Internal_ hosts	*	HTTP	Proxy
1001	In to zone	10.4.2.24	10.120.4.12	FTP	Pass
1002	In to zone	*	10.120.4.12	FTP	Drop
1003	In to zone	*	*	telnet	VPN

Rule #1000

The purpose of this rule, the first Administrator-created rule, is to permit certain hosts in the zone to have access to the Internet's World Wide Web. The following explains how the rule accomplishes this:

- **Direction**

The direction of the rule is *Out Of Zone*. This means the rule applies to packets sent by hosts inside the zone, seeking to establish sessions with external servers, as is the case with Internet access.

- **Source IP**

The address of the Source hosts is the host group *Internal_hosts*. A host group is a collection of IP addresses, in this case, the IP addresses of the internal hosts that will be permitted to connect to Web servers. The SMS or Group Administrator had to create this host group.

- **Destination IP**

The address of the Destination host., the hosts providing the Internet service, is given as an asterisk (*). The asterisk is the wildcard symbol that represents all hosts. It was used here because the intention of this rule is to allow internal users unrestricted access to the Web.

- **Service**

The service is *HTTP*. This is a service group that is provided with the SMS. It specifies that the protocol is 6 (TCP) and the destination port is 80, the standard protocol and port for HTTP services.

- **Action**

The action is *proxy*. This means that any HTTP sessions originated by hosts in the host group *Internal_hosts* will be forwarded to the proxy server, a host running the Lucent Proxy Agent application, which will establish the connection with the Web server.

Rules #1001 and 1002

The purpose of these rules is to allow one host outside the zone to access an FTP server inside the zone, and to deny access to all other hosts. The following explains how these rules accomplish this:

- *Rule #1001*
The purpose of rule #1001 is to allow one external host to access one internal FTP server.
Therefore, the direction of the rule is *In To Zone*, the Source IP is the address of the external host (10.4.2.24) and the Destination IP is the address of the internal server (10.120.4.12).
The service is the service group *FTP* (protocol = TCP, destination port = 21), and the action is *Pass*.
- *Rule #1002*
The purpose of rule #1002 is to deny all other requests for FTP services to that server. Therefore, the direction of the rule is *In To Zone*, the source IP is *, the destination IP is 10.120.4.12 again, the service is *FTP*, and the action is *Drop*.

Important! RULE ORDER

Rules #1001 and 1002 underscore the importance of the order of rules. The Brick processes rules in the order in which they appear in the zone's security policy.

Consequently, rules #1001 and 1002 would not work if they were in reverse order. If that were the case, the second rule would drop all FTP requests to server 10.120.4.12, including those emanating from host 10.4.2.24. These latter requests would be dropped before they ever reached rule #1002, which is intended to pass them.

Rule #1003

The purpose of this rule is to set up a client tunnel so that the sales representatives in the user group *sales_reps* can establish a secure telnet session with the hosts in the *internal_hosts* host group.

Therefore, the direction is *In To Zone*, the Source is * (asterisk), the Destination is * (asterisk), and the service is telnet. The action is VPN, because the purpose of this rule is to send telnet traffic through the tunnel.

□

Preventing Spoofing Attacks

Overview

Spoofing attacks are attacks in which the attacker uses a source address other than his/her own in order to penetrate the defenses and/or hide his/her identity.

Preventing the use of spoofing in the Alcatel-Lucent Virtual Firewall can be accomplished in one of two principal ways:

- Preventing spoofing in rules
- Through the Zone IP Addresses in the Brick Policy Assignment

Preventing spoofing in rules

One common way to prevent spoofing is to create a rule (typically, the first rule or two in the zone ruleset) which explicitly drops packets whose source addresses are invalid for the direction in which the packets are seen.

This method has two advantages:

- Violations are explicitly seen in the session log
- For complex networks where the list of hosts is long, searching in the rules is more efficient than the alternatives below

Through the zone IP addresses in the Brick Policy Assignment

Another way to prevent spoofing is to restrict the allowed packet addresses through the Policy Assignment tab in the Brick Editor.

There are three possible configurations for this:

1. A single assignment with the list of valid (for the Internet side) addresses in the Zone IP Addresses field
2. Two assignments, the first with the list of invalid (protected) addresses and the second being the normal packet processing zone. The first assignment has to have a “drop-all” ruleset. The second assignment can have an asterisk (*) wildcard character for the zone IP address and may or may not have an actual zone ruleset.
3. Two assignments, the first with the list of valid (for the Internet side) addresses and the second is an asterisk (*). The first assignment is the normal packet processing ruleset and can be blank. The second assignment must be a “drop-all” zone ruleset.

The advantage to Option #1 is that no session entries are created for spoofed packets, so performance and capacity may be better. When a spoofed packet is detected in this case, a log entry is created in the Administrative Events log with the text error in zone table configuration (error code 8).

Options #2 and #3 are functionally identical and really differ only in the list of addresses that are created (one being the list of all the addresses not in the other one). The advantage to this case is that the packets are logged in the Session log, which can be more easily searched and is the more natural place to expect it to be logged.

□

To Create a Brick Zone Ruleset

Overview

To create a Brick zone ruleset, you have to display the Brick Zone Ruleset Editor and enter a zone name. To take effect, the ruleset then has to be assigned to a port on a Brick device.

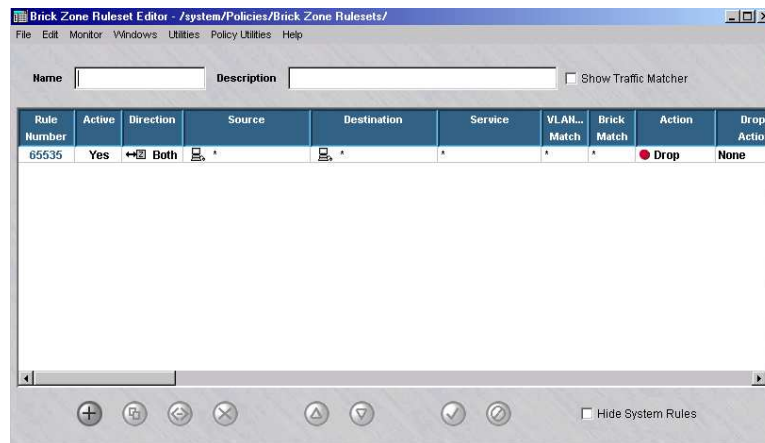
To display the Brick Zone Ruleset Editor

Complete the following steps to display the Brick Zone Ruleset Editor.

- 1 Open the folder of the group that will contain this Brick zone ruleset.
 - 2 Open the **Policies** folder.
 - 3 Right-click the Brick Zone Rulesets folder and select **New Brick Zone Ruleset** from the pop-up menu.
-

Result The Brick Zone Ruleset Editor is displayed ([Figure 1-2, “Brick Zone Ruleset Editor”](#) (p. 1-14)).

Figure 1-2 Brick Zone Ruleset Editor



END OF STEPS

To create the ruleset

Complete the following steps to create a Brick zone ruleset. Individual rules can be added using the procedure described in [“To Create a Security Rule”](#) (p. 1-16):

- 1 In the **Name** field, enter a unique name to identify this ruleset. The name can contain up to 44 characters. It can consist of lower case letters, numbers and certain special characters.
- 2 In the **Description** field, you can enter an optional description of the ruleset. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers and certain special characters.
- 3 Display the File menu and select one of the **Save** options.

END OF STEPS



To Create a Security Rule

Overview

One security rule is automatically added to the Brick zone ruleset when it is created. This is rule #65535 (the drop all rule), and it is always the last rule in the Brick zone ruleset.

To provide the Brick zone ruleset with a full-fledged security policy, you need to create additional rules. Then, for these rules to take effect, you have to assign the ruleset to a port on at least one Brick device. Refer to the procedure [“To Assign a Brick Zone Ruleset to a Physical Port”](#) (p. 1-46).

Task

Complete the following steps to create a security rule within a Brick zone ruleset.

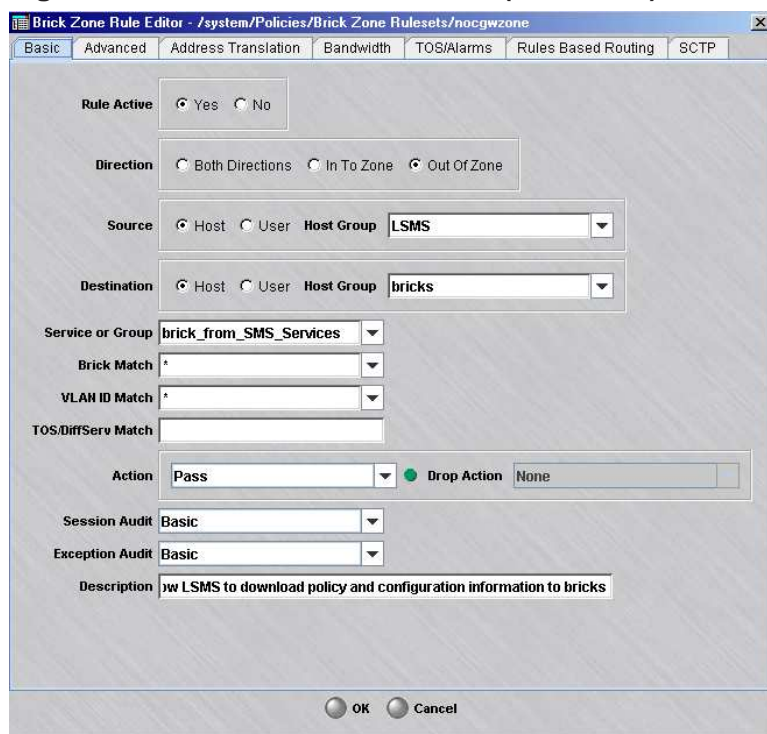
- 1 Access the Brick Zone Ruleset Editor as described in [“To display the Brick Zone Ruleset Editor”](#) (p. 1-14).

The Brick Zone Ruleset Editor is displayed (see [Figure 1-2, “Brick Zone Ruleset Editor”](#) (p. 1-14)).

- 2 With the Brick Zone Ruleset Editor displayed, do one of the following:
 - Right-click *an empty area* of the Rules Viewer and select **New** from the pop-up menu. When you save the new rule, it will automatically be numbered rule #1000, and all existing rules will be renumbered accordingly.
 - Right-click *an existing rule* in the Rules Viewer and select **New** from the pop-up menu. In this case, the new rule automatically takes the number of the rule you right-clicked, and the number of the existing rule increases by one number. All rules with higher numbers are renumbered accordingly.

The Brick Zone Rule Editor is displayed, initially showing the Basic tab (see [Figure 1-3, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-17)).

Figure 1-3 Brick Zone Rule Editor (Basic Tab)



By default, the rule is active, which means it takes effect as soon as the zone is assigned to a port and the ruleset is saved and applied to a Brick. To make this rule inactive (you can activate it later), click **No** in the **Rule Active** field.

- 3 In the **Direction** field, select the direction of the session. The choices are: **Both Directions**, **In To Zone**, and **Out of Zone**. **Both Directions** is the default.

The direction of a session refers to the session direction *vis a vis the zone and firewall*.

If the direction of a session is **In To Zone**, the session source is a host outside of the zone passing data traffic through the Brick device zone to the protected host.

If the direction of a session is **Out of Zone**, the session source is a protected host behind the Brick device passing data traffic through the Brick device zone to a router or other host on the other side of the Brick device.

Consequently, to be into the zone, a session source must be a host outside the zone, and to be out of the zone, a session source must be a host inside the zone.

In any rule except a VPN rule, you can also specify the direction as both. In VPN rules, the direction cannot be both.

-
- 4 In the **Source** and **Destination** fields, indicate the source and destination of the session. The source and destination can be hosts (the default) or users.

The following explains:

- **Hosts**

If the source or destination is to be a host (identified by its IP address), leave **Host** clicked (the default). Then, in the **Host Group** field, do one of the following:

- Leave the default asterisk in place. This means the source or destination is unrestricted (includes all hosts).
- Replace the asterisk with the IP address of a specific host. The source or destination will be only this host.

or

- Display the drop-down list and select either **Virtual Brick Addressor Browse**. If you select Browse, a Browse window is displayed and allows you to select a host group. The source or destination can only be assigned from the hosts in the selected host group.

- **Users**

If the source or destination is to be identified by the identity of the user, click **User**. Then, in the **User Group** field, display the drop-down list and do one of the following:

- Select one of the system-created user groups (**All_Users**, **Active_VPN_Users**, or **Active_VPN_UserTEPs**).

or

- Select **Browse** and use the Browse window to select an Administrator-created user group.

In either case, the source or destination will be the users in this user group.

-
- 5 In the **Service or Group** field, enter the service.

There are three ways to do this:

- Leave the default asterisk in place. This means the service is unrestricted (all services).
- Display the drop-down list and select a protocol from the drop-down list. The options include TCP, UDP, ICMP and most other IP-based protocols and services.
- Display the drop-down list and select **Browse**. A Browse window will appear and allow you to select a service group. You can select a service group you have created, or one that is provided with the SMS application.

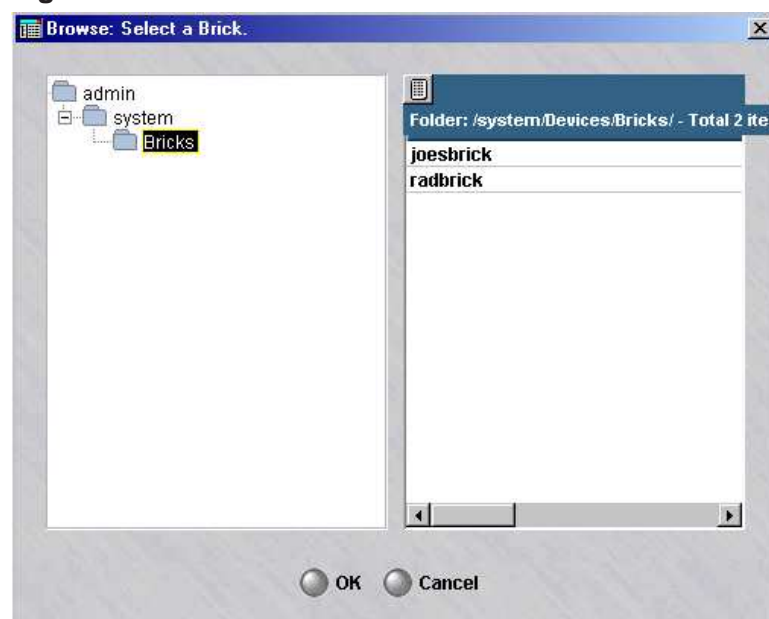
-
- 6 *Optional:* To make a rule specific to one or more Bricks, enter the name of a Brick device or click the down arrow next to the **Brick Match** field to display a drop-down list of choices.

The choices are:

- An asterisk (*), which is the default value. This means that this new rule is being defined for all Brick devices.

** *BROWSE* **. If this option is chosen, a Browse window is displayed to select the Brick device for which this rule is being defined ([Figure 1-4, “Browse: Select a Brick Window”](#) (p. 1-19)).

Figure 1-4 Browse: Select a Brick Window



Select a Brick device for which this rule is being created, and click **OK**. The system returns to the Brick Zone Rule Editor, with the **Brick Match** field populated with the name of the chosen Brick device.

-
- 7 In the **VLAN ID Match** field, enter one or more VLAN IDs that the Brick device will match when determining whether to pass, drop or proxy a session that matches all the other criteria in the rule.

The choices are:

- An asterisk (*), which is the default value. This means that the rule applies to all defined VLANs for this Brick device.
- Enter a single VLAN ID directly into the field, or a range of VLAN IDs separated by a dash. You can also enter a comma-separated list of single VLAN IDs or ranges of VLAN IDs (e.g., 1,2, 10-20, 30-40). This field can be used to specify additional matching criteria within the VLAN IDs that already match the zone assignment.
- **Port Default**, which means the default VLAN for this Brick device port.

8 In the **TOS/DiffServ Match** field, enter two hex digits or a range qualified by a matching mask. This specifies the TOS/DiffServ byte pattern, which serves as matching criteria for the rule to be triggered. The rule match is contingent on the second byte of the IP header having this value. (See [“To Add TOS/Alarm Capability to a Rule” \(p. 1-38\)](#) for additional details.)

9 In the **Action** field, select, from the drop-down list, the action the Brick device will take when a session matches the fields in this rule. The following explains each action.

Action	Explanation
Drop	<p>The default. The Brick device will discard the session and not allow it through the port.</p> <p>If you choose Drop, you can also specify whether or not to notify the initiator of the session (the source) that the session was dropped.</p> <p>This is done in the Drop Action field. The default is None (no notification). To send the source a message indicating the session was dropped, click Notify.</p>
Pass	<p>The Brick device will permit the session through the port.</p>
Proxy	<p>The Brick device will reflect the session to the host running the Lucent Proxy Agent application, which will determine whether the session is dropped or passed. Any session type can be proxied; however, the current version of the Lucent Proxy Agent supports only HTTP, SMTP and FTP.</p> <p><i>Note: the Lucent Proxy Agent application is no longer being supported and has been replaced by the Rules-Based Routing feature.</i></p>

Action	Explanation
VPN	<p>The Brick device encrypts or decrypts the session. If you choose this action, External will automatically be entered in the Virtual Private Network field on the Advanced tab.</p> <p>To use the VPN action, you must also set up the tunnel (see “Overview” (p. 10-1) and “Overview” (p. 11-1)).</p>
VPN Proxy	<p>This Brick encrypts or decrypts the session, and also reflects the session to the proxy host. This action causes External to be automatically entered in the Virtual Private Network field on the Advanced tab.</p> <p>To use the VPN Proxy action, you must also set up the tunnel and put an entry for the zone in the Proxy Table.</p>

Important! When you specify a drop action of **notify** or **reset**, it may be possible for an attacker to identify the actual IP addresses of hosts behind the Brick.

This may occur in two cases:

1. If you only specify this action for hosts that actually exist.
2. If you specify this action for hosts on the directly connected subnet in the zone on the protected side of the firewall. This will occur because the packet will only be forwarded to that zone if the Brick can determine how to route the packet to that zone, which requires that the host actually exists in order to resolve the ARP request.

-
- 10** In the **Drop Action** field, indicate the type of action to be taken when a session is dropped by the Brick. The default is **None**.

The other options include:

- ICMP All
- Reset TCP, Ignore Others (for dropped TCP sessions, the Brick spoofs an RST packet to the source only; for non-TCP sessions, no notification of any type is sent)
- Reset TCP, ICMP Others (for dropped TCP sessions, the Brick spoofs an RST packet to the source only; for non-TCP sessions, an ICMP 3/13 message is sent)

-
- 11** In the **Session Audit** field, indicate the type of auditing. The default is **Basic**. With basic auditing turned on, the Brick will audit all open and closed session packets (both dropped and passed packets).

For certain application layer protocols, you can perform more detailed auditing by selecting **Detailed** from the drop-down list. Detailed auditing includes basic auditing plus the additional auditing described in the table below:

Application Layer Protocol	Basic Auditing	Detailed Auditing	Exception Auditing
TFTP	Options: <ul style="list-style-type: none"> • Original session • Reverse session 	Options: <ul style="list-style-type: none"> • Dynamic port opened for reverse session 	Options: <ul style="list-style-type: none"> • Controls whether or not the Brick generates an audit message when the message type is not recognized
FTP	Options: <ul style="list-style-type: none"> • Control channel session • Data channel session • Any parsing errors • 3-Way FTP warning 	Options: <ul style="list-style-type: none"> • Dynamic ports opened for data channel • Commands sent over the FTP channel (as per the FTP logging enhancement feature) 	Options: <ul style="list-style-type: none"> • Not used.
SIP	Options: <ul style="list-style-type: none"> • SIP sessions • RTP/RTCP sessions 	Options: <ul style="list-style-type: none"> • Passed SIP requests and responses (the SIP Audit Session option Options tab must also be selected) 	Options: <ul style="list-style-type: none"> • Blocked SIP requests and responses (the SIP Audit Session option Options tab must also be selected)

Application Layer Protocol	Basic Auditing	Detailed Auditing	Exception Auditing
DNS	Options: <ul style="list-style-type: none"> • Original session 	Options: <ul style="list-style-type: none"> • DNS query packet • DNS response packet 	Options: <ul style="list-style-type: none"> • Block reason with packet type host name, rr type and rr class
GTP	Options: <ul style="list-style-type: none"> • Original session 	Options: <ul style="list-style-type: none"> • GTP message packet 	Options: <ul style="list-style-type: none"> • Block reason with message type and blocked information element type
H.323	Options: <ul style="list-style-type: none"> • H.323 session • H.245 session • RTP session • RTCP session • Any parsing errors 	Options: <ul style="list-style-type: none"> • Dynamic ports opened for H.245 • Dynamic ports opened for RTP • Dynamic ports opened for RTC 	Options: <ul style="list-style-type: none"> • Controls whether or not the Brick generates an audit message when certain internal message structures are invalid or unrecognized (Examples: invalid choice, unrecognized type, value out of range)
HTTP	Option: <ul style="list-style-type: none"> • Original session 	Option: <ul style="list-style-type: none"> • HTTP command protocol 	Option: <ul style="list-style-type: none"> • Controls whether or not the Brick generates an audit message when the packet is dropped due to a security violation

Application Layer Protocol	Basic Auditing	Detailed Auditing	Exception Auditing
SQL*Net	Option: <ul style="list-style-type: none"> • Original session • Redirected session • Any parsing errors 	Option: <ul style="list-style-type: none"> • Dynamic port opened for redirected session 	Option: <ul style="list-style-type: none"> • Controls whether or not the Brick generates an audit because: <ul style="list-style-type: none"> – - the redirect is to a different host – - the Brick cannot parse the message – - the message type is not recognized
NetBIOS	Option: <ul style="list-style-type: none"> • Original Session • Dynamic session 	Option: <ul style="list-style-type: none"> • Dynamic ports opened 	Option: <ul style="list-style-type: none"> • Not used.
RPC	Option: <ul style="list-style-type: none"> • Original Session 	Option: <ul style="list-style-type: none"> • Program number, version, procedure for each call • Dynamic ports opened for new sessions 	Option: <ul style="list-style-type: none"> • Controls whether or not the Brick generates an audit message when the packet is dropped due to a security violation

If you do not want to perform any auditing, select **None** from the drop-down list. It is recommended that you use the auditing feature carefully. The more sessions that are audited, the larger the log files become, the more space they occupy on your hard drive, and the longer reports take to run. Also, since detailed auditing can generate many more records than basic auditing, you should use this feature judiciously.

-
- 12** In the **Exception Audit** field, indicate the type of auditing. The default is **Basic**. The other alternatives are **None** or **Detailed**. When one of these options is selected, the Brick logs a specific exception condition that pertains to the particular connection. If

TCP validation is enabled, it will generate and transmit a new audit message, destined for the Session Log. This message will contain enough information regarding the type of validation failure to help you determine the source of the invalid packets. Sessions that fail other types of validation including application filter exceptions will be logged with auditing exceptions.

Basic exception auditing reports only the first error in a session for TCP validation. Detailed exception auditing reports one log message for each invalid packet in the session. Detailed auditing should only be turned on for troubleshooting your configuration.

-
- 13** In the **Description** field, you can enter an optional description of the rule. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers and certain special characters.

If the ruleset will contain many rules, the description can help you identify specific rules quickly.

-
- 14** Click **OK** to dismiss the Brick Zone Rule Editor and return to the Brick Zone Ruleset Editor.

The new rule is displayed in the Rules Viewer. If it is a Brick-specific rule, the rule is highlighted in yellow in the Rules Viewer.

-
- 15** Display the File menu and select one of the **Save** options.

END OF STEPS



To Add Time and Day Restrictions to a Rule

Overview

The SMS provides the ability to control the time and days that users may access specific network resources. For example, customers may want their employees to only have access to certain resources, such as the Internet or a company database server, during certain time periods or during certain days of the week. This allows the customer to set up and enforce a flexible enterprise-wide security policy.

Time-of-day restrictions are treated by the Brick device as any other matching criterion in a rule. In other words, if all the matching criteria in a rule, including the time and day, match the information in the header of a packet, the Brick device takes the appropriate action. If the time and day does not match the information in the header, no action is taken, even if the rest of the information in the packet matches the rule. In this case, Brick device processing proceeds to the next rule.

When creating rules, you can specify the time period, as well as the days of the week, to which the rule applies. You can also specify whether the time will be based on the actual SMS time or a time offset relative to the SMS time.

When a zone policy is applied to multiple Brick devices in different zones, this choice is particularly relevant. To have all Brick devices enforce the same policy at the same absolute time, use the SMS time in the policy. However, to have each Brick device change its policy at a particular local time, use the Brick device time in the policy.

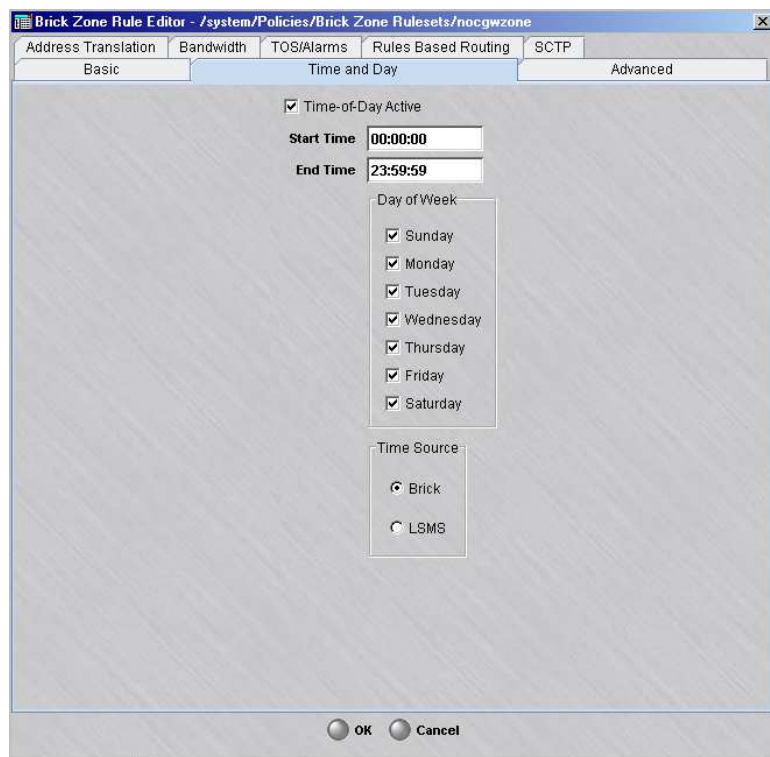
Important! Time-of-day restrictions can only be set for user-created rules. These restrictions cannot be set for SMS-generated rules. If the color of the Rule Number is blue on the Brick Zone Ruleset Editor, it is an SMS-generated rule; the Time and Day tab does not appear on the Brick Zone Rule Editor when you edit an SMS-generated rule.

Task

To add time and day restrictions to a rule, follow the steps below:

- 1 With the Brick Zone Rule Editor displayed (see [Figure 1-3, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-17)), click the **Time and Day** tab, and then click the **Time-of-Day Active** checkbox. The time and day fields are displayed (see [Figure 1-5, “Brick Zone Rule Editor \(Time and Day Tab\)”](#) (p. 1-27)).

Figure 1-5 Brick Zone Rule Editor (Time and Day Tab)



-
- 2 In the **Start Time** and **End Time** fields, enter the time period to which the rule applies. You must enter the start and end times in 24-hour time (such as 18:00:02 — 19:33:41).
 - 3 In the **Day of the Week** box, check the days of the week to which this rule will apply. By default, all the days of the week are checked. You can select any combination of days of the week.

END OF STEPS



To Add Advanced Features to a Rule

Overview

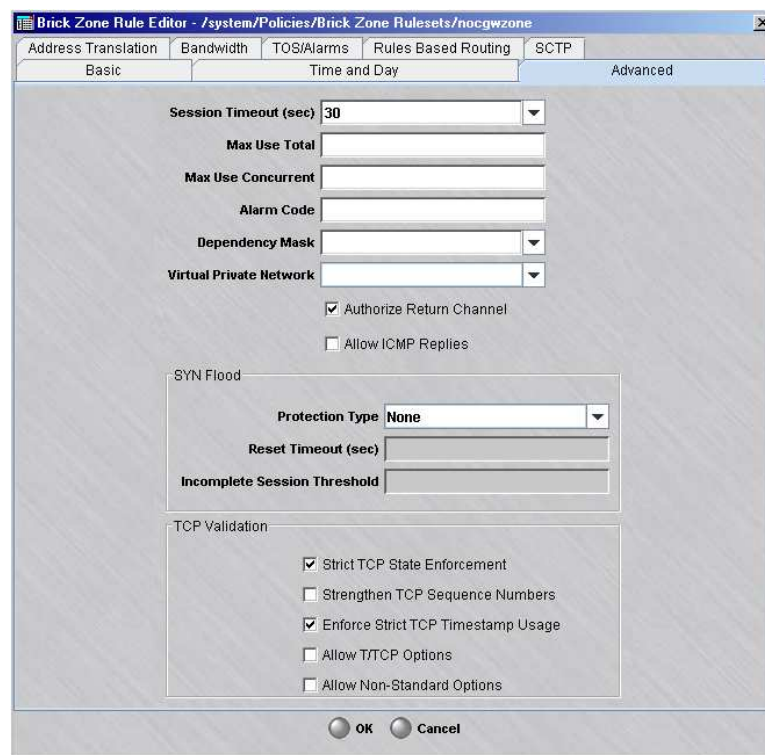
The SMS provides a number of advanced features that you can add to a security rule. These features allow you to change the session timeout, set usage limitations on the rule, associate an alarm with the rule, and set up SYN flood protection and perform strict TCP validation.

These features can be accessed by clicking **Advanced** on the Rule Editor to display the Advanced tab (Figure 1-6, “Brick Zone Rule Editor (Advanced Tab)” (p. 1-28)) .

You can add these features to a rule when you initially create the rule, or you can add them to a rule after the ruleset has been saved and applied. If you add them after the rule has been saved and applied, you will have to save and apply the ruleset again for the features to take effect.

Important! The **Dependency Mask** field allows you to associate a dependency mask with a rule. This feature is described in [Chapter 7, “Dependency Masks”](#).

Figure 1-6 Brick Zone Rule Editor (Advanced Tab)



Change the Session Timeout

The **Session Timeout** field refers to the length of time a particular packet remains in the session cache. The default timeout period for a rule with the action set to **Pass** varies according to the protocol:

- TCP = 300 seconds
- UDP = 30 second
- ICMP = 10 seconds

To change the default, enter the number of seconds in the **Session Timeout** field, or select a number from the drop-down list. The maximum number of seconds for any protocol is 99999.

If a rule uses a service group that contains more than one protocol, the longest default timeout period is the one used.

Important! If you set the Session Timeout to 0, every packet in the session will be audited. Use this only for troubleshooting your configuration.

Limit Total Usage

To limit the total number of times a rule can be invoked, enter a number in the **Max Use Total** field.

An important use of this feature is to create one-time-only rules. For example, you may want to allow an external client access to a server on your Intranet once, and no more. By setting the **Max Use Total** field to 1, you would accomplish this.

Although this rule (with the limit) remains in the Viewer Panel, it is no longer operative. If that client attempts to access the server a second time, this rule will not permit the connection to be made.

Important! If you create one-time-only rules, you may want to delete them after they have been used. This is because each time the security policy is reapplied or the Brick rebooted, the counter is reset to zero, and the rule becomes active again.

Important! The **Max Use Total** setting on the Advanced tab should not be modified for *administrativezone* and *nocgwzone* system rules 202, 205, 208, or 211.

Limit Concurrent Usage

To limit the number of active sessions that can be associated with a rule at a given time, enter a number in the **Max Use Concurrent** field.

In effect, this number determines how many sessions will be permitted through the Brick to or from a particular host (or host group) at one time.

This feature is particularly useful for limiting the load on overworked servers or as a protection against some types of denial-of-service attacks.

Associate an Alarm with a Rule

It is possible to associate an alarm with a rule, so that every time the rule is invoked by an inbound or outbound session an Administrator is notified.

To do this, you must first make sure this feature has been activated. Open the Configuration Assistant and click **Alarms** to display the Alarms parameter. Make sure the **Trigger Alarm Code** checkbox is checked (refer to the *Using the Configuration Assistant* chapter in the *SMS Administration Guide*).

Then, create an alarm trigger whose type is Alarm Code. When creating this trigger, enter the alarm code (any number between 1 - 65535). To associate this trigger with a rule, enter the same alarm code in the **Alarm Code** field on the Advanced tab. Then, when a session invokes this rule, the alarm associated with the alarm code will be triggered, and the Administrator will be notified. The method of notification (e.g., console message, email, page, etc.) is determined by the action type associated with the alarm.

Virtual Private Network

Whenever a rule is defined with an "Action" of "VPN", the user needs to determine how the encryption and decryption of IPSec packets will occur.

The possible actions that can be taken are:

- **EXTERNAL** - Outgoing packets will be encrypted, incoming packets will be decrypted. This is the default choice and should be used for Brick zone rulesets where the port is not directly connected to the Internet.
- **INTERNAL** - Outgoing packets will be decrypted, incoming packets will be encrypted. You may choose to use this in certain limited situations where the port protected by the ruleset is directly connected to the Internet.
- **BOTH** - In some networks, it may be necessary for a packet to travel through more than VPN tunnel. Rather than defining separate rules for External and Internal, you can select this option if the packet may be encrypted, decrypted and re-encrypted as it traverses the network.

The default ruleset *vpnzone* provides rules with this option set for both Internal and External. For more information on the *vpnzone*, please review *Appendix B* in the *SMS Policy Guide*.

Authorize Return Channel

The **Authorize Return Channel** checkbox is checked by default. This means you do not have to create a separate rule to permit bi-directional communication. The Brick is session-oriented, not packet oriented. The purpose of this feature is to allow packets to come back from an established session.

For example, if you create a rule allowing internal hosts to establish an HTTP session with an external server, you do not have to create a second rule allowing the reverse packets through the Brick.

Certain applications, such as SNMP traps, require only uni-directional communication. For these types of applications, uncheck this box.

Allow ICMP Replies

The **Allow ICMP Replies** checkbox is unchecked by default. If you check this checkbox, the Brick will allow any ICMP replies pertaining to this session issued by a device on the network to pass. One use of this feature is with traceroutes, because traceroutes use ICMP reply messages to track a packet's progress through the network.

Some systems use ICMP messages to pass along valuable error indications, such as "service unavailable". Unless this checkbox is checked, these packets will be dropped. When this box is checked, the validity of ICMP replies will be verified to the extent possible.

You must enable this option if you want ICMP messages, such as "Destination Unreachable," to be able to get back to the originating device when the Brick is NATing the original session.

Set Up SYN Flood Protection

A SYN flood attack is an attack in which the attacker sends a number of TCP synchronize (SYN) packets to a host, without the corresponding acknowledge packet (ACK) expected at the conclusion of the TCP 3-way handshake. Resources such as memory and processes are consumed on the host for each TCP connection setup, until the connection attempt times out. The attacker merely has to send a large number of TCP SYN packets to a host to cause that host, if unprepared, to run out of resources, which may cause the system to crash or require a reboot. Note that many servers have "hardened TCP stacks" providing a built-in SYN flood protection. Such servers do not require the Brick's syn flood protection. It is best in such cases to leave syn flood protection off so as to not interfere with the server's own protection.

For a more detailed discussion of the Brick's flood protection features, see [Appendix C, "Denial of Service Attacks"](#). To protect against a SYN flood attack, follow the steps below:

- 1 In the **Syn Flood** box on the Advanced tab, select **Send reset on timeout** from the drop-down list in the **Protection Type** field. This causes the other two fields in the box to become active.

2 The **Reset Timeout** field has a default value of 3 seconds. This value indicates how long the Brick will wait after the Incomplete Session Threshold has been reached to reset the connection to a given destination IP address. You may increase or decrease this value if you need to.

3 The **Incomplete Session Threshold** is set to 1000 incomplete sessions to a destination automatically. If the number of incomplete sessions to a particular destination IP address exceeds this number, a countdown timer is started. This timer counts down from the **Reset Timeout** value set in the previous field.

If the Brick does not observe all three packets of the TCP handshake plus a fourth data packet before the timer reaches zero, the Brick sends a TCP reset (RST) packet to the destination IP address and removes that session from its cache.

If the three-way handshake is completed through the Brick, the session is allowed to continue as usual, and the number of half-open sessions to that IP address is decreased by one.

END OF STEPS

Set up TCP Validation

The TCP Validation box contains the following checkboxes:

- **Strict TCP State Enforcement**
This box is checked by default for new rules in existing zones, all rules in new zones, and all rules in all zones for new installations. When it is checked for a particular rule, the Brick will ensure that all TCP-based sessions matching the rule begin with a valid TCP handshake, including a valid TCP checksum.
In addition, the Brick will consider packets with unusual TCP flag sets as being invalid (e.g., SYN+RST will be deemed invalid). When a TCP session fails to set up properly because of failure in the initial packet of the handshake, the end-session record will have a reason-code indicating that the session was dropped and the reason for dropping it.
- **Strengthen TCP Sequence Numbers**
This checkbox is not checked by default. When this checkbox is checked, the Brick will overwrite Initial Sequence Numbers (ISNs) with a known good pseudo-random sequence number in both directions. This allows the Brick to protect hosts that generate poor or predictable ISNs.
- **Allow T/TCP Options**
T/TCP options is a variant of TCP described in RFC 1644. If it is in use, then check this box.

- **Allow Non-Standard Options**

Some TCP versions use experimental TCP options. If you use one of these, then you must check this box.

- **Enforce Strict TCP Timestamp Usage**

If this option is unchecked, some TCP usage that is not strictly compliant with RFC 1323 will be accepted. For example, when the option is checked, connections that negotiate timestamp usage in the TCP handshake, must use it during data transfer. If the option is unchecked, then one or more data transfer segments may omit the timestamp option.



To Add Bandwidth Management to a Rule

Overview

This feature establishes the criteria for bandwidth for all sessions that use a given rule. These criteria can be set as any of the following guarantees or limits:

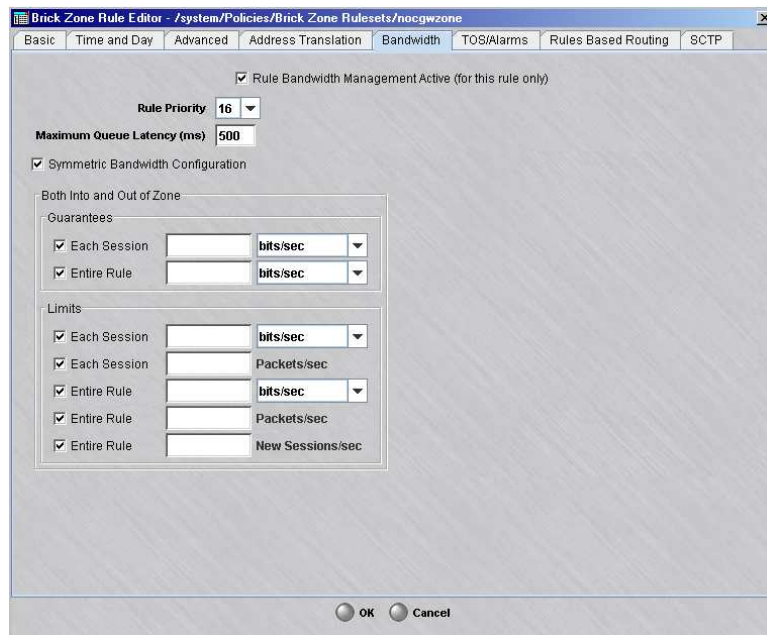
- Bit rate guarantees
- Bit rate limits
- Packet rate limits
- New session limits

Guarantee

A guarantee sets the minimum acceptable bandwidth when traffic must be reduced in one place to accommodate higher traffic demands elsewhere. A limit caps the allowable bandwidth, packet rate, or new session rate under high traffic loads. These limits can defend against denial of service attacks, for example.

To access the feature, follow the steps below:

-
- 1 In the Brick Zone Rule Editor ([Figure 1-2, “Brick Zone Ruleset Editor” \(p. 1-14\)](#)), click **Bandwidth** to display the Bandwidth tab, and then click the checkbox entitled **Rule Bandwidth Management Active (for this rule only)**. The options on the Bandwidth tab are displayed ([Figure 1-7, “Brick Zone Rule Editor \(Bandwidth Tab\)” \(p. 1-35\)](#)).

Figure 1-7 Brick Zone Rule Editor (Bandwidth Tab)

-
- 2 Set the **Rule Priority** parameter. In order to set this parameter, a few things need to be explained about how bandwidth allocation is performed.

For the purposes of this discussion, consider a class as some entity with bandwidth criteria placed on it. Classes can either be at the session, rule, zone, or physical port level. The physical port level is considered the highest level.

Rule priority controls which rule class(es) get any available excess bandwidth once their guarantees are satisfied. More generally, when two classes are competing for bandwidth, both are under their limits, and some higher level class is over its limit, the following are the rules:

- When both classes are under the guarantee, available bandwidth is allocated on a round robin basis.
- When one class is over guarantee and one is under guarantee, available bandwidth is allocated to the class that is under guarantee first.
- When both classes are over guarantee, the bandwidth is allocated to the class with the better priority (the lower value).

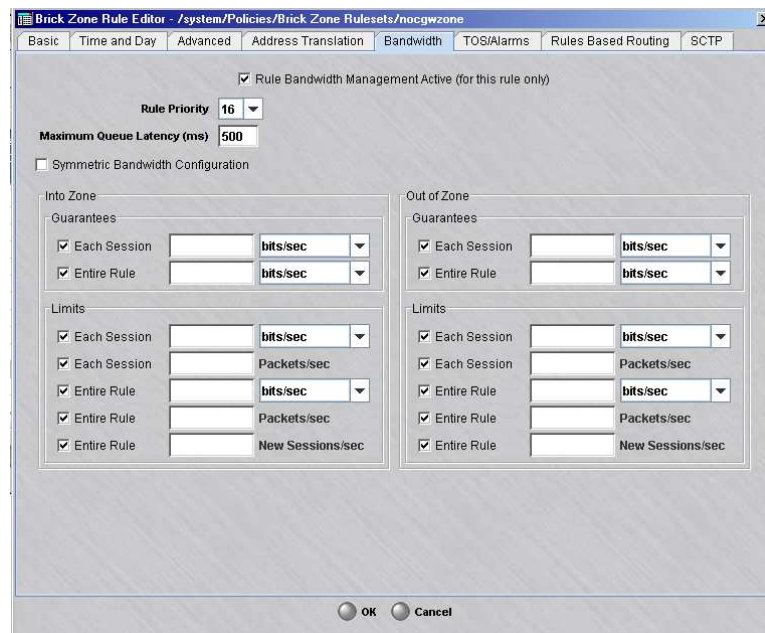
While rule priority is an important parameter for solving contention circumstances, it is much less important than the actual bandwidth parameters set below.

- 3 Set the **Maximum Queue Latency** parameter. This parameter controls the maximum amount of time a packet will be queued in the Brick, in milliseconds.

Setting this value appropriately prevents stale packets from being forwarded after they are no longer useful (i.e., lower values are usually better) and allows better bandwidth sharing (higher values are better). A value of 1 completely inhibits queuing which prevents most quality of service enforcement from occurring. Tuning this value is, at best, empirical in many applications (especially TCP-based ones). For applications such as H.323, this is fairly simple, because any latency over 50 msec is considered undesirable.

- 4 By default, all bandwidth guarantees and limits apply to packets going into, versus out of, the zone. While the Brick always enforces quality of service in a stateful manner, it can apply enforcement in an asymmetric fashion. Many applications, such as HTTP, have extremely asymmetric bandwidth requirements. If you want, however, you can set the bandwidth guarantees and limits independently in each direction. To do this, uncheck the **Symmetric Bandwidth Configuration** checkbox. This causes checkboxes on the right side of the screen to apply to "Into Zone" packets only, and a second set of checkboxes in the right hand side of the screen to apply to "Out of Zone" packets only. This display is shown in [Figure 1-8, "Bandwidth Tab \(Asymmetric Bandwidth Configuration\)"](#) (p. 1-36).

Figure 1-8 Bandwidth Tab (Asymmetric Bandwidth Configuration)



-
- 5** Set the **Guarantees** and **Limits** parameters. The check boxes next to each parameter enable or disable that parameter while preserving its configured value.

Note that the guaranteed bandwidth specified here for a rule is summed across all the rules in the zone. These guarantees must be no more than the guaranteed 'Transmit Bandwidth' and 'Receive Bandwidth' for the zone (as configured in the Policy Assignment tab of the Brick Editor).

END OF STEPS



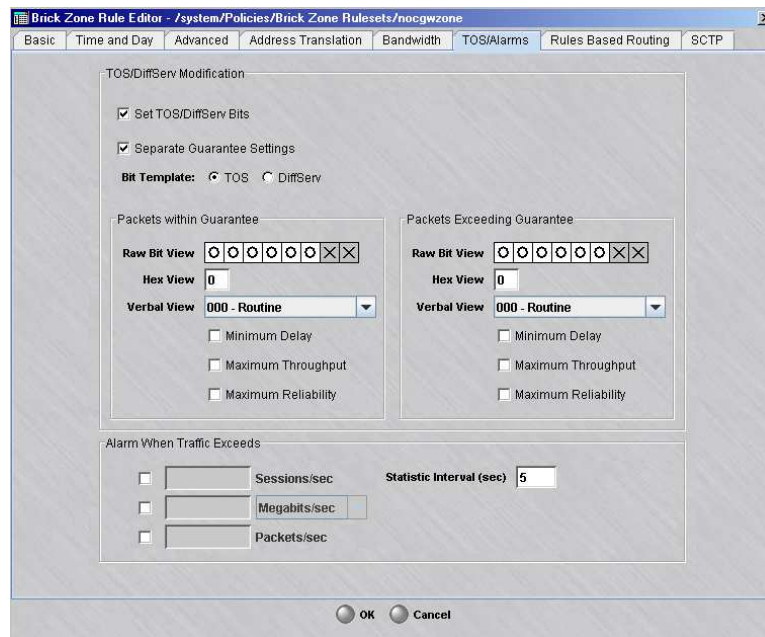
To Add TOS/Alarm Capability to a Rule

Overview

DiffServ is a method (defined in RFCs 2474/2475) that intermediate systems like routers can use to prioritize IP traffic, depending on the setting of the precedence portion of the ToS field (bits 0, 1, 2) and bits 3, 4, 5 of the ToS field.

To edit the TOS/DiffServ and QoS alarm parameters, click **TOS/Alarms** in the Brick Zone Rule Editor (Figure 1-2, “Brick Zone Ruleset Editor” (p. 1-14)) to display the TOS/Alarms tab (Figure 1-9, “Brick Zone Rule Editor (TOS/Alarms Tab)” (p. 1-38)).

Figure 1-9 Brick Zone Rule Editor (TOS/Alarms Tab)



These parameters inform the Brick device how to set the TOS/DiffServ bits in the second byte of the IP header, per firewall rule. Several equivalent representations of the ToS bit are presented for user convenience. There are separate settings for packets within the guarantee and those exceeding the guarantee. The Set **TOS/DiffServ Bits** feature can be disabled while still preserving the configuration.

If the TOS settings are also configured and enabled at the zone level (see the Policy Assignment tab on the Brick editor), the TOS values in the rules will be overridden by the values assigned for the zone.

One can also have a ToS/DiffServ byte pattern serve as matching criteria for the rule to be triggered. This is done in the **TOS/DiffServ Match** field in the Basic tab (see Step 8 on Step 8). The field takes two hex digits or a range and a mask. The rule match is contingent on the second byte of the IP header having this value.

You can create an alarm called **Rule Bandwidth Exceeded Alarm** (see the *Reports, Alarms and Logs Guide*). Part of its configuration is done here under the *Alarm When Traffic Exceeds* section (see [Figure 1-9, “Brick Zone Rule Editor \(TOS/Alarms Tab\)”](#) (p. 1-38)) and is triggered when traffic exceeds any of the rule limits specified here. These limits may be individually disabled. The **Statistic Interval** specifies how long, in seconds, the Brick accumulates traffic counts before it calculates the averages.



To Configure IP Address(es) for Rules-Based Routing

Overview

The Rules Based Routing tab of the Brick Zone Rule Editor is used to configure IP addresses for forward routing of packets and/or reverse routing of packets to a proxy server for content filtering if the Rules-Based Routing feature is being used and the packet matches the rule that is being applied. One or both IP address fields can be entered, or can be left blank (the default) if the Rules-Based Routing feature is not being used. A host group can be selected in place of an individual IP address to forward route packets to all/some of the hosts in the host group. The host group that is selected must have individual IP address entries (no address ranges) and cannot have nested host groups.

Complete the following steps to configure IP addresses for the Rules-Based Routing feature:

- 1 In the Brick Zone Rule Editor (Figure 1-2, “Brick Zone Ruleset Editor” (p. 1-14)), click **Rules Based Routing** to display the Rules Based Routing tab.

Result The Rules Based Routing tab of the Brick Zone Rule Editor is displayed (Figure 1-10, “Brick Zone Rule Editor (Rules Based Routing Tab)” (p. 1-40)).

Figure 1-10 Brick Zone Rule Editor (Rules Based Routing Tab)



2	To	Do This
	<p>Configure an IP address or host group for forward routing of packets that match the rule</p>	<p>Enter an IP address in the Forward Route Address or Host Group field.</p> <p><i>OR</i></p> <p>Click the down arrow next to the field to display a drop-down list and select a host group from the list. You can select /*BROWSE*/, which displays a Browse window and allows you to select a host group in a different group.</p> <p>If a host group is selected, select one of the options in the field below the Address or Host Group field:</p> <ul style="list-style-type: none"> • Always Route to All Hosts (the default) • Route to Hosts that Only Respond to Ping • Route to Hosts that Only Respond to SYN <p>If this option is chosen, enter a Source TCP Port and Destination TCP Port (valid values are 1 to 65535, inclusive (no default).</p>
	<p>Configure an IP address for reverse routing of packets that match the rule</p>	<p>Enter an IP address in the Reverse Route Address field.</p> <p>Optionally, click the Route Return Packets Back to Cached Source MAC/Media checkbox to invoke the Brick device to use the source MAC and media of the first packet as the destination MAC and media of the return packets.</p>

3 Click the **OK** button to apply the change(s) on the tab.

END OF STEPS



To Enable Sctp Processing on a Rule

When to use

Use this procedure to enable processing of Stream Control Transport Protocol (SCTP) packets for a rule within a Brick zone ruleset.

Task

Complete the following steps to enable Sctp processing on a rule.

- 1 In the Brick Zone Rule Editor (Figure 1-2, “Brick Zone Ruleset Editor” (p. 1-14)), click **SCTP** to display the Sctp tab.

Result The Sctp tab of the Brick Zone Rule Editor is displayed (Figure 1-11, “Brick Zone Rule Editor (Sctp Tab)” (p. 1-42)).

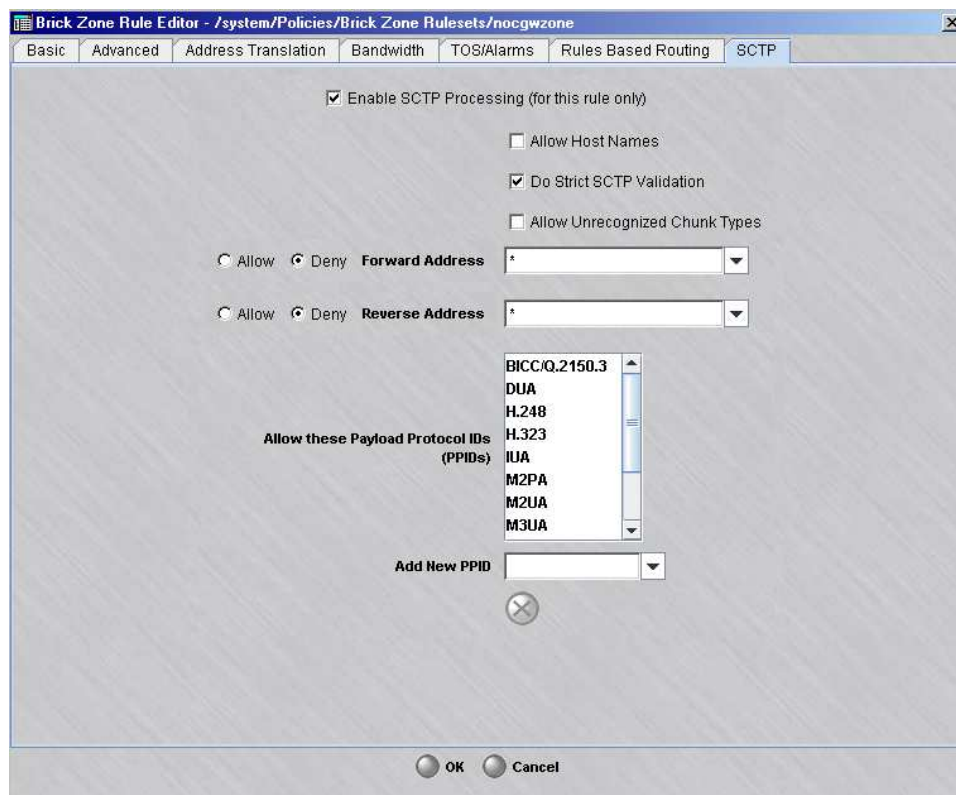
Figure 1-11 Brick Zone Rule Editor (Sctp Tab)



- 2 To enable Sctp processing, click the **Enable Sctp Processing (for this rule only)** checkbox.

Result The SCTP processing feature is enabled and related options are displayed on the SCTP tab (Figure 1-12, “Brick Zone Rule Editor (SCTP Feature Enabled With Processing Options)” (p. 1-43)).

Figure 1-12 Brick Zone Rule Editor (SCTP Feature Enabled With Processing Options)



- To permit selection of a specific host group as the source/destination address(es) from which to allow or deny SCTP traffic, click the **Allow Host Names** checkbox (this option is unchecked, by default).
- To perform basic SCTP validation, which validates the size/range of only certain “chunks” of the SCTP packet by the Brick device, as defined in RFC 2960 (SHUTDOWN, SHUTDOWN_ACK, COOKIE_ACK, SHUTDOWN_COMPLETE), click the **Do Strict SCTP Validation** checkbox to remove the check. To enforce the validation of all “chunk types” of the SCTP packet by the Brick device, as defined in RFC 2960, check the **Do Strict SCTP Validation** checkbox to place a check in it (this is the default).

-
- 5 To allow the Brick device to handle additional SCTP packet “chunk types” not covered in RFC 2960, click the **Allow Unrecognized Chunk Types** checkbox (this option is unchecked, by default).
-

6

To	Do This
Allow SCTP traffic from a specific source address or host group	Click the Allow radio button to the left of the Forward Address field and enter an asterisk (*) (which is the default value) or IP address in this field, or select a host group from the drop-down list
Deny SCTP traffic from a specific source address or host group	Click the Deny radio button to the left of the Forward Address field and enter an asterisk (*) (which is the default value) or IP address in this field, or select a host group from the drop-down list
Allow SCTP traffic to be delivered to a specific destination address or host group	Click the Allow radio button to the left of the Reverse Address field and enter an asterisk (*) (which is the default value) or IP address in this field, or select a host group from the drop-down list
Deny SCTP traffic from being delivered to a specific destination address or host group	Click the Deny radio button to the left of the Reverse Address field and enter an asterisk (*) (which is the default value) or an IP address in this field, or select a host group from the drop-down list

-
- 7 The **Allow these Payload Protocol IDs (PPIDs)** field lists the current SIGTRAN applications that will be allowed as payload in the SCTP packets.

To add a SIGTRAN application to this list, click the down arrow next to the **Add New PPID** field to display a drop-down list, and select the application ID.

To remove a SIGTRAN application, select the application ID from the **Allow these Payload Protocol IDs (PPIDs)** field list and click the Remove (⊗) button.

-
- 8** Click the **OK** button to apply the change(s) on the tab.

END OF STEPS



To Assign a Brick Zone Ruleset to a Physical Port

When to use

Use this task to assign a Brick zone ruleset to a physical port on the Brick. For the ruleset to take effect, you have to assign it to a port on a least one Brick device. This is done using the Policy Assignment tab of the Brick Editor

Task

Complete the following steps to assign a zone ruleset to a port.

- 1 Open the Devices folder and click the Bricks folder to display all configured Bricks.
-

- 2 Double-click the Brick that contains the port.

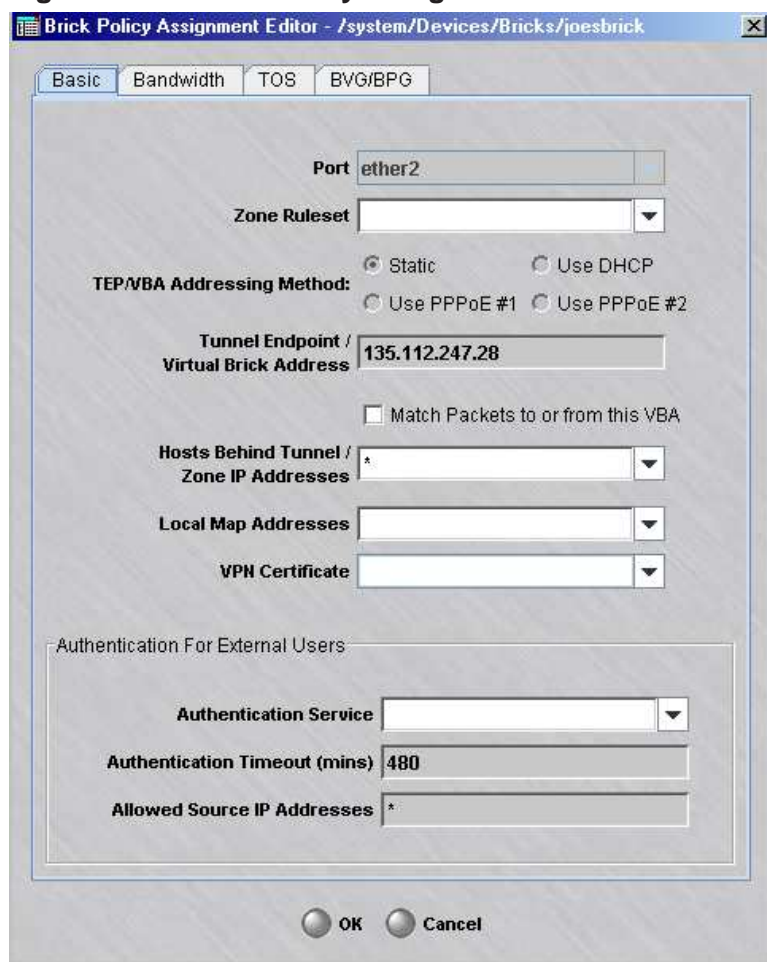
Result The Brick Editor is displayed.

- 3 Click **Policy Assignment** to display the Policy Assignment tab.
-

- 4 Double-click the port to be provisioned.

Result The Brick Policy Assignment Editor is displayed, with the port you selected entered in the **Port** field ([Figure 1-13, “Brick Policy Assignment Editor”](#) (p. 1-47)).

Figure 1-13 Brick Policy Assignment Editor



-
- 5 In the **Zone Ruleset** field, display the drop-down list and select a ruleset from the rulesets listed, or select **Browse**. A Browse window will appear. Double-click the Brick zone ruleset you want to assign to the port.
-
- 6 Select the **TEP/VBA Addressing Method** to be used - Static, Use DHCP, Use PPPoE#1, or Use PPPoE#2.

.....

7 If you want to match packets to or from the VBA, place a check in the checkbox for **Match Packets to or from this VBA**.

.....

8 If this port will be the endpoint of a LAN-LAN or client tunnel, enter an IP address in the **Tunnel Endpoint/Virtual Brick Address** field.

This address is also the Virtual Brick Address (VBA). If you will be using Network Address Translation, you may require a VBA (refer to [Chapter 6, “Network Address Translation”](#)).

.....

9 In the **Hosts Behind Tunnel/Zone IP Addresses** field, enter the addresses of the hosts to be protected by this Brick zone ruleset.

There are several ways to enter this information:

- Leave the default asterisk in place. This means every host connected to this port will be protected by the ruleset.
 - Replace the asterisk with an IP address, a range of IP addresses (in the format 1.1.1.1-1.1.1.10), or an IP address and subnet mask (for example, 10.10.10.0/24).
 - Display the drop-down list and select and select a host group from the host groups listed, or select **Browse**. A Browse window will appear and allow you to select a host group. For details about host groups, refer to [Chapter 2, “Host Groups”](#).
-

10 If you will be providing client tunnel users with an address on your local LAN, enter the address or range, or enter a host group from the drop-down list in the **Local Map Addresses** field. (See [Appendix A, “Local Presence”](#) for a more detailed discussion of local presence).

.....

11 If you will be using an X.509 certificate from Entrust or Verisign to authenticate client tunnel users, enter the name of the certificate in the **VPN Certificate** field. This must be the same name you used when setting up this certificate using the Certificate Manager (see [Chapter 9, “Digital Certificates”](#)).

.....

12 If this port will terminate a client tunnel, the users of that tunnel will have to be authenticated. If you will be using a database of user accounts that does *not* reside on the SMS host, you will have to enter an authentication service in the **Authentication Service for External Users** box.

If the port will not be terminating a client tunnel, or if the user database resides on the SMS host, you can skip this field. See [“Overview” \(p. 8-1\)](#) for a detailed explanation of user authentication.

-
- 13 Click **OK** to dismiss the Brick Policy Assignment Editor and return to the Policy Assignment tab of the Brick Editor. The new port assignment will be displayed.
-
- 14 Display the File menu and select **Save and Apply**, to apply the ruleset assignment to the Brick port.

Result The Apply Brick window is displayed, with the associated Brick device ([Figure 1-14, “Apply Brick Window” \(p. 1-49\)](#)).

Figure 1-14 Apply Brick Window



-
- 15 By default, the Brick session cache is used for comparison against the first packet in a session to determine whether to pass all subsequent packets. Optionally, click Clear cache to clear the Brick session cache. (For additional details, refer to [Figure 5-26, “Network Topology - Example 3” \(p. 5-52\)](#)).

.....
16 Click **OK** to apply the ruleset assignment to the selected Brick port.

.....
E N D O F S T E P S
.....



Maintaining Brick Zone Rulesets

Overview

Brick zone rulesets have to be maintained in order to stay up-to-date. This can involve editing the Brick zone ruleset (such as adding, changing or deleting rules), copying or moving the Brick zone ruleset, or deleting the Brick zone ruleset.

This section explains how to copy, move and delete Brick zone rulesets. It also explains how to apply a Brick zone ruleset. The section [“Maintaining Security Rules” \(p. 1-59\)](#) explains how to modify and delete rules in the Brick zone ruleset.

To view existing Brick zone rulesets

Complete the following steps to display all existing Brick zone rulesets in a given folder.

-
- 1 From the Navigator window, open the folder of the group that contains the Brick zone rulesets you want to view.
-
- 2 Open the Policies folder and click the Brick Zone Rulesets folder. All existing rulesets for that folder will be displayed in the Navigator window (see [Figure 1-15, “View Brick Zone Rulesets” \(p. 1-52\)](#)).

For each ruleset, the window shows the Administrator who created it, the date and time it was **Last Modified**, and a brief **Description**, if one was included when the ruleset was created.

Figure 1-15 View Brick Zone Rulesets

Alcatel-Lucent Security Management Server

Folder: /system/Policies/Brick Zone Rulesets/ - Total 7 item(s)

Name	Admin	Last Modified	Description
administrativezone	setup	2007-04-09 11:31:38	Zone protecting the LSMS
dhcpzone_on_inside_if	setup	2007-04-09 11:31:39	Default zone for mobile bri
dhcpzone_on_internet_if	setup	2007-04-09 11:31:39	Default zone for mobile bri
firewall	setup	2007-04-09 11:31:38	Zone protecting the brick (
nocgwzone	admin	2008-06-12 16:31:21	Zone protecting the LSMS
proxyzone	setup	2007-04-09 11:31:39	Zone protecting the proxie
vpnzone	setup	2007-04-09 11:31:38	An example rule set allowi

Brick zone rulesets

END OF STEPS

To use the traffic matcher tool

Complete the following steps to use the Traffic Matcher tool. The Traffic Matcher tool allows you to enter specifications for simulated traffic and then search for existing rules in a given Brick zone ruleset that would be triggered by the particular traffic pattern entered.

- 1 Select the **Brick Zone Rulesets** folder in the Folders panel.

Result A list of existing Brick zone rulesets is displayed in the Contents panel.

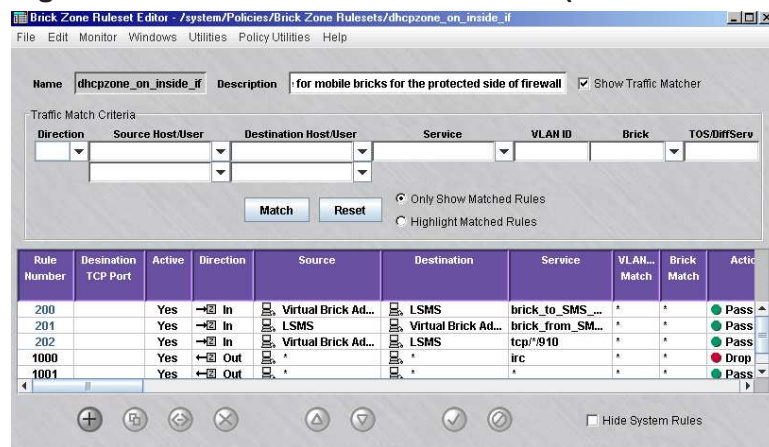
- 2 Double-click on a Brick zone ruleset in the Contents panel.

Result The Brick Zone Ruleset Editor is displayed with a table of existing rules in the ruleset.

- 3 Click the **Show Traffic Matcher** checkbox in the top portion of the Brick Zone Ruleset Editor.

Result The fields for the Traffic Matcher tool are displayed on the Brick Zone Ruleset Editor above the table of rules in a box entitled **Traffic Match Criteria** (Figure 1-16, “Brick Zone Ruleset Editor (with Traffic Matcher Tool Enabled)” (p. 1-53)).

Figure 1-16 Brick Zone Ruleset Editor (with Traffic Matcher Tool Enabled)



- 4 Enter the value(s) in one or more fields of the **Traffic Match Criteria** box on which to conduct a search for the matching rule(s).

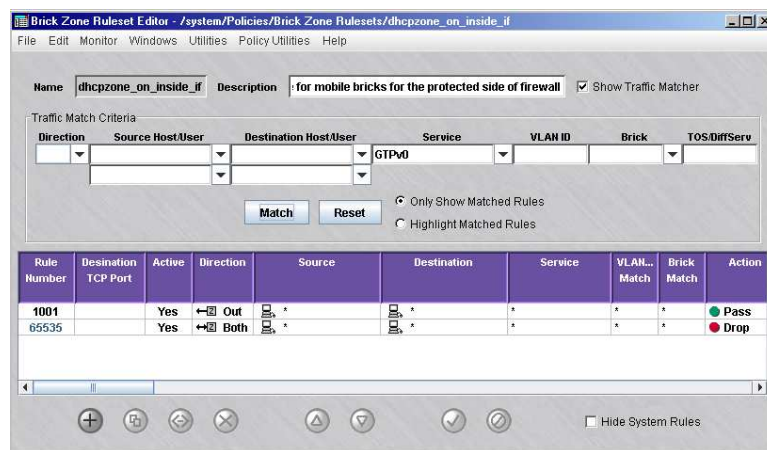
If you leave a field blank, a rule will not be excluded on the basis of that field and it will be ignored in the rule search. If all fields are left blank, all rules will be retrieved when a search is performed.

To clear the value(s) entered in the traffic fields, click the **Reset** button and re-enter the new value(s).

5 Click the **Match** button.

Result If one or more rules in the ruleset match the traffic pattern entered, and the **Only Show Matched Rules** radio button is selected (the default), only the rules that match the specified traffic criteria are displayed in a table on the Brick Zone Ruleset Editor. [Figure 1-17, “Brick Zone Ruleset Editor \(Traffic Match Search Performed\)”](#) (p. 1-54) shows an example of a search that was performed on a given ruleset using specific traffic criteria.

Figure 1-17 Brick Zone Ruleset Editor (Traffic Match Search Performed)



To display the entire set of rules in a ruleset and highlight the rules that match the entered criteria, select the **Highlight Matched Rules** radio button. The Editor window will re-display the entire list of rules in the ruleset with the rule(s) that match the traffic criteria entered highlighted in yellow.

6 When you have completed the search, close the Brick Zone Ruleset Editor.

END OF STEPS

To copy a Brick zone ruleset

Complete the following steps to copy a Brick zone ruleset. When you edit a Brick zone ruleset (using the **Save As** option to save the edited version with a new name) or copy a Brick zone ruleset, you have the option of copying components from the original ruleset (host groups, nested host groups, service groups, application filters), specifying how those components will be named and where they will be created and stored.

- 1 With the existing Brick zone rulesets displayed in the Navigator window, right-click on the Brick zone ruleset to be copied and select **Copy** from the pop-up menu.

Result A Copy window is displayed (Figure 1-18, “Copy Window” (p. 1-55)).

Figure 1-18 Copy Window



- 2 Complete the following fields, as applicable:
 - **Name**— Enter a new name for the ruleset copy (Example: **nocgwzone_1**). This field is required. The name entered must be different from the name of the original ruleset by at least one character if it is in the same group.
 - **Folder**— The ruleset copy will be created and stored in the **Brick Zone Rulesets** folder under the *system* group, by default. To specify a different folder, click the **Browse** button, which brings up a Select a Brick Zone Ruleset Folder window, and select the folder where the ruleset copy will be created and stored.
 - **Copy components when copying within a Group**— Click this checkbox (place a check in it) to copy the components from the original ruleset. If this checkbox is unchecked, and the copied ruleset is in the same group as the original one, the copied ruleset uses the same components as the original. If the ruleset is copied across groups, the components are copied regardless of this field setting.

- **Component Suffix**— If the **Copy components when copying within a Group** option is checked, enter a text string that will be appended to the name of each copied component (host group, nested host group, service group, application filter). If this field is left blank, the name of each copied component will be spelled the same way as the original one. Name collisions are possible, regardless of whether or not an entry is made in this field. If this occurs, the content of the conflicting component(s) is examined. If the content is the same, the duplicated ruleset uses the existing component(s) and the new copy(ies) is not made. If the content is different, a copy of the component(s) is created and stored under a unique name by appending its database ID to the name.

Note: the Component Suffix used must be unique and cannot be reused across zone rulesets.

- **Component Folder**— This field is used to specify a common folder name for the copied components. The folder name is relative to the top-level folder of each type of component. If this field is left blank, the copied components are stored in the top-level folder of that component. This folder is created as needed (if it does not already exist). A path string can be specified; / or \ delimiters can be used to delineate folder names.

-
- 3 Click **OK** to copy the ruleset and dismiss the Copy window.

Result The system returns to the Navigator. The copy of the ruleset appears in the Contents panel.

The ruleset (and components, if specified) are copied to the chosen folder(s).

.....
E N D O F S T E P S
.....

To move a ruleset

Complete the following steps to move a ruleset::

-
- 1 With the existing rulesets displayed in the Navigator window, right-click the ruleset to be moved and select **Move** from the pop-up menu. A Browse window will appear.
-
- 2 Select the folder you want to move the ruleset to, and click **OK** to dismiss the Browse window. The ruleset will be moved to the folder you indicated.

.....
E N D O F S T E P S
.....

Delete a Ruleset

To delete a ruleset, follow the steps below:

- 1 With the existing rulesets displayed in the Navigator window, right-click the ruleset you want to delete and select **Delete** from the pop-up menu. A confirmation window similar to the one shown in [Figure 1-19, “Confirmation Window \(Brick Zone Rulesets\)”](#) (p. 1-57) will appear.

Figure 1-19 Confirmation Window (Brick Zone Rulesets)



- 2 Click **Yes** to delete the ruleset. The ruleset will be removed from the Navigator window.

END OF STEPS

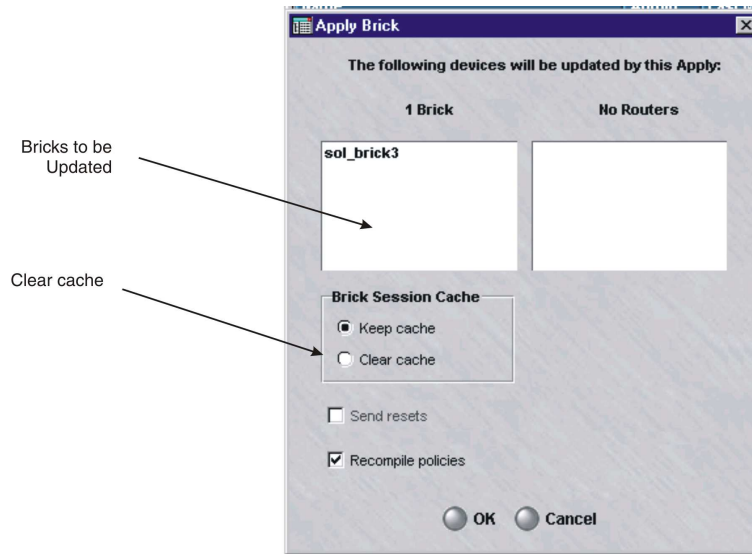
To apply a Brick zone ruleset

Whenever changes are made to a ruleset that has been assigned to a port on one or more Brick devices, the ruleset has to be applied to the Brick devices. There are two ways to apply a ruleset:

- Right-click the ruleset in the Navigator window and select **Apply** from the pop-up menu
— or —
- Select **Save and Apply** from the File menu in the Brick Zone Ruleset Editor after making the changes.

Regardless of which method you use, the Save Policy window will appear. This window, which is shown in [Figure 1-20, “Apply Policy Window”](#) (p. 1-58), indicates the Bricks that will be updated by this apply.

Figure 1-20 Apply Policy Window



Click **OK** to perform the apply.

You cannot select specific Bricks to be updated, and not others. If you only want to update one of the Bricks shown, you will have to do a device update on that Brick instead.

Important! *Brick SESSION CACHE*

When you select "clear cache", the "send resets" checkbox, which was grayed out, becomes active and is checked by default. This causes the Brick to send TCP resets to all TCP sessions for which strict TCP enforcement is enabled. The Brick sends two resets per session: one to each endpoint, using the other endpoint as the source address.

This makes it appear to each endpoint that the other endpoint has aborted the connection. If the resets are not sent, then the endpoint will continue to retransmit any outstanding (unacknowledged) packets it was sending for some number of minutes. If TCP strict enforcement is in effect, these retransmissions will be dropped.

By sending the resets, the Brick forestalls these retransmissions. Also, depending on the application, the end user may be given timely notification that the connection has been broken.

□

Maintaining Security Rules

Overview

Security rules have to be maintained in order to stay up-to-date. This can involve modifying rules that require updating, deleting rules that are no longer needed, and renumbering rules that have gotten out of order.

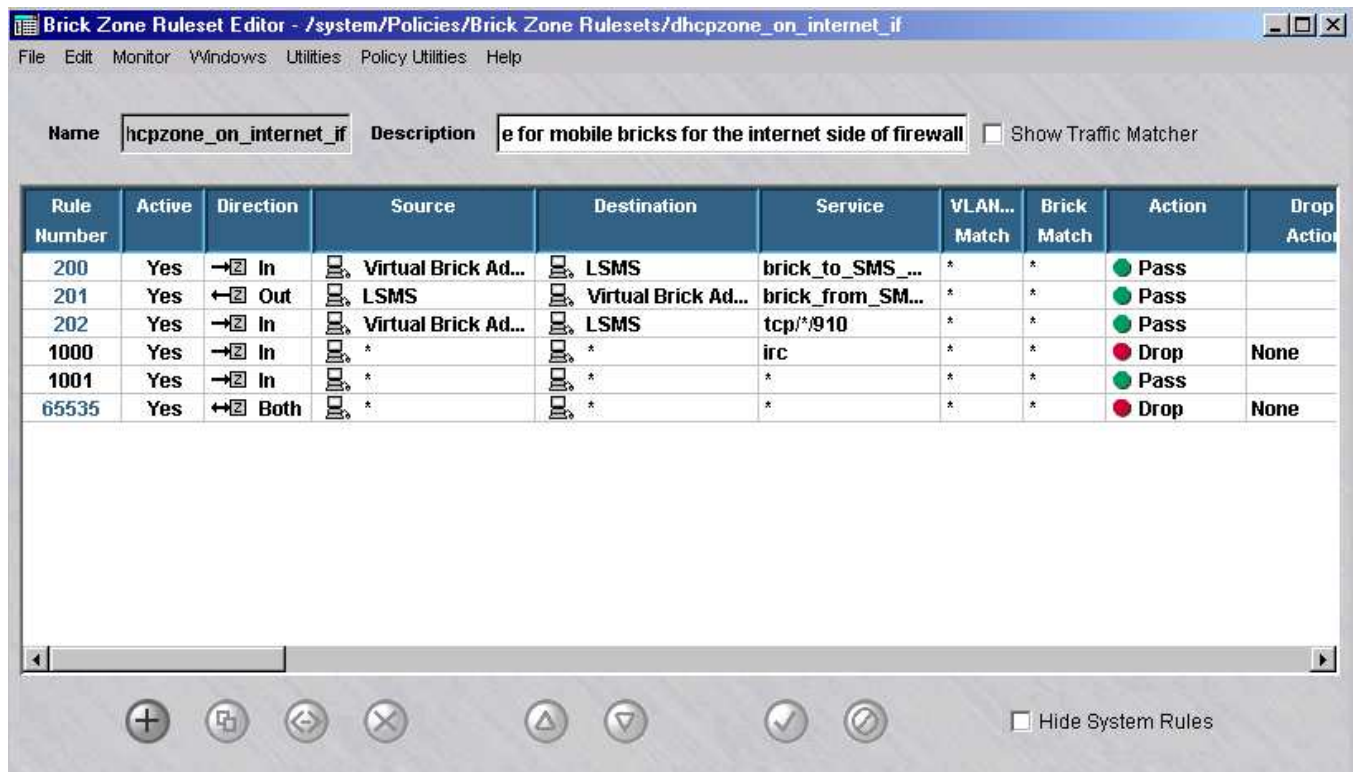
To view security rules

Complete the following steps to view the security rules in a Brick zone ruleset.

- 1 Access the Brick Zone Ruleset Editor as described in “To display the Brick Zone Ruleset Editor” (p. 1-14).

The Brick Zone Ruleset Editor (see Figure 1-21, “Brick Zone Ruleset Editor (Viewing Rules in a Zone Ruleset)” (p. 1-59)) shows all the security rules that have been created to date for that ruleset, in ascending order.

Figure 1-21 Brick Zone Ruleset Editor (Viewing Rules in a Zone Ruleset)













-
- 2 To filter the display of rules in the zone ruleset, use the Traffic Matcher tool. Refer to [“To use the traffic matcher tool” \(p. 1-53\)](#) for instructions on how to filter the display of rules using the Traffic Matcher tool.

For each rule, it shows all the fields for the Basic, Time of Day, Advanced, Address Translation, Bandwidth and TOS/Alarms tabs.

However, to view all the fields, you will have to use the horizontal scrollbar, since they do not all fit on the screen at one time.

ICONS

The Brick Zone Ruleset Editor uses the icons below to indicate the direction, source, destination and action of each rule:

	In to	Out of	Both		
Direction =					
		Host Group	User Group		
Source/Destination =					
Action =	Pass	Drop	Proxy	VPN	VPN Proxy
					
	(green)	(red)	(yellow)	(purple)	(purple/yellow)

END OF STEPS

To modify a rule

Complete the following steps to modify a security rule:

- 1 From the Navigator window, open the appropriate Group and Policies folders, and click the Brick Zone Rulesets folder to display all rulesets.
- 2 Double-click the Brick zone ruleset containing the rule you want to modify. The Brick Zone Ruleset Editor will appear (see [Figure 1-2, “Brick Zone Ruleset Editor”](#) (p. 1-14)).
- 3 Double-click the rule you want to modify. The Brick Zone Rule Editor will appear (see [Figure 1-3, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-17)) with the fields populated with the rule you selected.
- 4 Make any necessary changes to the information in the Basic, Time of Day, Advanced, Address Translation, Bandwidth and TOS/Alarms tabs.
- 5 When you have finished, click **OK** to dismiss the Brick Zone Rule Editor and return to the Brick Zone Rulesets Editor. The changes you just made will appear in the rule.
- 6 Display the File menu and select one of the **Save** options. You may want to choose **Save and Apply**, since the changes will not take effect until they are applied to the Brick(s) on which this ruleset is loaded (see [“To apply a Brick zone ruleset”](#) (p. 1-57) above).

END OF STEPS

To duplicate a rule

Complete the following steps to duplicate a security rule:

- 1 With the Brick Zone Ruleset Editor displayed, select the rule you want to duplicate, right-click the rule, and select the **Duplicate** option. The Brick Zone Rule Editor will appear (see [Figure 1-3, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-17)) with the fields populated with the rule you selected.

- 2 Make any necessary changes to the information in the Basic, Time of Day, Advanced, Address Translation, Bandwidth and TOS/Alarms tabs. If you do not make any changes, you will have two identical rules in the ruleset, which serves no purpose.
- 3 When you have finished, click **OK** to dismiss the Brick Zone Rule Editor and return to the Brick Zone Rulesets Editor. The new rule will appear one row above the rule you initially selected to duplicate.
- 4 Display the File menu and select one of the **Save** options.

END OF STEPS

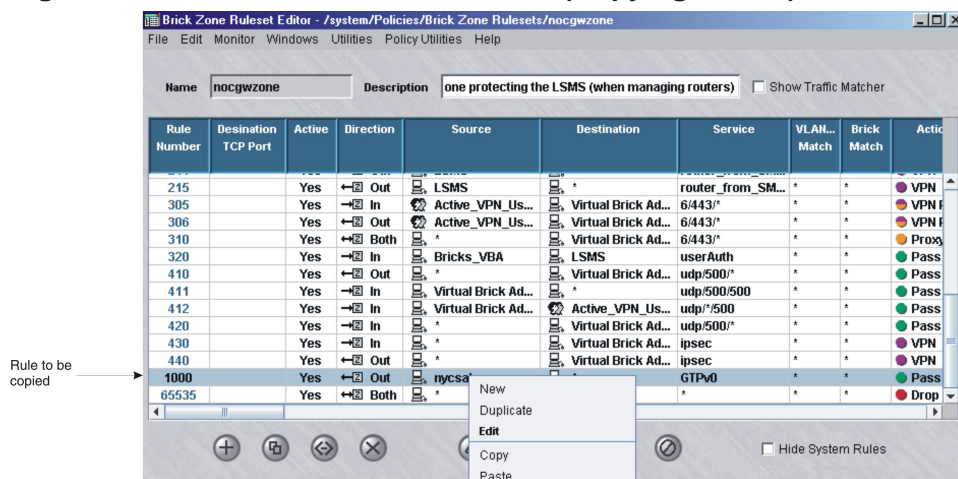
To copy/paste a rule

One or more rules can be copied at a time and then pasted into the same ruleset, or pasted into another ruleset in the same group. The pasted rule can be edited and changed, just like any other rule. Only user-defined rules can be copied and pasted.

Complete the following steps to copy/paste a rule:

- 1 With the Brick Zone Ruleset Editor open, right-click on the rule you wish to copy and select **Copy** from the pop-up menu (Figure 1-22, “Brick Zone Ruleset Editor (Copying a Rule)” (p. 1-62) shows an example).

Figure 1-22 Brick Zone Ruleset Editor (Copying a Rule)



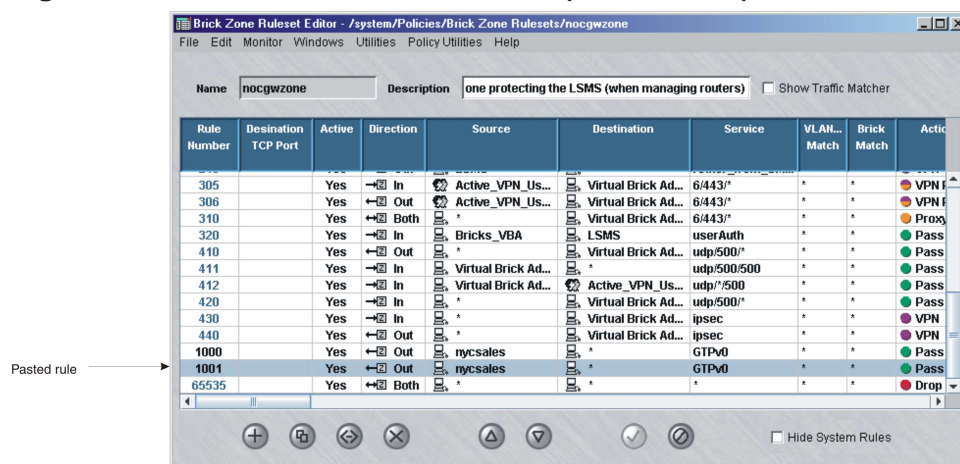
To select a range of rules to copy, click on the first rule to highlight it, then shift-click on the last rule to highlight it and all the intervening rules. Right-click somewhere in the selected range, and choose **Copy** from the pop-up menu.

To add single rules to the selection, press Ctrl-click on each rule to be copied, right-click on one of the rules in the selection, and choose **Copy** from the pop-up menu.

- 2 Select the point (rule) in the ruleset where the copied rule(s) will be pasted, right-click on the rule, and choose **Paste** from the pop-up menu.

Result The copied rule(s) is pasted in the ruleset above the rule selected, in the same relative order in which they were copied (Figure 1-23, “Brick Zone Ruleset Editor (Pasted Rule)” (p. 1-63) shows an example).

Figure 1-23 Brick Zone Ruleset Editor (Pasted Rule)



The pasted rule(s) is highlighted and renumbered according to its relative position in the ruleset.

To make changes to the copied/pasted rule, double-click on it to bring up the Brick Zone Rule Editor. You can change the Direction, Source, Destination, Service or Group, Action, and so forth.

Note: All changes to a Brick zone ruleset must be saved and applied to the Brick device to take effect.

END OF STEPS

To delete a rule

To delete a security rule, follow the steps below:

- 1 With the Brick Zone Ruleset Editor displayed, select the rule you want to delete and click the **DELETE** button.

A confirmation window similar to the one shown in [Figure 1-24, “Confirmation Window \(Rules\)”](#) (p. 1-64) will appear

Figure 1-24 Confirmation Window (Rules)



- 2 Click **Yes** to delete the rule and dismiss the Confirmation window. The rule will be removed from the Brick Zone Ruleset Editor.

END OF STEPS

Renumber a Rule

The order of the rules is important because the Brick applies the rules to incoming and outgoing packets in ascending numerical order, beginning with the first rule. This includes both the rules created automatically by the SMS and the rules created by Administrators.

If there is no match, the next rule is applied, and so forth, until the last, the default Drop All rule, is applied, and the packet is dropped.

To renumber a rule, select the rule and click the **UP** or **DOWN** buttons. Repeat as often as necessary until the rules are in the correct order. The rule will move up or down one row with each click.

Activate or Deactivate a Rule

You can activate and deactivate rules. When a rule is deactivated, it remains in the ruleset, but it is not applied to traffic through the Brick port until an Administrator activates the rule. You can tell if a rule is activated or deactivated by looking at the Active column in the Brick Zone Ruleset Editor (see [Figure 1-2, “Brick Zone Ruleset Editor”](#) (p. 1-14)).

To activate a rule, select the rule in the Brick Zone Ruleset Editor, right-click the rule, and select the **Activate** option.

To deactivate a rule, select the rule, right-click the rule, and select the **Deactivate** option.



Brick Zone Ruleset Templates Provided with the SMS Application for VoIP/NOE Traffic

Overview

A set of pre-defined Brick zone rulesets are provided with the SMS application when it is installed to make it easier to provision the Brick to monitor and protect call and data traffic in a VoIP network. Each pre-defined Brick zone ruleset is pre-configured with the required rules and other rule components (pre-defined host groups, service groups, application filters) which allow the Brick to secure the media and call control sessions at a specific location in the VoIP network.

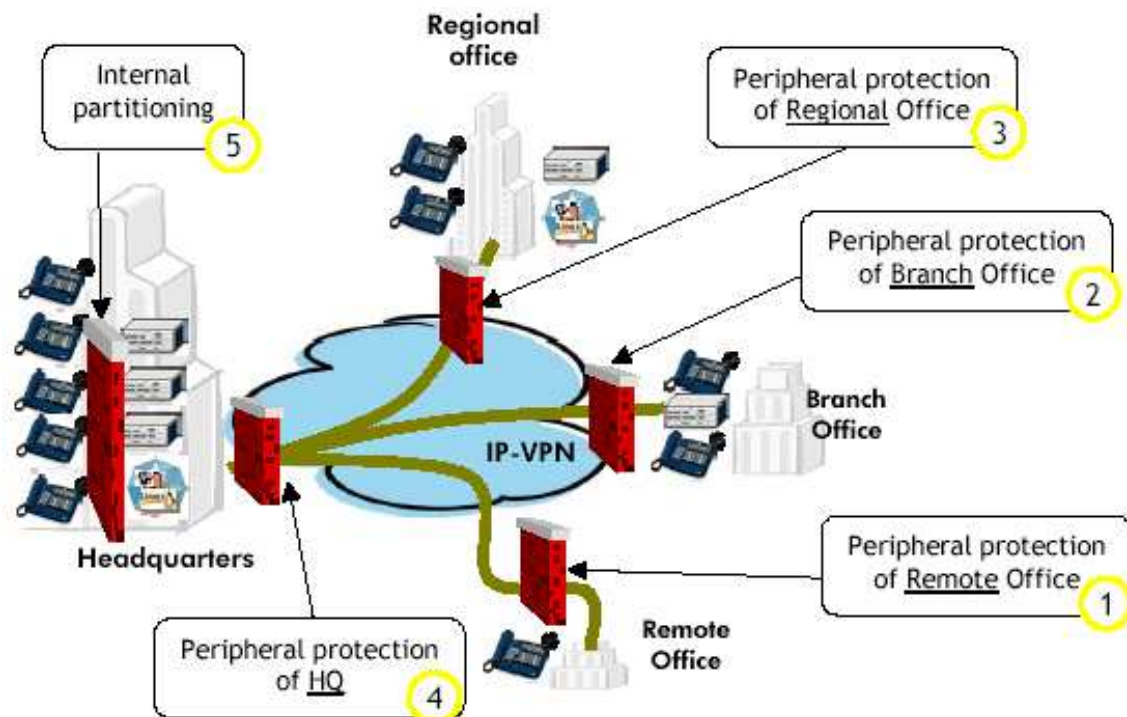
All required settings and parameters are pre-provisioned within these Brick zone ruleset templates for VoIP traffic. All that is required is for you to edit the host group templates called within the rules of the ruleset and add the IP addresses of the equipment (IP phones, servers, MGWs) that are behind the Brick zone being protected by the selected ruleset.

For details about editing host groups and adding IP addresses for equipment, including the pre-defined host group templates that are provided for VoIP traffic, refer to [Chapter 2, “Host Groups”](#).

Choosing the pre-defined Brick zone ruleset to be used

The name of the pre-defined Brick zone ruleset to be applied is based on the Brick's deployment location (refer to [Figure 1-25, “Deployment Sites for Brick Serving as Application Layer Gateway \(ALG\) in VoIP Network”](#) (p. 1-67)).

Figure 1-25 Deployment Sites for Brick Serving as Application Layer Gateway (ALG) in VoIP Network



A Brick can be deployed to serve as an Application Layer Gateway (ALG) in any of the following locations to secure and monitor VoIP call and data traffic:

- **Headquarters**— the Brick may be deployed to protect the Headquarters LAN at the periphery, including IP phones, call servers, and Media Gateway (MGW) units, or may be used to internally partition the headquarters LAN
- **Regional Office**— A remote location that is connected to the Headquarter site through IP-VPN or a local IPSec VPN where another call server is deployed. A MGW unit may also be present. The Brick deployed at a Regional Office site protects the data traffic, the IP phone call and data traffic, media call control messages to the MGWs, and call and signalling messages between call servers.
- **Branch Office**— A remote location that is connected to the Headquarters site through IP-VPN or a local IPSec VPN where a MGW is present along with IP phone devices. The Brick deployed at a Regional Office site protects the data traffic, the IP phone call and data traffic, media and call control messages to the MGWs.
- **Remote Office**— A remote location that is connected to the Headquarters site through IP-VPN or a local IPSec VPN. Only IP phone devices are being protected by the Brick at a Remote Office site; no call servers or MGW units are present.

Related information

[Chapter 2, “Host Groups”](#) describes the use of host groups in the SMS and the pre-defined host group templates that are provided with the SMS application to make it easier to configure the Brick for VoIP/NOE traffic.

[Chapter 4, “Service Groups”](#) describes the use of service groups in the SMS and the pre-defined service group templates that are provided with the SMS application to make it easier to configure the Brick for VoIP/NOE traffic.

[Chapter 5, “Application Filters”](#) describes the use of application filters in the SMS and the pre-defined application filter templates that are provided with the SMS application to make it easier to configure the Brick for VoIP/NOE traffic.

[Appendix E, “Configuring the Brick for VoIP/NOE Traffic Using Pre-Defined SMS Templates”](#) provides step-by-step instructions and illustrations on how to set up a Brick device to serve as an Application Layer Gateway (ALG) in various VoIP deployment scenarios (Headquarters site, Regional Office site, Branch Office site, Remote Office site, NATed configurations) using the pre-defined Brick zone ruleset, host group, service group, and application filter templates that are provided with the SMS application.

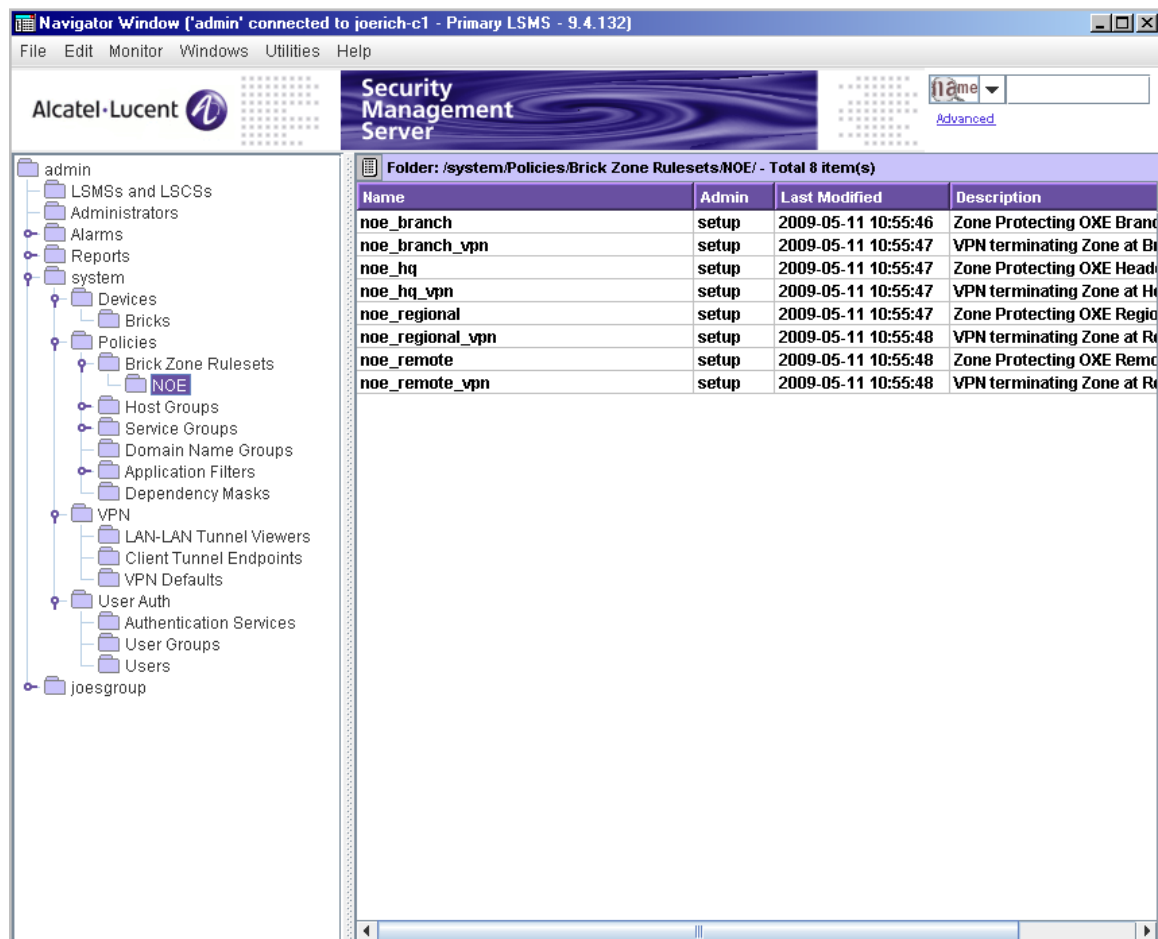
The section *Deployment of a Brick device as an application layer gateway (ALG) for VoIP/NOE phone communications* in the *SMS Administration Guide* provides general information about how a Brick device can be deployed as an ALG to monitor and secure communications between IP phones and other network elements in a VoIP network environment.

Downloading the NOE Brick zone ruleset templates

After the SMS application is installed, the NOE templates can be downloaded by right-clicking on a group folder (such as the *system* folder initially created when the SMS is installed) and choosing **Create NOE Template** from the pop-up menu.

When the NOE template set is created, the pre-defined zone ruleset templates are automatically downloaded to a subfolder named **NOE** under the **Brick Zone Rulesets** folder in the SMS Navigator (refer to [Figure 1-26, “Pre-Defined NOE Brick Zone Rulesets \(NOE Subfolder in SMS Navigator\)”](#) (p. 1-69)).

Figure 1-26 Pre-Defined NOE Brick Zone Rulesets (NOE Subfolder in SMS Navigator)



To remove the NOE zone ruleset templates and all component templates from the Policies folder, right-click on the group folder and choose **Delete NOE Template**. A confirmation dialog box is displayed, asking if you are sure that you want to delete the NOE policy template (which includes all of the zone ruleset templates are component templates) under the selected group folder. Choose **Yes**. The NOE templates are removed from the Policies folder under the selected group. If, however, one of the zone ruleset templates is in use on one or more Bricks, an error message is generated and the templates are not removed.

The following sections show the pre-defined Brick zone rulesets provided with the SMS application for VoIP/NOE traffic.

noe_hq

Defaults are used in this ruleset except where shown (specifically, the Detailed Audit and Sessions Timeout settings).

Brick Zone Ruleset Editor - /system/Policies/Brick Zone Rulesets/NOE/noe_hq

File Edit Monitor Windows Utilities PolicyUtilities Help

Name Description Show Traffic Matcher

Rule Number	Desination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action
1000		Yes	← Out	bricks	LSMS	brick_to_SMS...	*	*	Pass
1001		Yes	→ In	LSMS	bricks	brick_from_SM...	*	*	Pass
1002		Yes	← Out	*	LSMS	tcp/*910	*	*	Pass
1003		Yes	↔ Both	LSMS	LSMS	SMS_Redundancy	*	*	Pass
1004		Yes	→ In	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1005		Yes	→ In	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1006		Yes	→ In	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1007		Yes	← Out	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1008		Yes	← Out	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1009		Yes	← Out	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1010		Yes	← Out	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1011		Yes	← Out	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1012		Yes	← Out	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1013		Yes	→ In	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1014		Yes	→ In	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1015		Yes	→ In	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1016		Yes	← Out	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1017		Yes	← Out	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1018		Yes	← Out	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1019		Yes	→ In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1020		Yes	→ In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass

Hide System Rules

Rule Number	Destination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action
1017		Yes	Out	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1018		Yes	Out	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS	*	*	Pass
1019		Yes	In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1020		Yes	In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1021		Yes	In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS	*	*	Pass
1022		Yes	In	NOE_PCS	NOE_Call_Serv...	NOE_PCS_to_C...	*	*	Pass
1023		Yes	In	NOE_PCS	NOE_Call_Serv...	NOE_PCS_to_CS	*	*	Pass
1024		Yes	Out	NOE_Call_Serv...	NOE_PCS	NOE_CS_to_PC...	*	*	Pass
1025		Yes	Out	NOE_Call_Serv...	NOE_PCS	NOE_CS_to_PC...	*	*	Pass
1026		Yes	Out	NOE_Call_Serv...	NOE_PCS	NOE_CS_to_PCS	*	*	Pass
1027		Yes	Out	NOE_PHONES_...	NOE_TFTP_Ser...	NOE_inside_HQ	*	*	Pass
1028		Yes	In	NOE_PHONES_...	NOE_TFTP_Ser...	NOE_outside_HQ	*	*	Pass
1029		Yes	Out	NOE_PHONES_...	NOE_Presentati...	NOE_http_servi...	*	*	Pass
1030		Yes	In	*	NOE_TFTP_Ser...	sip	*	*	Pass
1031		Yes	Out	*	NOE_TFTP_Ser...	sip	*	*	Pass
1032		Yes	Out	NOE_Call_Serv...	NOE_Networke...	sip	*	*	Pass
1033		Yes	In	NOE_Networke...	NOE_Call_Serv...	sip	*	*	Pass
1034		Yes	In	NOE_Networke...	NOE_Call_Serv...	NOE_ABC_Signa...	*	*	Pass
1035		Yes	Out	NOE_Call_Serv...	NOE_Networke...	NOE_ABC_Signa...	*	*	Pass
1036		Yes	In	NOE_Networke...	NOE_MGWs_He...	H323_App	*	*	Pass
1037		Yes	Out	NOE_MGWs_He...	NOE_Networke...	H323_App	*	*	Pass
65535		Yes	Both	*	*	*	*	*	Drop

noe_hq_vpn

Defaults are used in this ruleset except where shown (specifically, the Detailed Audit and Sessions Timeout settings).

Brick Zone Ruleset Editor - /system/Polices/Brick Zone Rulesets/NOE/noe_hq_vpn

Name: noe_hq_vpn Description: VPN terminating Zone at Headquarters Site Show Traffic Matcher

Rule Number	Desination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action
190		Yes	In	*	Virtual Brick Ad...	UDP_Encapsulat...	*	*	Pass
191		Yes	Out	Virtual Brick Ad...	*	UDP_Encapsulat...	*	*	Pass
192		Yes	In	Virtual Brick Ad...	*	UDP_Encapsulat...	*	*	Pass
193		Yes	Out	*	Virtual Brick Ad...	UDP_Encapsulat...	*	*	Pass
305		Yes	In	Active_VPN_Us...	Virtual Brick Ad...	6/443/*	*	*	VPN Pro
306		Yes	Out	Active_VPN_Us...	Virtual Brick Ad...	6/443/*	*	*	VPN Pro
310		Yes	Both	*	Virtual Brick Ad...	6/443/*	*	*	Proxy
410		Yes	Out	*	Virtual Brick Ad...	udp/500/*	*	*	Pass
411		Yes	In	Virtual Brick Ad...	*	udp/500/500	*	*	Pass
412		Yes	In	Virtual Brick Ad...	Active_VPN_Us...	udp/*/500	*	*	Pass
420		Yes	In	*	Virtual Brick Ad...	udp/500/*	*	*	Pass
430		Yes	In	*	Virtual Brick Ad...	ipsec	*	*	VPN
440		Yes	Out	*	Virtual Brick Ad...	ipsec	*	*	VPN
1000		Yes	Out	bricks	LSMS	brick_to_SMS_...	*	*	Pass
1001		Yes	In	LSMS	bricks	brick_from_SM...	*	*	Pass
1002		Yes	Out	*	LSMS	tcp/*/910	*	*	Pass
1003		Yes	Both	LSMS	LSMS	SMS_Redundancy	*	*	Pass
1004		Yes	Out	NOE_VPN_Outs...	NOE_VPN_Outs...	*	*	*	VPN
1005		Yes	In	*	*	*	*	*	VPN
1006		Yes	Out	*	*	*	*	*	VPN
65535		Yes	Both	*	*	*	*	*	Drop

Hide System Rules

noe_regional

Defaults are used in this ruleset except where shown (specifically, the Detailed Audit and Sessions Timeout settings).

Brick Zone Ruleset Editor - /system/Policies/Brick Zone Rulesets/NOE/noe_regional

File Edit Monitor Windows Utilities Policy Utilities Help

Name: **noe_regional** Description: **Zone Protecting OXE Regional Office Site** Show Traffic Matcher

Rule Number	Destination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action
1000		Yes	In	bricks	LSMS	brick_to_SMS_...	*	*	Pass
1001		Yes	Out	LSMS	bricks	brick_from_SM...	*	*	Pass
1002		Yes	In	*	LSMS	tcp/*910	*	*	Pass
1003		Yes	Both	LSMS	LSMS	SMS_Redundancy	*	*	Pass
1004		Yes	In	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1005		Yes	In	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1006		Yes	In	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1007		Yes	Out	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1008		Yes	Out	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1009		Yes	Out	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1010		Yes	In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1011		Yes	In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1012		Yes	In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1013		Yes	Out	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1014		Yes	Out	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1015		Yes	Out	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1016		Yes	In	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1017		Yes	In	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1018		Yes	Out	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1019		Yes	Out	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass

Hide System Rules

The screenshot shows the 'Brick Zone Ruleset Editor' window for the policy 'noe_regional'. The description is 'Zone Protecting OXE Regional Office Site'. Below the header, there is a table of 20 rules (1011-1029 and 65535) with columns for Rule Number, Desination TCP Port, Active, Direction, Source, Destination, Service, VLAN... Match, Brick Match, and Action. Most rules are 'Pass' actions, while rule 65535 is a 'Drop' action.

Rule Number	Desination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action
1011		Yes	In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS...	*	*	Pass
1012		Yes	In	NOE_Call_Serv...	NOE_TFTP_Ser...	NOE_CS_to_CS	*	*	Pass
1013		Yes	Out	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1014		Yes	Out	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1015		Yes	Out	NOE_TFTP_Ser...	NOE_Call_Serv...	NOE_CS_to_CS	*	*	Pass
1016		Yes	In	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1017		Yes	In	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1018		Yes	Out	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1019		Yes	Out	NOE_Call_Serv...	NOE_Call_Serv...	NOE_CS_to_CS...	*	*	Pass
1020		Yes	In	NOE_PCS	NOE_Call_Serv...	NOE_PCS_to_C...	*	*	Pass
1021		Yes	In	NOE_PCS	NOE_Call_Serv...	NOE_PCS_to_CS	*	*	Pass
1022		Yes	Out	NOE_Call_Serv...	NOE_PCS	NOE_CS_to_PC...	*	*	Pass
1023		Yes	Out	NOE_Call_Serv...	NOE_PCS	NOE_CS_to_PC...	*	*	Pass
1024		Yes	Out	NOE_Call_Serv...	NOE_PCS	NOE_CS_to_PCS	*	*	Pass
1025		Yes	Out	NOE_PHONES_...	NOE_TFTP_Ser...	NOE_inside_Reg...	*	*	Pass
1026		Yes	In	NOE_PHONES_...	NOE_TFTP_Ser...	NOE_outside_R...	*	*	Pass
1027		Yes	Out	NOE_PHONES_...	NOE_Presentati...	NOE_http_service	*	*	Pass
1028		Yes	In	*	NOE_TFTP_Ser...	sip	*	*	Pass
1029		Yes	Out	*	NOE_TFTP_Ser...	sip	*	*	Pass
65535		Yes	Both	*	*	*	*	*	Drop

noe_regional_vpn

Defaults are used in this ruleset except as shown (specifically, the type of VPN as internal or both, and Detailed Audit setting).

Brick Zone Ruleset Editor - /system/Polices/Brick Zone Rulesets/NOE/noe_regional_vpn

Name: noe_regional_vpn Description: VPN terminating Zone at Regional Office Site Show Traffic Matcher

Rule Number	Destination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action
190		Yes	In	*	Virtual Brick Ad...	UDP_Encapsulat...	*	*	Pass
191		Yes	Out	Virtual Brick Ad...	*	UDP_Encapsulat...	*	*	Pass
192		Yes	In	Virtual Brick Ad...	*	UDP_Encapsulat...	*	*	Pass
193		Yes	Out	*	Virtual Brick Ad...	UDP_Encapsulat...	*	*	Pass
305		Yes	In	Active_VPN_Us...	Virtual Brick Ad...	6/443/*	*	*	VPN Proxy
306		Yes	Out	Active_VPN_Us...	Virtual Brick Ad...	6/443/*	*	*	VPN Proxy
310		Yes	Both	*	Virtual Brick Ad...	6/443/*	*	*	Proxy
410		Yes	Out	*	Virtual Brick Ad...	udp/500/*	*	*	Pass
411		Yes	In	Virtual Brick Ad...	*	udp/500/500	*	*	Pass
412		Yes	In	Virtual Brick Ad...	Active_VPN_Us...	udp/*/500	*	*	Pass
420		Yes	In	*	Virtual Brick Ad...	udp/500/*	*	*	Pass
430		Yes	In	*	Virtual Brick Ad...	ipsec	*	*	VPN
440		Yes	Out	*	Virtual Brick Ad...	ipsec	*	*	VPN
1000		Yes	Out	bricks	LSMS	brick_to_SMS_...	*	*	Pass
1001		Yes	In	LSMS	bricks	brick_from_SM...	*	*	Pass
1002		Yes	Out	*	LSMS	tcp/*/910	*	*	Pass
1003		Yes	Both	LSMS	LSMS	SMS_Redundancy	*	*	Pass
1004		Yes	In	*	*	*	*	*	VPN
1005		Yes	Out	*	*	*	*	*	VPN
65535		Yes	Both	*	*	*	*	*	Drop

Hide System Rules

noe_branch

Defaults are used in this ruleset except as shown (specifically, the Detailed Audit and Sessions Timeout settings).

Brick Zone Ruleset Editor - /system/Polices/Brick Zone Rulesets/NOE/noe_branch

Name: **noe_branch** Description: **Zone Protecting OXE Branch Office Site** Show Traffic Matcher

Rule Number	Desination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action
1000		Yes	Out	NOE_PHONES_...	NOE_TFTP_Ser...	NOE_inside_Bra...	*	*	Pass
1001		Yes	Out	*	NOE_TFTP_Ser...	sip	*	*	Pass
1002		Yes	Out	NOE_PCS	NOE_Call_Serv...	NOE_PCS_to_C...	*	*	Pass
1003		Yes	Out	NOE_PCS	NOE_Call_Serv...	NOE_PCS_to_CS	*	*	Pass
1004		Yes	In	NOE_Call_Serv...	NOE_PCS	NOE_CS_to_PC...	*	*	Pass
1005		Yes	In	NOE_Call_Serv...	NOE_PCS	NOE_CS_to_PC...	*	*	Pass
1006		Yes	In	NOE_Call_Serv...	NOE_PCS	NOE_CS_to_PCS	*	*	Pass
1007		Yes	In	*	NOE_PCS	NOE_PCS	*	*	Pass
1008		Yes	Out	NOE_PHONES_...	NOE_Presentati...	NOE_http_service	*	*	Pass
1009		Yes	Out	NOE_PHONES_...	NOE_TFTP_Ser...	NOE_Service-NA...	*	*	Pass
1010		Yes	Out	*	NOE_TFTP_Ser...	sip	*	*	Pass
1011		Yes	Out	NOE_PHONES_...	NOE_Presentati...	NOE_http_service	*	*	Pass
65535		Yes	Both	*	*	*	*	*	Drop

Hide System Rules

noe_branch_vpn

Defaults are used in this ruleset except as shown (specifically, the type of VPN as internal or both, and Sessions Timeout setting).

Brick Zone Ruleset Editor - /system/Polices/Brick Zone Rulesets/NOE/noe_branch_vpn

Name: **noe_branch_vpn** Description: **VPN terminating Zone at Branch Office Site** Show Traffic Matcher

Rule Number	Desination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action	Drop Action
190		Yes	In	*	Virtual Brick Ad...	UDP_Encapsulat...	*	*	Pass	
191		Yes	Out	Virtual Brick Ad...	*	UDP_Encapsulat...	*	*	Pass	
192		Yes	In	Virtual Brick Ad...	*	UDP_Encapsulat...	*	*	Pass	
193		Yes	Out	*	Virtual Brick Ad...	UDP_Encapsulat...	*	*	Pass	
305		Yes	In	Active_VPN_Us...	Virtual Brick Ad...	6443/*	*	*	VPN Proxy	
306		Yes	Out	Active_VPN_Us...	Virtual Brick Ad...	6443/*	*	*	VPN Proxy	
310		Yes	Both	*	Virtual Brick Ad...	6443/*	*	*	Proxy	
410		Yes	Out	*	Virtual Brick Ad...	udp/500/*	*	*	Pass	
411		Yes	In	Virtual Brick Ad...	*	udp/500/500	*	*	Pass	
412		Yes	In	Virtual Brick Ad...	Active_VPN_Us...	udp/*/500	*	*	Pass	
420		Yes	In	*	Virtual Brick Ad...	udp/500/*	*	*	Pass	
430		Yes	In	*	Virtual Brick Ad...	ipsec	*	*	VPN	
440		Yes	Out	*	Virtual Brick Ad...	ipsec	*	*	VPN	
1000		Yes	Out	bricks	LSMS	brick_to_SMS_...	*	*	Pass	
1001		Yes	In	LSMS	bricks	brick_from_SM...	*	*	Pass	
1002		Yes	Out	*	LSMS	tcp/*/910	*	*	Pass	
1003		Yes	In	*	*	*	*	*	VPN	
1004		Yes	Out	*	*	*	*	*	VPN	
65535		Yes	Both	*	*	*	*	*	Drop	None

Hide System Rules

noe_remote

Defaults are used in this ruleset except as shown (specifically, the Detailed Audit and Sessions Timeout settings).

Rule Number	Destination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action
1000		Yes	Out	NOE_PHONES_...	NOE_TFTP_Ser...	NOE_inside_Re...	*	*	Pass
1001		Yes	Out	*	NOE_TFTP_Ser...	sip	*	*	Pass
1002		Yes	Out	NOE_PHONES_...	NOE_Presentati...	NOE_http_service	*	*	Pass
1003		Yes	Out	NOE_PHONES_...	NOE_TFTP_Ser...	NOE_Service-NA...	*	*	Pass
1004		Yes	Out	*	NOE_TFTP_Ser...	sip	*	*	Pass
1005		Yes	Out	NOE_PHONES_...	NOE_Presentati...	NOE_http_service	*	*	Pass
65535		Yes	Both	*	*	*	*	*	Drop

noe_remote_vpn

Defaults are used in this ruleset except as shown (specifically, the type of VPN as internal or both and the Detailed Audit setting).

Rule Number	Destination TCP Port	Active	Direction	Source	Destination	Service	VLAN... Match	Brick Match	Action	Drop Action	Audit
190		Yes	In	*	Virtual Brick Ad...	UDP_Encapsulat...	*	*	Pass		Basic
191		Yes	Out	Virtual Brick Ad...	*	UDP_Encapsulat...	*	*	Pass		Basic
192		Yes	In	Virtual Brick Ad...	*	UDP_Encapsulat...	*	*	Pass		Basic
193		Yes	Out	*	Virtual Brick Ad...	UDP_Encapsulat...	*	*	Pass		Basic
305		Yes	In	Active_VPN_Us...	Virtual Brick Ad...	6/443/*	*	*	VPN Proxy		Basic
306		Yes	Out	Active_VPN_Us...	Virtual Brick Ad...	6/443/*	*	*	VPN Proxy		Basic
310		Yes	Both	*	Virtual Brick Ad...	6/443/*	*	*	Proxy		Basic
410		Yes	Out	*	Virtual Brick Ad...	udp/500/*	*	*	Pass		Basic
411		Yes	In	Virtual Brick Ad...	*	udp/500/500	*	*	Pass		Basic
412		Yes	In	Virtual Brick Ad...	Active_VPN_Us...	udp/*/500	*	*	Pass		Basic
420		Yes	In	*	Virtual Brick Ad...	udp/500/*	*	*	Pass		Basic
430		Yes	In	*	Virtual Brick Ad...	ipsec	*	*	VPN		Basic
440		Yes	Out	*	Virtual Brick Ad...	ipsec	*	*	VPN		Basic

2 Host Groups

Overview

Purpose

This chapter explains how to set up, use, and maintain host groups. It also describes the three host groups that are provided with the SMS application.

Contents

What is a Host Group?	2-2
To Set Up a Host Group	2-4
Global Host Groups	2-8
Dynamic Host Groups	2-10
Nested Host Groups	2-12
Maintaining Host Groups	2-15
Host Groups Provided with the SMS Application	2-22
Host Group Templates Provided with the SMS Application for VoIP/NOE Traffic	2-25



What is a Host Group?

Definition

A host group is a collection of IP addresses. It can consist of one IP address, multiple IP addresses, or one or more ranges of IP addresses. It can also contain the wildcard asterisk (*), which means the host group includes every IP address connected to the port to which it is applied.

There are two types of host groups — standard host groups, which can only be used in the group in which they were created, and global host groups, which can be used in every group.

Uses

The primary purpose of a host group is to enable you to enter more than one IP address when a source or destination address is required, or a list of hosts is needed. There are a number of instances where host groups can be used:

- To identify the source and destination hosts in a security rule in a Alcatel-Lucent *VPN Firewall Brick™* Security Appliance zone ruleset.
- To identify the local and remote hosts in a main or router tunnel ruleset.
- To specify the mapped source or destination address when performing network address translation.
- To enter the source and destination addresses in a dependency mask.
- To identify the hosts that Alcatel-Lucent IPsec Client users will be permitted to access through a client tunnel (also known as the Host Access List).
- To specify the hosts that will be used to provide Alcatel-Lucent IPsec Client users with a presence on the local LAN.
- To enter the hosts that users of a LAN-LAN tunnel will be able to access behind the other tunnel endpoint.

Benefits

In addition to allowing you to enter multiple IP addresses in a field, host groups greatly simplify the process of creating and managing Brick zone rulesets. For example, host groups:

- Help minimize the number of rules needed. When a new host is added to a zone, you can add a new address to the appropriate host group instead of creating a new rule.
- Allow you to intuitively identify the hosts impacted by a rule. A name such as "marketing_servers" in the **Source** or **Destination** field of a rule is much easier to identify than a string of IP addresses.
- Enable you to enter the same set of addresses or address ranges, in multiple rules, quickly and consistently.



To Set Up a Host Group

Overview

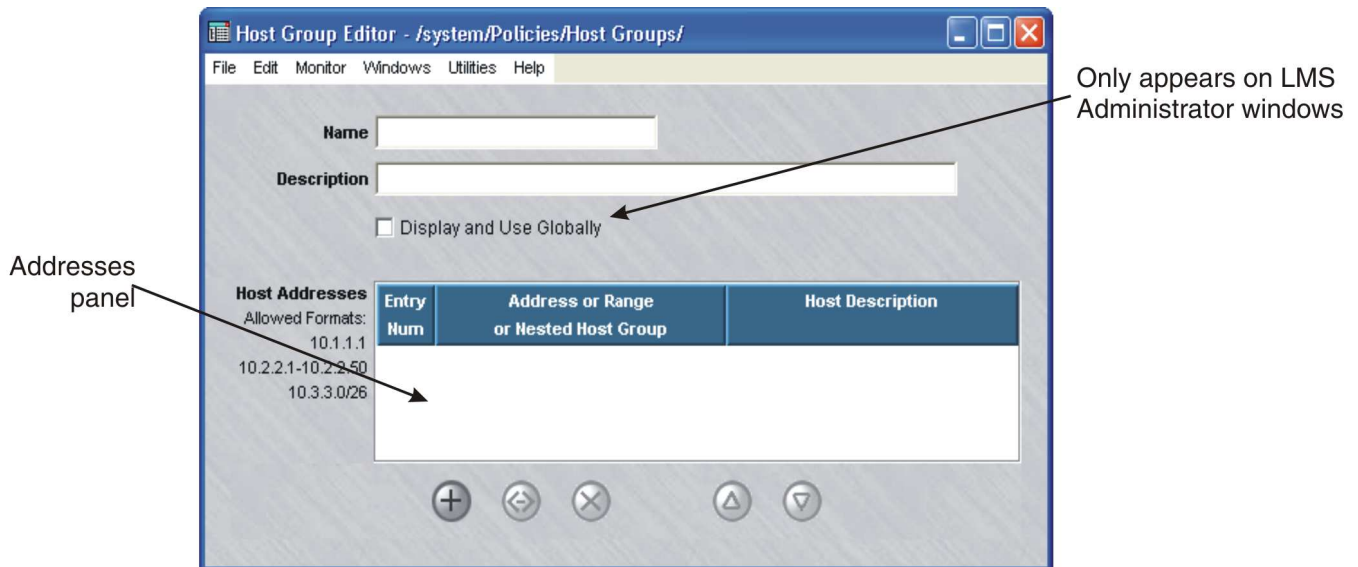
To set up a host group, you have to display the Host Groups Editor and enter the name of the host group and the IP address(es) to be included in the host group.

Display the Host Groups Editor

The easiest way to display the Host Groups Editor is to:

- 1 Open the folder of the group that will contain this host group.
- 2 Open the Policies folder.
- 3 Right-click the Host Groups folder and select **New Host Group** from the pop-up menu. The Host Group Editor is displayed (Figure 2-1, “Host Group Editor” (p. 2-4)).

Figure 2-1 Host Group Editor



END OF STEPS

To create a host group

Complete the following steps to create a host group.

- 1 In the **Name** field, enter a unique name to identify this host group. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers and certain special characters.
- 2 In the **Description** field, you can enter an optional description of the host group. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers and certain special characters.

This is the description that will appear on the Navigator window when you view existing host groups. See the section below entitled [“To view host groups”](#) (p. 2-15).

Important! If you are an SMS Administrator, you will see a checkbox entitled “Display and Use Globally” under the **Description** field. Click this checkbox if you want to make this a global host group. See [“Global Host Groups”](#) (p. 2-8) below for an explanation of global host groups.

- 3 Right-click in the addresses panel and select **New** from the pop-up menu (or click the **New** button) to display the Host Group Entry window, which is shown in [Figure 2-2, “Host Group Entry Window”](#) (p. 2-5).

Figure 2-2 Host Group Entry Window



Important! If you intend to incorporate another host group within this one, click the **Select Nested Host** checkbox. For an explanation of nested host groups, see [“Nested Host Groups” \(p. 2-12\)](#) below.

- 4 In the Host Group Entry window, do the following:
 - In the **Address of Range** field, enter a single IP address, multiple IP addresses, or a range of IP addresses. To enter multiple IP addresses, separate the addresses by a comma. To enter an address range, enter a beginning address, a dash, and an ending address.
 - In the **Host Description** field, enter a brief description of the address(es). This field is optional.
 - Click **OK** to dismiss the window. You will be returned to the Host Group Editor. The entry you just created will appear in the address panel.
 - An **Entry Num** value will be automatically assigned to this entry. This number is used to better manage potentially large numbers of entries in a host group and to stress and make explicit the fact that host group entries are ordered, which is important for DIRECT NAT. The feature will allow an administrator to sort entries by **Entry Num**, **Address**, or **Description**. Double-clicking on the column header will reverse the sort order. Sorting changes the order of presentation, but does not change the **Entry Num** order. The **Entry Num** can only be changed by right-clicking on the entry and using the **Up** and **Down** functions.

 - 5 To add more addresses to this host group, repeat [Step 3](#) and [Step 4](#) until all addresses have been entered. [Figure 2-3, “IP Addresses in a Host Group” \(p. 2-7\)](#) shows a host group with two address range entries.
-

Figure 2-3 IP Addresses in a Host Group

Host Addresses	Entry Num	Address or Range or Nested Host Group	Host Description
Allowed Formats: 10.1.1.1 10.2.2.1-10.2.2.50 10.3.3.0/26	1	10.1.1.1;10.1.1.10-10.1.1.19;10.2....	sales department
	2	10.1.1.2	finance department

- 6 Display the File menu and select one of the **Save** options.

Important! If the IP addresses are in another program, such as Microsoft Word or Excel, you can copy the addresses into the clipboard, and then paste them into the **Host Addresses** field using **[Ctrl-V]**.

END OF STEPS



Global Host Groups

Definition

A global host group is a host group that is created in one SMS group, but can be seen and used in every other group. Only SMS Administrators can create global host groups. Global host groups are dynamically linked, and will update everywhere when modified.

To create a global host group

If you are an SMS Administrator, you create a global host group by clicking the **Display and Use Globally** checkbox on the Host Group Editor (see [Figure 2-1, “Host Group Editor”](#) (p. 2-4)). You can do this when you create the host group, or you can do this after the host group has been created by editing the host group (see [“To modify a host group”](#) (p. 2-18) below).

When creating a global host group, make sure the name you give the host group is unique across all groups. If you attempt to give a global host group a name that is in use elsewhere, the host group will not be created, and you will get an error message indicating the save failed because there is an object in another group with the same name.

To view global host groups

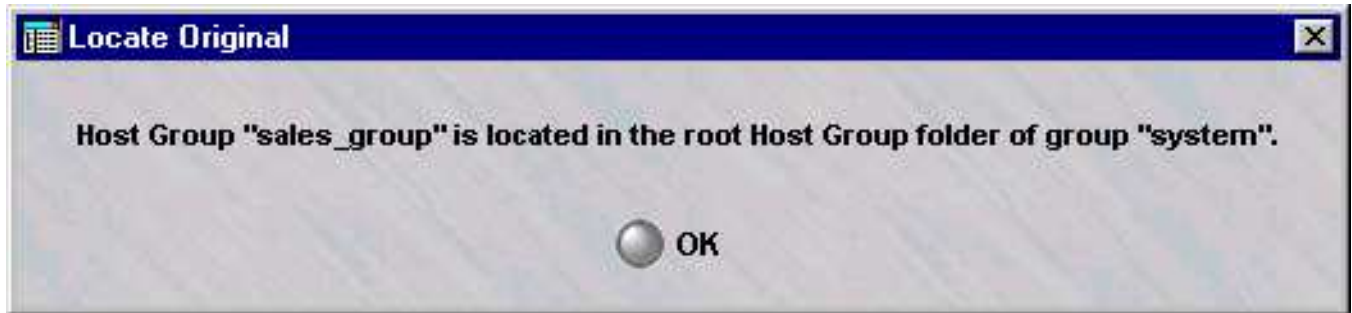
Global host groups appear in the Navigator window, just like standard host groups (see [“To view host groups”](#) (p. 2-15) below). When you view the host groups that have been created in a specific group, all the host groups *created in that group* appear the same, regardless of whether or not they are global.

However, any global host groups created in *a different group* can be identified by a globe icon that appears to the left of the entry, as shown below.

 sales_group	stan	2002-05-27 10:25:18	sales department host group
---	------	---------------------	-----------------------------

To determine the group in which the global host group was originally created, right-click the host group in the Navigator window and select **Original Location** from the pop-up menu. A window similar to the one shown in [Figure 2-4, “Locate Original Window”](#) (p. 2-9) will appear.

Figure 2-4 Locate Original Window



To remove the global status of a host group

Just as it is possible to make a non-global host group global by clicking the **Display and Use Globally** checkbox on the Host Group Editor (see [Figure 2-1, “Host Group Editor”](#) (p. 2-4)), it is possible to remove the global status of a host group by unchecking this checkbox. However, this can only be done if the host group is *not* in use globally. If the host group is in use in any group *other than the one in which it was created*, you cannot remove its global status.

It is also possible to delete a global host group — but only the original source of the host group. You cannot delete a global host group from any folder in which it appears except the folder in which it was originally created. No host group — global or standard — can be deleted if it is in use anywhere (see [“To delete a host group”](#) (p. 2-21) below).

Similarly, only the original source of a global host group can be moved (see [“To move a host group”](#) (p. 2-20) below)

Permissions

Permissions over global host groups are based on the group in which the host group was originally created. If an administrator has FULL policy permissions for that group, then that administrator has FULL permissions over all global host groups created in that group.

If an administrator has FULL permission over a global host group, the administrator can edit the host group in any of the groups in which it appears. An administrator can create a copy of a global host group as long as the administrator has FULL permissions over the *destination* group.

□

Dynamic Host Groups

Definition

The SMS provides dynamic host groups that obtain IP addresses dynamically from the DNS server for DHCP, PPPoE#1, or PPPoE#2 sessions on a Brick. These dynamic host groups are called:

- **dynamicDhcpDnsServers**
- **dynamicPppoe1DnsServers**
- **dynamicPppoe2DnsServers**

The availability of these dynamic host groups is not readily apparent when upgrading an SMS to the Release 9.4 software patch.

Task

To copy these host groups into the **system** Host Groups folder, follow the steps below:

- 1 Right-click on the **system** folder and select **New Group**.

The Group Editor window is displayed.

- 2 Create a "placeholder" group by entering a Group Name and Description in the respective fields.
-

- 3 Select **Save and Close** from the File menu.

The "placeholder" group is created.

- 4 Expand the new group in the explorer and click on the **Policies** folder to display the **Host Groups** folder.
-

- 5 Click on the **Host Groups** folder.

The host groups, including the dynamic host groups, are displayed in the Contents Panel.

-
- 6 Move or copy the dynamic host groups to the **Host Groups** folder under the **system** folder.

END OF STEPS



Nested Host Groups

Definition

A nested host group is a host group that has other host groups nested inside it. Nesting host groups allows you to create host groups composed of one or more other host groups. Including host groups within another host group is generally easier, and less prone to error, than entering the IP addresses of the internal host groups manually.

For example, suppose the administrator of a bank with several branch offices creates individual host groups for each branch office. The administrator could then create one large host group for the entire bank, and nest each of the branch office host groups within the overall bank host group. This is much easier than entering the IP addresses of each branch one at a time.


You can nest global host groups within other global host groups, or within other non-global host groups. When a global host group contains non-global host groups nested within it, the host group loaded to the Brick contains only the contents of the host groups visible in the same group as the Brick Zone Ruleset. You can create nested host groups with up to 32 levels of nesting.

Cycles are prohibited in nested host groups. For example, suppose host group A is a nested host group, and contains host group B within it. Further, suppose host group B is also a nested host group, and includes host group C within it. In this instance, host group C would not be permitted to be a nested host group containing host group A.

Task

To nest a host group within another host group, follow the steps below:

- 1 Create a host group as previously explained (see [“To create a host group” \(p. 2-5\)](#) above). Follow the steps until the Host Group Entry window appears ([Figure 2-2, “Host Group Entry Window” \(p. 2-5\)](#)).
- 2 Click the **Select Nested Host** checkbox in the Host Group Entry window (see [Figure 2-2, “Host Group Entry Window” \(p. 2-5\)](#)). The **Address or Range** field will become a **Nested Host Group** field, as shown in [Figure 2-5, “Host Group Entry Window \(Nested Host Groups\)” \(p. 2-13\)](#).

Figure 2-5 Host Group Entry Window (Nested Host Groups)The image shows a dialog box titled "Host Group Entry" with a close button (X) in the top right corner. At the top, there are two radio buttons: "Specify Host" (which is unselected) and "Select Nested Host" (which is selected). Below the radio buttons, there is a label "Nested Host Group" followed by a white rectangular drop-down menu. Underneath that is a label "Host Description" followed by a larger white rectangular text input field. At the bottom of the dialog box, there are two buttons: "OK" and "Cancel", each with a circular icon to its left.

-
- 3 In the **Nested Host Group** field, select a host group that will be nested in this host group from the drop-down list. If a description was entered for that host group, it will appear in the **Host Description** field.

 - 4 To nest additional host groups in this host group, repeat Steps 2 and 3 for each additional host group.

 - 5 When you have finished, click **OK**. [Figure 4-6, “Nested Service Groups” \(p. 4-12\)](#) shows an example of two host groups nested within a another host group.

Figure 2-6 Nested Host Groups

Host Addresses	Entry Num	Address or Range or Nested Host Group	Host Description
Allowed Formats: 10.1.1.1 10.2.2.1-10.2.2.50 10.3.3.0/26	1	atlanta_branch	atlanta sales representatives
	2	boston_branch	

END OF STEPS



Maintaining Host Groups

Overview

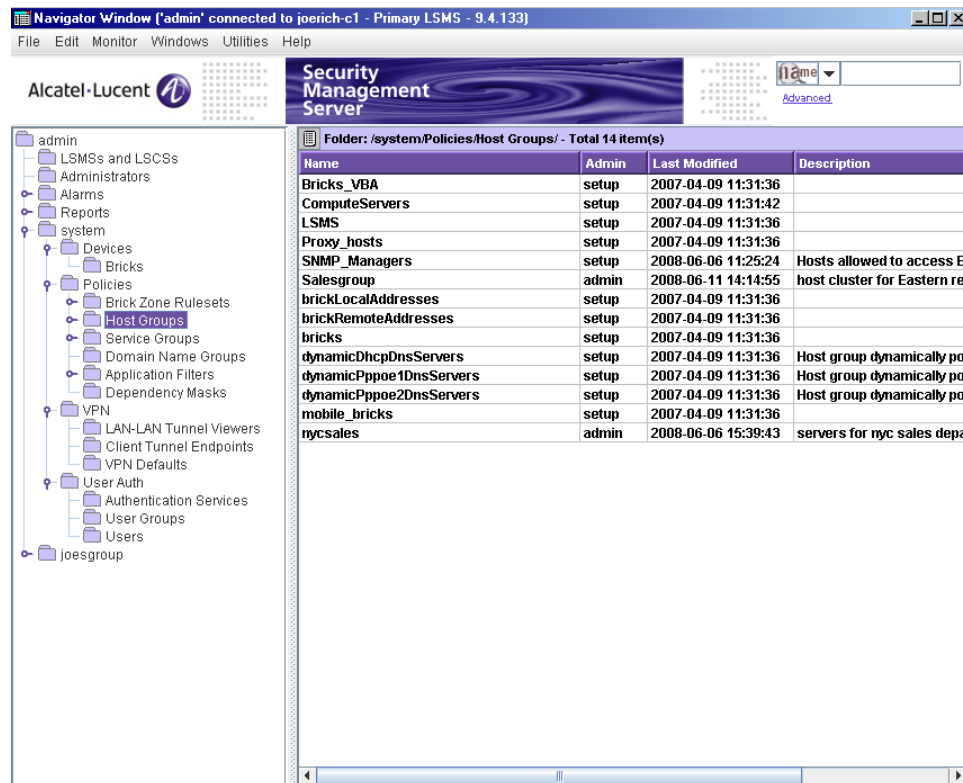
Once a host group has been created and saved, it can be viewed, modified, copied, or moved. If the host group is no longer needed, it can be deleted.

To view host groups

You must view the host groups before you can edit, move, copy, or delete a specific host group. To view all the host groups that have been created to date in a particular group, follow the steps below:

- 1 Open the appropriate group folder, and then open the Policies folder.
- 2 Click the Host Groups folder. All existing host groups will be displayed in the Navigator window, as shown in [Figure 2-7, “Navigator Window \(View Host Groups\)”](#) (p. 2-15).

Figure 2-7 Navigator Window (View Host Groups)



As you can see, for each host group, the Navigator window shows the Administrator who created the host group, the date and time the host group was created, and a brief description, if one was entered when the host group was created.

Important! As [Figure 2-7, “Navigator Window \(View Host Groups\)” \(p. 2-15\)](#) shows, the following host groups are provided with the SMS application:

- **Bricks_VBA**
- **ComputeServers**
- **LSMS**
- **Proxy_hosts**
- **SNMP_Managers**
- **brickLocalAddresses**
- **brickRemoteAddresses**
- **bricks**
- **dynamicDhcpDnsServers**
- **dynamicPppoe1DnsServers**
- **dynamicPppoe2DnsServers**
- **mobile_bricks**

The **bricks** host group only appears in the **system** group.

You can tell these host groups have been provided with the application because the Administrator who created them is *setup*. These host groups are explained in the section below entitled [“Permissions” \(p. 2-9\)](#).

.....
E N D O F S T E P S

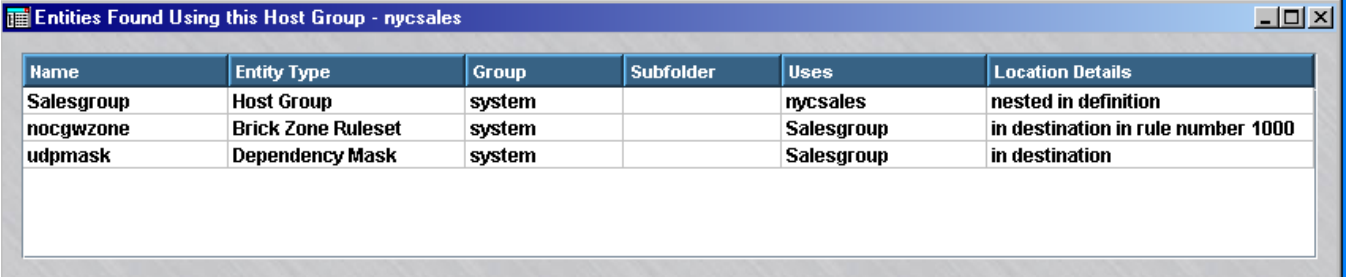
To find entities that use a host group

The SMS application allows you to perform a search for all entities (such as a zone ruleset, application filter, or dependency mask) that are using the selected host group, and, if it is nested in a host group, the parent host group and all entities that are using the parent host group.

- 1 With the host groups displayed in the Navigator window, right-click the desired host group and select **Find Entities Using this Host Group** from the pop-up menu.

Result An Entities Found Using this Host Group window is displayed (Figure 2-8, “Entities Found Using this Host Group Window” (p. 2-17)) showing the results of the search.

Figure 2-8 Entities Found Using this Host Group Window



Name	Entity Type	Group	Subfolder	Uses	Location Details
Salesgroup	Host Group	system		nycsales	nested in definition
nocgwzone	Brick Zone Ruleset	system		Salesgroup	in destination in rule number 1000
udpmask	Dependency Mask	system		Salesgroup	in destination

The display results on the window show the entity type that uses or contains the selected host group, the group and subfolder in which the entity is found, whether the entity uses the host group or the parent host group (if it is a nested host group), and details about the entity (if, for example, the entity is a Brick zone ruleset, the display lists the specific rules in which the host group is used).

To view additional details about the entity, right-click on the entry in the table and select **View** from the pop-up menu.

To edit the entity, right-click on the entry in the table and select **Edit** from the pop-up menu.

-
- 2 Click the Close () button at the top of the window to close the window.

END OF STEPS

To find overlapping IP addresses in a host group

A tool is provided to find IP address overlaps within a given host group. Complete the following steps to find overlapping IP addresses.

-
- 1 Click on the **Host Groups** folder.

Result A list of currently defined host groups is displayed in the Contents Panel.

-
- 2 Right-click on a host group to display a pop-up menu and select **Find Overlaps within this Host Group**.

Result An Overlaps found within Host Group window is displayed. Each case of overlapping IP addresses is shown on a separate row. Each row shows pairs of entry numbers that overlap each other, with the overlapping IP address or range of addresses. A host group is included in the list if the overlapping IP address exists in a nested host group.

.....
E N D O F S T E P S
.....

To modify a host group

You can modify any field in a host group. To modify a host group, follow the steps below:

-
- 1 With the host groups displayed in the Navigator window, double-click the host group you want to modify.

The Host Groups Editor will appear, as shown in [Figure 2-1, “Host Group Editor” \(p. 2-4\)](#).

-
- 2 Change any information in the **Description** and **Host Addresses** fields and the **Display and Use Globally** checkbox.

-
- 3 Display the File menu and select one of the **Save** options.

.....
E N D O F S T E P S
.....

To copy a host group

You can copy a host group to a different folder in the same group, or to a folder in another group, provided you have *Full* permission for that group. When you edit a host group (using the **Save As** option to save the edited version with a new name) or copy a host group, you have the option of copying components from the original host group (nested host groups, IP addresses), specifying how those components will be named and where they will be created and stored.

Complete the following steps to copy a host group:

-
- 1 With the host groups displayed in the Navigator window, right-click the host group to be copied and select **Copy** from the pop-up menu.

Result A Copy window is displayed (Figure 2-9, “Copy Window” (p. 2-19)).

Figure 2-9 Copy Window



2 Complete the following fields, as applicable:

- **Name**— Enter a new name for the host group copy (Example: **Salesgroup_2**). This field is required. The name entered must be different from the name of the original host group by at least one character if it is in the same group.
- **Folder**— The copy of the host group will be created and stored in the **Host Groups** folder under the *system* group, by default. To specify a different folder, click the **Browse** button, which brings up a Select a Host Group Folder window, and select the folder where the host group copy will be created and stored.
- **Copy components when copying within a Group**— Click this checkbox (place a check in it) to copy the components from the original host group. If this checkbox is unchecked, and the copied host group is in the same group as the original one, the copied host group uses the same components as the original. If the host group is copied across groups, the components are copied regardless of this field setting.
- **Component Suffix**— If the **Copy components when copying within a Group** option is checked, enter a text string that will be appended to the name of each copied component (nested host group). If this field is left blank, the name of each copied component will be spelled the same way as the original one. Name collisions are possible, regardless of whether or not an entry is made in this field. If this occurs, the content of the conflicting component(s) is examined. If the content is the same, the duplicated host group uses the existing component(s) and

the new copy(ies) is not made. If the content is different, a copy of the component(s) is created and stored under a unique name by appending its database ID to the name.

Note: the Component Suffix used must be unique and cannot be reused across zone rulesets.

- **Component Folder**— This field is used to specify a common folder name for the copied components. The folder name is relative to the top-level folder of each type of component. If this field is left blank, the copied components are stored in the top-level folder of that component. This folder is created as needed (if it does not already exist). A path string can be specified; / or \ delimiters can be used to delineate folder names.

-
- 3 Click **OK** to copy the host group and dismiss the Copy window.

Result The system returns to the Navigator window. The copy of the host group appears in the Contents panel.

The host group (and components, if specified) is copied to the chosen folder(s).

END OF STEPS

To move a host group

You can move a host group to a different folder in the same group, or to a folder in another group, provided you have *Full* permission for that folder. However, you can only move it to the folder labeled "Host Groups" or to a subfolder under that folder. If you try to move it to another folder, you will get an error message. Before moving a host group, make sure that it is not currently in use. If you attempt to move a host group that is in use, you will get an error message. If you receive an error message when attempting to move a host group, you can run the "find entity" utility to discover where it is in use (see "[To find entities that use a host group](#)" (p. 2-16) above). If a host group is moved to a new group, and the host group has other host groups nested within it, the original host group is moved to the new group, but the host groups nested within it are copied to the new group, not moved. If a host group with the same name and content already exists in the new group, the copy does not take place. If a host group with the same name but different content already exists, the copy is renamed.

To move a host group, follow the steps below:

- 1 With the host groups displayed in the Navigator window, right-click the host group you want to move and select **Move** from the pop-up menu. A Browse window will appear.

-
- 2 Select the folder you want to move this host group to. The host group will be moved to the folder you selected, and you will be returned to the Navigator window.

END OF STEPS

To delete a host group

Before deleting a host group, make sure that it is not currently in use. If you attempt to delete a host group that is in use, you will get an error message. If you receive an error message when attempting to delete a host group, you can run the "find entity" utility to discover where it is in use (see ["To find entities that use a host group" \(p. 2-16\)](#) above).

Complete the following steps to delete a host group.

-
- 1 With the host group displayed in the Navigator window, right-click the host group you want to delete and select **Delete** from the pop-up menu.

A confirmation window similar to the one shown in [Figure 2-10, "Confirmation Window \(Host Groups\)" \(p. 2-21\)](#) will appear.

Figure 2-10 Confirmation Window (Host Groups)



-
- 2 Click **Yes** to delete the host group and dismiss the Confirmation window. The host group will be removed from the Navigator window.

END OF STEPS

□

Host Groups Provided with the SMS Application

Overview

Pre-defined host groups are automatically created when the SMS application is installed. The following explains the purpose of each of these host groups.

Bricks

The *Bricks* host group only appears in the **system** group. It contains the IP address of every Brick that has been configured in every group. Every time an Administrator configures a Brick, its IP address is automatically added to this host group. If the Brick is deleted, the IP address is automatically removed from the host group.

The *Bricks* host group is used in three of the preconfigured Brick zone rulesets that are provided with the SMS application. The following explains:

- *administrativezone*
The purpose of the *administrativezone* ruleset is to protect an SMS that is managing Bricks, but not routers (or is managing routers in the clear). It is assigned to the port connecting the SMS to the Brick, and it only allows the SMS to communicate with that Brick and any other Bricks that have been configured, unless, of course, an Administrator adds new rules to the zone's default ruleset. The *Bricks* host group is the source of the rule allowing traffic from the Bricks to the SMS (Rule #203), and the destination of the rule allowing traffic to the Bricks from the SMS (Rule #204).
Using this host group as the source and destination of these rules ensures that any time a new Brick is configured, it can automatically communicate with the SMS.
- *nocgwzone*
The purpose of the *nocgwzone* ruleset is to protect an SMS that is managing routers through an encrypted management tunnel. As in *administrativezone*, the *Bricks* host group can be found in Rules #203 and #204.
- *proxyzone*
The purpose of the *proxyzone* ruleset is to protect the host(s) running the Proxy Agent software, or other content security software, and to permit proxied SMTP and HTTP sessions to be reflected from the Brick to the proxy hosts, and back to the Brick if necessary.
Note: the Proxy Agent software is no longer supported and has been replaced by the Rules-Based routing feature.
The *Bricks* host group is used as the destination in Rule #234. This rule allows the proxy host(s) to query a Brick about reflected sessions. Using this host group guarantees that the proxy host(s) will be able to query any Brick that has been configured.

mobile_bricks

The *mobile_bricks* host group is used for Bricks that are in a dynamic environment. Bricks that are coming from behind a NAT box or a DHCP or PPPoE environment are coming from an IP (public address) that can in theory change at any given time. The Brick coming from a dynamic environment will register this public address with the SMS on first contact. This host group can be used to manually place the public addresses in this host group instead of *, which is what will most likely be configured initially since the SMS, in theory, does not know the public IP of these Bricks coming from a dynamic environment.

Bricks_VBA

The *Bricks_VBA* host group appears in every group. It contains the Virtual Brick Address (VBA) of every zone that has been assigned to a port on a Brick. Every time an Administrator assigns a zone a VBA, the VBA is automatically added to this host group. If the VBA or the zone assignment is deleted, the VBA is automatically removed from the host group.

The *Bricks_VBA* host group is used as the destination in Rule #232 in the Proxy Zone ruleset. The purpose of this rule is to allow proxied sessions using dual-reflection to be reflected out of the Proxy Zone to a Brick. Using this host group guarantees that the reflected sessions can reach the Brick.

LSMS

The *LSMS* host group is global and, therefore, appears in every group. This group contains the IP address of a single SMS (in a standalone configuration) or two SMSs (in a redundant configuration). The public IP addresses of these SMSs, if assigned, are also in this host group. The group is used in the default "system" rules for the administrativezone, firewall, nocgwzone and vpnzone.

ComputeServers

The *ComputeServers* host group is global and is nested in the *SMS* host group. This group contains the IP address of every Compute Server (CS) configured in the system.

Proxy_hosts

The *Proxy_hosts* host group appears in every group. It contains the IP address(es) of the proxy host(s) in the Proxy Zone. These are the hosts running the Alcatel-Lucent Proxy Agent software, or other third-party content security software.

The *Proxy_hosts* host group is used as the destination in Rule #231 in the Proxy Zone ruleset. The purpose of this rule is to allow sessions reflected by the Brick to enter the Proxy Zone.

brickLocalAddresses and brickRemoteAddresses

The *brickLocalAddresses* and *brickRemoteAddresses* host groups appear in every group. They are used when a Brick is set up with private (non-registered) IP addresses. They are only populated (manually) if an administrator needs to conserve public IP addresses. There are default "system" rules already setup that use these host groups in the administrativezone and nocgwzone (For more information on this scenario, please review *Appendix B* in the *SMS Administration Guide*.)

The *brickLocalAddresses* group should be populated with the "public" address that the SMS needs to respond to in order to reach the "private" IP address of the Brick.

The *brickRemoteAddresses* host group should be populated with the "private" IP addresses of the Bricks with non-registered IP addresses.



Host Group Templates Provided with the SMS Application for VoIP/NOE Traffic

Overview

A set of pre-defined host group templates are provided with the SMS application when it is installed to make it easier to configure the Brick to handle VoIP/NOE traffic. These pre-defined host group templates are already pre-selected in the appropriate Brick zone ruleset, NOE application filter, or client tunnel endpoint, as required. The name of the host group template indicates the location and type of equipment that is being protected by the Brick in the VoIP network environment. For example, the host group template **NOE_PHONES_Headquarters** would contain the IP addresses of IP phone devices positioned behind the Brick zone at a corporate Headquarters site.

To configure these pre-defined host groups, all that is required is for you to enter the IP addresses of the equipment (such as the IP hardphones/softphones, call servers, MGW GA/GD boards, Presentation Server, Passive Communication Server), as indicated by the host group name.

In some cases, pre-defined host groups are nested within another pre-defined host group template that is assigned to the application filter, zone rule, or tunnel endpoint.

Related information

Review the following sections in this Guide for related information on use of the pre-defined templates provided with the SMS application for NOE/VoIP traffic:

- [“Brick Zone Ruleset Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 1-66)
- [“Service Group Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 4-27)
- [“TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-123)
- [“NOE Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-68)
- [Appendix E, “Configuring the Brick for VoIP/NOE Traffic Using Pre-Defined SMS Templates”](#)

Downloading the NOE host group templates

After the SMS application is installed, the NOE templates can be downloaded by right-clicking on a group folder (such as the *system* folder initially created when the SMS is installed) and choosing **Create NOE Template** from the pop-up menu.

When the NOE template set is created, the pre-defined host group templates are automatically downloaded in a subfolder named **NOE** under the **Host Groups** folder in the SMS Navigator (refer to [Figure 2-11, “Pre-Defined NOE Host Group Templates \(NOE Subfolder in Navigator\)”](#) (p. 2-26)).

Figure 2-11 Pre-Defined NOE Host Group Templates (NOE Subfolder in Navigator)

The screenshot shows the Alcatel-Lucent Security Management Server interface. The left pane displays a navigation tree with the following structure:

- admin
 - LSMSs and LSCSs
 - Administrators
 - Alarms
 - Reports
 - system
 - Devices
 - Bricks
 - Policies
 - Brick Zone Rulesets
 - Host Groups
 - NOE** (highlighted)
 - Service Groups
 - Domain Name Groups
 - Application Filters
 - Dependency Masks
 - VPN
 - LAN-LAN Tunnel Viewers
 - Client Tunnel Endpoints
 - VPN Defaults
 - User Auth
 - Authentication Services
 - User Groups
 - Users
 - joesgroup

The main pane displays a table of templates under the folder `/system/Policies/Host Groups/NOE/` (Total 41 item(s)).

Name	Admin	Last Modified	Description
NOE_Call_Server_HQ	setup	2009-05-11 10:55:43	Call Server at Headquar
NOE_Call_Server_Regional	setup	2009-05-11 10:55:43	Call Server at Regional S
NOE_Call_Servers	setup	2009-05-11 10:55:43	Nested Host Groups of C
NOE_GA_IPs_Branch_Office	setup	2009-05-11 10:55:43	GA IP Addresses for Bra
NOE_GA_IPs_Headquarters	setup	2009-05-11 10:55:43	GA IP Addresses for Hea
NOE_GA_IPs_Outside_HQ	setup	2009-05-11 10:55:43	GA IP Addresses for Out
NOE_GA_IPs_Outside_Regional	setup	2009-05-11 10:55:43	GA IP Addresses for Out
NOE_GA_IPs_Regional_Office	setup	2009-05-11 10:55:43	GA IP Addresses for Reg
NOE_GA_Public_IPs_Branch_Office	setup	2009-05-11 10:55:43	GA Public IP Addresses
NOE_GD_IPs_Branch_Office	setup	2009-05-11 10:55:43	GD IP Addresses for Bra
NOE_GD_IPs_Headquarters	setup	2009-05-11 10:55:43	GD IP Addresses for Hea
NOE_GD_IPs_Outside_HQ	setup	2009-05-11 10:55:44	GD IP Addresses for Out
NOE_GD_IPs_Outside_Regional	setup	2009-05-11 10:55:44	GD IP Addresses for Out
NOE_GD_IPs_Regional_Office	setup	2009-05-11 10:55:44	GD IP Addresses for Reg
NOE_GD_Public_IPs_Branch_Office	setup	2009-05-11 10:55:44	GD Public IP Addresses
NOE_MGWs_Branch_Office	setup	2009-05-11 10:55:44	MGWs of Branch Office
NOE_MGWs_Headquarters	setup	2009-05-11 10:55:44	MGWs of Headquarters
NOE_MGWs_Regional_Office	setup	2009-05-11 10:55:44	MGWs of Regional Office
NOE_Networked_Call_Servers	setup	2009-05-11 10:55:44	Networked Call Servers
NOE_Networked_MGWs	setup	2009-05-11 10:55:44	Networked MGWs
NOE_OTUC_Server	setup	2009-05-11 10:55:44	OTUC Server
NOE_Omni_Vista_4760_Server	setup	2009-05-11 10:55:44	Omni Vista 4760 Server
NOE_PCS	setup	2009-05-11 10:55:44	Passive Communication
NOE_PHONES_Branch_Office	setup	2009-05-11 10:55:44	NOE Phones of Branch C
NOE_PHONES_Headquarters	setup	2009-05-11 10:55:44	NOE Phones of Headqua
NOE_PHONES_MGW_inside_Branch	setup	2009-05-11 10:55:44	Nested Host Groups of E
NOE_PHONES_MGW_inside_HQ	setup	2009-05-11 10:55:44	Nested Host Groups of H
NOE_PHONES_MGW_inside_Regional	setup	2009-05-11 10:55:44	Nested Host Groups of F
NOE_PHONES_MGW_outside_HQ	setup	2009-05-11 10:55:44	Nested Host Groups for
NOE_PHONES_MGW_outside_Regional	setup	2009-05-11 10:55:44	Nested Host Groups for
NOE_PHONES_Regional_Office	setup	2009-05-11 10:55:44	NOE Phones of Regional

To remove the NOE zone ruleset templates and all component templates from the Policies folder, right-click on the group folder and choose **Delete NOE Template**. A confirmation dialog box is displayed, asking if you are sure that you want to delete the NOE policy template (which includes all of the zone ruleset templates are component templates) under the selected group folder. Choose **Yes**. The NOE templates are removed from the Policies folder in the selected group. If, however, one of the zone ruleset templates is in use on one or more Bricks, an error message is generated and the templates are not removed.

Summary description of pre-defined host groups for VoIP/NOE traffic

The following list summarizes the pre-defined host group templates provided with the SMS application for VoIP/NOE traffic. The list indicates whether a host group is nested within another host group.

- **NOE_Call_Server_HQ**— a blank host group that must be filled in with the IP address(es) of the call server(s) at the Headquarters site
- **NOE_Call_Server_Regional**— a blank host group that must be filled in with the IP address(es) of the call server(s) at the Regional site
- **NOE_Call_Servers**— a nested host group that contains the **NOE_Call_Server_HQ** and **NOE_Call_Server_Regional** host groups
- **NOE_GA_IPs_Branch_Office**— a blank host group that must be filled in with the private IP addresses assigned to the GA boards of the Media Gateway located behind the Brick at a Branch Office site. This host group is pre-selected in one of the pre-defined NOE application filter templates applied in the Brick zone ruleset for VoIP/NOE traffic.
- **NOE_GA_IPs_Headquarters**— a blank host group that must be filled in with the private IP addresses assigned to the GA boards of the Media Gateway located behind the Brick at a Headquarters site. This host group is pre-selected in one of the pre-defined NOE application filter templates applied in the Brick zone ruleset for VoIP/NOE traffic.
- **NOE_GA_IPs_Regional_Office**— a blank host group that must be filled in with the private IP addresses assigned to the GA boards of the Media Gateway located behind the Brick at a Regional Office site. This host group is pre-selected in one of the pre-defined NOE application filter templates applied in the Brick zone ruleset for VoIP/NOE traffic.
- **NOE_GA_IPs_Outside_HQ**— a nested host group that contains the **NOE_GA_IPs_Regional_Office** and **NOE_GA_IPs_Branch_Office** host groups
- **NOE_GA_IPs_Outside_Regional**— a nested host group that contains the **NOE_GA_IPs_Headquarters** and **NOE_GA_IPs_Branch_Office** host groups
- **NOE_GA_Public_IPs_Branch_Office**— a host group that must be filled in with the public IP addresses that will be assigned to the GA boards of the Media Gateway located at a Branch Office site when the Brick performs NAT during MGW-CS IPlink protocol communications. This host group is pre-selected in one of the pre-defined NOE application filter templates applied in the Brick zone ruleset for VoIP/NOE traffic.
- **NOE_GD_IPs_Branch_Office**— a host group that must be filled in with the private IP addresses of the GD board of the Media Gateway that are converted to one of the public IP addresses from the **NOE_GA_Public_IPs_Branch_Office** host group when the Brick performs NAT during MGW-CS IPlink protocol communications. This host group is pre-selected in one of the pre-defined NOE application filter templates applied in the Brick zone ruleset for VoIP/NOE traffic.

- **NOE_GD_IPs_Headquarters**— a host group that must be filled in with the private IP addresses of the GD board of the Media Gateway located behind the Brick at a Headquarters site. This host group is pre-selected in one of the pre-defined NOE application filter templates applied in the Brick zone ruleset for VoIP/NOE traffic.
- **NOE_GD_IPs_Regional_Office**— a host group that must be filled in with the private IP addresses of the GD board of the Media Gateway located behind the Brick at a Regional Office site. This host group is pre-selected in one of the pre-defined NOE application filter templates applied in the Brick zone ruleset for VoIP/NOE traffic.
- **NOE_GD_IPs_Outside_HQ**— a nested host group that contains the **NOE_GD_IPs_Regional_Office** and **NOE_GD_IPs_Branch_Office** host groups
- **NOE_GD_IPs_Outside_Regional**— a nested host group that contains the **NOE_GD_IPs_Headquarters** and **NOE_GD_IPs_Branch_Office** host groups
- **NOE_GD_Public_IPs_Branch_Office**— a host group that contains the public IP addresses that are assigned to the GD board of the Media Gateway at a Branch Office site when the Brick performs NAT during MGW-CS IPlink protocol communications. This host group is pre-selected in one of the pre-defined NOE application filter templates applied in the Brick zone ruleset for VoIP/NOE traffic.
- **NOE_MGWs_Branch_Office**— a nested host group that contains the **NOE_GA_IPs_Branch_Office** and **NOE_GD_IPs_Branch_Office** host groups
- **NOE_MGWs_Headquarters**— a nested host group that contains the **NOE_GA_IPs_Headquarters** and **NOE_GD_IPs_Headquarters** host groups
- **NOE_MGWs_Regional_Office**— a nested host group that contains the **NOE_GA_IPs_Regional_Office** and **NOE_GD_IPs_Regional_Office** host groups
- **NOE_PHONES_MGW_inside_Branch**— a nested host group that contains the **NOE_PHONES_Branch_Office** and **NOE_MGWs_Branch_Office** host groups.
- **NOE_PHONES_MGW_inside_HQ**— a nested host group that contains the **NOE_PHONES_Headquarters** and **NOE_MGW_Headquarters** host groups.
- **NOE_PHONES_MGW_inside_Regional**— a nested host group that contains the **NOE_PHONES_Regional_Office** and **NOE_MGWs_Regional_Office** host groups.
- **NOE_PHONES_MGW_outside_HQ**— a nested host group that contains the **NOE_PHONES_Branch_Office**, **NOE_MGWs_Branch_Office**, **NOE_PHONES_Regional_Office**, **NOE_MGWs_Regional_Office**, and **NOE_PHONES_Remote_Office** host groups.
- **NOE_PHONES_MGW_outside_Regional**— a nested host group that contains the **NOE_PHONES_Branch_Office**, **NOE_MGWs_Branch_Office**, **NOE_PHONES_Headquarters**, **NOE_MGWs_Headquarters**, and **NOE_PHONES_Remote_Office** host groups.
- **NOE_PHONES_Branch_Office**— a blank host group that must be filled in with the IP addresses of the IP phone devices located at the Branch Office site.

- **NOE_PHONES_Headquarters**— a blank host group that must be filled in with the IP addresses of the IP phone devices located at the Headquarters site.
- **NOE_PHONES_Regional_Office**— a blank host group that must be filled in with the IP addresses of the IP phone devices located at the Regional Office site.
- **NOE_PHONES_Remote_Office**— a blank host group that must be filled in with the IP addresses of the IP phone devices located at the Remote Office site.
- **NOE_VPN_Branch**— a nested host group that contains the **NOE_PHONES_Branch_Office** and **NOE_MGWs_Branch_Office** host groups.
- **NOE_VPN_HQ_HUB**— a nested host group that contains the **NOE_PHONES_Branch_Office**, **NOE_MGWs_Branch_Office**, **NOE_PHONES_Regional_Office**, **NOE_MGWs_Regional_Office**, **NOE_PHONES_Remote_Office**, **NOE_PHONES_Headquarters**, **NOE_MGWs_Headquarters**, **NOE_Call_Server_HQ**, and **NOE_TFTP_Server_HQ** host groups.
- **NOE_VPN_Remote**— a nested host group that contains the **NOE_PHONES_Remote_Office** host group.
- **NOE_VPN_Outside_HQ**— a nested host group that contains the **NOE_PHONES_Branch_Office**, **NOE_MGWs_Branch_Office**, **NOE_PHONES_Regional_Office**, **NOE_MGWs_Regional_Office**, **NOE_PHONES_Remote_Office**, **NOE_Call_Server_Regional**, and **NOE_TFTP_Server_Regional** host groups.
- **NOE_Networked_Call_Servers**— a blank host group that must be filled in with the IP addresses of the networked call servers, typically positioned behind the Brick at a Headquarters site and carrying VoIP traffic over SIP trunks.
- **NOE_Networked_MGWs**— a blank host group that must be filled in with the IP addresses of the GA and GD boards of the Media Gateway, typically positioned behind the Brick at a Headquarters site and converting TDM voice to VoIP packets.
- **NOE_TFTP_Server_HQ**— a blank host group that must be filled in with the IP address of the TFTP server at a Headquarter site.
- **NOE_TFTP_Server_Regional**— a blank host group that must be filled in with the IP address of the TFTP server at a Regional Office site.
- **NOE_TFTP_Servers**— a nested host group that contains the **NOE_TFTP_Server_HQ** and **NOE_TFTP_Server_Regional** host groups.
- **NOE_Presentation_Server**— a blank host group that must be filled in with the IP address of the Presentation Server.
- **NOE_OTUC_Server**— a blank host group that must be filled in with the IP address of the OTUC Server.
- **NOE_Omni_Vista_4760_Server**— a blank host group that must be filled in with the IP address of the OmniVista 4760 Server.
- **NOE_PCS**— a blank host group that must be filled in with the IP address of the Passive Communication Server (PCS).



3 Domain Name Groups

Overview

Purpose

This chapter explains how to set up, use, and maintain domain name groups.

Contents

Domain Name Groups	3-2
To Set Up a Domain Name Group	3-3
Global Domain Name Groups	3-7
Nested Domain Name Groups	3-10
Maintaining Domain Name Groups	3-12



Domain Name Groups

Definition

A domain name group is a collection of domain names. There are two types of domain name groups — standard domain name groups, which can only be used in the group in which they were created, and global domain name groups, which can be used in every group.

Uses

Creating a domain name group allows you to conveniently define a list of domain names once and reference it as an entity from within application filters. The idea is to define a domain name group and then protect those domains for a particular application protocol. In Release 8.0, this feature is limited to DNS application filters, but in future releases you will be able to reference it from others.



To Set Up a Domain Name Group

Overview

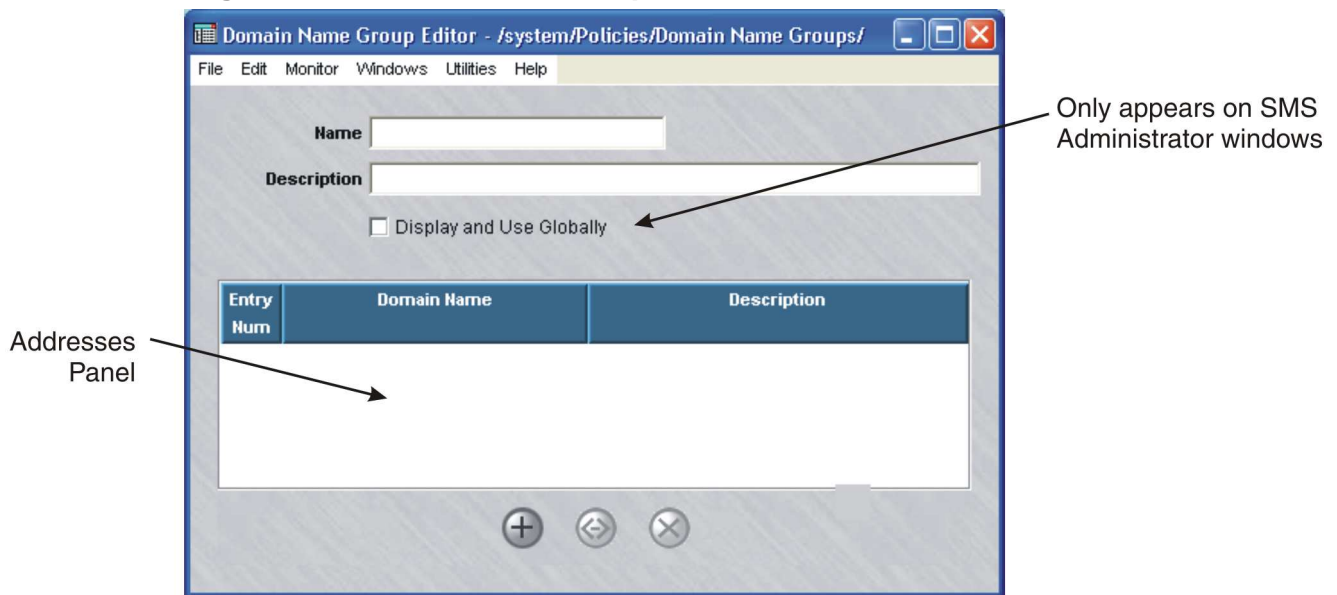
To set up a domain name group, you have to display the Domain Name Groups Editor and enter the name of the domain name group and an optional description.

To display the Domain Name Groups Editor

The easiest way to display the Domain Name Group Editor is to:

- 1 Open the folder of the group that will contain this domain name group.
- 2 Open the Policies folder.
- 3 Right-click the Domain Name Group folder and select **New Domain Name Group** from the pop-up menu. The Domain Name Group Editor is displayed (Figure 3-1, “Domain Name Group Editor” (p. 3-3)).

Figure 3-1 Domain Name Group Editor



END OF STEPS

To create a domain name group

Complete the following steps to create a domain name group.

- 1 In the **Name** field, enter a unique name to identify this domain name group. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers and certain special characters.
- 2 In the **Description** field, you can enter an optional description of the domain name group. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers and certain special characters.

This is the description that will appear on the Navigator window when you view existing domain name groups. Refer to the section [“Overview” \(p. 3-12\)](#).

Important! If you are an SMS Administrator, you will see a checkbox entitled “Display and Use Globally” under the **Description** field. Click this checkbox if you want to make this a global domain name group. Refer to the section [“Global Domain Name Groups” \(p. 3-7\)](#) for an explanation of global domain name groups.

- 3 Right-click in the addresses panel and select **New** from the pop-up menu (or click the **New** button) to display the Domain Name Group Entry window, which is shown in [Figure 3-2, “Domain Name Group Entry Window” \(p. 3-4\)](#).

Figure 3-2 Domain Name Group Entry Window



Important! If you intend to incorporate another domain name group within this one, click the **Select Nested Domain Name** checkbox. For an explanation of nested domain name groups, refer to the section [“Nested Domain Name Groups”](#) (p. 3-10).

- 4 In the Domain Name Group Entry window, do the following:
 - In the **Domain Name** field, enter a unique name to identify this domain name group. The name can contain up to 73 characters. It can consist of upper and lower case letters, numbers and certain special characters.
 - In the **Description** field, you can enter an optional description of the domain name. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers and certain special characters.
 - Click **OK** to dismiss the window. You will be returned to the Domain Name Group Editor. The entry you just created will appear in the address panel.
- 5 To add more groups to this domain name group, repeat [Step 3](#) and [Step 4](#) until all addresses have been entered. [Figure 3-3, “Domain names in a Domain Name Group”](#) (p. 3-5) shows a domain name group editor screen with two domain names entered.

Figure 3-3 Domain names in a Domain Name Group



- 6 Display the File menu and select one of the **Save** options.

Important! If the domain names or descriptions are in another program, such as Microsoft Word or Excel, you can copy the names and descriptions into the clipboard, and then paste them into the **Domain Name** and **Description** fields using **[Ctrl-V]**.

.....
E N D O F S T E P S



Global Domain Name Groups

Definition

A global domain name group is a domain name group that is created in one SMS group, but can be seen and used in every other group. Only SMS Administrators can create global domain name groups. Global domain name groups are dynamically linked, and will update everywhere when modified.

Create a Global Domain Name Group

If you are an SMS Administrator, you create a global domain name group by clicking the **Display and Use Globally** checkbox on the Domain Name Group Editor (see [Figure 3-1, “Domain Name Group Editor” \(p. 3-3\)](#)). You can do this when you create the domain name group, or you can do this after the domain name group has been created by editing the domain name group (see “Modify a Domain Name Group” below).

When creating a global domain name group, make sure the name you give the domain name group is unique across all groups.

View Global Domain Name Groups

Global domain name groups appear in the Navigator window, just like standard domain name groups (refer to the section [“Overview” \(p. 3-12\)](#)). When you view the domain name groups that have been created in a specific group, all the domain name groups *created in that group* appear the same, regardless of whether or not they are global.

However, any global domain name groups created in *a different group* can be identified by a globe icon that appears to the left of the entry, as shown below.

 XYZDomains	admin	2004-03-12 07:41:40	XYZ Corp.
--	-------	---------------------	-----------

To determine the group in which the global domain name group was originally created, right-click the domain name group in the Navigator window and select **Original Location** from the pop-up menu. A window similar to the one shown in [Figure 3-4](#), “[Locate Original Window](#)” (p. 3-8) is displayed.

Figure 3-4 Locate Original Window



Removing the Global Status of a Domain Name Group

Just as it is possible to make a non-global domain name group global by clicking the **Display and Use Globally** checkbox on the Domain Name Group Editor (see [Figure 3-1](#), “[Domain Name Group Editor](#)” (p. 3-3)), it is possible to remove the global status of a domain name group by unchecking this checkbox. However, this can only be done if the domain name group is *not* in use globally. If the domain name group is in use in any group *other than the one in which it was created*, you cannot remove its global status.

It is also possible to delete a global domain name group — but only the original source of the domain name group. You cannot delete a global domain name group from any folder in which it appears except the folder in which it was originally created. No domain name group — global or standard — can be deleted if it is in use anywhere (refer to the section “[To delete a domain name group](#)” (p. 3-17)).

Similarly, only the original source of a global domain name group can be moved (refer to the section “[To move a domain name group](#)” (p. 3-16)).

Permissions

Permissions over global domain name groups are based on the group in which the domain name group was originally created. If an administrator has FULL policy permissions for that group, then that administrator has FULL permissions over all global domain name groups created in that group.

If an administrator has FULL permission over a global domain name group, the administrator can edit the domain name group in any of the groups in which it appears. An administrator can create a copy of a global domain name group as long as the administrator has FULL permissions over the *destination* group.



Nested Domain Name Groups

Definition

A nested domain name group is a domain name group that has other domain name groups nested inside it. Nesting domain name groups allows you to create domain name groups composed of one or more other domain name groups. Including domain name groups within another domain name group is generally easier, and less prone to error, than entering the domain name of the internal domain name groups manually.

For example, suppose the administrator of XYZ Corp. with several divisions creates individual domain name groups for each division. The administrator could then create one large domain name group for the corporation, and nest each of the division domain name groups within the overall corporate domain name group. This way of organizing the domain names might be easier to manage than one large list.

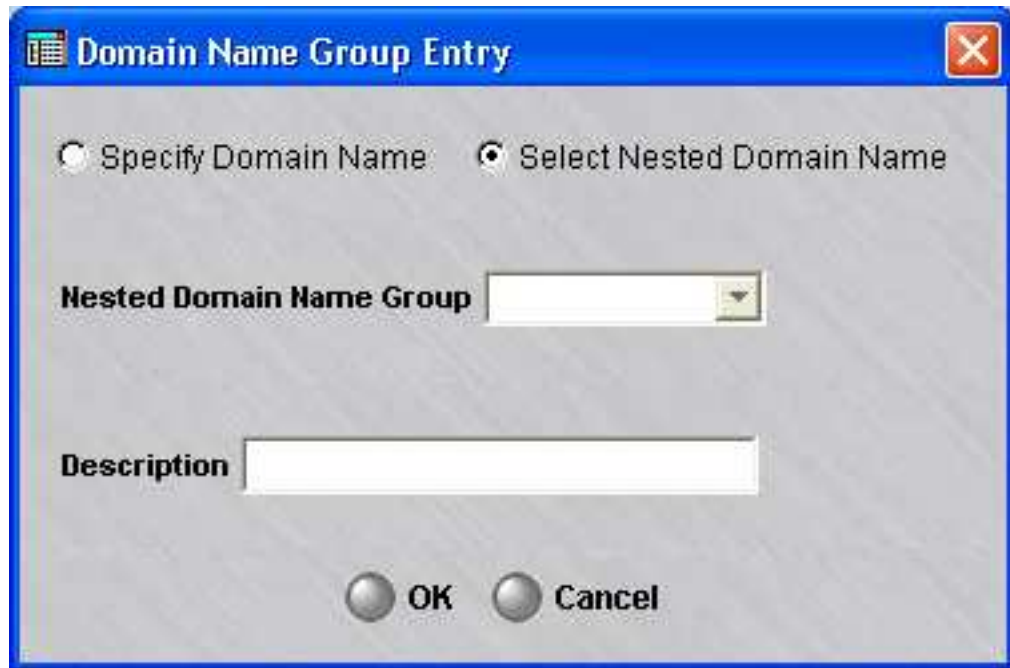
You can nest global domain name groups within other global domain name groups, or within other non-global domain name groups. Administrators will be warned if a global domain name group contains non-global domain name groups nested within it, because the names and contents of these non-global domain name groups will be exposed to administrators who manage other groups. You can create nested domain name groups with up to 32 levels of nesting.

Cycles are prohibited in nested domain name groups. For example, suppose domain name group A is a nested domain name group, and contains domain name group B within it. Further, suppose domain name group B is also a nested domain name group, and includes domain name group C within it. In this instance, domain name group C would not be permitted to be a nested domain name group containing domain name group A.

Task

To nest a domain name group within another domain name group, follow the steps below:

- 1 Create a domain name group as previously explained (see "Create a Domain Name Group" above). Follow the steps until the Domain Name Group Entry window appears (Figure 3-2, "Domain Name Group Entry Window" (p. 3-4)).
- 2 Click the **Select Nested Domain Name** checkbox in the Domain Name Group Entry window (see Figure 3-2, "Domain Name Group Entry Window" (p. 3-4)). The **Domain Name** field will become a **Nested Domain Name Group** field, as shown in Figure 4-5.

Figure 3-5 Domain Name Group Entry Window (Nested Domain Name Groups)

-
- 3** In the **Nested Domain Name Group** field, select a domain name group that will be nested in this group from the drop-down list. If a description was entered for that group, it will appear in the **Domain Name Description** field.

 - 4** To nest additional groups in this domain name group, repeat Steps 2 and 3 for each additional group.

 - 5** When you have finished, click **OK**.

END OF STEPS



Maintaining Domain Name Groups

Overview

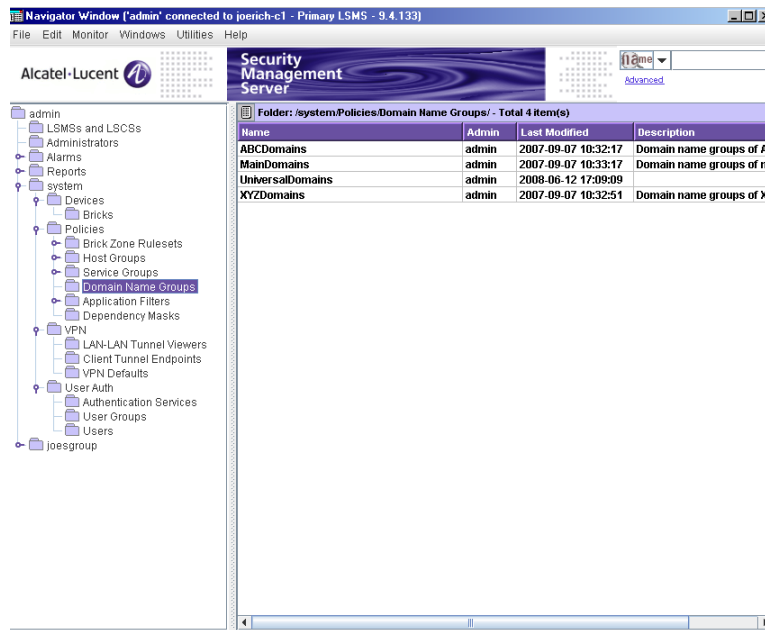
Once a domain name group has been created and saved, it can be viewed, modified, copied, or moved. If the domain name group is no longer needed, it can be deleted.

To view domain name groups

You must view the domain name groups before you can edit, move, copy, or delete a specific domain name group. To view all the groups that have been created to date in a particular group, follow the steps below:

- 1 Open the appropriate group folder, and then open the Policies folder.
- 2 Click the Domain Name Groups folder. All existing domain name groups will be displayed in the Navigator window, as shown in [Figure 3-6, “Navigator Window \(View Domain Name Groups\)”](#) (p. 3-12).

Figure 3-6 Navigator Window (View Domain Name Groups)



As shown, for each domain name group, the Navigator window shows the Administrator who created the domain name group, the date and time the group was created, and a brief description, if one was entered when the group was created.

END OF STEPS

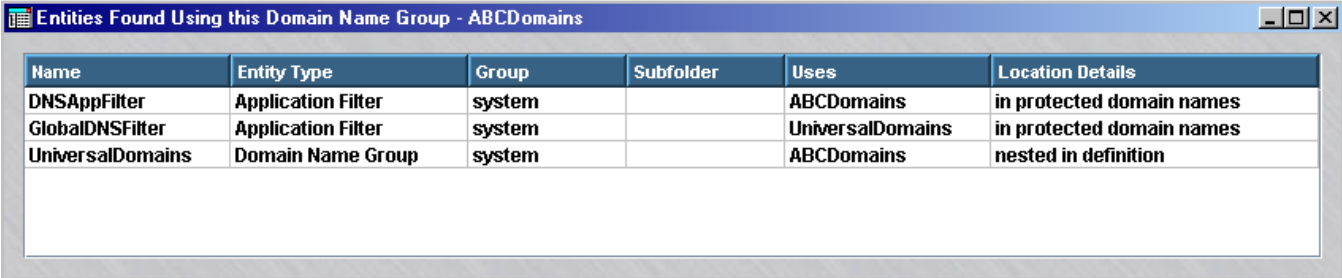
To find entities that use a domain name group

The SMS application allows you to perform a search for all entities (such as a zone ruleset, application filter, or dependency mask) that are using the selected domain name group, and, if it is nested within another domain name group, all entities that are using the parent domain name group.

- 1 With the domain name groups displayed in the Navigator window, right-click the desired domain name group you want and select **Find Entities Using this Domain Group** from the pop-up menu.

Result An Entities Found Using this Domain Name Group window is displayed (Figure 3-7, “Entities Found Window” (p. 3-13)) showing the results of the search.

Figure 3-7 Entities Found Window




Name	Entity Type	Group	Subfolder	Uses	Location Details
DNSAppFilter	Application Filter	system		ABCDomains	in protected domain names
GlobalDNSFilter	Application Filter	system		UniversalDomains	in protected domain names
UniversalDomains	Domain Name Group	system		ABCDomains	nested in definition

The display results on the window show the entity type that uses or contains the selected domain name group, the group and subfolder in which the entity is found, whether the entity uses the domain name group or the parent domain name group (if it is a nested domain name group), and details about the entity.

To view additional details about the entity, right-click on the entry in the table and select **View** from the pop-up menu.

To edit the entity, right-click on the entry in the table and select **Edit** from the pop-up menu.

- 2 Click the Close () button at the top of the window to close the window.

END OF STEPS

To modify a domain name group

You can modify any field in a domain name group. To modify a domain name group, follow the steps below:

- 1 With the domain name groups displayed in the Navigator window, double-click the domain name group you want to modify.

The Domain Name Groups Editor will appear, as shown in [Figure 3-3, “Domain names in a Domain Name Group”](#) (p. 3-5).

- 2 The **Name** and **Description** fields and the **Display and Use Globally** checkbox can be modified.

- 3 Display the File menu and select one of the **Save** options.

END OF STEPS

To copy a domain name group

You can copy a domain name group to a different folder in the same group, or to a folder in another group, provided you have *Full* permission for that group. When you edit a domain name group or (using the **Save As** option to save the edited version with a new name) or copy a domain name group, you have the option of copying components from the original domain name group (nested domain name group), specifying how those components will be named and where they will be created and stored.

Complete the following steps to copy a domain name group:

- 1 With the domain name groups displayed in the Navigator window, right-click the domain name group you want to copy and select **Copy** from the pop-up menu.

Result A Copy window is displayed (Figure 3-8, “Copy Window” (p. 3-15)).

Figure 3-8 Copy Window



-
- 2 Complete the following fields, as applicable:
- **Name**— Enter a new name for the domain name group copy (Example: **ABCDomains_2**). This field is required. The name entered must be different from the name of the original domain name group by at least one character if it is in the same group.
 - **Folder**— The copy of the domain name group will be created and stored in the **Domain Name Groups** folder under the *system* group, by default. To specify a different folder, click the **Browse** button, which brings up a Select a Domain Name Group Folder window, and select the folder where the domain name group copy will be created and stored.
 - **Copy components when copying within a Group**— Click this checkbox (place a check in it) to copy the components from the original domain name group. If this checkbox is unchecked, and the copied domain name group is in the same group as the original one, the copied domain name group uses the same components as the original. If the domain name group is copied across groups, the components are copied regardless of this field setting.
 - **Component Suffix**— If the **Copy components when copying within a Group** option is checked, enter a text string that will be appended to the name of each copied component (nested domain name group). If this field is left blank, the name of each copied component will be spelled the same way as the original one. Name collisions are possible, regardless of whether or not an entry is made in this field. If this occurs, the content of the conflicting component(s) is examined. If the content is the same, the duplicated domain name group uses the existing

component(s) and the new copy(ies) is not made. If the content is different, a copy of the component(s) is created and stored under a unique name by appending its database ID to the name.

Note: the Component Suffix used must be unique and cannot be reused across zone rulesets.

- **Component Folder**— This field is used to specify a common folder name for the copied components. The folder name is relative to the top-level folder of each type of component. If this field is left blank, the copied components are stored in the top-level folder of that component. This folder is created as needed (if it does not already exist). A path string can be specified; / or \ delimiters can be used to delineate folder names.

-
- 3 Click **OK** to copy the domain name group and dismiss the Copy window.

Result The system returns to the Navigator window. The copy of the domain name group appears in the Contents panel.

The domain name group (and components, if specified) are copied to the chosen folder(s).

END OF STEPS

To move a domain name group

You can move a domain name group to a different folder in the same group, or to a folder in another group, provided you have *Full* permission for that folder. However, you can only move it to the folder labeled "Domain Name Groups" or to a subfolder under that folder. If you try to move it to another folder, you will get an error message. Before moving a domain name group, make sure that it is not currently in use. If you attempt to move a domain name group that is in use, you will get an error message. If you receive an error message when attempting to move a domain name group, you can run the "find entity" utility to discover where it is in use (see "Find Entities Using a Domain Name Group" above). If a domain name group is moved to a new group, and the domain name group has other domain name groups nested within it, the original domain name group is moved to the new group, but the domain name groups nested within it are copied to the new group, not moved. If a domain name group with the same name and content already exists in the new group, the copy does not take place. If a domain name group with the same name but different content already exists, the copy is renamed.

To move a domain name group, follow the steps below:

- 1 With the domain name groups displayed in the Navigator window, right-click the group you want to move and select **Move** from the pop-up menu. A Browse window will appear.
- 2 Select the folder you want to move this domain name group to. The domain name group will be moved to the folder you selected, and you will be returned to the Navigator window.

END OF STEPS

To delete a domain name group

Before deleting a domain name group, make sure that it is not currently in use. If you attempt to delete a domain name group that is in use, you will get an error message. If you receive an error message when attempting to delete a domain name group, you can run the "find entity" utility to discover where it is in use (see "Find Entities Using a Domain Name Group" above).

To delete a domain name group, follow the steps below:

- 1 With the domain name group displayed in the Navigator window, right-click the domain name group you want to delete and select **Delete** from the pop-up menu.

Result A confirmation window similar to the one shown in [Figure 3-9](#), "Confirmation Window (Domain Name Groups)" (p. 3-17) is displayed.

Figure 3-9 Confirmation Window (Domain Name Groups)



-
- 2 Click **Yes** to delete the domain name group and dismiss the Confirmation window.

Result The domain name group is removed from the Navigator window.

END OF STEPS



4 Service Groups

Overview

Purpose

This chapter explains how to set up, use and maintain service groups. It also describes the service groups that are provided with the SMS application.

Contents

What is a Service Group?	4-2
To Set Up a Service Group	4-4
Global Service Groups	4-9
Nested Service Groups	4-11
Maintaining Service Groups	4-13
Service Groups Provided with the SMS Application	4-23
Service Group Templates Provided with the SMS Application for VoIP/NOE Traffic	4-27



What is a Service Group?

Definition

A service group defines a service or collection of services. A service consists of a protocol, destination port, and source port. A service group can also contain the wildcard asterisk (*) in place of a specific protocol, destination port or source port. The asterisk means that this service group applies to every protocol, destination port or source port (whichever has the asterisk).

There are two types of service groups — standard service groups, which can only be used in the group in which they were created, and global service groups, which can be used in every group.

Uses

You can create service groups of your own, or you can apply one of the service groups that have been provided with the SMS application. These “canned” service groups cover most of the important services used by Internet applications today.

A set of pre-defined service group templates are also provided with the SMS application when it is installed to make it easier to configure the Brick to filter and inspect call and data traffic between IP phone devices and the other network elements (such as the Media Gateway unit and call servers) in a VoIP network environment. These pre-defined service group have been specifically set up for VoIP/NOE traffic and are already assigned to rules in pre-defined Brick zone rulesets for this purpose. For more details, refer to the section

There are two reasons why you may need to create your own service groups:

- To handle applications that are not addressed by any of the service groups provided with the SMS
and
- To create a service group that contain multiple services.

You can use the service groups you create, as well as the service groups provided, in any security rule in a Alcatel-Lucent Firewall *VPN Firewall Brick*TM Security Appliance zone ruleset, and in any dependency mask you create.

Benefits

Service groups simplify the process of creating and managing a zone's security policy. For example:

- They enable you to enter the same set of services or protocols in multiple rules quickly and consistently.
- They allow you to identify intuitively the services impacted by a rule. A name such as "AllHTTP" is much easier to identify than a string of Web ports.
- They permit you to keep the number of rules to a minimum by letting you add new services when needed, instead of new rules.



To Set Up a Service Group

Overview

To set up a service group, you have to display the Service Groups Editor and enter the name of the service group and the services (protocol/destination port/source port) to be included in the group.

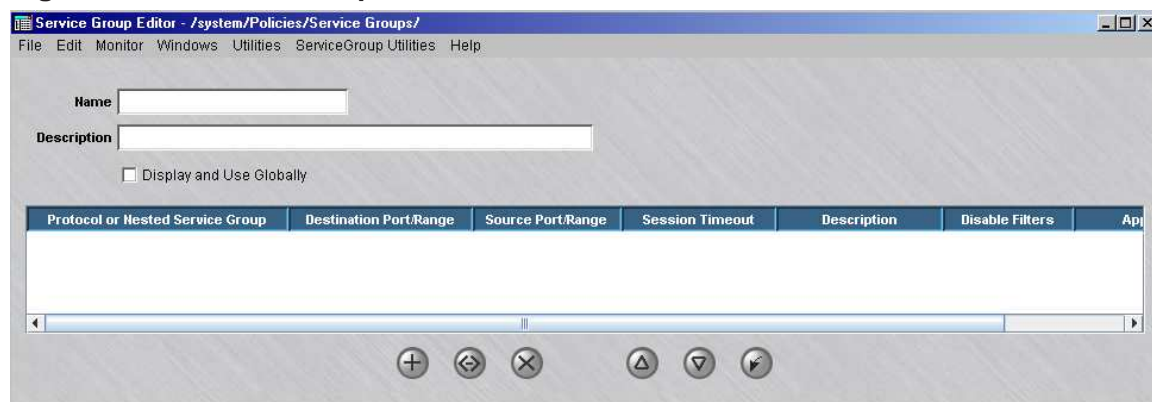
Display the Service Group Editor

The easiest way to display the Service Groups Editor is to:

- 1 Open the folder of the group that will contain this service group.
- 2 Open the Policies folder.
- 3 Right-click the Service Groups folder and select **New Service Group** from the pop-up menu.

Result The Service Groups Editor is displayed (Figure 4-1, “Service Group Editor” (p. 4-4)).

Figure 4-1 Service Group Editor



END OF STEPS

To create a service group

Complete the following steps to create a service group.

- 1 In the **Name** field, enter a unique name to identify this service group. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers and certain special characters.
-

- 2 In the **Description** field, you can enter an optional description of the service group. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers and certain special characters.

This is the description that will appear on the Navigator window when you view existing service groups. See the section below entitled [“To view a list of existing service groups”](#) (p. 4-13).

Important! If you are an SMS Administrator, you will see a checkbox entitled **Display and Use Globally** under the **Description** field. Click this checkbox if you want to make this a global service group. See [“Global Service Groups”](#) (p. 4-9) below for an explanation of global service groups.

- 3 Right-click the **Services** panel and select **New** from the pop-up menu (or click the **New** button).

Result The Service Editor is displayed (Figure 4-2, “Service Editor” (p. 4-6)).

Figure 4-2 Service Editor

Service Editor - /system/Policies/Service Groups/

Description [Text Field]

Specify Service Select Nested Service Group

Protocol * [Dropdown Menu]

Destination Port or Range (TCP/UDP) * [Text Field] - [Text Field]
Type (ICMP)

Source Port or Range (TCP/UDP) * [Text Field] - [Text Field]
Code (ICMP)

Session Timeout (seconds) [Text Field]

Disable all filtering including built-in

Application Filter [Dropdown Menu]

OK Cancel

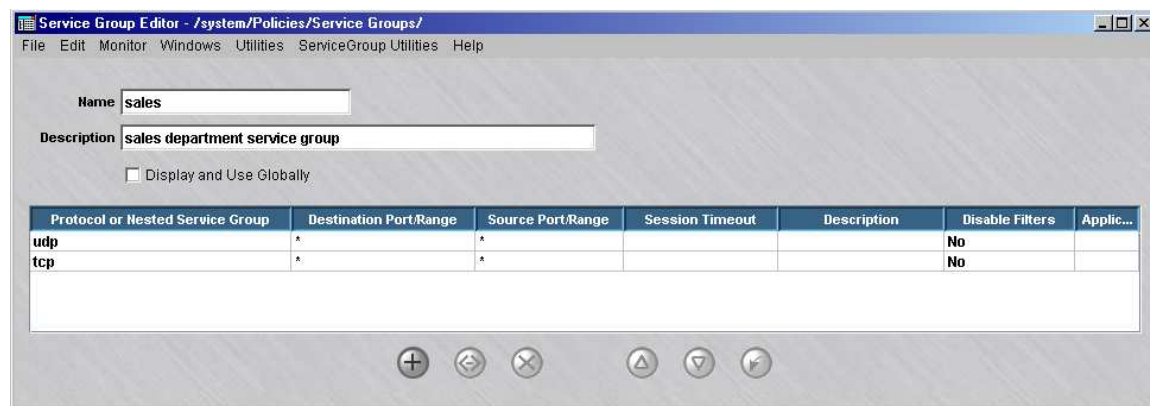
-
- 4 The purpose of the Service Editor is to allow you to add services to the service group.

With this screen displayed, do the following:

1. In the **Description** field, enter a brief description (up to 80 characters) of this service. This is optional.
If you choose to enter a description, the text you enter will appear in the Description column in the **Services** panel of the Service Group Editor (see [Figure 4-1, “Service Group Editor”](#) (p. 4-4)).
If you intend to incorporate another service group within this one, click the **Select Nested Service Group** checkbox. For an explanation of nested service groups, see [“Nested Service Groups”](#) (p. 4-11).
2. In the **Protocol** field, select a protocol from the drop-down list. This field is required. The choices are:
 - Asterisk (any protocol)
 - ICMP
 - SCTP
 - TCP
 - UDP
3. In the **Destination Port or Range** field, enter the number, or range of numbers, of the destination port. If you do not want to specify a destination port, leave the default asterisk in place.
4. In the **Source Port or Range** field, enter the number, or range of numbers, of the source port. If you do not want to specify a source port, leave the default asterisk in place.
5. Optionally, in the **Session Timeout (seconds)** field, you can specify a session timeout for this service group which overrides the rule session timeout setting. This field defines the time that a session of the specified service type may be idle before timing out. This setting can simplify the rule structure by eliminating the need to have different rules for different timeouts.
6. To disable any application filtering of traffic by the Brick for the specified protocol, including the built-in application filtering normally performed by the Brick, click the **Disable all filtering including built-in** checkbox to place a check in it (enable this option). When this option is checked, the **Application Filter** field is cleared and grayed out, and the Brick will not perform any application filtering on the specified protocol traffic. This option is disabled, by default. This option applies to the FTP, TFTP, IKE (udp) and EUA (udp) traffic protocols.
Built-in application filtering of traffic is normally performed by the Brick on the following protocols/ports:
 - FTP - tcp/21
 - TFTP - udp/69
 - IKE - udp/500 and udp/4500
 - EUA - udp/911/9020 (for example, dstport=911 and srcport=9020)

7. To use a configured application filter, select the filter from the drop-down list. To create a new application filter, follow the procedure that is explained in [Chapter 5, “Application Filters”](#).
8. Click **OK** to dismiss the Service Editor and return to the Service Group Editor. The service you just created will appear in the Services panel.
9. If you need to add other services to this service group, repeat 1 - 8 for each additional service. [Figure 4-3, “Service Group with Two Services”](#) (p. 4-8) shows a service group called *sales* that contains two services (UDP and TCP).

Figure 4-3 Service Group with Two Services



5. Display the File menu and select one of the **Save** options.

Important! *Import Services*

If the service group you want to create will contain services that already have their own service groups, such as HTTP or SMTP, there is a shortcut you can take.

Instead of selecting **New** when you right-click in the Service Group Editor (see [“To create a service group”](#) (p. 4-5) above), select **Import Services** instead. A browse window will appear instead of the Service Editor ([Figure 4-2, “Service Editor”](#) (p. 4-6)) and allow you to select an existing service group and copy its services directly into the Service Group Editor ([Figure 4-1, “Service Group Editor”](#) (p. 4-4)). The name of the service group from which you copied the services will appear in the **Description** field. Repeat this step for each service to be copied, and then save the new service group.

There is no linkage between the two service groups. If you change the original service group at a later time, this will not affect the new service group.

END OF STEPS



Global Service Groups

Definition

A global service group is a service group that is created in one group, but can be seen and used in every other group. Only SMS Administrators can create global service groups.

To create a global service group


If you are an SMS Administrator, you create a global service group by clicking the **Display and Use Globally** checkbox on the Service Group Editor (see [Figure 4-1, “Service Group Editor”](#) (p. 4-4)). You can do this when you create the service group, or you can do this after the service group has been created by editing the service group (see [“To modify a service group”](#) (p. 4-18) below).

When creating a global service group, make sure the name you give the service group is unique across all groups. If you attempt to give a global service group a name that is in use elsewhere, the service group will not be created, and you will get an error message indicating the save failed because there is an object in another group with the same name.

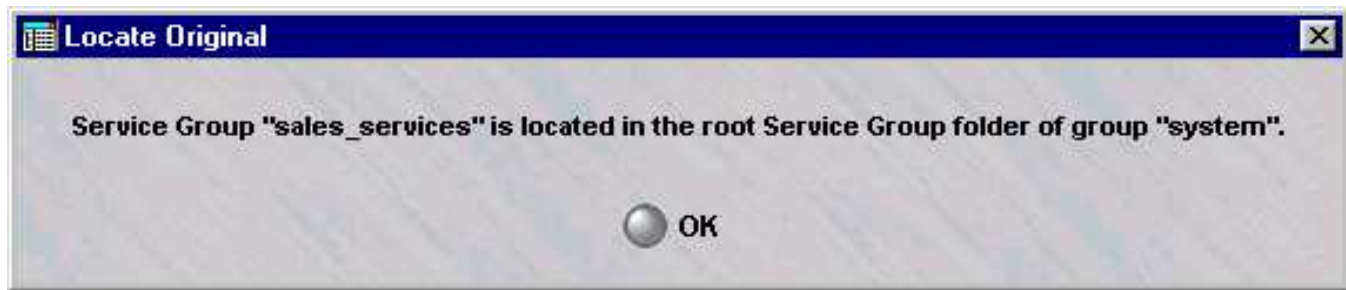
To view global service groups

Global service groups appear in the Navigator window, just like standard service groups (see [“To view a list of existing service groups”](#) (p. 4-13) below). When you view the service groups that have been created in a specific group, all the service groups *created in that group* appear the same, regardless of whether or not they are global.

However, any global service groups created in *a different group* can be identified by a globe icon that appears to the left of the entry, as shown below.

 sales_services	stan	2002-05-27 14:34:18	sales department services
--	------	---------------------	---------------------------

To determine the group in which the global service group was originally created, right-click the service group in the Navigator window and select **Original Location** from the pop-up menu. A window similar to the one shown in [Figure 4-4, “Locate Original Window \(Service Groups\)”](#) (p. 4-10) will appear.

Figure 4-4 Locate Original Window (Service Groups)

To remove the global status of a service group

Just as it is possible to make a non-global service group global by clicking the **Display and Use Globally** checkbox on the Service Group Editor (see [Figure 4-1, “Service Group Editor”](#) (p. 4-4)), it is possible to remove the global status of a service group by unchecking this checkbox. However, this can only be done if the service group is *not* in use globally. If the service group is in use in any group *other than the one in which it was created*, you cannot remove its global status.

It is also possible to delete a global service group — but only the original source of the service group. You cannot delete a global service group from any folder in which it appears except the folder in which it was originally created. No service group — global or standard — can be deleted if it is in use anywhere (see [“To delete a service group”](#) (p. 4-21) below).

Similarly, only the original source of a global service group can be moved (see [“To move a service group”](#) (p. 4-20) below)

Permissions

Permissions over global service groups are based on the group in which the service group was originally created. If an administrator has FULL policy permissions for that group, then that administrator has FULL permissions over all global service groups created in that group.

If an administrator has FULL permission over a global service group, the administrator can edit the service group in any of the groups in which it appears. An administrator can create a copy of a global service group as long as the administrator has FULL permissions over the *destination* group.

Nested Service Groups

Definition

A nested service group is a service group that has other service groups nested inside it. Nesting service groups allows you to create service groups composed of one or more other service groups. Including service groups within another service group is generally easier, and less prone to error, than entering the protocols and ports of the service groups manually.

For example, suppose an administrator wants to create a service group with the FTP, HTTP and SMTP services in it. Since service groups for these services already exist, the easiest way to create this service group is to simply nest the three existing service groups in the larger service group. This is much easier than entering the protocols and ports used by each service one at a time.

You can nest global service groups within other global service groups, or within other non-global service groups. When a global service group contains non-global service groups nested within it, the service group loaded to the Brick will contain only the contents of the service groups visible in the same group as the Brick Zone Ruleset. You can create nested service groups with up to 32 levels of nesting.

Cycles are prohibited in nested service groups. For example, suppose service group A is a nested service group, and contains service group B within it. Further, suppose service group B is also a nested service group, and includes service group C within it. In this instance, service group C would not be permitted to be a nested service group containing service group A.

Task

To nest a service group within another service group, follow the steps below:

- 1 Create a service group as previously explained (see [“To create a service group”](#) (p. 4-5) above). Follow the steps until the Service Editor appears ([Figure 4-2, “Service Editor”](#) (p. 4-6)).

- 2 Click the **Select Nested Service Group** checkbox in the Service Editor (see [Figure 4-2, “Service Editor”](#) (p. 4-6)). The protocol and port fields will become a **Nested Service Group** field, as shown in [Figure 4-5, “Service Editor \(Nested Service Groups\)”](#) (p. 4-12).

Figure 4-5 Service Editor (Nested Service Groups)



- 3 In the **Nested Service Group** field, select a service group that will be nested in this service group from the drop-down list.
- 4 To nest additional service groups in this service group, repeat Steps 2 and 3 for each additional service group.
- 5 When you have finished, click **OK**. [Figure 4-6, “Nested Service Groups” \(p. 4-12\)](#) shows an example of two service groups nested within a another service group.

Figure 4-6 Nested Service Groups

Protocol or Nested Service Group	Destination Port/Range	Source Port/Range	Session Timeout	Description	Disable Filters	A...
GTPv0						
ftp						
sctp						

END OF STEPS



Maintaining Service Groups

When to use

Once a service group has been created and saved, it can be viewed, modified, copied, or moved. If the service group is no longer needed it can be deleted.

To view a list of existing service groups

You must view the service groups before you can edit, move, copy, or delete a specific service group. To view all the service groups that have been created to date in a particular group, follow the steps below:

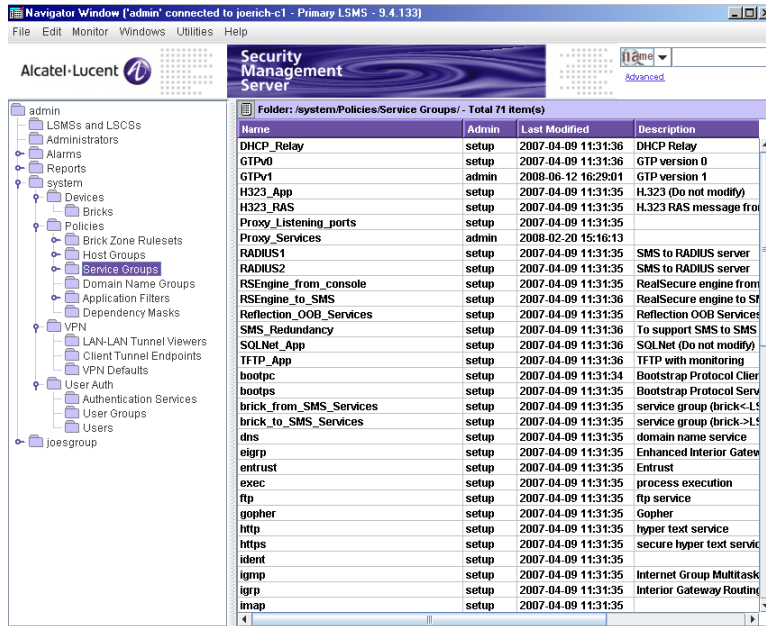
- 1 Open the appropriate group folder, and then open the Policy Components folder.
- 2 Click the Service Groups folder. All existing service groups, including those provided with the SMS application, will be displayed in the Navigator window (see [Figure 4-7, “Navigator Window \(View Service Groups\)”](#) (p. 4-14)).

For each service group, the Navigator window shows the Administrator who created the service group, the date and time the service group was created, and a brief description, if one was entered when the service group was created.

If the service group is one of the more than 60 provided with the SMS application, the word “setup” will appear in the Admin column. Use the scrollbar to view the service groups that do not fit in this window.

To view the contents of a particular service group, you have to display the service group in the Service Group Editor. See the section below entitled “[To modify a service group](#)” (p. 4-18) for instructions.

Figure 4-7 Navigator Window (View Service Groups)



END OF STEPS

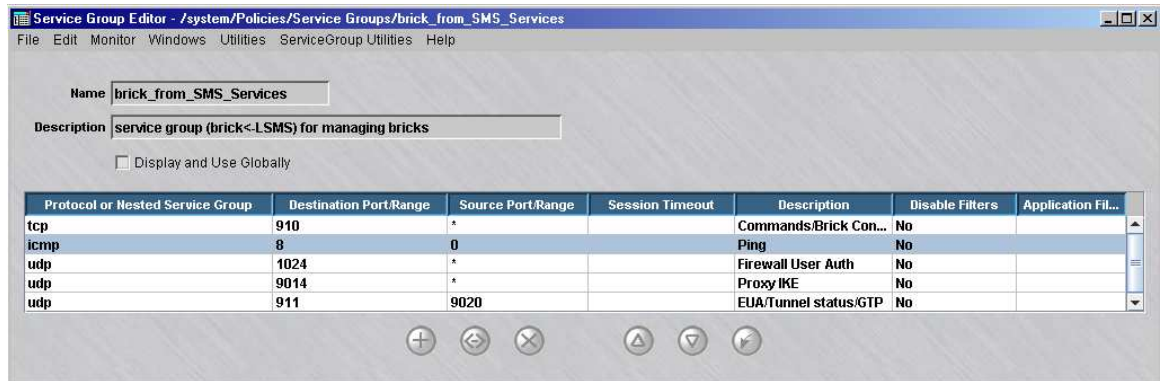
To view details of an existing service group

Complete the following steps to view the service details of an existing service group.

- 1 Double-click on a service group listed in the Contents panel of the Navigator.

Result A read-only view of the Service Group Editor is displayed, showing a tabular listing of the service details for the selected group (Figure 4-8, “Service Group Editor (Service Group Details)” (p. 4-15)).

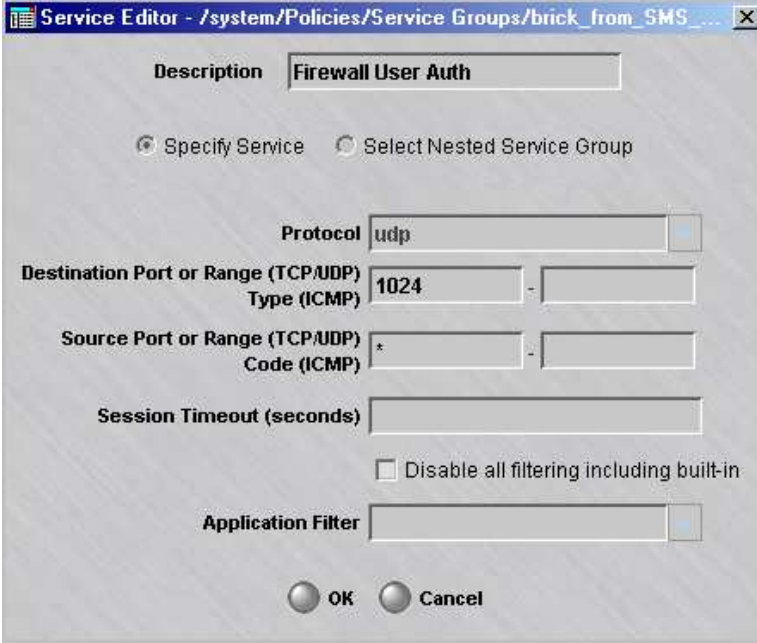
Figure 4-8 Service Group Editor (Service Group Details)



- 2 To view individual details of a protocol/service within a service group, right-click on the protocol in the **Protocol or Nested Service Group** column of the Service Group Editor and select **View** from the pop-up menu.

Result A read-only version of the Service Editor is displayed, showing the details of the selected protocol/service (Figure 4-9, “Service Editor (View Protocol/Service Details)” (p. 4-16)).

Figure 4-9 Service Editor (View Protocol/Service Details)



The screenshot shows a dialog box titled "Service Editor - /system/Policies/Service Groups/brick_from_SMS...". The "Description" field contains "Firewall User Auth". There are two radio buttons: "Specify Service" (selected) and "Select Nested Service Group". The "Protocol" field is set to "udp". The "Destination Port or Range (TCP/UDP)" field is set to "1024" and the "Type (ICMP)" field is empty. The "Source Port or Range (TCP/UDP)" field is set to "*" and the "Code (ICMP)" field is empty. The "Session Timeout (seconds)" field is empty. There is a checkbox labeled "Disable all filtering including built-in" which is unchecked. The "Application Filter" field is empty. At the bottom, there are "OK" and "Cancel" buttons.

END OF STEPS

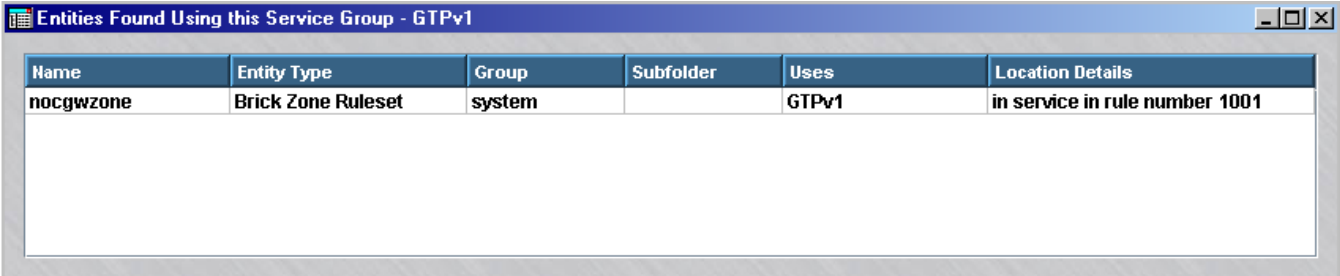
To find entities using a service group

The SMS application allows you to perform a search for all entities (such as a zone ruleset, application filter, or dependency mask) that are using the selected service group and, if it is applied within a parent entity, the entities that are using the parent entity.

- 1 With the service groups displayed in the Navigator window, right-click the desired service group and select **Find Entities Using this Service Group** from the pop-up menu.

Result An Entities Found Using this Service Group window is displayed (Figure 4-10, “Entities Found Using this Service Group Window” (p. 4-17) showing the search display results.

Figure 4-10 Entities Found Using this Service Group Window



Name	Entity Type	Group	Subfolder	Uses	Location Details
nocgwzone	Brick Zone Ruleset	system		GTPv1	in service in rule number 1001

The display results on the window show the entity type that uses or contains the selected service group, the group and subfolder in which the entity is found, whether the entity uses the service group or a parent entity, and details about the entity (if, for example, the entity is a Brick zone ruleset, the display lists the specific rules in which the service group is used).

To view additional details about the entity, right-click on the entry in the table and select **View** from the pop-up menu.

To edit the entity, right-click on the entry in the table and select **Edit** from the pop-up menu.

- 2 Click the Close (✕) button at the top of the window to close the window.

END OF STEPS

To modify a service group

You can modify any field in a service group except the **Name** field. Complete the following steps to modify a service group:

- 1 With the service groups displayed in the Navigator window, double-click the service group you want to modify.

The Service Group Editor will appear, with the **Name** field greyed-out and the other fields active, as shown in [Figure 4-11, “Service Group Editor \(Edit Mode\)”](#) (p. 4-18).

Figure 4-11 Service Group Editor (Edit Mode)



- 2 Do any of the following:
 - To modify a service, double-click the service in the Services panel. The Service Editor ([Figure 4-2, “Service Editor”](#) (p. 4-6)) will appear, with the fields populated. Make any necessary changes to the protocol, ports or description. Then, click **OK** to dismiss the Service Editor and return to the Service Group Editor.
 - To add a new service, right-click in the Services panel and select **New** from the pop-up menu. The Service Editor ([Figure 4-2, “Service Editor”](#) (p. 4-6)) will appear. Enter the information requested and click **OK** to dismiss the Service Editor and return to the Service Group Editor.
 - To delete a service, right-click the service and select **Delete** from the pop-up menu. A confirmation window will appear. Click **Yes** to confirm the deletion and return to the Service Group Editor.
 - To reorder the services, right-click a service and select **Up** or **Down** from the pop-up menu. Repeat until the services are in the correct order.
 - To import a service from another service group, right-click in the Services panel and select **Import Services**. Select the service group to be imported from the browse window that appears, and click **OK**.

Important! You can use the buttons at the bottom of the Service Group Editor instead of the right-click menu in each of the above actions, if you prefer.

-
- 3 Display the File menu and select one of the **Save** options.

END OF STEPS

To copy a service group

You can copy a service group to a different folder in the same group, or to a folder in another group, provided you have *Full* permission for that group. When you edit a service group (using the **Save As** option to save the edited version with a new name) or copy a service group, you have the option of copying components from the original service group (nested service groups, application filters), specifying how those components will be named and where they will be created and stored.

Complete the following steps to copy a service group:

-
- 1 With the service groups displayed in the Navigator window, right-click the service group you want to copy and select **Copy** from the pop-up menu.

Result A Copy window is displayed ([Figure 4-12, “Copy Window”](#) (p. 4-19)).

Figure 4-12 Copy Window



-
- 2 Complete the following fields, as applicable:
 - **Name**— Enter a new name for the service group copy (Example: **TFTP_Appx**). This field is required. The name entered must be different from the name of the original host group by at least one character if it is in the same group.
 - **Folder**— The copy of the service group will be created and stored in the **Service Groups** folder under the *system* group, by default. To specify a different folder, click the **Browse** button, which brings up a Select a Service Group Folder window, and select the folder where the service group copy will be created and stored.

- **Copy components when copying within a Group**— Click this checkbox (place a check in it) to copy the components from the original service group. If this checkbox is unchecked, and the copied service group is in the same group as the original one, the copied service group uses the same components as the original. If the service group is copied across groups, the components are copied regardless of this field setting.
- **Component Suffix**— If the **Copy components when copying within a Group** option is checked, enter a text string that will be appended to the name of each copied component (nested service group, application filter). If this field is left blank, the name of each copied component will be spelled the same way as the original one. Name collisions are possible, regardless of whether or not an entry is made in this field. If this occurs, the content of the conflicting component(s) is examined. If the content is the same, the duplicated service group uses the existing component(s) and the new copy(ies) is not made. If the content is different, a copy of the component(s) is created and stored under a unique name by appending its database ID to the name.
Note: the Component Suffix used must be unique and cannot be reused across zone rulesets.
- **Component Folder**— This field is used to specify a common folder name for the copied components. The folder name is relative to the top-level folder of each type of component. If this field is left blank, the copied components are stored in the top-level folder of that component. This folder is created as needed (if it does not already exist). A path string can be specified; / or \ delimiters can be used to delineate folder names.

3 Click **OK** to copy the service group and dismiss the Copy window.

Result The system returns to the Navigator window. The copy of the service group appears in the Contents panel.

The service group (and components, if specified) is copied to the chosen folder(s).

END OF STEPS

To move a service group

You can move a service group to a different folder in the same group, or to a folder in another group, provided you have *Full* permission for that folder. However, you can only move it to the folder labeled "service Groups" or to a subfolder under that folder. If you try to move it to another folder, you will get an error message. Before moving a service group, make sure that it is not currently in use. If you attempt to move a service group that is in use, you will get an error message. If you receive an error message when attempting to move a service group, you can run the "find entity" utility to discover where it is in use (see ["To find entities using a service group"](#) (p. 4-17))

above). If a service group is moved to a new group, and the service group has other service groups nested within it, the original service group is moved to the new group, but the service groups nested within it are copied to the new group, not moved. If a service group with the same name and content already exists in the new group, the copy does not take place. If a service group with the same name but different content already exists, the copy is renamed.

To move a service group, follow the steps below:

- 1 With the service groups displayed in the Navigator window, right-click the service group you want to move and select **Move** from the pop-up menu. A Browse window will appear.
- 2 Select the folder you want to move this service group to. The service group will be moved to the folder you selected, and you will be returned to the Navigator window.

END OF STEPS

To delete a service group

Before deleting a service group, make sure that it is not currently in use. If you attempt to delete a service group that is in use, you will get an error message. If you receive an error message when attempting to delete a service group, you can run the "find entity" utility to discover where it is in use (see ["To find entities using a service group"](#) (p. 4-17) above).

To delete a service group, follow the steps below:

- 1 With the service groups displayed in the Navigator window, right-click the service group you want to delete and select **Delete** from the pop-up menu.

A confirmation window similar to the one shown in [Figure 4-13, "Confirmation Window \(Service Groups\)"](#) (p. 4-21) will appear.

Figure 4-13 Confirmation Window (Service Groups)



-
- 2 Click **Yes** to delete the service group and dismiss the Confirmation window. The service group will be removed from the Navigator window.

END OF STEPS



Service Groups Provided with the SMS Application

Overview

More than 60 service groups are automatically created when the SMS application is installed. The service groups are described in the table below. The service groups are listed in alphabetical order, with the service groups that begin with uppercase letters listed first.

Two of the service groups — **H.323_App** and **SQLNET_App** — have entries in the Application-Layer Monitoring field.

Service Group Name	Description	Protocol	Dest Port	Source Port
DHCP_Relay	DHCP Relay	UDP	67-68	67-68
GTPv0	GTP version 0	UDP	3386	*
GTPv1	GTP version 1	UDP	2123	*
		UDP	2152	*
H.323_App	H.323 (do not modify)	TCP	1720	*
H.323_RAS	H.323 RAS message from endpoint to gatekeeper	UDP	1719	*
Proxy_Listening_ports		ICMP	8	0
RADIUS1	SMS to RADIUS server	UDP	1645-1646	*
RADIUS2	SMS to RADIUS server	UDP	1812-1813	*
RSEngine_from_console	RealSecure engine from console	TCP	25	*
		UDP	162	*
		TCP	2998	*
		TCP	901	*
RSEngine_to_SMS	RealSecure engine to SMS	TCP	9005	*
Reflection_OOB_Services	Reflection OOB services for the brick proxy	UDP	1024	*
SQLNet_App	SQLNet (do not modify)	TCP	1521	*
TFTP_App	TFTP w/monitoring	UDP	69	*
bootpc	Bootstrap Protocol Client	UDP	68	*

Service Group Name	Description	Protocol	Dest Port	Source Port
bootps	Bootstrap Protocol Server	UDP	67	*
Brick_from_SMS_services	service group (brick<-SMS) for managing Brick devices	TCP	910	*
		ICMP	8	0
		UDP	911	9020
		UDP	1024	*
		UDP	9014	*
Brick_to_SMS_Services	service group (brick->SMS) for managing Brick devices	TCP	9000	*
		UDP	9014	*
		UDP	9020	911
		TCP	900	*
dns	domain name service	UDP	53	*
		TCP	53	*
eigrp	Enhanced Interior Gateway Routing Protocol (EIGRP)	88	*	*
entrust	Entrust	TCP	389	*
		TCP	709	*
		TCP	829	*
exec	process execution	TCP	512	*
ftp	ftp service	TCP	21	*
gopher	Gopher	TCP	70	*
http	Hyper text service	TCP	80	*
https	secure hyper text service	TCP	443	*
ident		TCP	113	*
igmp	Internet Group Multitasking Protocol (IGMP)	2	*	*
igrp	Internet Gateway Routing Protocol (IGRP)	9	*	*
imap		TCP	143	*
ipsec	IPSEC (ESP only)	51	*	*
		50	*	*

Service Group Name	Description	Protocol	Dest Port	Source Port
irc	Internet Relay Chat protocol	TCP	194	*
		UDP	194	*
kerberos	Kerberos	TCP	88	*
ldap	Lightweight Directory Access Protocol	TCP	389	*
login	Login Host Protocol	TCP	49	*
netbios-gm	NetBIOS Datagram Service	UDP	138	*
netbios-ns	NetBIOS Name Service	UDP	137	*
netbios-ssn	NetBIOS Session Service	TCP	139	*
netstat	Netstat	TCP	15	*
nntp	Network News Transfer Protocol	TCP	119	*
ntp	Network Time Protocol	UDP	123	*
ospf	Open Shortest Path First (OSPF)	89	*	*
ping_request	Ping service	ICMP	8	0
ping_resp	Ping service	ICMP	0	0
pop3		TCP	110	*
realaudio_control	Real Audio Control	TCP	7070	*
realaudio_data	Real Audio Data	UDP	6970-7170	*
rip/rip2	Routing Information Protocol (RIP/RIP2)	UDP	520	520
rlogin	Remote Login Service	TCP	513	*
router_from_SMS_other	service group (udp & icmp) (router <-SMS) for managing router	UDP	9014	*
		UDP	69	
router_from_SMS_tcp	service group (router <-SMS) for managing router	TCP	9017-9018	*

Service Group Name	Description	Protocol	Dest Port	Source Port
router_to_SMS_other	service group (udp & icmp) (router ->SMS) for managing router	ICMP	*	*
		UDP	69	*
		UDP	9014	*
router_to_SMS_tcp	Service group (tcp) (router -> SMS) for managing router	TCP	9017-9018	*
		TCP	23	*
rsh	Remote shell	TCP	514	*
securID	SMS to ACE Server	UDP	5500	*
secure_remote_admin_from_SMS	Remote administration from SMS to Client	TCP	*	7000
secure_remote_admin_to_SMS	Remote administration from Client to SMS	TCP	443	*
		TCP	7000	*
		TCP	9041	*
sip	Session Initiation Protocol	TCP	5060	*
		UDP	5060	*
smtp	SMTP	TCP	25	*
snmp	SNMP Request	UDP	161	*
snmp_trap	SNMP Trap	UDP	162	*
ssh	Secure Shell	TCP	22	*
syslog	Syslog	UDP	514	*
telnet	telnet service	TCP	23	*
traceroute_icmp	icmp traceroute ports	ICMP	8	0
traceroute_udp	udp traceroute ports	UDP	32000 - 53000	*
userAuth	VBA to SMS for user authentication	TCP	9010-9011	*
whois		TCP	63	*
x11	X-Window System	TCP	6000-6063	*
		UDP	177	*



Service Group Templates Provided with the SMS Application for VoIP/NOE Traffic

Overview

A set of pre-defined service group templates are provided with the SMS application when it is installed to make it easier to configure the Brick to handle VoIP/NOE traffic. These service group templates are automatically pre-selected in rules that have been set up in pre-defined Brick zone rulesets that can be assigned to a port on a Brick that is protecting IP phones, call servers, or a Media Gateway (MGW) unit in a portion of the VoIP network. Some of these pre-defined service groups call one of the pre-defined application filters that are also provided with the SMS application when it is installed that have already been properly configured to check call and data traffic from the equipment at a specific location.

The name of the pre-defined service group template indicates the type of equipment for which the services are being checked and/or the deployment location of the Brick serving as an ALG, and whether the Brick is also performing NAT for MGW-CS IPlink protocol messaging in that location. For example, the pre-defined service group **NOE_Service-NAT_Branch** is already configured in a rule to inspect UDP traffic between IP phone devices located behind the Brick zone at a Branch office site and call servers. This pre-defined service group calls the pre-defined application filter **TFTP_with_NOE_NAT_Branch** that has all of the proper field settings already defined and the IP addresses of the IP phone devices pre-selected in a host group. It is also set up to allow the Brick to perform NAT for IPlink protocol messaging between the MGW at a Branch office site and call servers.

Related information

Review the following sections in this Guide for related information on use of the pre-defined templates provided with the SMS application for VoIP traffic:

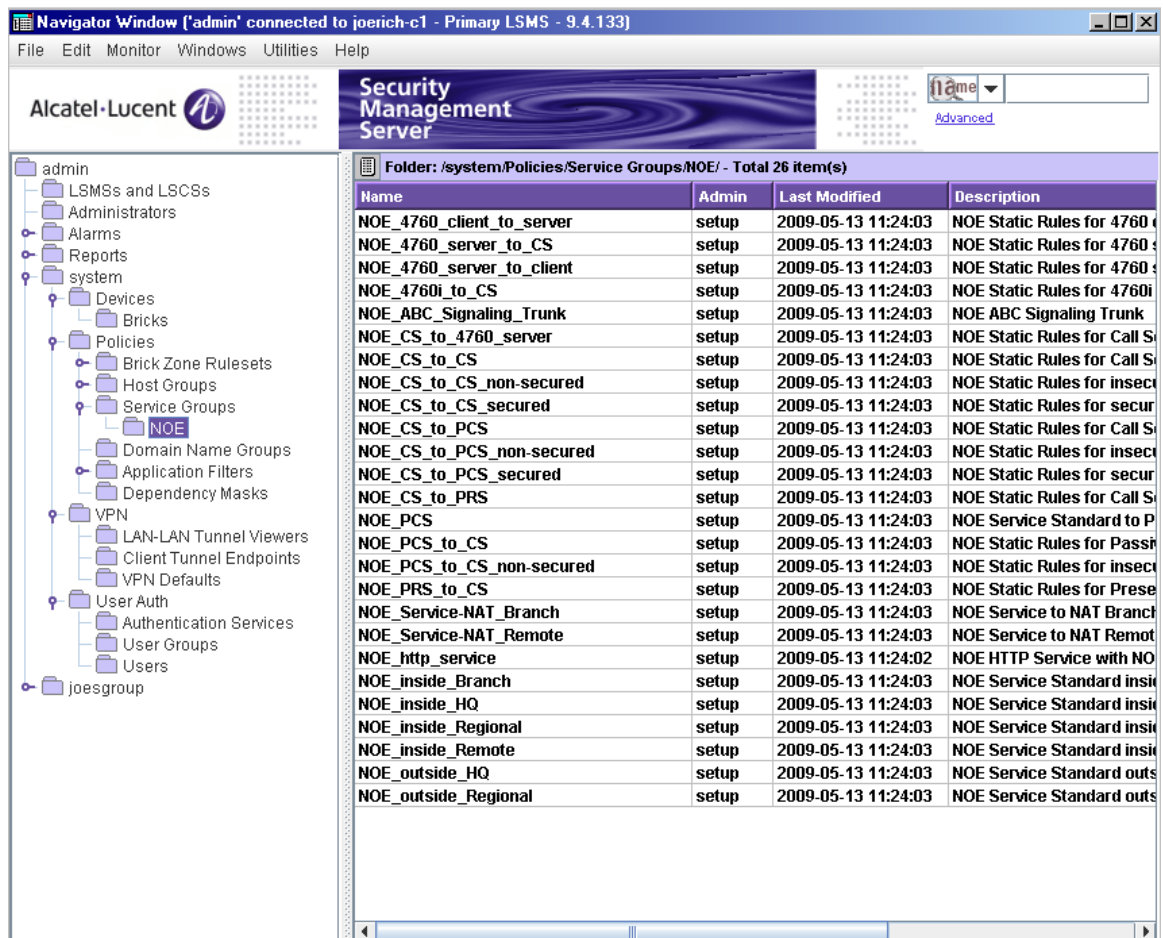
- [“Brick Zone Ruleset Templates Provided with the SMS Application for VoIP/NOE Traffic” \(p. 1-66\)](#)
- [“Host Group Templates Provided with the SMS Application for VoIP/NOE Traffic” \(p. 2-25\)](#)
- [“TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic” \(p. 5-123\)](#)
- [“NOE Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic” \(p. 5-68\)](#)
- [Appendix E, “Configuring the Brick for VoIP/NOE Traffic Using Pre-Defined SMS Templates”](#)

Downloading the NOE service group templates

After the SMS application is installed, the NOE templates can be downloaded by right-clicking on a group folder (such as the *system* folder initially created when the SMS is installed) and choosing **Create NOE Template** from the pop-up menu.

The pre-defined service group templates are automatically downloaded to a subfolder named **NOE** under the **Service Groups** folder in the SMS Navigator (refer to Figure 4-14, “Pre-Defined NOE Service Group Templates (NOE subfolder in Navigator)” (p. 4-28)).

Figure 4-14 Pre-Defined NOE Service Group Templates (NOE subfolder in Navigator)



To remove the NOE zone ruleset templates and all component templates from the Policies folder, right-click on the group folder and choose **Delete NOE Template**. A confirmation dialog box is displayed, asking if you are sure that you want to delete the NOE policy template (which includes all of the zone ruleset templates are component templates) under the selected group folder. Choose **Yes**. The NOE templates are

removed from the Policies folder under the selected group. If, however, one of zone ruleset templates is in use on one or more Bricks, an error message is generated and the templates are not removed.

Summary description of pre-defined service group templates for VoIP/NOE traffic

Pre-defined Service Group Name	Protocol	Dest Port	Source Port	Pre-defined Application Filter
NOE_4760_client_to_47640_server (contains nested service group http_service)	UDP	88	*	None
	UDP	500	*	
	UDP	30014	*	
	TCP	3389	*	
	TCP	30014	30019	
	TCP	30022	30025	
	TCP	30100-30149	*	
NOE_4760_server_to_4760_client	TCP	30500-30509	*	None
NOE_4760_server_to_CS (contains nested service groups ping_request,ftp,telnet,ssh)	TCP	2535	*	None
	UDP	2556	*	
NOE_4760i_to_CS (contains nested service groups ftp,telnet,ssh,http_service,https)	TCP	5540	*	None
NOE_CS_to_4760_server	LDAP			None
NOE_ABC_Signaling_Trunk	UDP	2571	2571	None
NOE_CS_to_CS (contains nested service groups ping_request,sip,H323_RAS)	UDP	32640	32640	None
	UDP	2556	2556	
	UDP	2562-2564	*	
	UDP	2567	*	
	TCP	113	*	
	TCP	513-1023	*	
	TCP	2562	*	
	TCP	2535	*	
	TCP	2555	*	

Pre-defined Service Group Name	Protocol	Dest Port	Source Port	Pre-defined Application Filter
NOE_CS_to_CS_non_secured (contains nested service group ftp)	TCP	113	*	None
	TCP	513-1023	*	
NOE_CS_to_CS_secured (contains the nested service group ssh)	SSH			None
NOE_CS_to_PCS (contains the nested service group ssh)	UDP	32640	3260	None
NOE_CS_to_PCS_non-secured	TCP	113	*	None
	TCP	513-1023	*	
NOE_CS_to_PCS_secured (contains the nested service group ssh)	SSH			None
NOE_CS_to_PRS	UDP	2000	2570	None
NOE_Service-NAT_Branch	UDP	69	*	TFTP_with_NOE_NAT_Branch
NOE_Service-NAT_Remote	UDP	69	*	TFTP_with_NOE_NAT_Remote
NOE_inside_Branch	UDP	69	*	TFTP_with_NOE_inside_Branch
NOE_inside_HQ	UDP	69	*	TFTP_with_NOE_inside_HQ
NOE_inside_Regional	UDP	69	*	TFTP_with_NOE_inside_Regional
NOE_inside_Remote	UDP	69	*	TFTP_with_NOE_inside_Remote
NOE_outside_HQ	UDP	69	*	TFTP_with_NOE_outside_HQ
NOE_outside_Regional	UDP	69	*	TFTP_with_NOE_outside_Regional
NOE_PCS	UDP	69	*	TFTP_with_NOE_PCS
NOE_PCS_to_CS (contains nested service group TFTP_App)	TCP	514	*	None
NOE_PCS_to_CS_non-secured	TCP	113	*	None
	TCP	513-1023	*	

Pre-defined Service Group Name	Protocol	Dest Port	Source Port	Pre-defined Application Filter
NOE_PRS_to_CS	UDP	2570	2000	None
NOE_http_service	TCP	80	*	NOE_HTTP



5 Application Filters

Overview

Purpose

This chapter explains how to use application filters. Its purpose is to allow additional application layer validation, inspection and access control directly on the Alcatel-Lucent *VPN Firewall Brick™* Security Appliance.

After defining an application filter, the filter may need to be added to a new or existing service group and then incorporated into a Brick zone ruleset. In this release, the administrator may define application filters for the following scenarios:

- DHCP Relay - A DHCP client request for an IP address can be passed through a Brick to the appropriate DHCP server.
- DNS - Domain Name Service security filters.
- ESP - Allows ESP NULL encapsulated data through the Brick and validates the proper format of the packets.
- GTP - General Packet Radio Service Tunneling Protocol packet filtering through a Brick.
- H.323 VoIP - Used for Voice over IP, can be used with H.245, RTP, or RTCP.
- HTTP - User defined URI strings can be blocked and much more.
A pre-defined HTTP application filter template is provided with the SMS application when it is installed to filter HTTP URLs passed between IP phone devices and a presentation server in a VoIP network. For additional details, refer to the section [“Pre-defined HTTP application filter for VoIP/NOE traffic”](#) (p. 5-60).
- NOE - Application filter configured specifically to inspect and validate voice and signalling traffic between IP phone devices and other network elements (call servers, Media Gateway) in a VoIP call network. Can also be configured on the Brick to decode and protect IPlink signalling traffic between the Media Gateway and call server(s) in a VoIP call network. An NOE application filter is created and

used within a TFTP application filter.

A set of pre-defined NOE application filter templates are provided with the SMS application when it is installed to make it easier to configure the Brick to handle VoIP/NOE traffic. For details, refer to the section [“NOE Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-68).

- RPC - The Brick can selectively permit particular RPC calls and can dynamically open ports through the Brick, based upon the results of the port mapper call.
- RTSP - Real Time Streaming Protocol (RTSP) application filter secures communications between RTSP clients and multimedia streaming servers. Restricts valid port ranges and source/destination addresses for dynamic opening of the exact RTP/RTCP pinhole by the Brick device for RTSP media streams, prevents filename-based attacks by inspecting RTSP commands in client/server exchanges, and protects against buffer overflow attacks.
- SIP - Session Initiation Protocol (SIP) specific packet filtering through the Brick.
- FTP - File Transfer Protocol filtering to prevent unauthorized ftp server connections at the packet and port level.
- SMTP - Simple Mail Transfer Protocol filtering of e-mail provides protection against various forms of attacks.
- SQL*Net - Oracle Corporation’s remote data access software that allows databases and their applications residing on different computers to communicate as peer applications. When this application filter is applied to a service group on a rule, the Brick looks at the appropriate string(s) and determines the exact port (pinhole) for relaying messages between client and server
- TFTP - Trivial File Transfer Protocol filtering of file transfers from a server by the Brick.

A set of pre-defined TFTP application filter templates are provided with the SMS application when it is installed to make it easier to configure the Brick to handle VoIP/NOE traffic. For details, refer to the section [“TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-123).

The FTP and TFTP application filters support the use of ports other than the default TCP port 21 for FTP and UDP port 69 for TFTP.

Contents

Process Flow	5-4
DHCP Relay Application Filter	5-5
DNS Application Filters	5-8
ESP Application Filter	5-18
FTP Application Filter	5-21
GTP Application Filter	5-31

H.323 Application Filters	5-47
HTTP Application Filter	5-53
NOE Application Filter	5-62
NOE Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic	5-68
RPC Application Filters	5-74
RTSP Application Filter	5-79
SIP Application Filters	5-87
SMTP Application Filters	5-105
SQL*Net Application Filter	5-114
TFTP Application Filter	5-117
TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic	5-123
Global Application Filters	5-127
To Copy an Application Filter	5-129



Process Flow

Application filter processing

Up to three steps are required to activate an application filter:

- Create the application filter
- If necessary, add the filter to a new or existing service group.
Most of the application filters available in this release are already included into "canned" service groups with a predefined protocol and port number.
- After you have defined or customized your filter, simply include the service group in a rule in a Brick Zone Ruleset and "apply" that rule update to the Brick.



DHCP Relay Application Filter

Overview

When a DHCP client wishes to obtain an IP address from a DHCP server, it broadcasts a request on the local LAN. If there is a DHCP server on the same LAN, then it responds to the request.

However, if a DHCP server is outside the LANs protected by a Brick, the client request needs to be forwarded through the Brick to the server. This filter enables the DHCP client request to be forwarded by the Brick to a server located in another network. The request can be passed through the Brick in the clear or through a LAN - LAN tunnel.

To provide DHCP relay in the clear (not through a tunnel), the Brick zone policy ruleset closest to the client must be created with a rule passing a service group with a "DHCP_Relay" application filter assigned to it. A DHCP relay in the clear only works if the Virtual Brick Address (VBA) is set to the interface address.


In the case of a LAN-LAN tunnel in which the DHCP client is behind one Brick and the DHCP server is behind the other Brick—and there is more than one zone on the client-side virtual interface (physical port + VLAN)—make sure the "Route multicast packets to first matching zone" checkbox is checked and the client zone is assigned the 0.0.0.0 address, in addition to any other valid zone address in the Policy Assignment tab.

Brick to non-Brick tunnels must follow the same client-side procedures as documented above but also ensure that the Brick interface IP address is in the client-side "Hosts behind Tunnel" list.

Task

Complete the following steps to configure this application filter.

- 1 From the SMS Navigator, open the Policies folder in the desired group, and click **Application Filters**.
- 2 In the right-hand column, double-click the DHCP_Relay filter. The window shown in [Figure 5-1, "DHCP Application Filter Editor"](#) (p. 5-6) is displayed.

Figure 5-1 DHCP Application Filter Editor

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. See [“RPC Application Filters” \(p. 5-74\)](#) below for an explanation of global application filters.

-
- 3** In the **DHCP Server IP Address** field, enter the IP address of the DHCP server.
Optionally, you can specify IP addresses for two DHCP servers, separated by a comma. If two DHCP servers are specified, the Brick sends all broadcast DHCP requests to both servers. This supports redundancy, and hence, greater reliability.
-
- 4** Open the File menu and select one of the **Save** options.

-
- 5 Integrate the DHCP_Relay service group into a rule in a Brick zone ruleset. Then, save and apply the ruleset change to the desired Brick.

END OF STEPS



DNS Application Filters

Definition

Domain Name Service (DNS) is a fundamental component of the network, trusted implicitly by all applications but with no (deployed) authentication. Thus it is crucial that outside data not be able to contaminate internal DNS databases, while still providing full information about outside machines.

The DNS application filter is designed to avoid a DNS spoof attack. The Brick device will ensure that an outside DNS server will not be able to obtain the address (or some other attribute) of an inside host. The filter will check the correctness of each DNS packet and block unwanted information selectively. This filter will also prevent many attacks, such as DNS poisoning and transaction ID prediction.

Configure a DNS Application Filter

A DNS application filter called `dnsDefault` is created automatically in the system group for a newly installed SMS, or for each pre-existing group for an upgraded SMS. For a newly installed SMS and for newly created groups, `dnsDefault` is additionally assigned to the UDP entry of the `dns` service group. You may decide just to further configure the `dnsDefault` application filter or create a new one. Note: The initial settings of a newly created DNS application filter are identical to the `dnsDefault` application filter "out of the box."

To create a new DNS application filter, follow the steps below:

- 1 From the SMS Navigator, open the Policies folder in the desired group, and click **Application Filters**.
- 2 In the right-hand column, double-click the `dnsDefault` filter. The DNS application filter editor screen shown in [Figure 5-2, "DNS Application Filter Editor \(Protected Domain Names Tab\)"](#) (p. 5-9) will appear.

Figure 5-2 DNS Application Filter Editor (Protected Domain Names Tab)

The screenshot shows the 'Application Filter Editor' window for a filter named 'dnsDefault'. The 'Protected Domain Names' tab is active, displaying several configuration options:

- Name:** dnsDefault
- Description:** (empty text field)
- Display and Use Globally
- Type:** DNS
- Audit DNS Commands (Detailed Session Audit)
- Audit DNS Exceptions (Basic Exception Audit)
- Protected Domain Names:**
 - Allow Root Node Queries
 - Protected Domain Name Group:**
 - Only Allow Domain Names in this Group
 - Only Disallow Domain Names in this Group
 - Ignore Domain Names in this Group
 - Disallow Queries that return an Address in this Host Group

- 3 In the **Name** and **Description** fields, enter a name (default is dnsDefault) for the application filter and a brief description. The name is required, but the description is optional.

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to [“Global Application Filters”](#) (p. 5-127) for an explanation of global application filters.

4 The editor screen also contains the following checkboxes:

- **Audit DNS Commands (Detailed Session Audit)**

If this option is checked and if the "Session Audit" parameter in a rule is set to "Detailed", the Session Log will record and provide details for all DNS sessions that use this filter. If unchecked, the session log does not provide detailed data.

- **Audit DNS Exceptions (Basic Exception Audit)**

If this option is checked and the "Exception Audit" parameter in a rule is set to a value other than "None", the Session log will provide details on blocked DNS sessions that use this filter.

5 The **Protected Domain Names** tab contains these fields:

- **Allow Root Node Queries**

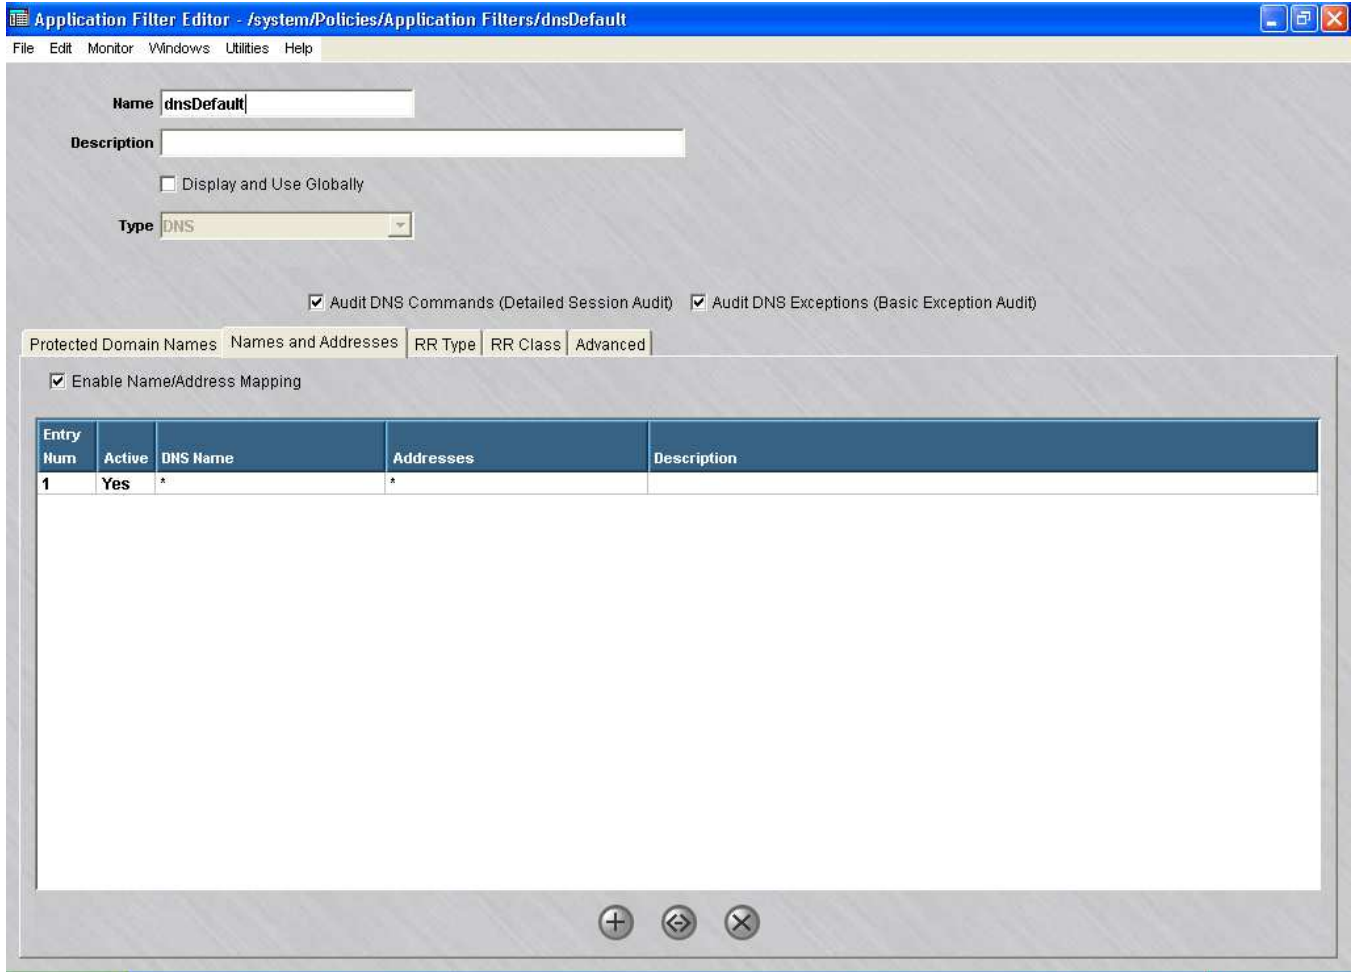
If this check box is checked, queries of the root node of the DNS database are permitted, otherwise they are blocked.

- **Protected Domain Name Group**

The DNS application filter can be configured so that only certain domain names will be allowed to have DNS queries run on them or it can be configured in a negative sense, whereby all DNS queries are permitted except for certain domain names. The filter will try to match domain names against the **Name** field of a RR or any other field inside a RR where a domain name may appear. These domain names are defined by domain name groups which must be created before they will show up as selections in the pull-down (see Chapter 5 - *Domain Names*). Choosing the option to Ignore Domain Names in this Group effectively disables this feature, even if a domain name group has been selected.

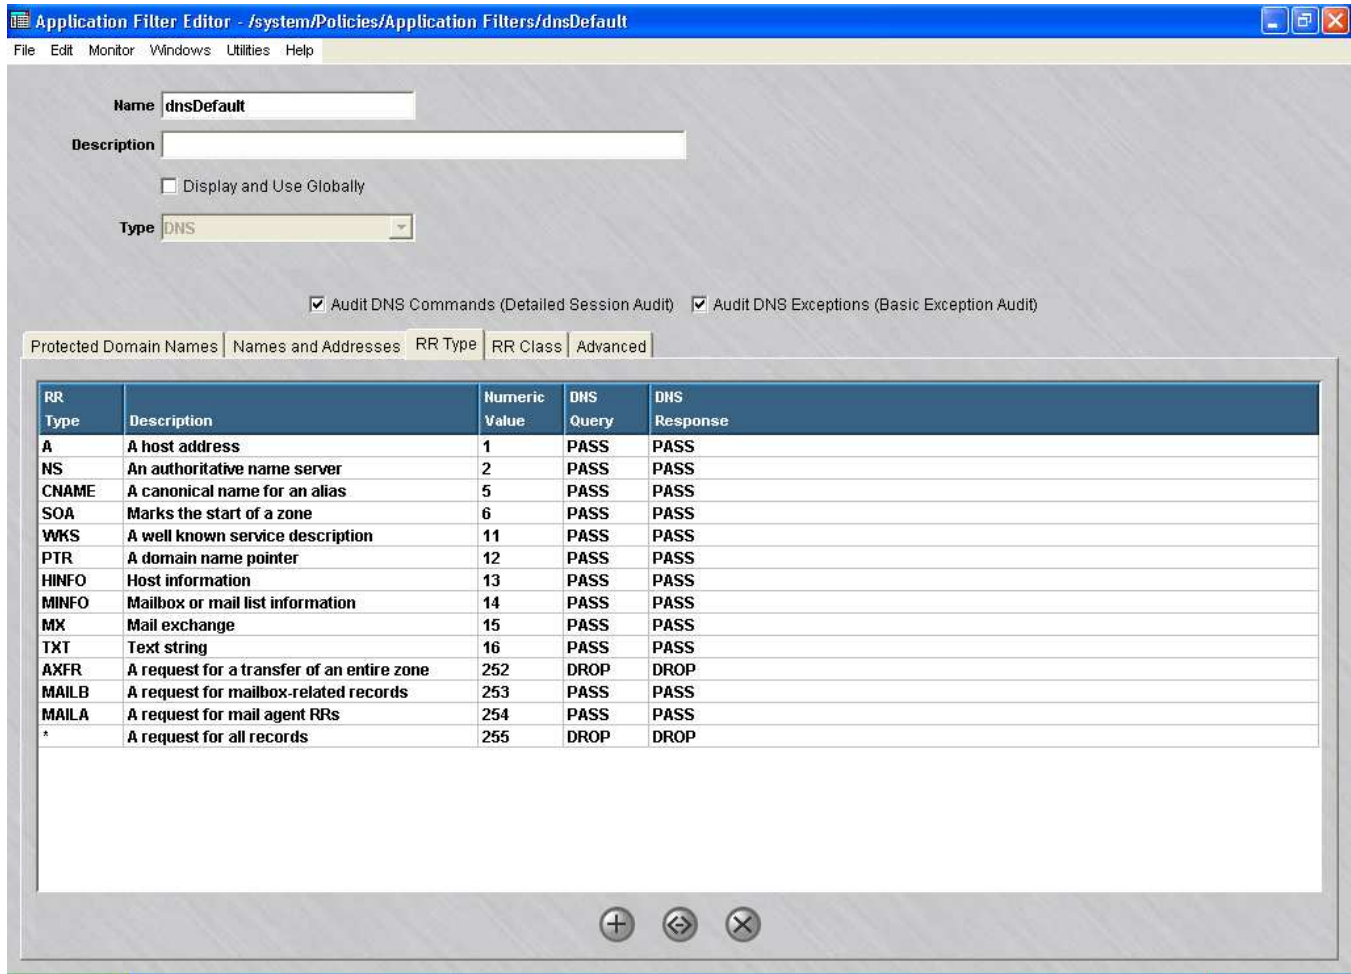
6 The **Names and Addresses** tab (see [Figure 5-3, "DNS Application Filter Editor Screen \(Names and Addresses Tab\)"](#) (p. 5-11)) is used to enter the name and address of the domain. The responses to queries received by the DNS application filter will be compared to all entries in the names and address table. For every matching name, the address must fall within the specified range. For every matching address, the name must fall within the specified range. If either of these conditions are not met, the response will be blocked. If there are no matches with either the name or address, the response will be passed. Names can be Fully Qualified Domain Names or a range of names such as *.alcatel-lucent.com. The address can be ranges of IPs or a specific IP. Host groups can be used in the usual manner.

Figure 5-3 DNS Application Filter Editor Screen (Names and Addresses Tab)



- 7 The **RR Type** tab shows the table in [Figure 5-4, “DNS Application Filter Editor Screen \(RR Type Tab\)”](#) (p. 5-12):

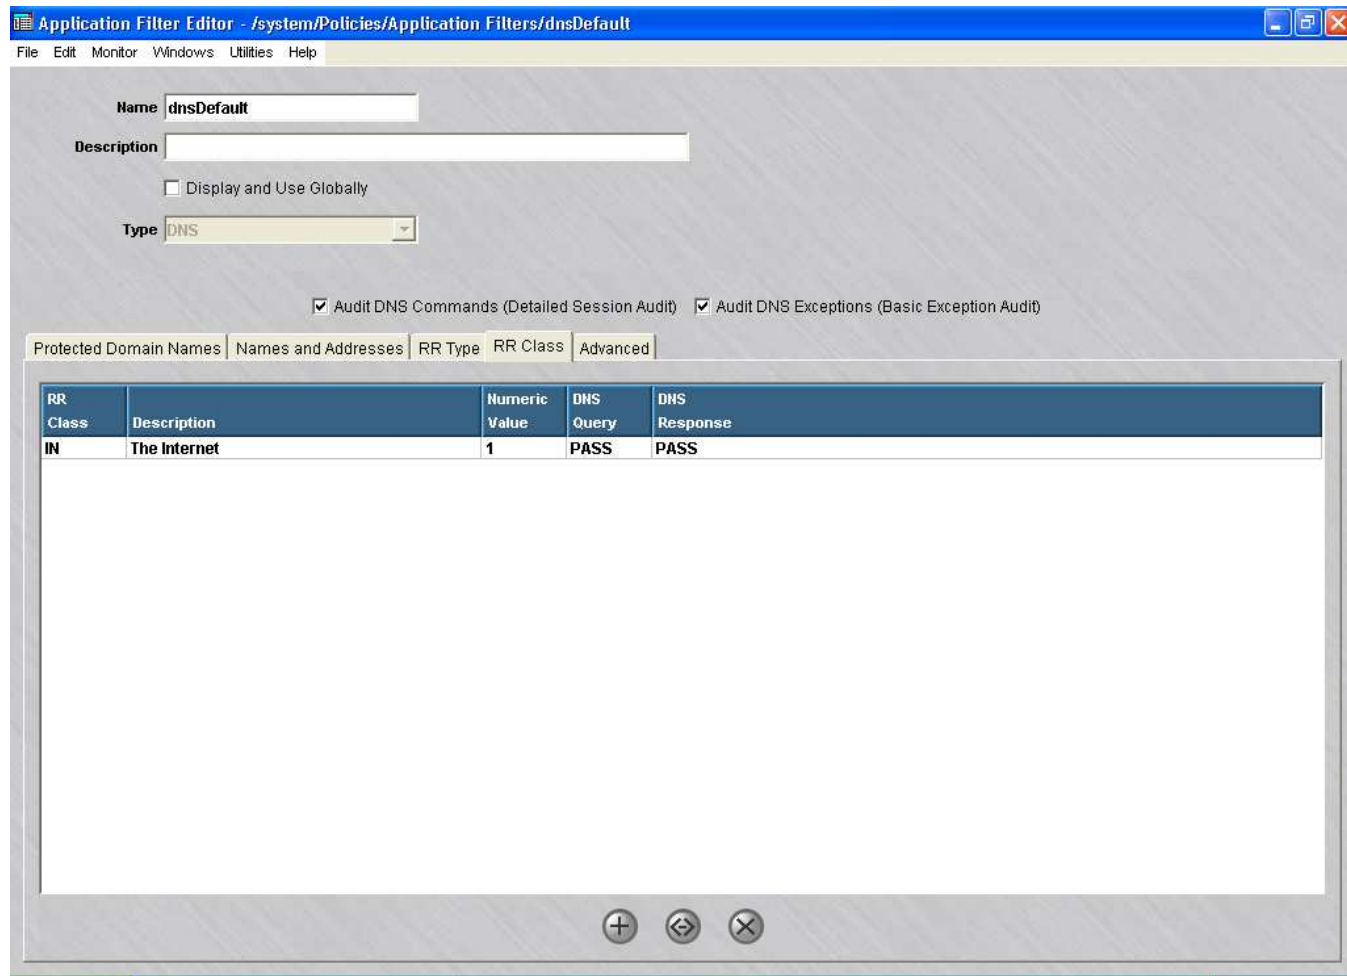
Figure 5-4 DNS Application Filter Editor Screen (RR Type Tab)



This part of the DNS application filter lets the administrator configure how resource record (RR) types are handled. Both the DNS query and response for each RR type can be passed or dropped. Entries may be added to the table (for new RR types) by right-clicking within the table and selecting New or deleted by choosing Delete. Entries can be sorted in various ways by left-clicking on a column header. Note that if a query is dropped, then the setting for its response is ignored.

8 The **RR Class** tab shows this table:

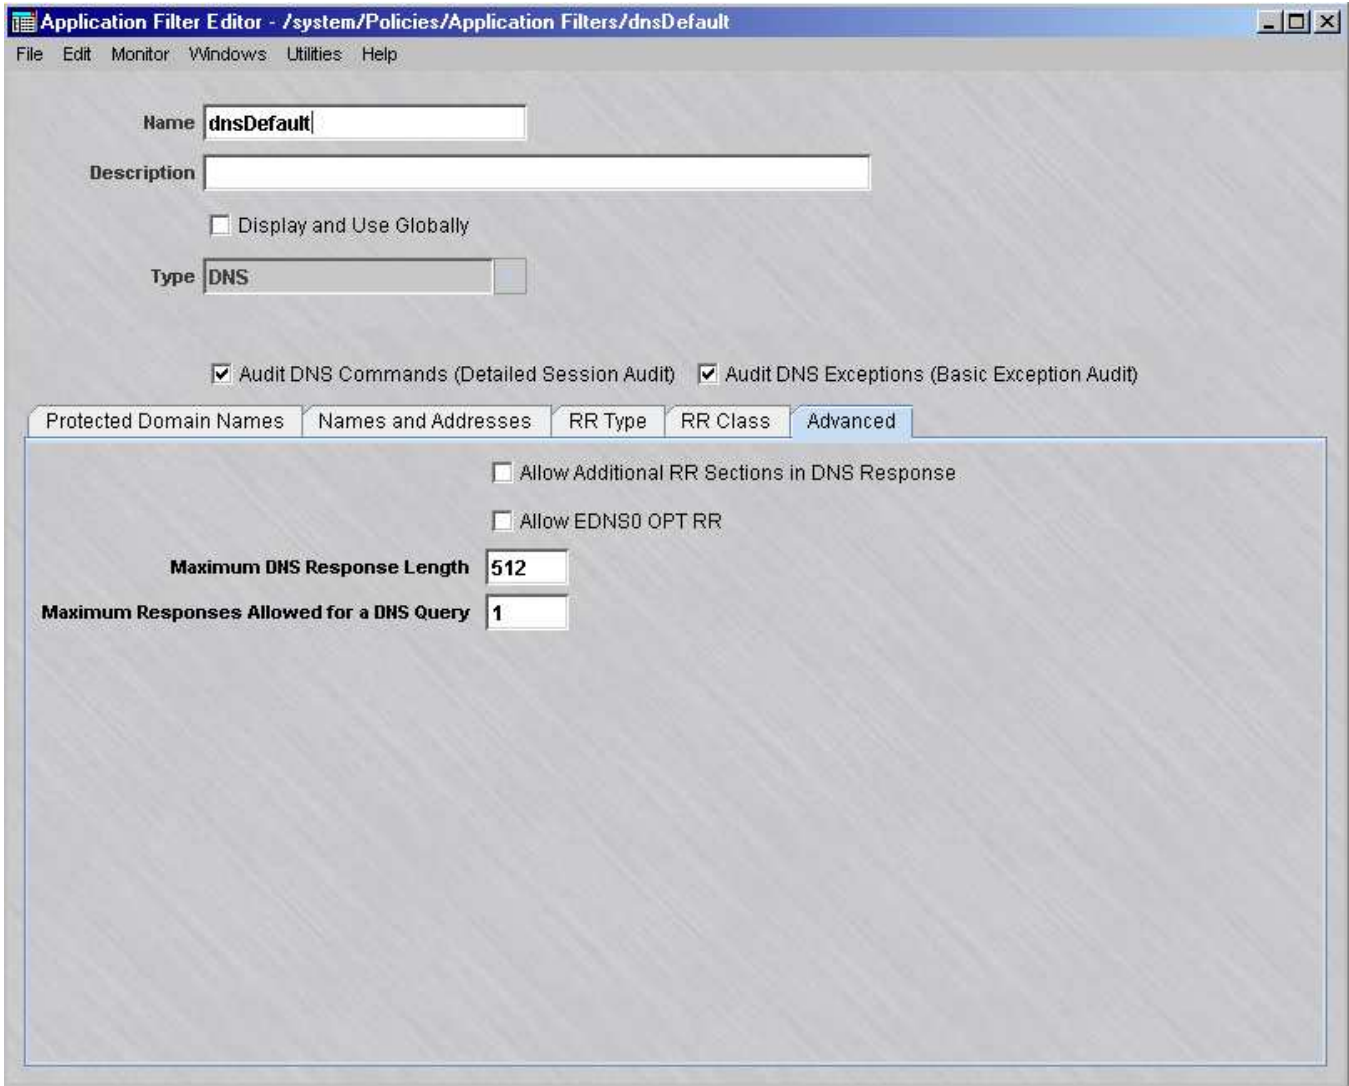
Figure 5-5 DNS Application Filter Editor (RR Class Tab)



This part of the DNS application filter lets the administrator configure resource record (RR) classes. Editing of this table is accomplished the same as the RR Type tab above.

- 9 The **Advanced** tab contains these fields:

Figure 5-6 DNS Application Filter Editor (Advanced Tab)



This tab contains the following fields:

- **Allow Additional RR Sections in DNS Response**
Additional RR sections, which appear at the end of some DNS responses, have been used in DNS spoof attacks to provide unsolicited information that a client may cache. If this option is not checked, all additional RR sections will be filtered out (discarded) by the Brick.
- **Allow EDNS0 OPT RR**
OPT pseudo-RR as defined by EDNS0 (RFC2671) is primarily used to facilitate establishment of a larger UDP payload size for DNS packet transfer. When this option is enabled (checkbox is checked), the DNS application filter will allow DNS Query and Response packets to have one OPT pseudo-RR included in the additional RR section. The DNS application filter checks the correctness on various fields of the OPT pseudo-RR before passing the packet.
- **Maximum DNS Response Length**
This size limits the maximum number of bytes permitted in a single DNS response.
- **Maximum Responses Allowed for a DNS Query**
This count limits how many DNS responses are allowed for a single DNS query, which is usually one.

-
- 10** After you have entered all the necessary data, open the **File** menu and select one of the **Save** options.

.....
 END OF STEPS

Features and Behaviors

-
- 1** The following is pertinent information regarding other DNS application filter features and behaviors:
- For DNS query packets, only the combination of a single query (QDCONT=1) and zero counts for all the other sections (ANCOUNT=0, NSCOUNT=0, ARCOUNT=0) are permitted. This prevents unsolicited information from being piggy-backed on legitimate queries.
 - The Brick ensures that the ID of a DNS response matches the ID of an outstanding DNS request, otherwise it drops the response.

- The Brick replaces the ID field of DNS queries with a randomly generated ID to avoid any predictable ID sequence. It restores the original ID before passing it back to the client.
- In order to pass as much legitimate information as possible, the Brick will not drop the entire response packet unless the answering section fails to pass the filter criteria. An RR in any other section of a DNS response will be removed from the packet if it fails to pass the filter criteria. An exception audit will be logged for each removed RR under this scenario.

Important! The filter only works on DNS over the UDP protocol. Most DNS queries and responses use UDP; however, zone transfer uses the TCP protocol. In order to block zone transfer, you must either delete the entry in the default DNS service group with TCP protocol on port 53 or not define such an entry in a newly created DNS service group.

.....
 END OF STEPS

To add the filter to a service group

In some cases, the dnsDefault application filter will automatically be added to the DNS service group. If not, you will need to add it or another DNS application filter to the DNS Service Group. To add this filter to the DNS Service Group, follow the steps below:

-
- 1 In the Navigator window, open the appropriate Service Group folder and double-click the DNS Service Group.
.....
 - 2 Highlight and double-click the entry for UDP port 53. DNS application filters are not supported for TCP.
.....
 - 3 Click the pulldown for the application filter, select your DNS application filter, and click **OK**.
.....
 - 4 Open the **File** menu and select one of the **Save** options.

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information on Brick Zone Rulesets, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). For additional information on service groups, see [Chapter 4, “Service Groups”](#).



ESP Application Filter

Overview

The ESP application filter allows ESP NULL encapsulated data through the Brick and validates the proper format of the packets.

To configure an ESP application filter

Complete the following steps to configure an ESP application filter.

- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.


- 2 In the right hand column, right-click the mouse and select **New Application Filter** from the pop-up menu.

Result The Application Editor is displayed.

- 3 Click the down arrow next to the **Type** field to display a drop-down list and select **ESP**.

Result The Application Filter Editor is re-displayed with the parameters for setting up an ESP application filter ([Figure 5-7, “ESP Application Filter” \(p. 5-18\)](#)).

Figure 5-7 ESP Application Filter



-
- 4 In the **Name** and **Description** fields, enter a name for the application filter and a brief description. The **Description** field is optional.

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the [“Global Application Filters”](#) (p. 5-127) section for an explanation of global application filters.

- 5 From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.

END OF STEPS

To add the filter to a service group

To add this application filter to a service group, follow the steps below:

- 1 In the Navigator window, open the appropriate Service Group folder and double-click on the service group or create a new one.

Result The Service Group Editor window is displayed.

- 2 Right-click and select **New** from the pop-up menu if you are creating a new service group.

Result The Service Editor window is displayed.

- 3 In the **Protocol** field, enter **50**.
-

- 4 Click the pulldown for the application filter, select the name of the application filter from the pull-down menu, and click **OK**.
-

- 5 Open the File menu and select one of the **Save** options.

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information on Brick Zone Rulesets, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). For additional information on service groups, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). [Chapter 4, “Service Groups”](#).



FTP Application Filter

Overview

FTP is one of the oldest protocols used over the Internet. In the last few years, virtually every ftp server has been hit with a variety of problems ranging from remote root hacks to denial of service attacks. The LVF can be used to strengthen the security for FTP servers and clients located in the secure zone.

FTP is among the three most commonly used applications on the Internet and is known to be highly vulnerable. Not only ftp servers and clients are at risk; the security weaknesses in the FTP protocol are used to launch attacks on other services. Considering the large scale deployment of FTP, it is not possible to revise the protocol for better security. Instead, network firewalls can help overcome the security weaknesses of the protocol.

In addition to supporting the Network Address Translation (NAT) and Port Address Translation (PAT) features in SMS and the blocking of FTP proxy transfers, the FTP application filters provide the following security measures:

- Filters FTP commands
- Checks for connection stealing
- Restricts dynamic ports
- Suppresses invalid user name responses
- Prevents brute force password guessing
- Checks protocol field lengths
- Checks for invalid characters
- Prohibits FTP on select user accounts

Configuring an FTP Application Filter

An FTP application filter called *ftpDefault* is created automatically in the system group for a newly installed SMS and, once installed, is created automatically for newly created groups. Similarly, an *ftpDefault* filter is created when an SMS upgrade is done for each of your existing groups. However, *ftpDefault* is not automatically assigned to any service group after installation or upgrade.

You may decide to just reconfigure the *ftpDefault* application filter or create a new one.

To configure an FTP application filter

To configure an FTP application filter, follow the steps below:

- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.

- 2 In the right hand column, right-click the mouse and select **New Application Filter** from the pop-up menu. After the Application Filter Editor appears, click the Type drop-down menu and select **FTP**.

The Application Filter Editor is displayed with the parameters for setting up an FTP application filter. [Figure 5-8, “FTP Application Filter Editor \(Commands Filter Tab \)” \(p. 5-23\)](#) shows a sample of the screen.

Note: The initial settings of a new FTP application filter are identical to the ftpDefault application filter that is installed with the product.

Figure 5-8 FTP Application Filter Editor (Commands Filter Tab)

The screenshot shows the 'Application Filter Editor' window with the 'Commands Filter' tab selected. The window title is '/system/Policies/Application Filters/'. The menu bar includes File, Edit, Monitor, Windows, Utilities, and Help. The main area contains the following elements:

- Name:** An empty text input field.
- Description:** An empty text input field.
- Display and Use Globally
- Type:** A dropdown menu set to 'FTP'.
- Audit FTP Commands (Detailed Session Audit) Audit FTP Exceptions (Basic Exception Audit)
- Commands Filter Tab:** Contains two sub-tabs: 'Pass List' and 'Block List'.
 - Pass List:** A list box containing the following commands: mkd, port, user, quit, stat, stor, type, pasv.
 - Block List:** A list box containing the command: site.
- Add:** Two empty text input fields, one under each list box, with a circular 'X' button below them.

- 3 In the **Name** and **Description** fields, enter a name for the application filter and a brief description. The **Description** field is optional.

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the [“Global Application Filters”](#) (p. 5-127) section for an explanation of global application filters.

-
- 4 Click the **Audit FTP Commands (Detailed Session Audit)** checkbox to monitor and log FTP commands issued when an FTP connection is established.

Click the **Audit FTP Exceptions (Basic Exception Audit)** checkbox to monitor and log application filter violations and the FTP traffic.

- 5 The Commands Filter tab contains these fields:

- **Pass List / Block List**

Command filtering is accomplished in one of two separate modes: pass or block. If you select the **Pass List** radio button, only those commands in the pass list are allowed. If you select the **Block List** radio button, commands in the block list are prohibited.

Commands can be added to either list by typing the command into the **Add** field under the respective list and pressing the Enter key. Only one list can be active at a time for updating.

A command in either list can be deleted by highlighting it and then clicking the **Remove selected** button, or by right-clicking on the command and choosing **Delete**. You can select multiple entries for deletion by clicking the left mouse button while pressing the Ctrl key, or by clicking the left mouse button while pressing the Shift key to select a range of commands in the list.

- 6 The Protocol Anomaly Check tab contains these fields ([Figure 5-9, “FTP Application Editor \(Protocol Anomaly Check Tab \)”](#) (p. 5-25) shows a sample screen):

Figure 5-9 FTP Application Editor (Protocol Anomaly Check Tab)

The screenshot shows the 'Application Filter Editor' window for the 'ftpDefault' filter. The 'Protocol Anomaly Check' tab is selected. The configuration includes:

- Name:** ftpDefault
- Description:** (empty text box)
- Display and Use Globally
- Type:** FTP
- Audit FTP Commands (Detailed Session Audit)
- Audit FTP Exceptions (Basic Exception Audit)
- Protocol Anomaly Check Tab:**
 - Enable Protocol Anomaly Checking
 - Maximum Path/File Name Length:** 255
 - Maximum Command Length:** 4
 - Maximum Parameter Length:** 255
 - Maximum Username/Password Length:** 32

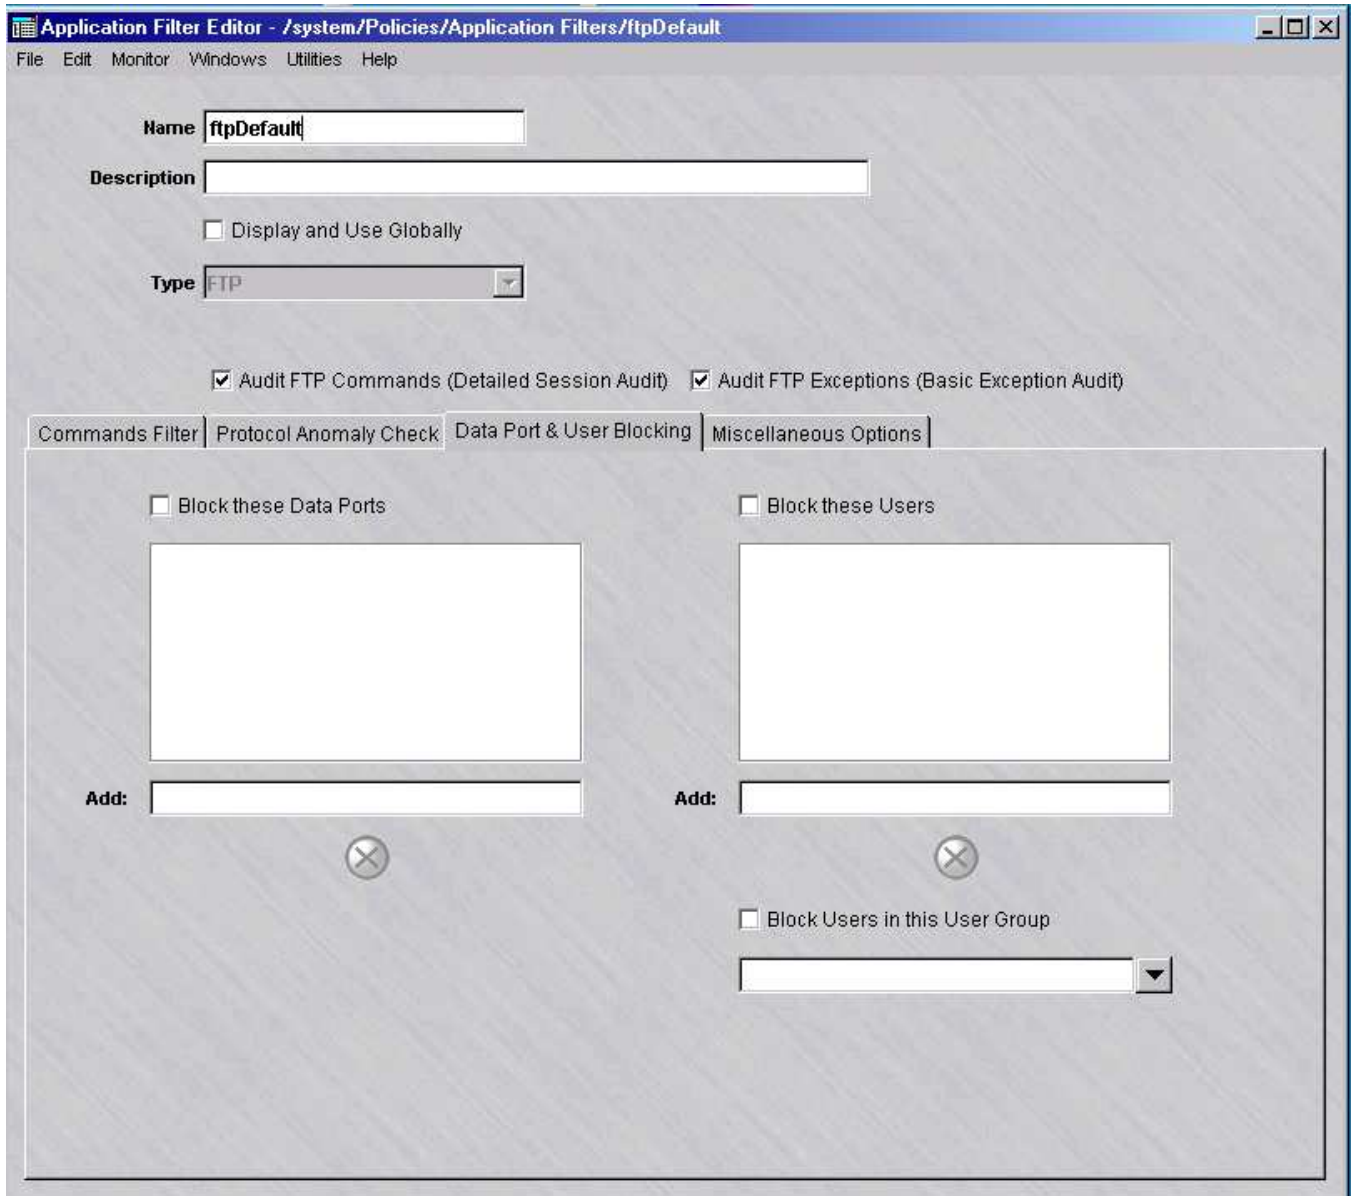
The following explains these fields:

- **Enable Protocol Anomaly Checking** - This checkbox is checked by default to enable Protocol Anomaly Checking, which parses all significant protocol fields to determine field lengths and protect from buffer overrun. To disable protocol anomaly checks, click the checkbox to remove the check.
- **Maximum Path/File Name Length** - If Protocol Anomaly Checking is enabled, this field is used to specify the maximum length for path / filenames. The default value is 255 characters. Change the field value for this parameter, as needed.



- **Maximum Command Length** - If Protocol Anomaly Checking is enabled, this field is used to specify the maximum length for commands. The default value is four characters. Change the field value for this parameter, as needed.
 - **Maximum Parameter Length** - If Protocol Anomaly Checking is enabled, this field is used to specify the maximum length of all parameters in a command line. The default value is 255 characters. Change the field value for this parameter, as needed.
 - **Maximum Username/Password Length** - If Protocol Anomaly Checking is enabled, this field is used to specify the maximum length of usernames and passwords. The default value is 32 characters. Change the field value for this parameter, as needed.
-

- 7 The Data Port & User Blocking tab contains these fields ([Figure 5-10, “FTP Application Filter Editor \(Data Port & User Blocking Tab\)”](#) (p. 5-27) shows a sample screen):

Figure 5-10 FTP Application Filter Editor (Data Port & User Blocking Tab)

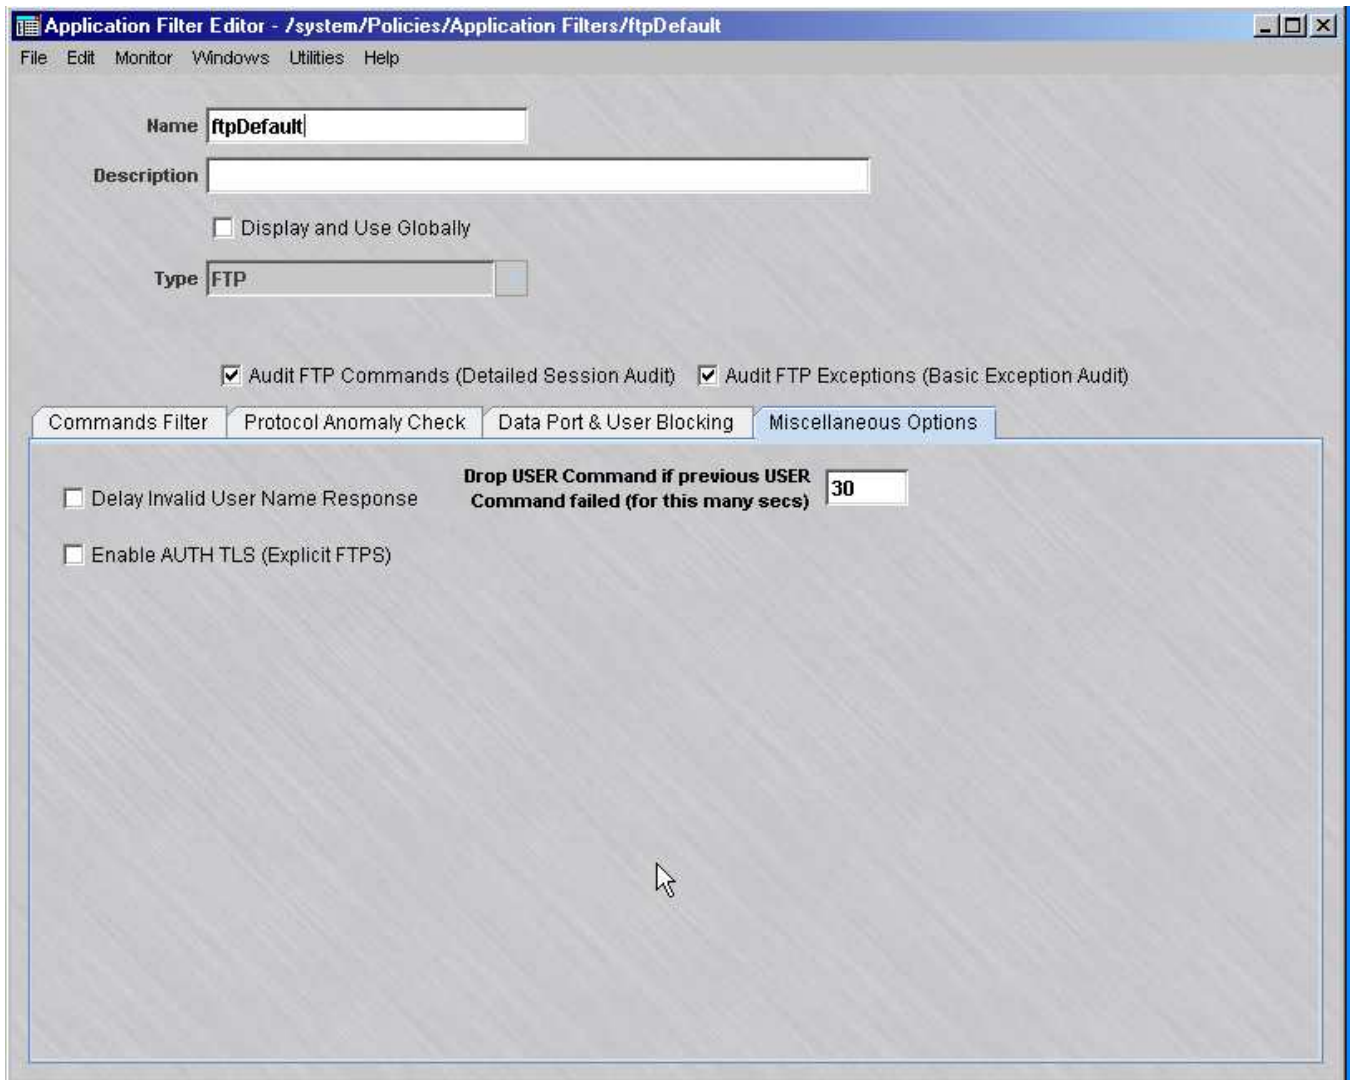


The following explains these fields:

- **Blocked Data Ports List** - To enable this feature, click the Block these Data Ports checkbox to place a check in the box. To add a port to the blocked ports list, enter the port number in the Add field below the list window and press the Enter key. Valid values are 1024 - 65353. To remove a port from the blocked ports list, select the port in the list window and click the Remove () button, or simply right-click on the port in the list window and choose **Delete**.
- **Blocked Users List** - To enable this feature, click the Block these Users checkbox to place a check in the box. To add a user to the blocked users list, enter the user ID (login) or user group name in the Add field below the list window and press the Enter key. To remove a user from the blocked users list, select the user ID in the list window and click the Remove () button, or simply right-click on the user ID in the list and choose **Delete**. You can also delete all users that been configured in a user group. To enable this feature, click the Block Users in this User Group checkbox to place a check in the box, click the down arrow next to the field below the checkbox to display the Select a User Group window, and select the user group.

-
- 8 The Miscellaneous Options tab contains these fields ([Figure 5-11, “FTP Application Editor \(Miscellaneous Options Tab\)”](#) (p. 5-29) shows a sample screen):

Figure 5-11 FTP Application Editor (Miscellaneous Options Tab)



The following explains these fields:

- **Delay Invalid User Name Response** - Click this checkbox to activate a delay in response to an invalid user name attempting to establish an FTP session.
- **Drop USER Command if previous USER Command failed (for this many secs)** - Enter the duration, in seconds, to block the USER command if the previous five consecutive USER command from the same host failed. The default value for this parameter is 30 seconds.
- **Enable AUTH TLS (Explicit FTPS)** - Click this checkbox to stop performing any validations on the FTP application layer after the SMS has accepted the AUTH TLS command with a 234 response. Strict TCP validation continues to be performed if this option is enabled. This allows for the confidentiality and integrity of TLS but removes a layer of protocol enforcement. This checkbox is unchecked, by default.

-
- 9 From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.

END OF STEPS

To add the filter to a service group

To add this application filter to a service group, follow the steps below:

-
- 1 In the Navigator window, open the appropriate Service Group folder and double-click on the service group or create a new one.
 - 2 Edit the tcp protocol entry, specifying the destination port for incoming messages using this application filter.
 - 3 Click the pulldown for the application filter, select the name of the application filter from the pull-down menu, and click **OK**.
 - 4 Open the File menu and select one of the **Save** options..

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information on Brick Zone Rulesets, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). For additional information on service groups, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). [Chapter 4, “Service Groups”](#).

END OF STEPS



GTP Application Filter

Definition

General Packet Radio Service Tunneling Protocol (GTP) handles the flow of user packet data and signaling information between the SGSN (Serving GPRS Support Node) and GGSN (Gateway GPRS Support Node) in a GPRS network. GTP is defined on both the Gn and Gp interfaces of GPRS and UMTS networks.

The Gp interface is the logical connection between Public Land Mobile Networks (PLMNs) that is used to support mobile data users. GTP is used to establish a connection between a local SGSN and the user's home GGSN.

The GTP application filter is used to filter the packets going through the Gp interface to provide security checking and protection.

GTP application filter (stateful)

GTP Application Filtering - Stateful provides the next higher level of security to GTP traffic. It is designed to:

- Reduce the effects of denial-of-service attacks
- Reduce the opportunities for user-session hijacking
- Enforce PLMN policies related to roaming and traffic volume
- Counter several known vulnerabilities related to theft of traffic

This level of application filtering adds to the GTP filtering options that require the retention of GTP requests, roaming procedure, and PDP context state information.

For stateful GTP application filtering, the state and context of GTP requests and responses are collected and examined by the SMS and can be used to apply a filtering method or action to be taken based on that context or state.

Additional security mechanisms can be provided using stateful GTP application filtering, including:

- Sending an appropriate GTP response to the sender of a GTP request that has been dropped due to protocol or policy violations
- Version enforcement for GTP request messages
- Setting a limit on the total number of packets that may be passed by a tunnel.
- Throttling of GTP echo request messages on each direction of a UDP session that exceed a pre-defined interval
- Throttling of GTP error indication messages on each direction of a UDP session that exceed a pre-defined interval

GTP PDP Context Deletion Monitor

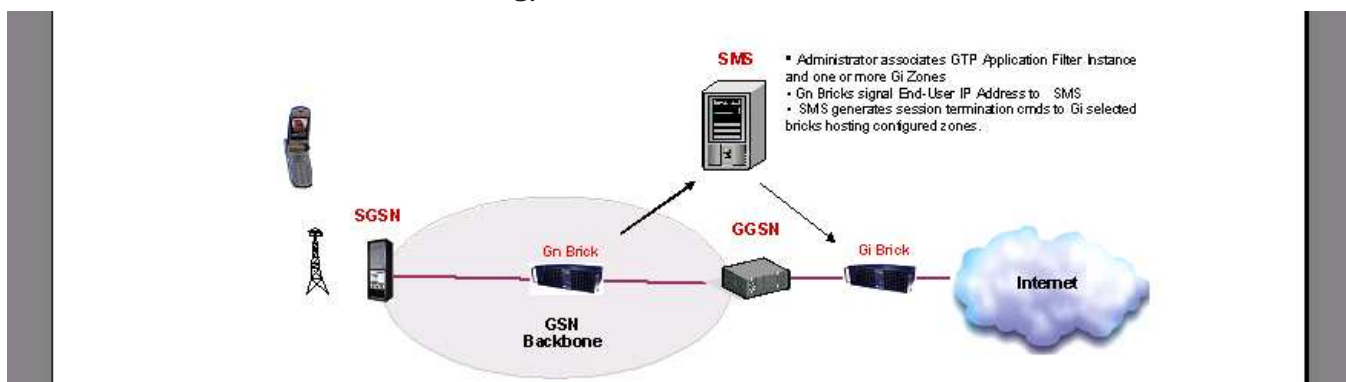
The increased security mechanisms provided by stateful GTP application filtering are further enhanced by a built-in GTP PDP Context Deletion Monitor that coordinates the termination of firewall sessions on the Gi interface of the Brick with the termination of the associated GTP PDP Contexts on the GPRS/UMTS backbone. This addresses vulnerabilities with known theft-of-service and over-billing exploits, as well as providing the means to tighten security related to GPRS/UMTS traffic on the non-backbone area of a service provider network.

This feature monitors the deletion of GTPv0 and GTPv1 GTP PDP Contexts at one or more GTP monitoring points on the GTP backbone and communicates these events to the SMS to trigger the deletion of firewall sessions that use the associated IP addresses. These GTP monitoring points may include one or more ruleset zones and control all Bricks that host the specified zones. Each GTP monitoring point is associated with a specific GTP PDP Context Deletion Monitor. However, a firewall control point may be a member of multiple GTP PDP Context Deletion Monitor instances, and the SMS may support multiple GTP PDP Context Deletion monitors.

When the last PDP Context associated with an IP address is deleted, the SGSN sends a message to the SMS indicating the deletion of the last PDP context and release of the associated IP address. The SMS generates session termination commands to the Brick hosting the zones that are configured with the Deletion Monitor feature.

The GTP PDP Context Deletion process is illustrated in [Figure 5-12, “GTP Application Filter -Stateful \(GTP PDP Context Deletion Monitoring\)”](#) (p. 5-32).

Figure 5-12 GTP Application Filter -Stateful (GTP PDP Context Deletion Monitoring)



This feature may be used while a Brick is configured for inline filtering of GTP traffic or while a Brick is passively monitoring GTP traffic out-of-line, on one or more unidirectional network “taps”. A Brick can simultaneously filter GTP traffic inline on some ports and passively monitor GTP traffic on other ports.

The combination of inline/passive monitoring points, multiple monitoring points, and zone-based definition of monitoring points allow the use of this feature to be scaled to support large GPRS/UMTS networks with complex topologies and multiple firewall layers.

The GTP PDP Context Deletion Monitor feature can be enabled or disabled on the Tunnel Management tab category of the Stateful GTP application filter tab. For additional instructions on enabling this feature, refer to [Step 11](#) in the procedure “[To configure a GTP application filter](#)” (p. 5-33).

GTP R6 information element (IE) removal and GTP R6 support

The GTP application filter now supports all of the messages and IEs defined in Release 6 (R6) of 3GPP TS 29.060.v6. The newly defined MBMS messages of R6 are blocked by the GTP application filter by default, so the additional R6 messages will be transparent to R5 equipment users after an SMS upgrade. GTP R6 users have to explicitly configure the GTP application filter to pass those messages.

To facilitate the inter-operability between GTP R5 and R6 equipment, R6 IEs can be removed from GTP R6 messages. The application filter software is pre-populated with a default list of GTP messages from which R6 IEs are automatically removed and a default list of R6 IEs that are stripped from these messages. Messages and IEs can be added to these pre-populated lists as needed through the SMS Version 1 GTP Application Filter Editor screens. Refer to [Step 4](#) and [Step 12](#) of the procedure “[To configure a GTP application filter](#)” (p. 5-33) for instructions on how to enable selective removal of R6 IEs from GTP messages and how to modify the list of GTP messages from which IEs are removed.

To configure a GTP application filter

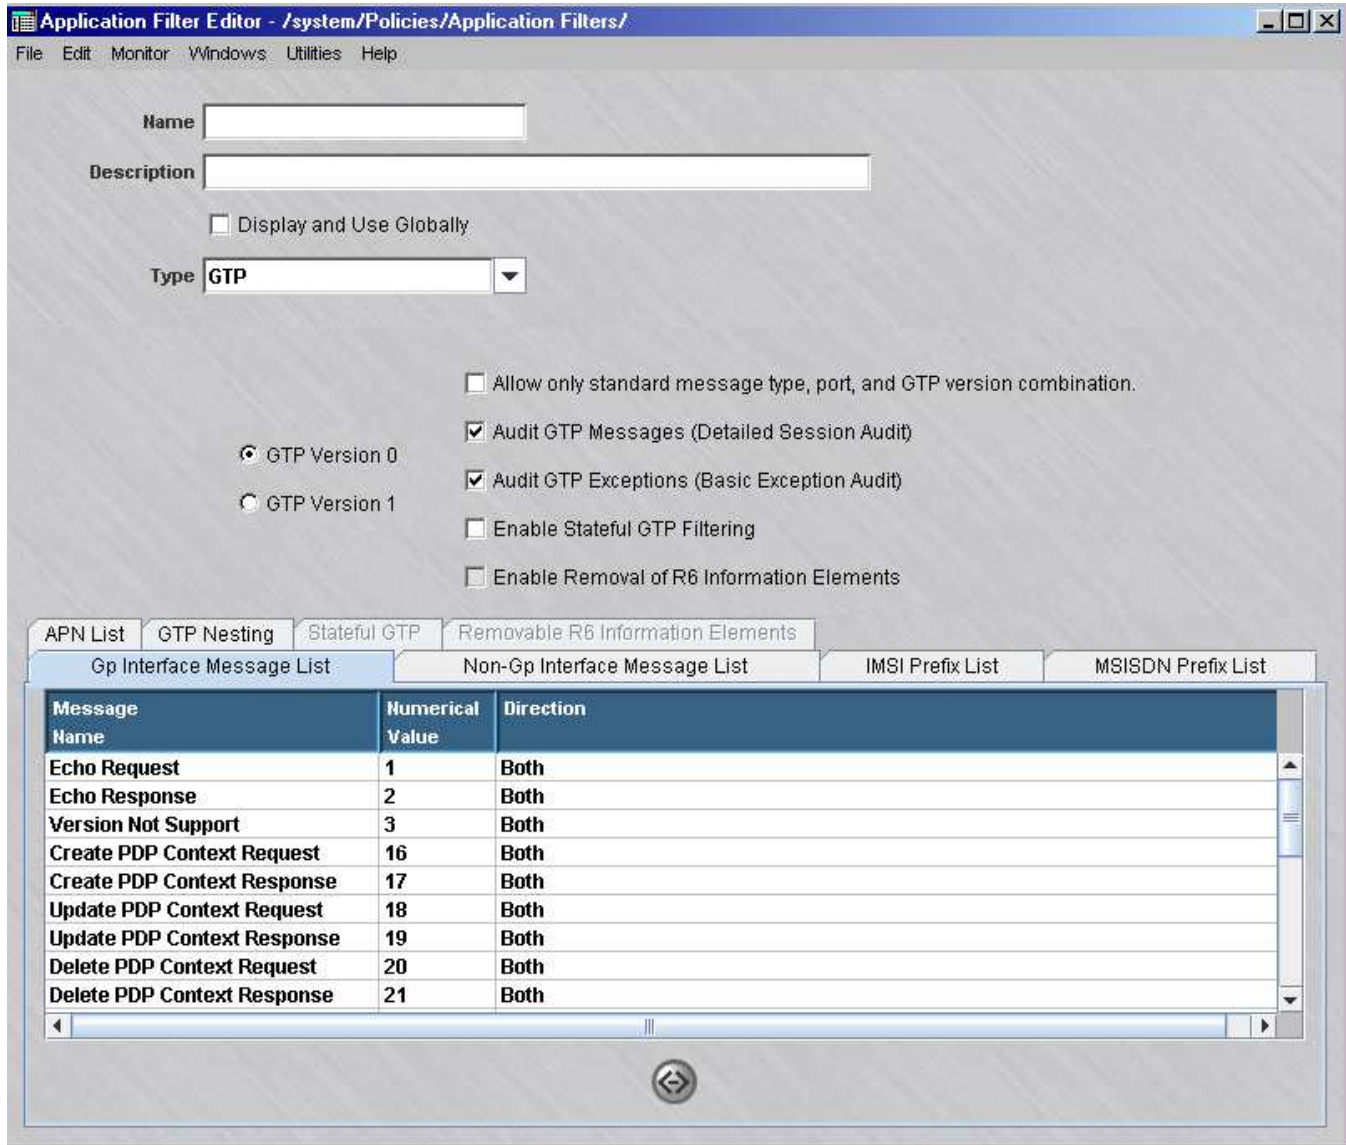
Complete the following steps to configure a GTP application filter.

- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.
- 2 In the right hand column, right-click the mouse and select **New Application Filter** from the pop-up menu. After the Application Filter Editor appears, click the Type drop-down menu and select **GTP**.

An alternate method is to click on the **Application Filters** folder in the Navigator to display a list of existing application filters, and select **gtp_v0** (for the GTP Version 0 application filter parameters) or **gtp_v1** (for the GTP Version 1 application filter parameters).

Result The GTP Application Filter window is displayed ([Figure 5-13, “GTP Application Filter Editor \(GTP Version 0 Parameters\)”](#) (p. 5-34)).

Figure 5-13 GTP Application Filter Editor (GTP Version 0 Parameters)



If you selected **GTP** from the **New Application Filter** pop-up menu, the GTP Version 0 application filter parameters are displayed on the Application Editor screen.

If you selected **gtp_v0** or **gtp_v1** from the Contents panel of application filters, the GTP Version 0 or GTP Version 1 application parameters are displayed, depending on which version was selected.

-
- 3 In the **Name** and **Description** fields, enter a name for the application filter and a brief description. The name is required, but the description is optional.

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the “[Global Application Filters](#)” (p. 5-127) section for an explanation of global application filters.


- 4 The area above the seven tabs contains these radio buttons and checkboxes:

- **GTP Version 0**
Selecting this option will set the application filter to GTP version zero provisional attributes. Once the application filter is saved this option cannot be changed.
- **GTP Version 1**
Selecting this option will set the application filter to GTP version one provisional attributes. Once the application filter is saved this option cannot be changed.
- **Allow only standard message type, port, and GTP version combination.**
If this option is checked, only the standard provisional attributes will be allowed along with the standard ports (GTP Version 0 Port: 3386 or GTP Version 1 Ports: 2123, 2152) depending on which GTP version is chosen. Service Groups attempting to use a GTP application filter with this option checked will ensure that the appropriate ports are used.
- **Audit GTP Messages (Detailed Session Audit)**
If this option is checked and if the “Session Audit” parameter in a rule is set to “Detailed”, the Session Log will record and provide details for all GTP sessions that use this filter. If unchecked, the session log does not provide detailed data.
- **Audit GTP Exceptions (Basic Exception Audit)**
If this option is checked and the “Exception Audit” parameter in a rule is set to a value other than “None”, the Session log will provide details on blocked GTP sessions that use this filter.
- **Enable Stateful GTP Filtering**
If this option is checked, stateful GTP filtering is activated and GTP requests and responses are filtered based on the message context and state that has been collected and retained by the SMS. Additional details about configuring Stateful GTP Filtering are provided in [Step 11](#) of this procedure.
- **Enable Removal of R6 Information Elements** *Note: this option is only available on the GTP Version 1 application filter; it is grayed out on the GTP Version 0 application filter screen. To access this option, select the GTP Version 1 radio button on the GTP Application Filter Editor screen (if it is not already selected). This option will already be available if you selected gtp_v1 from the Contents panel in Step 2).* If this option is checked, it enables selective removal of GTP R6

Information Elements (IEs) from certain GTP messages. When this option is enabled, the **Removable R6 Information Elements** tab is activated, which provides access to pre-populated lists of GTP messages and IEs that can be modified by the administrator. Additional details about configuring this option are provided in [Step 12](#) of this procedure.

- 5 The Gp Interface Message List tab contains the following fields ([Figure 5-14, “GTP Application Filter Editor \(GP Interface Message List Tab\)”](#) (p. 5-36) shows a sample screen):

Figure 5-14 GTP Application Filter Editor (GP Interface Message List Tab)



Message Name	Numerical Value	Direction
Echo Request	1	Both
Echo Response	2	Both
Version Not Support	3	Both
Create PDP Context Request	16	Both
Create PDP Context Response	17	Both
Update PDP Context Request	18	Both
Update PDP Context Response	19	Both
Delete PDP Context Request	20	Both
Delete PDP Context Response	21	Both

The following explains these fields:

- **Gp Interface Message List**

Each entry in the list corresponds to a standard Gp interface message and contains its Message Name, Numerical Value, and a Direction having a value that determines when it will be processed, "In" or "Out" of the zone, "Both" in and out of the zone, and "None" for neither in or out of the zone. The default direction is "Both" and messages may not be added to this standard list.

The MBMS messages defined in GTP R6 have their direction set to **None** by default so they can be ignored by existing R5 equipment users. However, in order to allow these newly added (for R6) GTP messages to be passed by the Brick for GTP R6 equipment users, the direction of the message(s) must be manually edited. To change the direction of a new R6 GTP message on the **Gp Interface Message List** table, right-click on the message and select **Edit**. An Edit Message Parameters window is displayed with the current Direction setting for that message. Change the Direction setting (to **In To Zone**, **Out Of Zone**, or **Both Directions**), and click **OK**. The system returns to the **Gp Interface Message List** table, showing the modified Direction setting of the GTP message.

- 6 The Non-Gp Interface Message List tab contains the following fields ([Figure 5-15](#), “GTP Application Filter Editor (Non-Gp Interface Message Tab)” (p. 5-37) shows a sample screen):

Figure 5-15 GTP Application Filter Editor (Non-Gp Interface Message Tab)

Message Name	Numerical Value	Direction
Node Alive Request	4	None
Node Alive Response	5	None
Redirection Request	6	None
Redirection Response	7	None
Send Routing Information for GP...	32	None
Send Routing Information for GP...	33	None
Failure Report Request	34	None
Failure Report Response	35	None
Note MS GPRS Present Request	36	None

The following explains these fields:

- **Non-Gp Interface Message List**

Each entry in the list corresponds to a Gp interface message and contains its Message Name, Numerical Value, and a Direction having a value that determines when it will be processed, “In” or “Out” of the zone, “Both” in and out of the zone, and “None” for neither in or out of the zone. The default direction is “None” and new messages with non-standard numerical values may be added and edited further.

- 7 The IMSI Prefix List tab contains the following fields and functions ([Figure 5-16](#), “GTP Application Filter Editor (IMSI) Prefix list Tab)” (p. 5-38) shows a sample screen):

Figure 5-16 GTP Application Filter Editor (IMSI) Prefix list Tab)

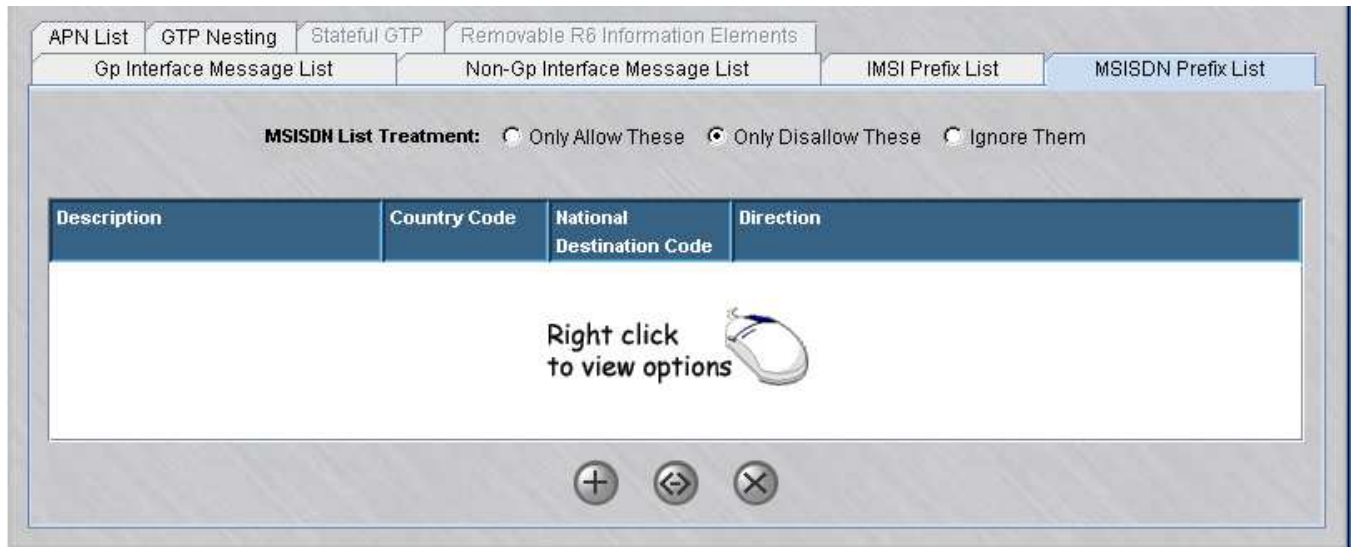
The following explains these fields:

- **IMSI Prefix List**

The International Mobile Subscriber Identity prefix list is blank for new GTP application filters and the user may define new IMSI entries for GTP packet processing. The IMSI list can be configured by choosing to "Only Allow These", "Only Disallow These", or "Ignore Them" for the current IMSI entries. A new list entry can be created by clicking the right mouse button and selecting **New** from the pop-up menu. The Create New IMSI Prefix window is displayed. In the **Description** field, enter a description for the entry (optional). In the **Mobile Country Code** field, enter a three digit code. In the **Mobile Network Code** field, enter a 2-3 digit code. Choose **In To Zone**, **Out of Zone**, or **Both Directions** for the filtering direction. Click the **OK** button. The new list entry is displayed on the tab listing. To edit an entry, right click on the entry in the list and choose **Edit** from the pop-up menu. The Edit IMSI Prefix window is displayed to change the entry parameters.

-
- 8 The MSISDN Prefix List tab contains the following fields and functions ([Figure 5-17, "GTP Application Filter \(MSISDN Prefix List Tab\)"](#) (p. 5-39) shows a sample screen):

Figure 5-17 GTP Application Filter (MSISDN Prefix List Tab)



The following explains these fields:

- **MSISDN Prefix List**

The Mobile Station Integrated Services Digital Network prefix list is blank for new GTP application filters and users may define new MSISDN entries. The MSISDN list can be configured by choosing to "Only Allow These", "Only Disallow These", or "Ignore Them" for the current IMSI entries. A new list entry can be created by clicking the right mouse button and selecting **New** from the pop-up menu. The Create New MSISDN Prefix window is displayed. In the **Description** field, enter a description for the entry (optional). In the **Country Code** field, enter a 1-3 digit code. In the **National Destination Code** field, enter a variable length code. Choose **In To Zone**, **Out of Zone**, or **Both Directions** for the filtering direction. Click the **OK** button. The new list entry is displayed on the tab listing. To edit an entry, right click on the entry in the list and choose **Edit** from the pop-up menu. The Edit MSISDN Prefix window is displayed to change the entry parameters.

-
- 9 The APN List tab contains the following fields and functions (Figure 5-18, "GTP Application Filter Editor(APN List Tab)" (p. 5-40) shows a sample screen):

Figure 5-18 GTP Application Filter Editor(APN List Tab)



The following explains these fields:

- **APN List**

The Access Point Names (APN) list is blank for new GTP application filters and users may define new APN entries for GTP packet processing. The APN list can be configured by choosing to "Only Allow These", "Only Disallow These", or "Ignore Them" for the current APN list entries. The **Selection Mode** pull-down list specifies the type of GTP packets that are allowed. A new list entry can be created by clicking the right mouse button and selecting **New** from the pop-up menu. The Create New APN window is displayed. In the **Description** field, enter a description for the entry (optional). In the **Access Point Name** field, enter a partially qualified domain name (such as *.alcatel-lucent.com) or a fully qualified domain name (such as 127.0.0.1 or alcatel-lucent.com). Choose **In To Zone**, **Out of Zone**, or **Both Directions** for the filtering direction. Click the **OK** button. The new list entry is displayed on the tab listing. To edit an entry, right click on the entry in the list and choose **Edit** from the pop-up menu. The Edit APN window is displayed to change the entry parameters.

-
- 10 The GTP Nesting tab contains the following field ([Figure 5-19, "GTP Application Filter Editor \(GTP Nesting Tab\)"](#) (p. 5-41) shows a sample screen):

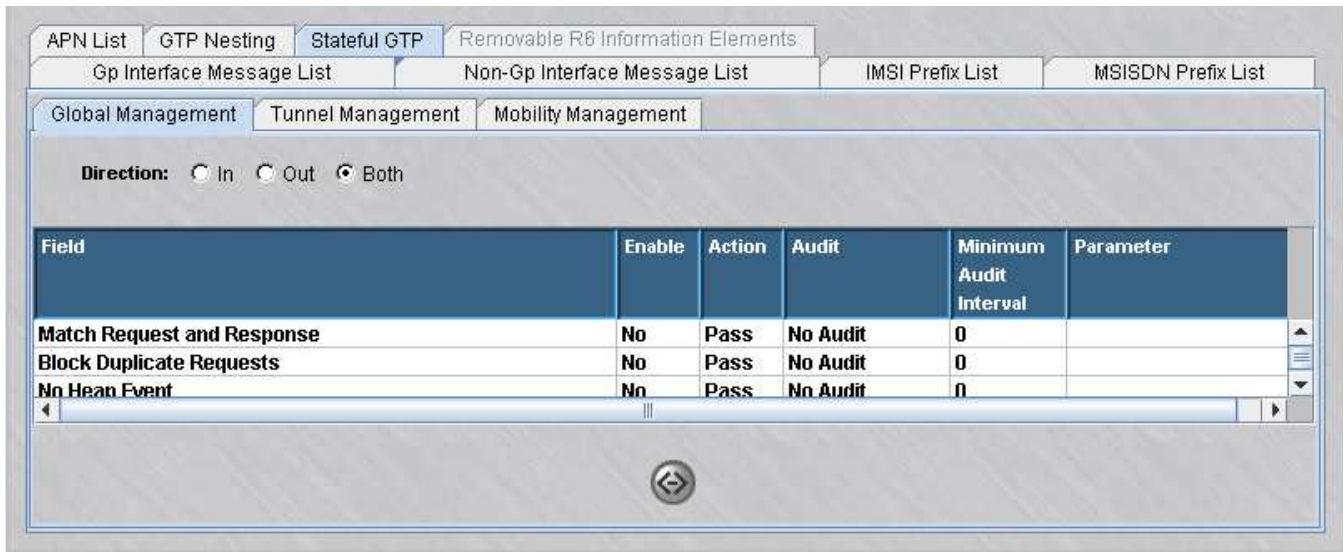
Figure 5-19 GTP Application Filter Editor (GTP Nesting Tab)

The following explains these fields:

- **Do not allow GTP within GTP Tunnel**
If this option is checked, GTP traffic will not be allowed within a GTP Tunnel session using this application filter.

-
- 11** The Stateful GTP tab is enabled when the **Enable Stateful GTP Filtering** checkbox in the top portion of the GTP Application Filter Editor screen is checked. This tab is divided into three tab categories ([Figure 5-20, “GTP Application Filter Editor \(Stateful GTP Tab\)”](#) (p. 5-42) shows a sample screen):

Figure 5-20 GTP Application Filter Editor (Stateful GTP Tab)



The three tabbed category lists are blank for new GTP application filters and users may define new Stateful GTP Filtering list entries for GTP packets coming "In" or "Out" of the zone, or in both directions.

The three tab categories for this filtering method are:

- Global Management
- Tunnel Management
- Mobile Management

which allow you to set global parameters for stateful GTP application filtering, parameters related to SGSN-to-GGSN tunnels, enable or disable the GTP PDP Context Deletion Monitor (in the Tunnel Management tab category), and set parameters related to mobile GTP data packets.

The filtering parameters in each category are disabled, by default. To configure a filtering parameter and change its settings, click on the related category tab to display the related parameter listing, right-click on the parameter to be modified, and select **Edit** from the displayed pop-up menu. The editing screen is displayed. For each filtering parameter, the following fields can be configured as follows:

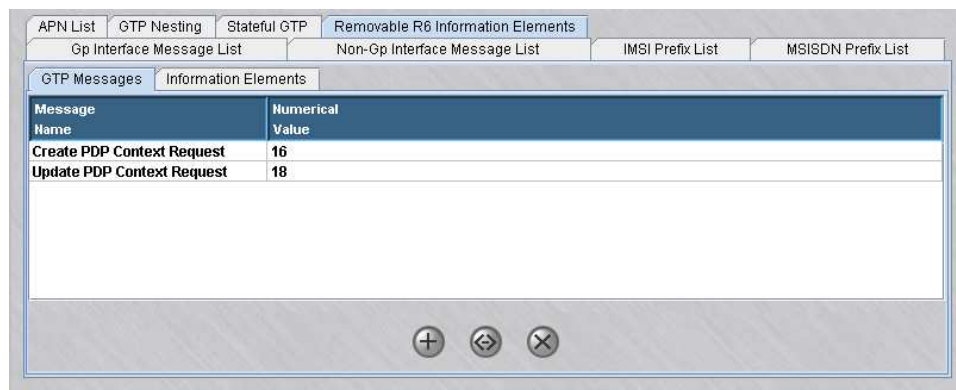
- **Enable/Disable** - click this checkbox to enable the filtering parameter, uncheck the checkbox to disable it.
- **Action** - click the down-arrow next to this field to display a drop-down list and select the action to be taken on the GTP packet if the state of the GTP request/response applies to this filtering parameter. The choices are: **Pass** or **Drop**.

- **Audit** - click the down-arrow next to this field to display a drop-down list and select the Audit setting for incoming/outgoing GTP packets. The choices are: **No Audit**, **Audit on Drop**, or **Audit Always**. The default setting is **No Audit**.
- **Minimum Audit Interval** - enter the minimum interval permissible between audit messages for this parameter. The default value is **0**.
- **Parameter (Optional)** - enter the numeric value for this parameter (if applicable).

After editing the value(s) of the selected filtering parameter, click the **OK** button. The edited parameter entry is displayed on the tab listing.

- 12 The Removable R6 Information Elements tab is activated when the **Enable Removal of R6 Information Elements** checkbox on the GTP Version 1 Application Filter Editor screen is checked (Figure 5-21, “GTP Application Filter Editor (Removable R6 Information Elements Tab) (Only Available on GTP Version 1)” (p. 5-43)).

Figure 5-21 GTP Application Filter Editor (Removable R6 Information Elements Tab) (Only Available on GTP Version 1)



This tab has two sub-tabular listings that contain:

- **GTP Messages**—A default list of GTP messages with R6 IEs that need to be stripped (removed) (Figure 5-22, “Removal R6 Information Elements (GTP Messages Sub-Tab)” (p. 5-44)).
- **Information Elements**—A default list of R6 IEs to be removed from the GTP messages (Figure 5-23, “Removable R6 Information Elements (Information Elements Sub-Tab)” (p. 5-44)).

Figure 5-22 Removal R6 Information Elements (GTP Messages Sub-Tab)

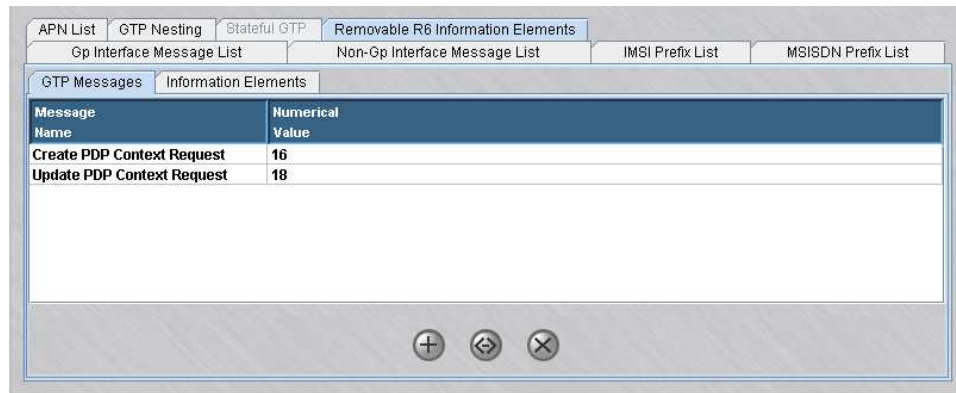
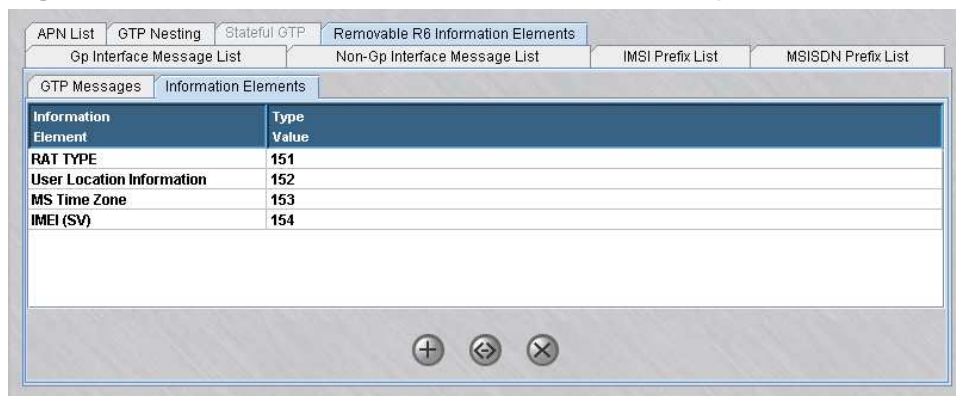


Figure 5-23 Removable R6 Information Elements (Information Elements Sub-Tab)



To	Do This
<p>Add a new GTP message to the list of messages from which R6 IEs will be removed</p>	<p>Right-click on one of the existing messages on the GTP Messages sub-tab and select New from the pop-up menu. A New Messages Parameters window is displayed. Enter the new Message Name (for example, Create MBMS Context Request) and the Message's Numerical Value (in this example, 102) and click OK. The GTP message will be added to the list.</p>

To	Do This
Modify an existing GTP message entry on the list	Right-click on the GTP message to be modified on the GTP Messages sub-tab and select Edit from the pop-up menu. An Edit Messages Parameters window is displayed with the current GTP message and numerical value (if shown). Make any modifications needed and click OK . The change(s) will be reflected on the list.
Delete a GTP message entry from the list.	Right-click on the GTP message to be deleted on the GTP Messages sub-tab and select Delete from the pop-up menu. A Delete Message Parameters window is displayed with the GTP entry shown. Click OK to confirm the deletion. The GTP message is deleted from the list.
Add an Information Element (IE) to the default list of IEs to be stripped from GTP messages	Click on the Information Elements sub-tab, right-click on one of the existing IEs in the list, and select New from the pop-up menu. A New Messages Parameters window is displayed. Enter the IE Name in the Message Name field (for example, MBMS Service Area) and the IE Type Value in the Message's Numerical Value field (in this example, 160), and click OK . The IE is added to the IE list.
Modify an existing IE entry on the list	Click on the Information Elements sub-tab, right-click on one of the IEs to be modified in the list, and select Edit from the pop-up menu. An Edit Messages Parameters window is displayed with the current IE and IE Type Value (if shown). Make any modifications needed and click OK . The change(s) will be reflected on the list.
Delete an IE from the list.	Click on the Information Elements sub-tab, right-click on one of the IEs to be deleted from the list, and select Delete from the pop-up menu. A Delete Message Parameters window is displayed with the IE entry shown. Click OK to confirm the deletion. The IE is deleted from the list.

-
- 13 From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.

.....
E N D O F S T E P S
.....

To add the filter to a service group

To add this application filter to a service group, follow the steps below:

-
- 1 In the Navigator window, open the appropriate Service Group folder and double-click on the service group or create a new one.
 - 2 Edit the tcp protocol entry, specifying the destination port for incoming messages using this application filter.
 - 3 Click the pulldown for the application filter, select the name of the application filter from the pull-down menu, and click **OK**.
 - 4 Open the File menu and select one of the **Save** options.
-

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information on Brick Zone Rulesets, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). For additional information on service groups, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). [Chapter 4, “Service Groups”](#).

.....
E N D O F S T E P S
.....



H.323 Application Filters

Overview

The LVF handles 1:1 NAT in addition to many-to-one NAT. In the latter case, the LVF must perform not only NAT, but Port Address Translation (PAT) as well. The combination is sometimes called NPAT. With NPAT, the private network presents the appearance of a single address to the public network, but there may be any number of actual (private) addresses and machines in the private network.

The H.323 protocol suite consists of several session and application layer protocols for exchanging "multimedia" traffic across an IP network. The endpoints typically exchange messages using a combination of TCP and UDP channels, most of which are negotiated dynamically. Generally, for RAS (Registration, Administration and Status) between an endpoint and a gatekeeper, a fixed port is used (UDP/1719). This channel corresponds to the "H323_RAS" Application Filter and Service Group.

A common scenario is one where there are two TCP channels open, and four unidirectional UDP streams flowing. The first TCP channel is bound to a default static port (TCP port 1720), when the call is made directly using an IP address, or a port negotiated in RAS admission Request/Confirm. This channel corresponds to the "H323_VoIP" Application Filter Group and to the "H323_App" Service Group. Messages passed over this channel may include H.323, Q.931 and H.225 protocols.

Inside these messages, the IP address and port number of a second TCP (H.245) channel is exchanged. One of the endpoints then initiates a new TCP channel connection to the newly negotiated port, which is used for H.245 messages. Then, over the second TCP channel, a pair of UDP ports on either endpoint is negotiated. These UDP channels are used for RTP and RTCP messages, in both directions.

This second channel corresponds to the "H323_VOIP" Application Filter. When opening the filter, note that there are four tabs available - for H.245, RTP, RTCP and Address Translation. The values under any of these tabs may be modified as needed.

The following is an explanation of "H.323 VoIP" application filter fields for the H.245, RTP, and RTCP (all the tabs contain the same fields):

- **"Max Number of Dynamic Channels:** Maximum number of session per call that can be created from a single H.323 rule.
- **"Unused Rule Timeout:** Dynamic rule lifetime meaning how long a dynamic rule can be left unused before it's deleted.
- **"Idle Session Timeout:** Session unused timeout.
- **"Allow Destination Ports:** Dynamic ports to be allowed to open.

The following is an explanation of the "H323 VoIP" application filter parameters for the Address Translation tab:

- **"Unaligned Sequence Encoding:** Some H.323 stacks encode sequences differently from the way the Brick expects them by default, Sonus equipment is one example of this. In the event that the Brick does not seem to be properly processing some messages or there are parsing error messages in the admin events log, you can try checking this box and applying the rules to the firewall to see if this will clear up the problem.
- **"NAT to VBA:** This is the list of all addresses that will be translated to the VBA. Addresses not included are left unaltered in the message. This field takes the form of any address, address range, address subnet or hostgroup.
- **Use Ports:** Establishes a list of the range of ports to be used to perform Port Address Translation -PAT. . The ports must be in the range of 1-65535. The default value will be "32768-65535". This field must contain a minimum of one valid port number.

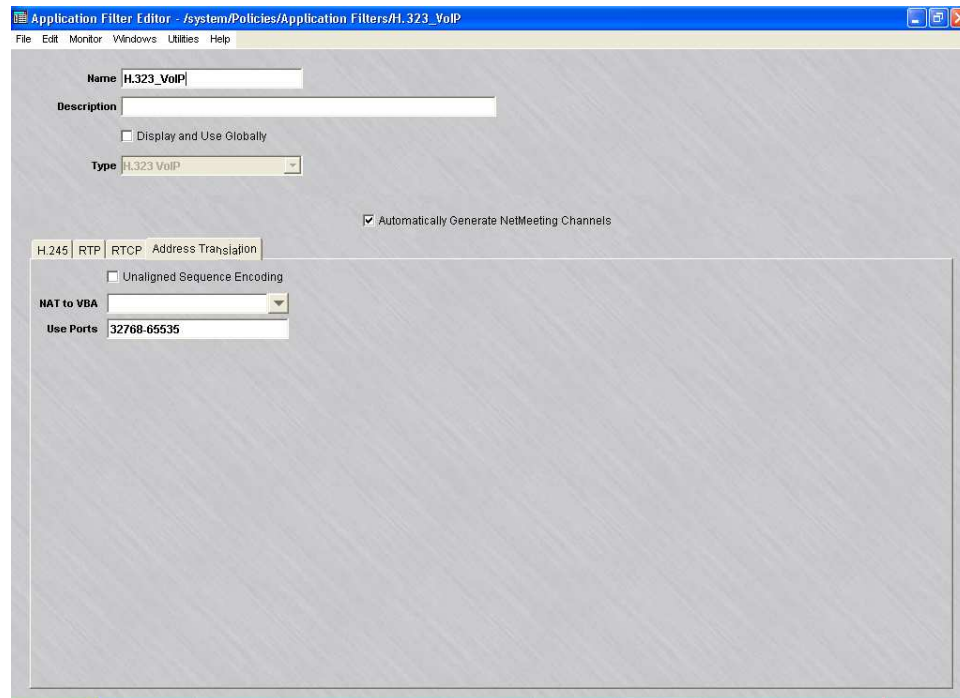
Important! The PAT feature is disabled when the NAT to VBA field is left blank.

The Brick will NAT to the VBA (and PAT as needed) whenever a packet is processed that matched a rule that used one of the H.323 application filters and that application filter contained a "NAT to VBA" entry that contains the address in the packet.

Whenever the Brick performs this translation (NAT to VBA), it will simultaneously perform a port translation (as needed) to a port that falls within the range specified in the application filter. (When a NAT is done using the existing capabilities, no port translation is done or should be done).

Important! NPAT will be performed only on those host whose IPs are included in "NAT to VBA" and the port number will be selected from port ranges specified in "Use Ports" field.

Figure 5-24, "H.323_VOIP Application Filter (Address Translation Tab)" (p. 5-49) shows the H.323_VOIP Application Filter tab for the Port Address Translation (PAT) supports.

Figure 5-24 H.323_VOIP Application Filter (Address Translation Tab)

The following is an explanation of the “H323 RAS” application filter fields for Address Translation tab:

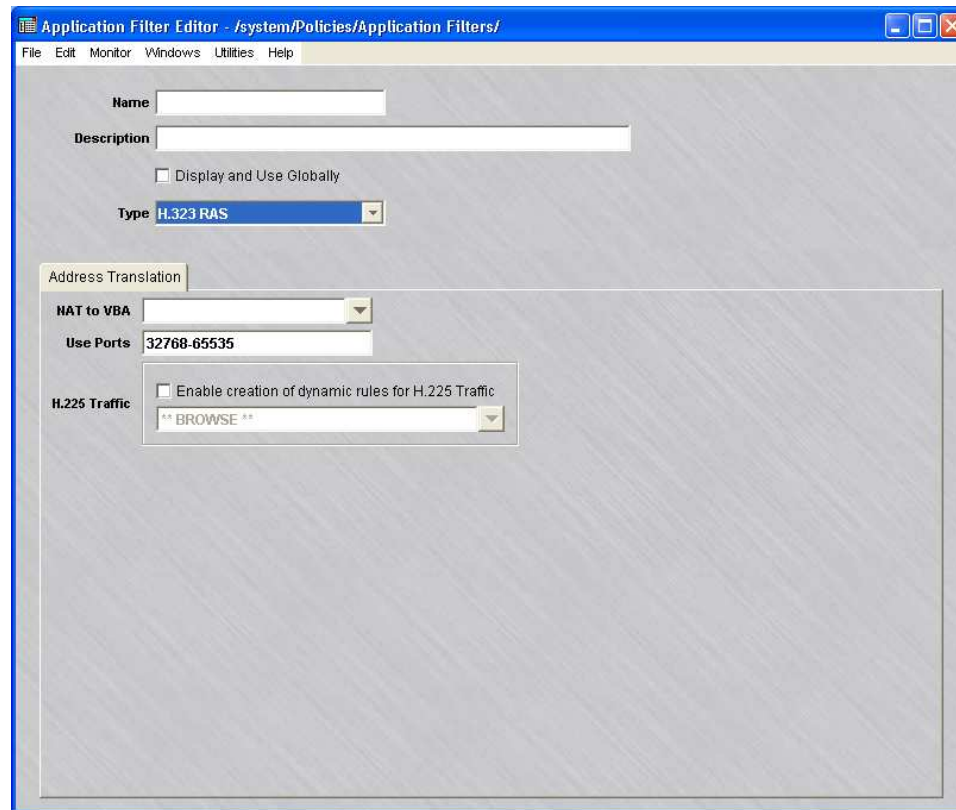
- **NAT to VBA:** This is the list of all addresses that will be translated to the VBA. Addresses not included are left unaltered in the message. This field takes the form of any address, address range, address subnet or hostgroup.

Important! The PAT feature is disabled when the NAT to VBA field is left blank.

Additional fields are:

- **“Use Ports:** Sets the range of ports to be used to perform Port Address Translation -PAT. The ports must be in the range of 1-65535. The default value will be “32768-65535”. This field must contain a minimum of one valid port number.
- **“Enable creation of dynamic rule for H.225 traffic:** When using NPAT, this button should be selected and the H.323 Application Filter used in the same zone, as the H.323 RAS Application Filter must be selected. This will enable the Brick to perform many to one NAT by means of PAT.

Figure 5-25, “H.323_RAS Application Filter” (p. 5-50) shows the H.323_RAS Application Filter.

Figure 5-25 H.323_RAS Application Filter

After making any changes to a H323 VoIP Application Filter, integrate the H323_RAS or H323_App Service Groups into rules in a Brick zone ruleset and save and apply the ruleset change(s) to the desired Brick.

If one to one NAT is required, the LVF must perform NAT only. If a Gatekeeper is used to address calls, it is important that NAT be performed with a static one-to-one mapping using "Direct NAT". Moreover, if the Gatekeeper does routed-mode calls, this NAT must be configured one-to-one statically using Host Groups.

For additional information about H.323 and application filters, see Chapter 1 VPN Firewall Brick in the SMS Technical Overview.

If many to one NAT are required, the LVF must perform not only NAT, but Port Address Translation (PAT) as well. Therefore, rules with H.323 RAS and H.323 VoIP application filters should be added to the private zones that contain the private hosts specified in "NAT to VBA" fields specified in the H.323 RAS and H.323 VoIP application filters.

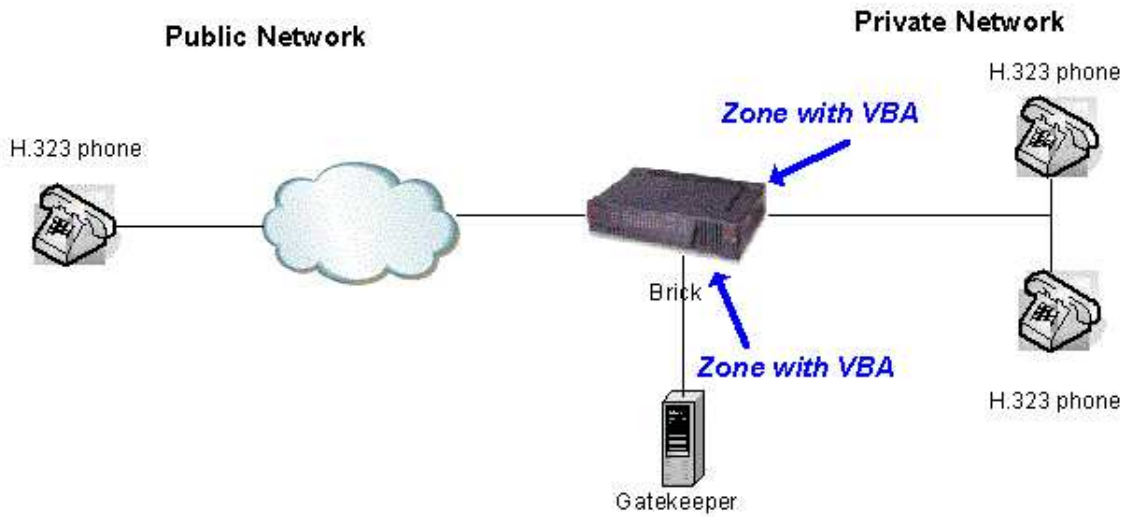
Important! H.323 application filters using NPAT must be applied to zones in which the private addresses are IN the zone (packets coming from the private addresses are coming OUT of the zone.)

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the [“Global Application Filters” \(p. 5-127\)](#) section for an explanation of global application filters.

The following are typical network topologies supported by NPAT using H.323 application filter on the Brick:

1. The Gatekeeper in the private network. In this case, the Gatekeeper’s public address is the VBA and the Brick will do destination NAT for incoming gatekeeper connections just as it does today.
2. The Gatekeeper on the public side of the Brick device. The LVF will support both the case in which the Gatekeeper handles the private addresses of the phones and the case in which it does not. In the latter case, the LVF will monitor and translate the registration requests and generate dynamic rules as needed for handling subsequent requests
3. The Gatekeeper in the DMZ ([Figure 5-26, “Network Topology - Example 3” \(p. 5-52\)](#)). In this case, the Gatekeeper may have either its own public address, or may have a private address with the NAT being done by the LVF. This is by far the most interesting case, as the NAT has to be done in the Gatekeeper’s zone and then bounced off to the other zone.

Figure 5-26 Network Topology - Example 3



□

HTTP Application Filter

Overview

An HTTP expression presented in a browser may consist of two elements — the Universal Resource Locator (URL) and the Universal Resource Indicator (URI) string. A typical URL has a protocol and a domain name defined, such as:

`http://www.alcatel-lucent.com`

If you are looking for information within a site, all of the data in the string following the URL is considered to be part of the URI string.

You can use this filter to define URI strings to be blocked (dropped) by the Brick device. Examples might include:

- *Strings that include extensions such as "mp3", "game", "exe", etc.*
- *Strings that contain a known virus.*

To configure an HTTP Application Filter

Complete the following steps to configure an HTTP application filter.

-
- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.
-
- 2 In the right-hand column, right-click the mouse and select **New Application Filter** from the pop-up menu.

Result The Application Filter Editor is displayed (Figure 5-27, “Application Filter Editor” (p. 5-54)).

Figure 5-27 Application Filter Editor

- 3 In the **Name** and **Description** fields, enter a name for the application filter and a brief description. The name is required, but the description is optional.

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the “[Global Application Filters](#)” (p. 5-127) section for an explanation of global application filters.

4 The **Commands** tab contains these fields:

- **Audit HTTP Commands (Detailed Session Audit)**

If this option is checked and if the "Session Audit" parameter in a rule is set to "Detailed", the Session Log will record and provide details for all HTTP sessions that use this filter. If unchecked, the session log does not provide detailed data.

- **Audit HTTP Exceptions (Basic Exception Audit)**

If this option is checked and the "Exception Audit" parameter in the ruleset is set to a value other than "None", the Session log will provide details on blocked HTTP sessions that use this filter.

- **Block Directory Traversal Above Server Root**

If checked, the Brick will prevent a URI from referencing a directory above or adjacent to the root directory provided by the server.

- **Treat Backslash ("/") As File Path Delimiter**

If checked, the Brick will treat the backslash character as a file path delimiter within URIs, otherwise it is treated as normal text. This setting applies if directory traversal is enabled.

- **Do Out-Of-Band URL Filtering**

If checked, the Brick will query the LPA to determine if web pages should be passed or blocked. This is an alternative option for doing URL filtering that should only be used if neither virus scanning nor proxy forwarding are enabled in the LPA. Refer to the LPA Guide for an overview of the option and to understand the level of LPA configuration needed for it.

The easiest way to explain how to configure and enable Out-of-Band URL filtering is to configure URL filtering the normal way, both on the SMS and on the LPA, and then:

- check this box
- use this application filter in the HTTP service group (or equivalent)
- use that service group in the rule with the PROXY action
- change that rule's action from PROXY to PASS (if you leave it as PROXY, the Brick will do the Out-of-Band URL filtering to the LPA *followed by* the normal reflection proxying to the LPA, which will be slow and undesired)
- in the Proxy_Listening_ports service group, make it so that UDP can be passed to the proxy port, if it's not set that way already

- **Commands**

HTTP commands such as GET, HEAD, and POST may be entered here and either specifically blocked or passed, depending on the setting of the radio button to the right of it.

- **Maximum Header Length**

This is the maximum allowable size, in bytes, of an HTTP request header. A value must be supplied here.

- **Maximum Header Length**

This is the maximum allowable size, in bytes, of an HTTP request body.

- **Enable Strict HTTP Syntax Checking**

If checked, the Brick will strictly enforce the HTTP protocol. In particular, it will:

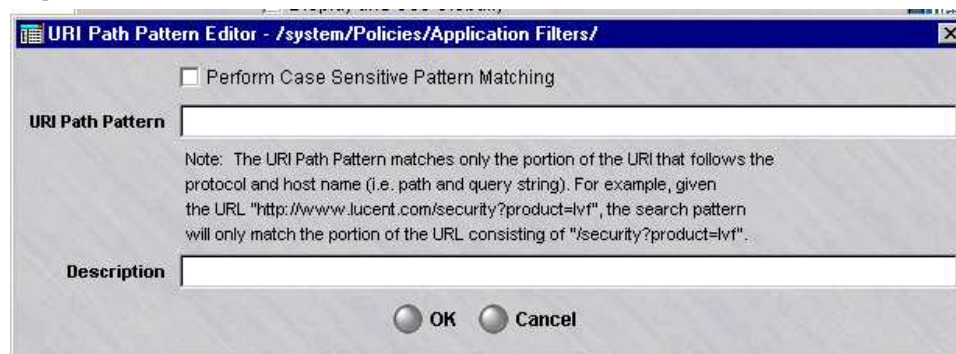
1. Ensure that the HTTP/m.n string is present in the command line and, if not, then the HTTP 0.9 is assumed and only GET is allowed.
2. Verify that unexpected spaces are not present in the request header.
3. Ensure that all header lines contain a ":".
4. Check character set validation: i.e., only printable characters can be used for keywords.
5. Verify that there are no more than three consecutive slashes.
6. Ensure that line delimiters can only be carriage return and new line.
7. Verify that line folding is not used.
8. Ensure that if there is a scheme, it must be HTTP.
9. Verify that the Content-Length occurs only once.

- **Drop re-connect attempts after a violation (for this many seconds).**

This setting applies when a command is blocked due to a violation discovered by Strict HTTP Syntax Checking. If a number, in seconds, is supplied for this field, the connection that was just denied will not be able to re-connect until this many seconds has elapsed.

-
- 5 Click **URI** to display the URI tab. You can use this tab to define the actual URI strings to be blocked. Simply right-click in the window to display the URI Path Pattern Editor. It is shown in [Figure 5-28, "URI Path Pattern Editor"](#) (p. 5-56).

Figure 5-28 URI Path Pattern Editor



Note that the URL path pattern does not match the host name. The maximum for each URI string is 4096 characters.

-
- 6 Fill in the desired URI string and click **OK** when you are done. Repeat as often as necessary. You may enter as many strings as you like for the filter.

The following is an example of a URI string:

```
/lsms  
/lsms/login.html
```

Note that you can use wildcards. The wildcard character "?" matches a single character (including "0") and the wildcard character "*" matches multiple characters. It is recommended that you add the wildcard character "*" to the beginning and the end of the URI string to be more inclusive.

If you specify just *lsms*, it will NOT block */lsms* or */lsms/login.html*. If you specify **lsms*, it will block */lsms* but NOT */lsms/login.html*. If you specify **lsms**, it will block */lsms* and */lsms/login.html*.

Every time an URI string is added, changed or deleted, you need to apply the Brick for the action to take effect.

The **Perform Case Sensitive Pattern Matching** checkbox must be checked to enable URI Path Pattern matching.

Please note that the SMS supports alternate encoding methods for the URI, including:

- %xx hex escaping
- 2 byte Unicode characters
- %Uxxxx *Microsoft*[®] encoding (where 'x' is a hexadecimal digit)

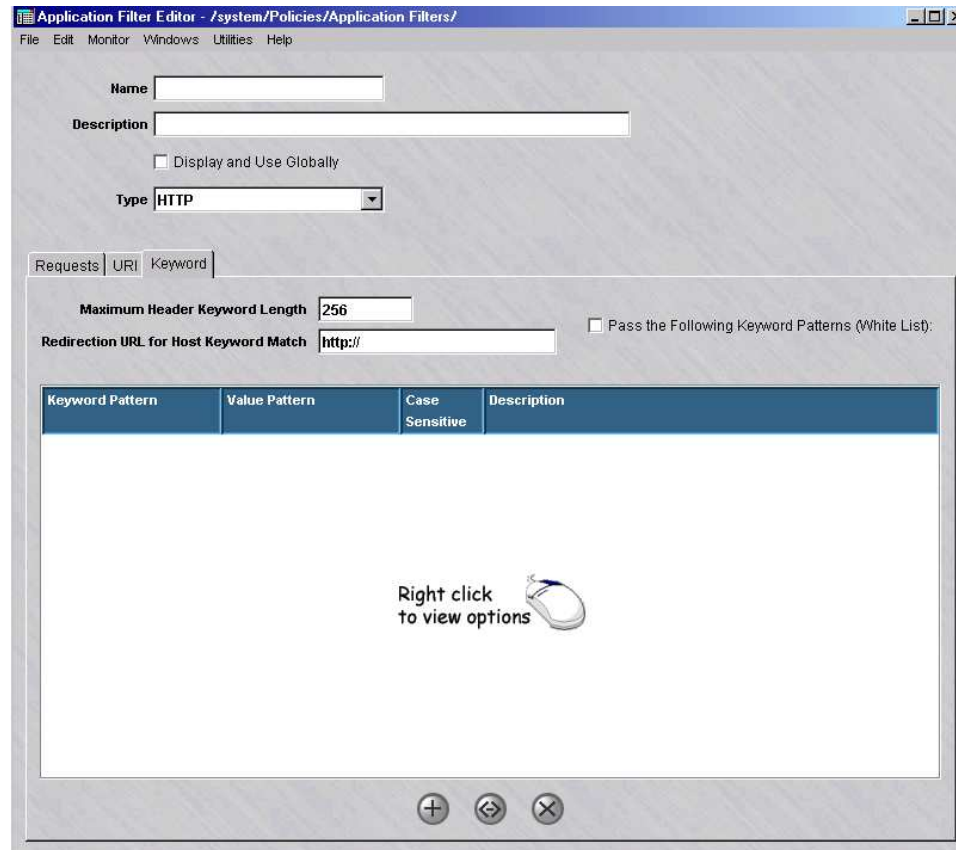
Regardless of encoding method, the Brick will fully decode the URI before attempting to perform pattern matching.

-
- 7 Click **Keyword** to display the **Keyword** tab.

Result

The Keyword tab of the Application Filter Editor is displayed ([Figure 5-29, "Application Filter Editor \(Keyword Tab\)"](#) (p. 5-58)).

Figure 5-29 Application Filter Editor (Keyword Tab)



8 Complete the tab fields as follows:

- In the **Maximum Header Keyword Length** field, you may enter the maximum length that a keyword can have in a request header.
- To use the Redirect URL feature, specify a URL in the **Redirection URL for Host Keyword Match** field. If the **Pass the Following Keyword Patterns (White List)** checkbox is unchecked, and the user is blocked as a result of a specified host keyword pattern, the user will be redirected to the specified URL. If the **Pass the Following Keyword Patterns (White List)** is checked, and the user is blocked as a result of a specified host keyword pattern that is *not* matched, the user will be redirected to the specified URL.

- To block an HTTP command matching a keyword and its value, enter patterns for both. Click the right mouse button and choose **New** to create a new entry. The Keyword Matching Pattern Editor is displayed.
You may enter as many patterns as you like for the filter. As with URI patterns, '*' and '?' are available as metacharacters. If the pattern case is significant, check the **Make Pattern Matching Case Sensitive** checkbox. Enter a textual description of the keyword in the **Description** field (optional). After you have completed entry of the keyword characteristics, click **OK**.
- To configure the keyword patterns list as white lists, click the **Pass the Following Keyword Patterns (White List)** checkbox.
By default, only the defined keyword patterns are blocked by the HTTP application filter.

-
- 9 After you have entered all the necessary data, open the File menu and select one of the **Save** options.

.....
E N D O F S T E P S
.....

To add the filter to a service group

To add this application filter to the HTTP Service Group or another existing service group, follow the steps below:

-
- 1 In the Navigator window, open the appropriate Service Group folder and double-click the HTTP service group.
 - 2 Highlight and double click the entry for tcp port 80.
 - 3 Click the pulldown for the application filter, select your HTTP application filter., and click **OK**.
 - 4 Open the File menu and select one of the **Save** options.

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information on Brick Zone Rulesets, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). For additional information on service groups, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). [Chapter 4, “Service Groups”](#).

.....
E N D O F S T E P S
.....

Pre-defined HTTP application filter for VoIP/NOE traffic

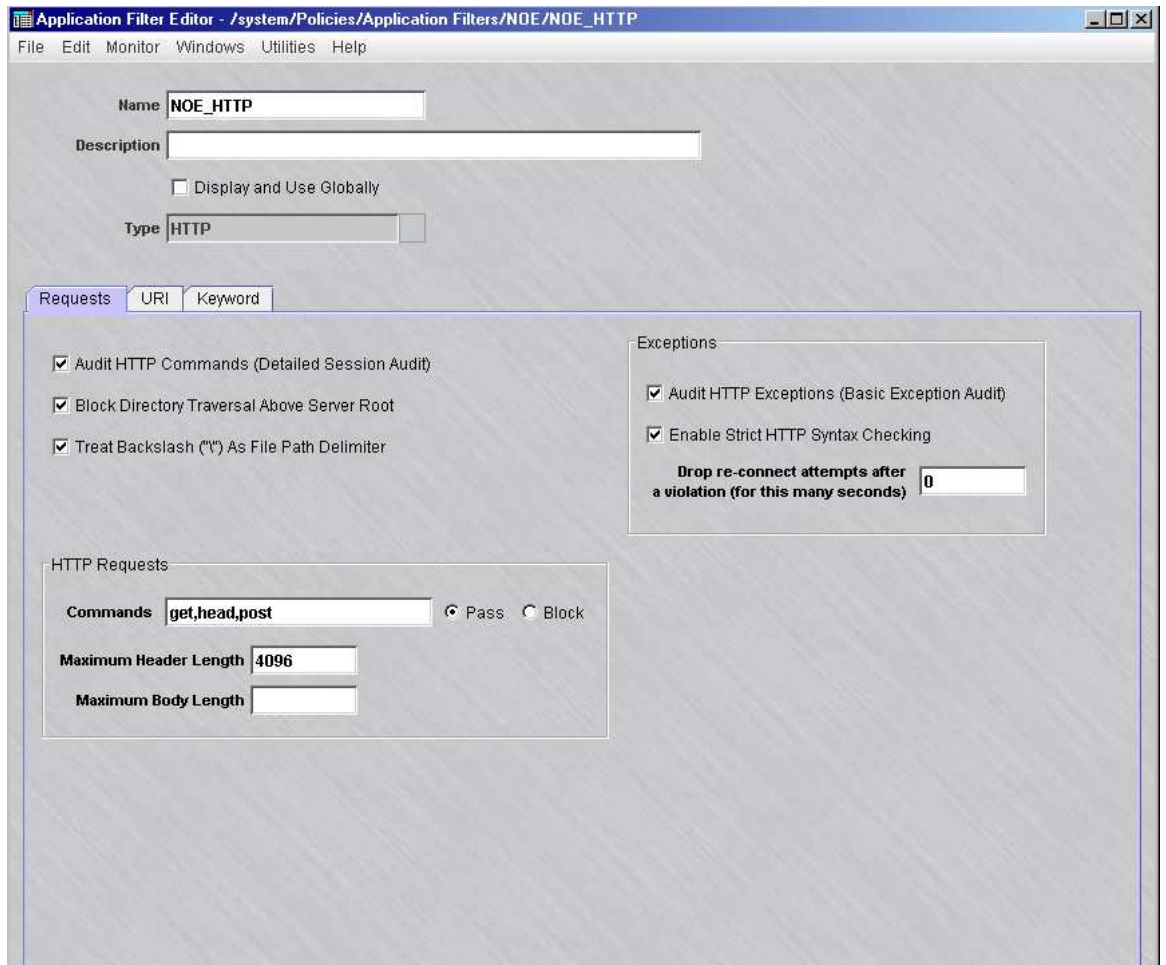
A pre-defined application filter named **NOE_HTTP** is provided with the SMS application filter when it is installed to make it easier to configure the Brick to inspect and pass HTTP URLs to/from IP phones and a presentation server in a VoIP network environment. The pre-defined **NOE_HTTP** application filter is pre-selected within a pre-defined **http_service** service group, which is, in turn, assigned to a rule in pre-defined Brick zone rulesets that can be assigned to a Brick interface.

After the SMS application is installed, the NOE templates can be downloaded by right-clicking on a group folder (such as the *system* folder initially created when the SMS is installed) and choosing **Create NOE Template** from the pop-up menu.

When the NOE template set is created, the **NOE_HTTP** application filter is automatically downloaded to an **NOE** subfolder under the **Application Filters** folder in the SMS Navigator.

This pre-defined application filter has all of the default settings of an HTTP application filter pre-selected (refer to [Figure 5-30, “Pre-Defined NOE_HTTP application filter for VoIP/NOE Traffic”](#) (p. 5-61) for an example of the Application Filter Editor window that shows the **NOE_HTTP** application filter).

Figure 5-30 Pre-Defined NOE_HTTP application filter for VoIP/NOE Traffic



NOE Application Filter

Overview

An NOE application filter is useful when a Brick device has been deployed to serve as an Application Layer Gateway (ALG) in a VoIP network. It can check, validate, and manage the bandwidth of Real Time Protocol (RTP) sessions between an IP phone device (like the NOE) and other network elements during VoIP communications.

This application filter also allows the Brick to secure IPlink signalling sessions between a Media Gateway (MGW) and call servers (CSs) in a VoIP network. This filter can be set up for the Brick to perform Network Address Translation (NAT) during establishment of the data link between the MGW and CS, replacing the private IP addresses of the GD and GA boards of the MGW behind the Brick zone with public IP addresses from a selected host group in outgoing/incoming messages to/from the CS.

After it is configured, the NOE application filter is selected and called within a TFTP application filter, which is, in turn, assigned to a service group and applied to a rule within the zone ruleset of the Brick that is protecting the call server(s) handling the VoIP voice and data traffic.

The NOE application filter parameters include an RTP session timer, a maximum activity timer for telnet debug sessions, and Quality of Service (Qos) settings for RTP sessions.

A set of pre-defined NOE application filter templates is provided with the SMS application when it is installed to make the configuration of the Brick as an Application Layer Gateway (ALG) to handle VoIP/NOE traffic easier. For more information about these templates, refer to the section [“NOE Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-68).

Additional information about the Application Layer Gateway (ALG)/NOE feature is provided in the section *Deployment of a Brick device as an application layer gateway (ALG) for VoIP/NOE phone communications* in the *SMS Administration Guide*.

Before you begin

Before you configure an NOE application filter, you can create a host group that contains the physical IP addresses of the call servers (one or two call servers can be defined in this type of host group) if redundant call servers are being used to handle VoIP traffic between IP phones in the network. Refer to [Chapter 2, “Host Groups”](#) in this Guide for details about creating a host group.

Alternatively, you can just enter the IP address(es) of the call server(s), instead of using a host group, in the **Call Servers** field of the NOE Application Filter window (refer to [Step 5](#) of the procedure).

If the Brick is being used to firewall MGW-CS communications, you must create two host groups of private IP addresses for the GD and GA boards in the MGW and two host groups of public IP addresses that the Brick can assign to the GD and GA boards when it is a NATed IPLink session (refer to [Step 13](#) and [Step 14](#) of the procedure).

To configure an NOE application filter

Complete the following steps to configure an NOE application filter.

- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.
- 2 In the right-hand column, right-click the mouse and select **New Application Filter** from the pop-up menu.
Result The Application Filter Editor is displayed.
- 3 Click the down arrow next to the **Type** field to display a drop-down menu and select **NOE**.

Result The NOE Application Filter Editor window is displayed (Figure 5-31, “NOE Application Filter Editor” (p. 5-64)).

Figure 5-31 NOE Application Filter Editor

-
- 4 In the **Name** and **Description** fields, enter a name for the application filter and a textual description, respectively. **Name** is a required field. The **Description** field is optional.

A checkbox labeled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the “[Global Application Filters](#)” (p. 5-127) section for an explanation of global application filters.

-
- 5 In the **Call Servers** field, enter one or two physical IP addresses of the call server to be used (each separated by a comma), or click the down arrow next to the field to display a list of host groups, and select the call server host group to be used (one or two call server IP addresses can be defined in this type of host group).
-
- 6 In the **Base Port** field, enter the base port of the call server to be used to pass UA signalling traffic between the call server and IP phones. The default port value is supplied initially.
-
- 7 In the **RTP/RTCP Session Lifetime (min)** field, specify the time duration of an RTP session and/or Remote TCP (RTCP) reporting activity resulting from a telnet debug session, in minutes. The default value is 5 minutes.
-
- 8 In the **Telnet Debug Max Timer (min)**, specify the maximum time limit for a telnet debugging session of an IP phone from a remote terminal, in minutes. The default value is 60 minutes (1 hour).
-
- 9 The **NOE Audit** portion of the window has the following options:
- **Session**— If this option is enabled (checkbox is checked) and the **Session Audit** parameter in a rule is set to **Basic** or **Detailed** (this parameter is set for a rule on the Brick Zone Rule Editor), the Session Log will record and provide details for all RTP sessions that use this filter. If this option is disabled (checkbox is unchecked), the Session Log does not provide any RTP session details. This option is enabled by default.
 - **Exception**— If this option is enabled (checkbox is checked) and the **Exception Audit** parameter in a rule is set to a value other than **None** (this parameter is set for a rule on the Brick Zone Rule Editor), the Session Log will record and provide details on blocked RTP sessions that use this filter.
-
- 10 To configure Real Time Protocol (RTP) Quality of Service (QoS) settings for RTP sessions, click on the **Enable QoS for RTP Traffic** checkbox.

The **Enable QoS for RTP Traffic** option is enabled by default. The related RTP QoS parameters are also enabled by default.

- 11 If the QoS for RTP Traffic option was enabled in [Step 10](#), enter values for the following fields:
- **RTP Stream QoS Priority**—Click the down arrow next to this field to display a drop-down list and select the priority for RTP data packets over other traffic.
 - **Maximum Queue Latency (ms)**—enter the time limit (in milliseconds) for the Brick device to queue congested RTP packets before discarding them (typically set to 40-60 ms).
 - **Enable In/Out Guarantee for all RTP Streams**—to set a minimum guaranteed QoS level for all incoming and outgoing RTP traffic, click this checkbox, enter a value in the field to the right of the checkbox, and select the unit measurement from the drop-down list.
Note: If the **Enable In/Out Guarantee for all RTP Streams** checkbox is unchecked, the configured value and unit measurement fields are greyed out and the option is disabled. The last configured value and unit measurement settings are retained so they can be reapplied if this option is reactivated at some later point.
 - **Enable In/Out Limit Per RTP Stream**—to set a maximum limit for each RTP stream, click this checkbox, enter a value in the field to the right of the checkbox, and select the unit measurement from the drop-down list.
Note: If the **In/Out Limit Per RTP Stream** checkbox is unchecked, the configured value and unit measurement fields are greyed out and the option is disabled. The last configured value and unit measurement settings are retained so they can be reapplied if this option is reactivated at some later point.
-

- 12 To enable the Security Optimization feature, click on the **Optimize Pin-holes** checkbox in the **Restrict pin-holes to RTP flows crossing Brick** portion of the window. This option is disabled, by default.

When this option is enabled, the Brick does not dynamically open pinholes for RTP traffic between IP hardphones/softphones that are behind the Brick. The Brick will only open pinholes for RTP traffic between IP hardphones/softphones that traverses the Brick.

Additionally, when the Security Optimization option is enabled, the **Hosts Behind Brick** field is activated, which allows you to select a host group (if configured) that contains a list of IP addresses for IP hardphones/softphones. If a phone call is placed between two IP phone devices defined in this host group, the Brick will not dynamically open pinholes for this RTP session. If no host group is selected, and a call is placed between two IP phone devices that are in a host group behind the Brick zone where the rule/application filter is being applied, the Brick will not dynamically open pinholes for the RTP session.

-
- 13** If the Brick is protecting IPlink signalling traffic between the Media Gateway (MGW) and call server(s), click the down arrow next to the **MGW GD IPs** field to display a list of host groups, and select the host group that contains the private IP address(es) of the MGW/GD board. If it is a NATed IPlink connection supported by the Brick between the MGW and call server(s), then an entry must also be made in the **MGW-GD Public IPs** field (see [Step 14](#)).

Then, click the down arrow next to the **MGW GA IPs** field to display a list of host groups, and select the host group that contains the private IP address(es) of the MGW/GA board. If it is a NATed IPlink connection supported by the Brick between the MGW and call server(s), then an entry must also be made in the **MGW-GA Public IPs** field (see [Step 14](#)).

-
- 14** If there is a NATed IPlink connection supported by the Brick between the Media Gateway (MGW) and call server(s), click the down arrow next to the **MGW-GD Public IPs** field to display a list of host groups, and select the host group that contains the pool of public IP addresses to be mapped to the private IP addresses of the MGW GD boards in the host group selected in [Step 13](#).

Then, click the down arrow next to the **MGW-GA Public IPs** field to display a list of host groups, and select the host group that contains the pool of public IP addresses to be mapped to the private IP addresses of the MGW GA boards in the host group selected in [Step 13](#).

IPlink protocol signalling messages are only exchanged between the GD board of a MGW unit and the CS(s). However, RTP/RTCP messages can be exchanged between GA and GD boards in the MGW and other network elements (call servers, IP hardphones/softphones).

-
- 15** From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.

END OF STEPS

To add the filter to a service group

The NOE application filter is called within a TFTP application filter, which is then added to a service group. For instructions on how to add a TFTP application filter to a service group, refer to the [“TFTP Application Filter”](#) (p. 5-117) section.

□

NOE Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic

Overview

A set of pre-defined NOE application filter templates are provided with the SMS application when it is installed to make it easier to configure the Brick to handle VoIP/NOE traffic. These NOE application filter templates are automatically pre-selected in the appropriate Brick zone ruleset. The name of the application filter template indicates where the Brick is being deployed to serve as an ALG in the VoIP network and whether this NOE application filter is pre-configured for the Brick to perform NAT between the TFTP server (MGW device) behind the Brick zone and call server(s) for IPlink signaling protocol messages.

Some of the pre-defined NOE application filters described in this section are also called within pre-defined TFTP application filters that have been set up to make it easier to configure the Brick to handle VoIP/NOE traffic. For details about these pre-defined TFTP application filters, refer to the section [“TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-123).

Related information

Review the following sections in this Guide for related information on use of the pre-defined templates provided with the SMS application for VoIP traffic:

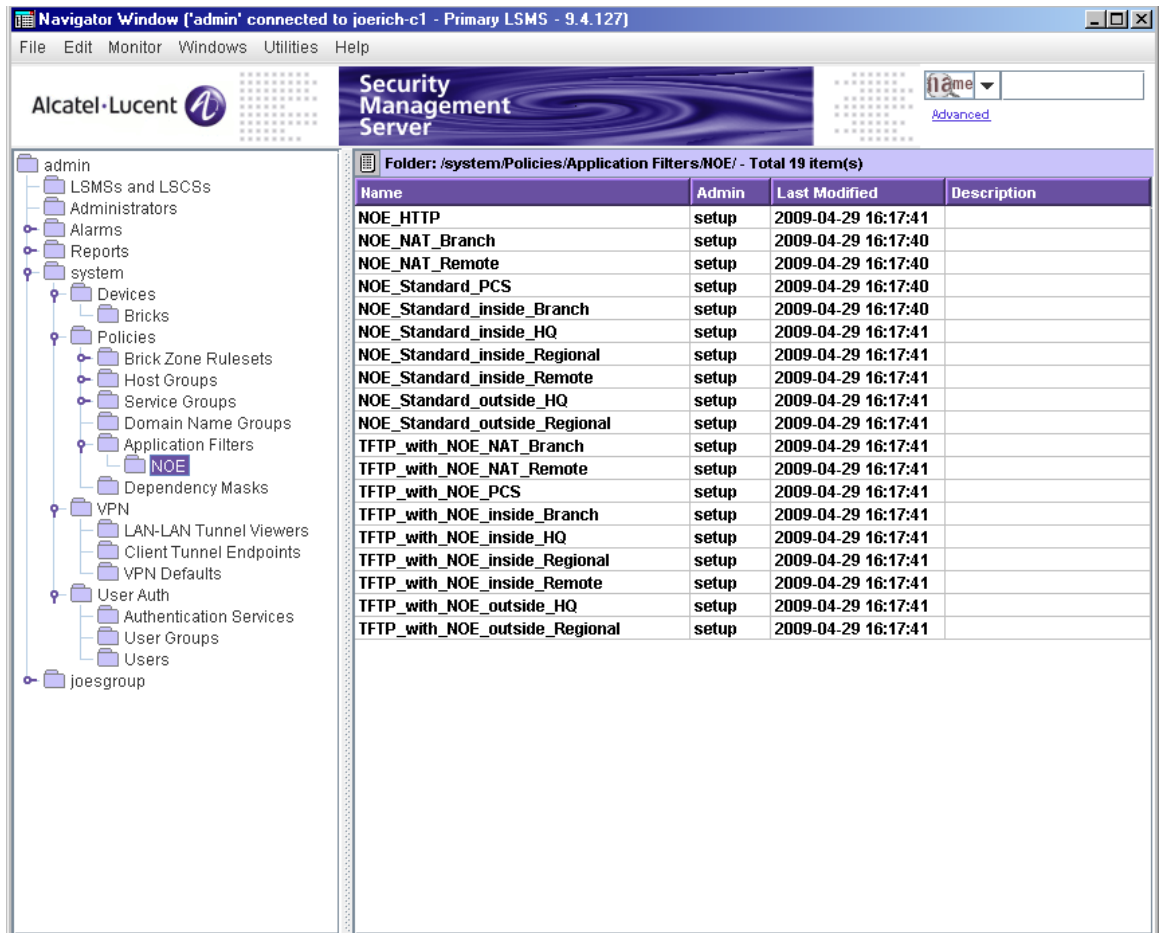
- [“Brick Zone Ruleset Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 1-66)
- [“Host Group Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 2-25)
- [“Service Group Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 4-27)
- [“TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-123)
- Appendix E, “Configuring the Brick for VoIP/NOE Traffic Using Pre-Defined SMS Templates”

Downloading the NOE application filter templates for VoIP/NOE traffic

After the SMS application is installed, the NOE templates can be downloaded by right-clicking on a group folder (such as the *system* folder initially created when the SMS is installed) and choosing **Create NOE Template** from the pop-up menu.

When the NOE template set is created, the pre-defined NOE application filter templates are automatically downloaded to a subfolder named **NOE** under the **Application Filters** folder in the SMS Navigator (refer to [Figure 5-31, “NOE Application Filter Editor”](#) (p. 5-64)) for a sample NOE Application Filter Editor window).

Figure 5-32 Pre-Defined NOE Application Filter Templates (NOE Subfolder in Navigator)



To remove the NOE zone ruleset templates and all component templates from the Policies folder, right-click on the group folder and choose **Delete NOE Template**. A confirmation dialog box is displayed, asking if you are sure that you want to delete the NOE policy template (which includes all of the zone ruleset templates are component templates) under the selected group folder. Choose **Yes**. The NOE templates are removed from the Policies folder under the selected group. If, however, one of zone ruleset templates is in use on one or more Bricks, an error message is generated and the templates are not removed.

The following sections describe the default settings within each of the pre-defined NOE application filter templates (refer to [Figure 5-31, “NOE Application Filter Editor”](#) (p. 5-64)).

NOE_NAT_Branch application filter

The **NOE_NAT_Branch** application filter has the standard NOE default settings with other fields configured as follows:

- In the **Call Servers** field, the pre-defined host group **Call Servers** is already selected.
- The RTP Security Optimization feature is enabled (checkbox is checked).
- In the **Hosts Behind Brick** field, the pre-defined host group **NOE_PHONES_MGW_Inside_Branch** is already selected.
- In the **MGW-GD IPs** field, the pre-defined host group **GD_IPs_Branch_Office** is already selected.
- In the **MGW-GA IPs** field, the pre-defined host group **GA_IPs_Branch_Office** is already selected.
- In the **MGW-GD Public IPs** field, the pre-defined host group **GD_Public_IPs_Branch_Office** is already selected.
- In the **MGW-GA Public IPs** field, the pre-defined host group **GA_Public_IPs_Branch_Office** is already selected.

NOE_NAT_Remote

The **NOE_NAT_Remote** application filter has the standard NOE default settings with other fields configured as follows:

- In the **Call Servers** field, the pre-defined host group **Call Servers** is already selected.
- The RTP Security Optimization feature is enabled (checkbox is checked).
- In the **Hosts Behind Brick** field, the pre-defined host group **NOE_Phones_Remote_Office** is already selected.
- The **MGW-GD IPs**, **MGW-GA IPs**, **MGW-GD Public IPs**, and **MGW-GA Public IPs** fields are left blank.

NOE_Standard_PCS

The **NOE_Standard_PCS** application filter has the standard NOE default settings with other fields configured as follows:

- In the **Call Servers** field, the pre-defined host group **Call Servers** is already selected.
- The RTP Security Optimization feature is disabled (checkbox is not checked).
- The **Hosts Behind Brick** field is left blank.
- The **MGW-GD IPs**, **MGW-GA IPs**, **MGW-GD Public IPs**, and **MGW-GA Public IPs** fields are left blank.

NOE_Standard_Inside_Branch

The **NOE_Standard_Inside_Branch** application filter has the standard NOE default settings with other fields configured as follows:

- In the **Call Servers** field, the pre-defined **Call Servers** host group is already selected.
- The RTP Security Optimization feature is enabled (checkbox is checked).
- In the **Hosts Behind Brick** field, the pre-defined host group **NOE_PHONES_MGW_Inside_Branch** is already selected.
- In the **MGW-GD IPs** field, the pre-defined host group **GD_IPs_Branch_Office** is already selected.
- In the **MGW-GA IPs** field, the pre-defined host group **GA_IPs_Branch_Office** is already selected.
- The **MGW-GD Public IPs** and **MGW-GA Public IPs** fields are left blank.

NOE_Standard_Inside_HQ

The **NOE_Standard_Inside_HQ** application filter has the standard NOE default settings with other fields configured as follows:

- In the **Call Servers** field, the pre-defined host group **Call Servers** is already selected.
- The RTP Security Optimization feature is enabled (checkbox is checked).
- In the **Hosts Behind Brick** field, the pre-defined host group **NOE_PHONES_MGW_inside_HQ** is already selected.
- In the **MGW-GD IPs** field, the pre-defined host group **GD_IPs_Headquarters** is already selected.
- In the **MGW-GA IPs** field, the pre-defined host group **GA_IPs_Headquarters** is already selected.
- The **MGW-GD Public IPs** and **MGW-GA Public IPs** fields are left blank.

NOE_Standard_inside_Regional

The **NOE_Standard_inside_Regional** application filter has the standard NOE default settings with other fields configured as follows:

- In the **Call Servers** field, the pre-defined host group **Call Servers** is already selected.
- The RTP Security Optimization feature is enabled (checkbox is checked).
- In the **Hosts Behind Brick** field, the pre-defined host group **NOE_PHONES_MGW_inside_Regional** is already selected.
- In the **MGW-GD IPs** field, the pre-defined host group **GD_IPs_Regional_Office** is already selected.

- In the **MGW-GA IPs** field, the pre-defined host group **GA_IPs_Regional_Office** is already selected.
- The **MGW-GD Public IPs** and **MGW-GA Public IPs** fields are left blank.

NOE_Standard_inside_Remote

The **NOE_Standard_inside_Remote** application filter has the standard NOE default settings with other fields configured as follows:

- In the **Call Servers** field, the pre-defined host group **Call Servers** is already selected.
- The RTP Security Optimization feature is enabled (checkbox is checked).
- In the **Hosts Behind Brick** field, the pre-defined host group **NOE_Phones_Remote_Office** is already selected.
- The **MGW-GD IPs**, **MGW-GA IPs**, **MGW-GD Public IPs**, and **MGW-GA Public IPs** fields are left blank.

NOE_Standard_outside_HQ

The **NOE_Standard_outside_HQ** application filter has the standard NOE default settings with other fields configured as follows:

- In the **Call Servers** field, the pre-defined host group **Call Servers** is already selected.
- The RTP Security Optimization feature is enabled (checkbox is checked).
- In the **Hosts Behind Brick** field, the pre-defined host group **NOE_PHONES_MG-W_outside_HQ** is already selected.
- In the **MGW-GD IPs** field, the pre-defined host group **GD_IPs_Outside_HQ** is already selected.
- In the **MGW-GA IPs** field, the pre-defined host group **GA_IPs_Outside_HQ** is already selected.
- The **MGW-GD Public IPs** and **MGW-GA Public IPs** fields are left blank.

NOE_Standard_outside_Regional

The **NOE_Standard_outside_Regional** application filter has the standard NOE default settings with other fields configured as follows:

- In the **Call Servers** field, the pre-defined host group **Call Servers** is already selected.
- The RTP Security Optimization feature is enabled (checkbox is checked).
- In the **Hosts Behind Brick** field, the pre-defined host group **NOE_PHONES_MG-W_outside_Regional** is already selected.
- In the **MGW-GD IPs** field, the pre-defined host group **GD_IPs_Outside_Regional** is already selected.

- In the **MGW-GA IPs** field, the pre-defined host group **GA_IPs_Outside_Regional** is already selected.
- The **MGW-GD Public IPs** and **MGW-GA Public IPs** fields are left blank.



RPC Application Filters

Overview

The Brick can inspect a Remote Procedure Call (RPC) stream for three things:

1. Message format violations
 2. Allowed program, procedure, and versions
 3. Port mapper responses that need a separate dynamic port to be opened
- The first of these capabilities is enabled by simply assigning a service group with the RPC Application Filter to a rule.
The others are controlled by the two different tabs on the RPC Application Filter screens.

To configure an RPC application filter

To configure an RPC application filter, follow the steps below:

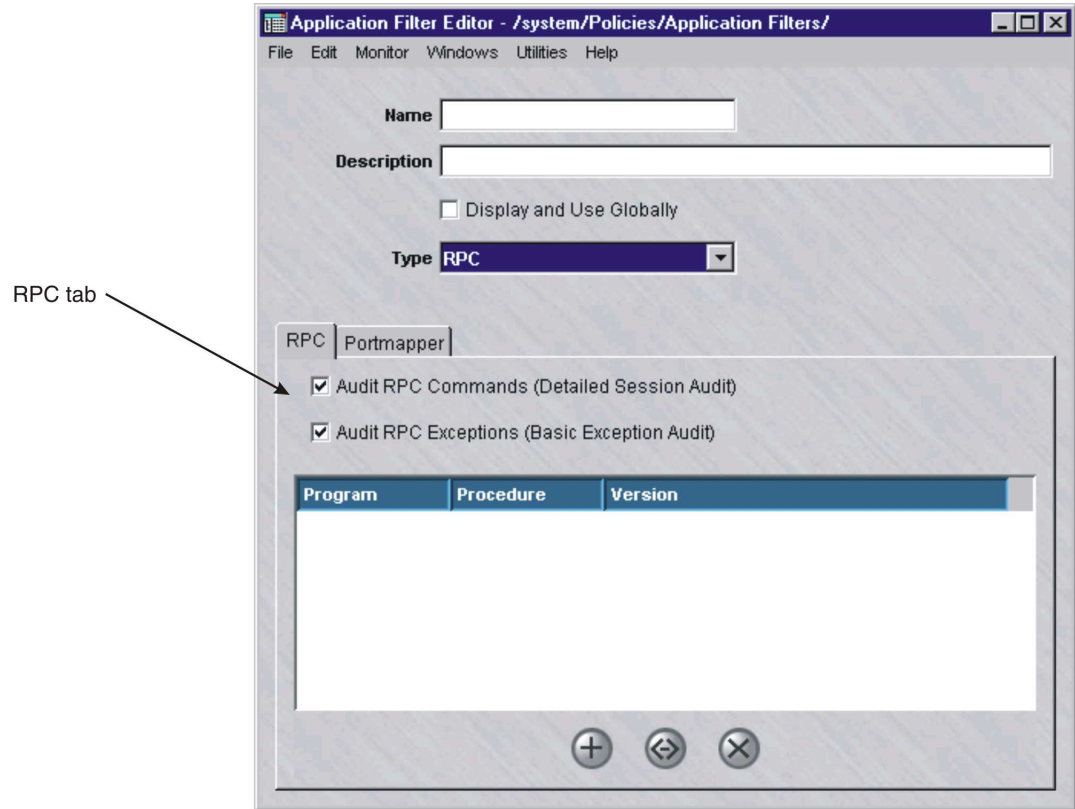
-
- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.

 - 2 Right-click the mouse and select **New Application Filter** from the pop-up menu. The Application Filter Editor will appear with HTTP as the default application filter.

 - 3 Select **RPC** from the **Type** pull-down.

Result The Application Filter Editor is displayed with the RPC tab (Figure 5-33, “Application Filter Editor (RPC Tab)” (p. 5-75)).

Figure 5-33 Application Filter Editor (RPC Tab)



- 4 In the **Name** and **Description** fields, enter a name for the application filter and a brief description. The name is required, but the description is optional.

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the “[Global Application Filters](#)” (p. 5-127) section for an explanation of global application filters.

5 The **RPC** tab contains these four fields:

- **Audit RPC Commands (Detailed Session Audit)**

If this option is checked and if the "Session Audit" parameter in a rule is set to "Detailed", the Session Log will record and provide details for all RPC calls that use this filter. If unchecked, the session log does not provide detailed data.

- **Audit RPC Exceptions (Basic Exception Audit)**

If this option is checked and the "Exception Audit" parameter in the ruleset is set to a value other than "None", the Session log will provide details on blocked RPC calls that use this filter.

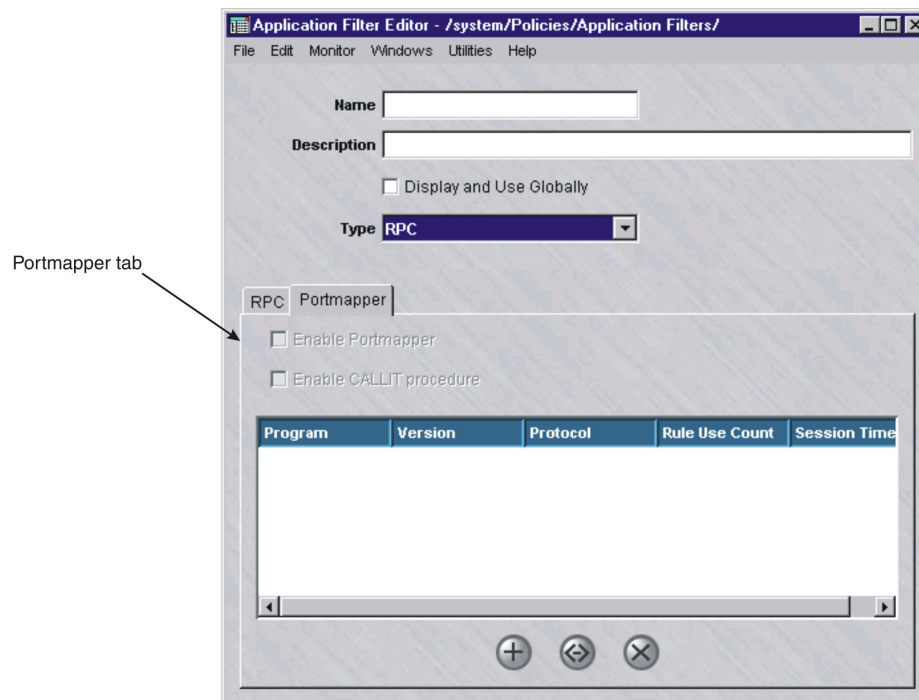
- **A table of all of the Program/Procedure/Version combinations that are allowed to pass through this filter.**

6 The **Portmapper** tab controls the ability of the firewall to dynamically create dynamic open ports in the firewall as well as to filter the programs that are requested from the portmapper.

Complete the tab as follows:

- Click Portmapper to display the portmapper tab.

Figure 5-34 Application Filter Editor (Portmapper Tab)



Do the following:

- To enable the firewall to inspect port mapper requests and to create dynamic open ports, click on Enable Portmapper. When enabled, the firewall will block all port mapper requests that do not match the list of allowed Program/Version/Protocol combinations in the table below. For those that match this triplet, Rule Use Count specifies how many sessions can use that rule and Session Timeout specifies how long a dormant session remains open, in seconds.
- To enable the caller to use the "callit" function of the portmapper, click on Enable CALLIT procedure. This function allows an RPC client to do a Remote Procedure Call to a named program/version/procedure in one step without having to come back to the actual registered port. The procedure so invoked must match the list of Program/Procedure/Version combinations on the RPC tab before the firewall will allow it.
- In the table on the bottom of the screen, enter the combinations of Program/Version/Protocol that may be requested of the portmapper through the firewall. A '*' in the protocol field means both UDP and TCP.

Important! When a port is dynamically opened via the portmapper, the secondary channel is assigned the same application filter and hence is restricted to matching the Program/Procedure/Version list in the RPC tab. It is also assigned the same QOS, VPN, and IP Address NAT as the original session.

.....
 END OF STEPS

To add the filter to a service group

To add this application filter to a service group, follow the steps below:

-
- 1 In the Navigator window, open the appropriate Service Group folder and double-click on the service group or create a new one (perhaps called rpc or portmapper).
 - 2 For portmapper, create an entry for udp, destination port 111.
 - 3 Click the pulldown for the application filter, select your RPC application filter, and click **OK**.
 - 4 Open the File menu and select one of the **Save** options.

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information on Brick Zone Rulesets, see [Chapter 1, "Alcatel-Lucent VPN Firewall Brick™ Security"](#)

[Appliance Zone Rulesets](#). For additional information on service groups, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). [Chapter 4, “Service Groups”](#).

END OF STEPS



RTSP Application Filter

Overview

Real-Time Streaming Protocol (RTSP) is a multimedia streaming protocol that is used by some applications, such as *Apple® QuickTime™* and *RealNetworks® RealPlayer®*, for the delivery of multimedia content over IP networks. It is typically based on TCP and is somewhat similar in operation and syntax to HTTP.

In client/server exchanges, the RTSP protocol is used by the client to communicate information such as the media file being requested, the application type being used, the file mechanism for the delivery (unicast, multicast, UDP, TCP), and other control information to the multimedia streaming server.

A typical exchange between a client and server in an RTSP session would be as follows:

1. The client establishes a TCP connection with the multimedia server(s), usually on TCP port 554.
2. The client issues a series of RTSP commands (SETUP, DESCRIBE) in establishing the TCP connection to the server, which provide session requirements, such as RTSP version, session transport, and UDP/TCP port information.
3. The server acknowledges and responds to the request and the session parameters with a Session ID to be used by the client in further exchanges.
4. Once the negotiation of transport parameters has been completed, the client issues a RTSP PLAY command to instruct the server to start delivery of the RTP data stream.
5. When the client decides to stop delivery of the media stream, it issues a TEARDOWN command, along with the assigned Session ID, instructing the server to stop the RTP delivery associated with that ID.

When an RTSP application filter is configured and assigned to a rule within a zone ruleset on the Brick, allowing RTSP traffic to pass, the Brick monitors and inspects all RTSP traffic between the client and server. The Brick receives RTSP messages on a configurable destination port (default TCP port 554). The application filter that is assigned to the rule can specify two host groups, which contain a list of allowed media source addresses and client destination addresses, respectively, for the RTSP media streams. Redirect options can be enabled within the application filter to allow RTSP media streams to be sent to different client IP addresses, or, if the media server is different from the RTSP server, to originate from a different server IP address. A different media server IP address may be defined within a source host group specified by the application filter.

Additional dynamic rules are created by the Brick when it sees the server response to the SETUP command. These rules define dynamic RTP pinholes (ports) for passing RTSP media streams back and forth between the client and server.

When the client transmits a TEARDOWN command to the server to terminate the media transfer session, the Brick decodes this message and forces the media sessions to idle out quickly.

As part of the filtering process, the Brick inspects the lengths of all parts of the RTSP message to ensure that it conforms with the message limits set in the application filter, to protect against buffer overflow attacks.

Other session controls, such as a maximum session idle timeout and number of unanswered RTSP client requests allowed, can be provisioned within the application filter.

Important! The RTSP application filter does not currently support Network Address Translation (NAT).

Related information

For more information about the Rules-Based Routing feature, refer to the “[Rules-based routing](#)” (p. 1-4) section in [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#).

To configure an RTSP application filter

Complete the following steps to configure an RTSP application filter:

- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.
- 2 In the right-hand column, right-click the mouse and select **New Application Filter** from the pop-up menu.
Result The Application Filter Editor is displayed.
- 3 Click the down arrow next to the **Type** field to display a drop-down menu and select **RTSP**.

Result The RTSP Application Filter Editor is displayed, initially showing the Options tab (Figure 5-35, “RTSP Application Filter (Options Tab)” (p. 5-81)).

Figure 5-35 RTSP Application Filter (Options Tab)

The screenshot shows the 'Application Filter Editor' window with the 'Options' tab selected. The configuration is as follows:

- Name:** [Empty text box]
- Description:** [Empty text box]
- Display and Use Globally
- Type:** RTSP (dropdown menu)
- Options Tab:**
 - Maximum Lengths:**
 - Entire Message: [Empty text box]
 - Message Header Part: [Empty text box]
 - Message Body Part: [Empty text box]
 - Header Line: 500
 - Request/Response Line: 500
 - Keyword: 200
 - Control:**
 - Strict Syntax Checking
 - Allow Client Address to Change
 - Block Directory Traversal Above Server Root
 - Maximum Session Idle Timeout: 300
 - Pipeline Limit: 4
 - Media:**
 - Media should use Rules Based Routes from the RTSP rule (if any)
 - Client can specify a destination multicast address
 - Transport:** UDP Interleaved TCP TCP
 - Dynamic Ports Range: 1024-65535
 - Maximum Multicast Layers: 1
 - Maximum Multicast TTL: [Empty text box]
 - Allowed media destination addresses (client): [Empty dropdown menu]
 - Allowed media source addresses (server): [Empty dropdown menu]
 - RTSP Audit:** Session Exception
 - Drop re-connect attempts after a violation (for this many seconds): [Empty text box]

- 4 In the **Name** and **Description** fields, enter a name for the application filter and a textual description, respectively. **Name** is a required field. The **Description** field is optional.

A checkbox labeled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the “[Global Application Filters](#)” (p. 5-127) section for an explanation of global application filters.

- 5 In the **Maximum Lengths** portion of the window, specify the maximum character length (limit) for each part of a RTSP message received, as follows:
- **Entire Message**— this field limits the maximum length of a message. The valid range is 500-999,999 characters. If this field is left blank, no limit is imposed.
 - **Message Header Part**— this field limits the maximum length of the header part of the message. The valid range is 500-999,999 characters. If this field is left blank, no maximum length is imposed.
 - **Message Body Part**— this field limits the maximum length of the body part of the message. This field also limits the length of interleaved media (if any). The valid range is 500-999,999 characters. If this field is left blank, no maximum length is imposed.
 - **Header Line**— this field limits the maximum length of each header in the message. The valid range is 100-9,999 characters. If this field is left blank, no maximum length is imposed.
 - **Request/Response Line**— this field limits the maximum length of the first line of the message. The valid range is 100-9,999 characters. If this field is left blank, no maximum length is imposed.
 - **Keyword**— this field limits the maximum length of each keyword of the message. The valid range is 100-9,999 characters. If this field is left blank, no maximum length is imposed.
-
- 6 In the **Control** portion of the window, configure the following RTSP options:
- **Strict Syntax Checking**— when this option is enabled (checkbox is checked), RTSP messages are checked more strictly for general conformance to the protocol specification (RFC 2326). This option is enabled, by default.
 - **Allow Client Address to Change**— when this option is enabled (checkbox is checked), it allows client RTSP requests to originate from different source IP addresses (as defined in the host group). This option is disabled, by default.
 - **Block Directory Traversal Above Server Root**— when this option is enabled (checkbox is checked), the application filter ensures that there are no more “previous directory” sequences (“../”) before the directory path in RTSP commands. This option is enabled, by default.

- **Maximum Session Idle Timeout**— this field allows you to specify the maximum idle timeout value (in seconds) that the server may specify. The valid range is 1-99999. If this field is left blank, the Session Timeout value is taken from the policy rule on which the application filter is applied.
 - **Pipeline Limit**— this field allows you to specify the maximum number of unanswered (outstanding) RTSP client messages. The valid range is 1-32. The default value is 4. If it is set to 1, then a response must be received before the next message can be sent.
-

7 The **RTSP Audit** portion of the window has the following options:

- **Session**— if this option is enabled (checkbox is checked) and the **Session Audit** parameter in a rule is set to **Basic** or **Detailed** (this parameter is set for a rule on the Brick Zone Rule Editor), the Session Log will record and provide details for all RTSP sessions that use this filter. If this option is disabled (checkbox is unchecked), the Session Log does not provide any RTSP session details. This option is enabled by default.
 - **Exception**— if this option is enabled (checkbox is checked) and the **Exception Audit** parameter in a rule is set to a value other than **None** (this parameter is set for a rule on the Brick Zone Rule Editor), the Session Log will record and provide details on blocked RTSP sessions that use this filter.
-

8 In the **Drop reconnect attempts after a violation (for this many seconds)** field, enter the number of seconds to block additional RTSP traffic from the specified client source address(es) to any destination address when a message is blocked.

Important! Care should be taken in configuring this option, since a spoofed message could cause traffic from a legitimate IP address to be blocked.

9 In the **Media** portion of the window, configure the following options:

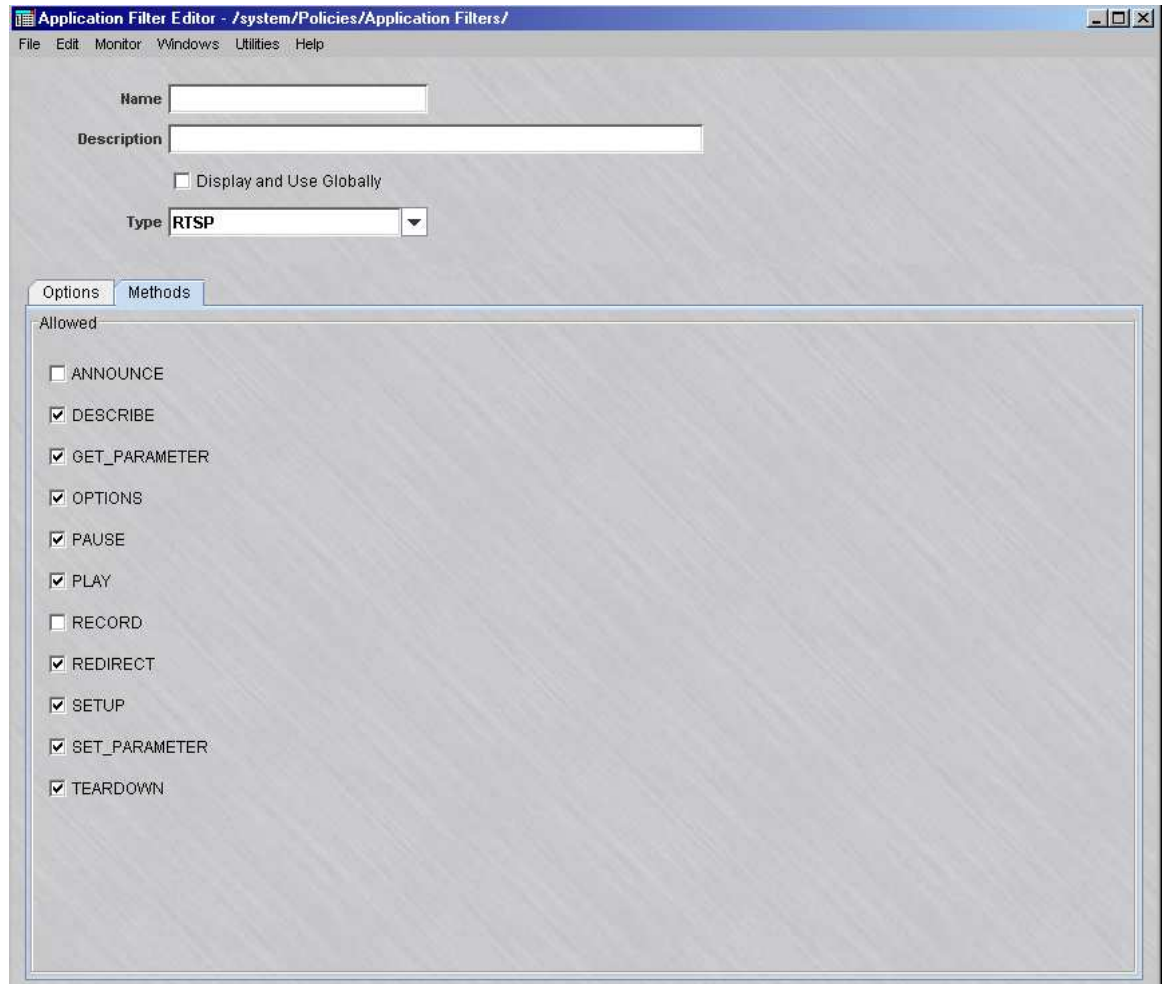
- **Media should use Rules Based Routes from the RTSP rule (if any)**— if this option is enabled (checkbox is checked), and Rules Based Routing is used on the RTSP rule to which this application filter is applied, the same Rules Based Routing is applied to the media pinhole rules dynamically generated by the Brick. This option is disabled, by default.
 - **Client can specify a multicast destination address**— if this option is enabled (checkbox is checked), the client is allowed to specify the destination address for multicast media. This option is disabled, by default.
 - **Transport** — choose one or more valid transport layer protocols for the media streams controlled by the RTSP session: **UDP**, **Interleaved TCP**, **TCP**. **Interleaved TCP** media is sent on the same TCP connection as the RTSP protocol. **UDP** and **Interleaved TCP** are the default protocols.
-

- **Dynamic Ports Range**— all media ports must fall within this specified range. The valid port range is 0-65535. The default port range is 1024-65535.
 - **Maximum Multicast Layers**— when media is sent to a multicast address, this field limits the number of sequential addresses that may be used for media transport.
 - **Allowed media destination addresses (client)**— this field is used to specify the alternative client IP address(es) for receiving RTSP media streams. Valid entries are an individual IP address, IP address range, or click the down arrow next to the field to display a drop-down menu, select **Browse**, which brings up a Select a Host Group window, and select a host group of valid client destination IP addresses. This field must contain one or more multicast addresses if multicast media is used. If this field is left blank, the RTSP media stream is sent to the client's IP address (the default).
 - **Allowed media source addresses (server)**— this field is used to specify the alternative media source IP address(es) if RTSP responses and media streams can be sent from another source address. Valid entries are an individual IP address, IP address range, or click the down arrow next to the field to display a drop-down menu, select **Browse**, which brings up a Select a Host Group window, and select a host group of valid media server (source) IP addresses. If this field is left blank, the RTSP media stream originates from the RTSP server's IP address (the default).
-

10 Click on the **Methods** tab.

Result The Methods tab of the RTSP Application Filter Editor is displayed (Figure 5-36, “RTSP Application Filter Editor (Methods Tab)” (p. 5-85)).

Figure 5-36 RTSP Application Filter Editor (Methods Tab)



- 11 Select the RTSP messages (commands) that can be exchanged between the client and server. If the **Strict Syntax Checking** option is not checked on the Options tab, any other command not shown on this tab is permitted.
- 12 From the **File** menu, select **Save and Close** to save the filter and close the Application Filter Editor window.

END OF STEPS

To add the filter to a service group

Complete the following steps to add this application filter to an existing service group. For complete details about creating a service group, refer to [Chapter 4, “Service Groups”](#).

- 1 In the Navigator window, open the Service Group folder and double-click on an existing service group.

Result The Service Group Editor is displayed.

- 2 In the **Name** and **Description** fields, enter a new name and description, respectively, for this service group, or leave the default entries to use the existing service group.

- 3 Click the pulldown for the application filter, select your RTSP application filter, and click **OK**.

- 4 Open the **File** menu and select one of the **Save** options.

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information on Brick Zone Rulesets, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). For additional information on service groups, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). [Chapter 4, “Service Groups”](#).

END OF STEPS



SIP Application Filters

About Firewalling SIP

This section discusses security aspects of the SIP protocol. More detailed configuration guidance is given in subsequent sections. The **Session Initiation Protocol (SIP)** is a signalling protocol used in applications such as Voice over IP (VoIP) and Instant Messaging services. SIP users have SIP IDs that function much like real-time email ID's. For instance, <sip:user1@example.com>. In a typical VoIP application, user1 registers his/her current "location", for example his/her PC in the corporate network user1-pc.branch2.example.com. A SIP Proxy server in the "example.com" domain has access to the database of current locations. Other users may now "call" user1 using the SIP URI sip:user1@example.com. The caller may use his/her own (outbound) SIP proxy to locate the (inbound) SIP proxy of the person he is calling. The software that implements the SIP part of the phone or soft phone is known as a User Agent.

There are three basic kinds of flows in a SIP call:

1. Initial requests and responses (INVITE)
2. Requests and responses once a dialog has been established (Example: BYE)
3. RTP media streams

The endpoints in each of the above SIP flows must be the endpoints of the initial request or otherwise allowed in the **Names and Other Addresses** table (tab) on the SIP Application Filter Editor as **Any** or **Media**.

With respect to firewalls, SIP calls are more complicated than most protocols. An ordinary firewall rule authorizes one address from the allowed source addresses to exchange data with another address from the allowed destinations. Typically the data is sent from some arbitrary "ephemeral" port number to some well-known port number. For example, accessing a Web page on the Internet permitted by a single rule allowing access to the web server at port 80. In fact, the access may be done with many TCP streams (connections), but with respect to the firewall aspect, that is irrelevant. More complicated protocols like FTP have an initial TCP stream for control and separate streams for each file transferred. The IP addresses are the same in the data stream, but the port numbers will be different and the role of source and destination may be reversed. The FTP control protocol contains the information necessary for the firewall to "open up a pinhole through the firewall" for each file transfer. The firewall simply has to read the protocol to learn the port numbers that will be used, and create a temporary rule allowing it. The case with SIP is more involved from a security standpoint. A SIP *dialog* or call extends over several UDP or TCP streams. The firewall must be able to keep track of them as a unit. The initial stream may be between an SIP proxy in the zone and some SIP proxy in the Internet. Subsequent SIP messages in the dialog may be passed directly from the User Agent in the zone to one

in the Internet or perhaps some other SIP proxy. However, the firewall cannot simply assume that it is OK when it sees a new address as an endpoint. Your security policy must authorize that address.

The SIP message contains headers showing:

- the SIP proxies that have forwarded the message so far,
- the final destination and
- an optional list of specific SIP proxies the message should be routed to on its way to the final destination. (application level routing)

Replies to this message will visit all the listed SIP proxies in opposite order. Proxies can be specified by name or address or both. Even the port number can be specified. This flexibility gives a would-be attacker a lot to play with. For example, he can create a SIP message that claims to be from a user in your domain, even though it is originating in the Internet. (While SIP user authentication can be enabled to prevent this, it is a good idea to put a separate check in the firewall.) The attacker can also route SIP messages to one of your internal http servers (perhaps as part of a Denial of Service attack). The SIP Application Filter is designed to enforce a basic security policy. By default, routing messages to other than port 5060 is disabled by the filter. Starting from this base, you can add an entry that allows messages from your domain to leave your security zone, and prevent any message from entering your zone while claiming to be from it. This is done by adding entries to the **Names and Other Addresses** table that name your domain and configuring the **Zone Side** and **Allow As** fields as described in the following sections. In the context of a SIP Application Filter, the term **Outside the Zone** refers to any SIP User Agent or Proxy located outside of the security zone and **Inside the Zone** to those on the inside of the security zone. For example, "outside" user agents maybe prevented from issuing an INVITE ("make a call") command.

The default **Names and Other Addresses** table contains an entry for **Rule Match** that permits any of the addresses in the original rule. Other addresses must be explicitly added to the table. For example, if the media is sent to a media gateway, that gateway must be allowed by the **Names and Other Addresses** table as **Any** or **Media**.

All addresses mentioned in the SIP message in the SIP request URI, From, To, Contact, Via, Route and Record-Routes headers are checked. The ones for which no pinhole is necessary are allowed as **SIP Routed** or **Any**.

There is a default table entry that allows all addresses with the default port (5060) as **SIP Routed**; the administrator typically does not have to modify these addresses. However, a Via that specifies port 80 (for example), would not be allowed by the default table.

Pinholes for the media are opened once the Brick device receives the SDP answer and recognizes both the source and destination address.

In most cases, the zone should have the policy to protect the User Agents and Proxies for which you are responsible. However it is possible to invert the roles played by "Outside" and "Inside". You can, for example, apply a SIP Application Filter on a public interface with Internet access. In this case, Internet hosts reside "Inside" the public zone. The "Outside" hosts are "behind" the Brick on other ports.

Basic configuration

The Zone Administrator enables SIP calls by creating a SIP Application Filter, a Service Group that uses the filter (see chapter on Service Groups), and configuring a pair of rules with this group. The rule allowing calls to phones in the zone has the direction *In to* the zone, and a well-known destination port - which is normally port 5060. The destination host group is a list of all addresses that may receive SIP messages. If all calls initially go through an inbound SIP proxy, then that is the only address that should be configured. If calls may go directly to the User Agent in the zone, then the addresses of the phones must be entered as destination addresses as well. The source address field can be used to constrain who can make calls to phones inside the zone. Similarly, if the phones are allowed to make outbound calls, then the administrator must create a rule for the *Out of* the zone direction and the SIP destination port. In this direction, the host group of SIP proxies and/or phones is placed in source column, and restrictions on who may be called, if any, are placed in the destination column. Port 5060 is the default for sip. Port 5061 is the default for sips. When sips is used, the entire message is encrypted, and the firewall cannot read its contents. Therefore the SIP Application Filter is not useful. If sips secured messages are allowed through the firewall, care should be taken that appropriate security policy is enforced in the application.

As a call is being setup, the source and / or destination of the signaling messages may change. For example, the destination may progress from the initial proxy to one "closer" to the user and finally to the address of the phone where the user is currently located. Whether this happens or not depends whether the proxies decide to "stay in the loop" or "drop out" once they have determined the downstream destination address. If the addresses can change, then the additional addresses must be configured in the **Names and Other Addresses** of the SIP Application Filter.

Important! The nature of the SIP protocol requires the Firewall to open up many pinholes based on the type of transaction being performed: voice, video, conference, presence, instant messaging, etc. Therefore, in order for a SIP call or transaction to properly complete a SIP Application Filter will be required on every zone that a SIP message is required to traverse. This is the most secure way to handle SIP calls.

Network Address /Port Address Translation (NAT/PAT) configuration

The following explains aspects of NAT/PAT configuration:

- The typical NAT scenario is a zone configured with private addresses, where the phones in the zone must be able to make calls to and receive calls from phones with public addresses. In this case, inbound calls should be directed to the Virtual Brick Address (VBA) and port 5060 or some other well-known port number. This address should be translated using the normal destination NAT feature of the rule to the private address of the proxy. Alternatively the proxy may have its own public address and reside in a "Demilitarized Zone" with other public servers. All private phones should be registered with this proxy. As the call progresses, the Brick will dynamically translate additional addresses in the SIP dialog to the Brick VBA and some dynamic port. This mapping exists only for the duration of the call. If the rule itself did not specify NAT' ting (in other words, if the proxy has a public address) then mapping must be specified in the SIP Application Filter. A zone must have a VBA whenever a rule with a SIP Application Filter specifies NAT. All NAT' ting must be to/from the VBA. When NAT' ting is used for a call originating in the private zone, the Brick essentially appears as the source User Agent; applications will see SIP messages with the Via and Contact of the Brick VBA.

The following are typical SIP network topologies supported by SIP Application Filter for the current software release:

- The SIP Proxy in the private network ([Figure 5-37, "SIP Proxy in the private network" \(p. 5-91\)](#)). In this case, the Proxy public address is the VBA and the Brick does destination NAT for incoming SIP messages.
- The Proxy on the public side of the Brick ([Figure 5-38, "SIP Proxy on public side of the Brick" \(p. 5-92\)](#))
- The Proxy in the DMZ ([Figure 5-39, "SIP Proxy in the DMZ" \(p. 5-92\)](#)). In this case, the Proxy may have either its own public address or may have a private address with the NAT being done by the LVF.

Figure 5-37 SIP Proxy in the private network




Figure 5-38 SIP Proxy on public side of the Brick





Figure 5-39 SIP Proxy in the DMZ



To create a SIP application filter

Complete the following steps to configure a SIP application filter:

- 1 From the SMS Navigator, open the Policies folder in the desired group and click Application Filters;
- 2 In the right hand column, right-click the mouse and select New Application Filter from the pop-up menu. After the Application Filter Editor appears, click the Type drop-down menu and select SIP. The window shown in [Figure 5-40, “SIP Application Filter Editor \(Options Default Tab\)”](#) (p. 5-93) will appear.

Figure 5-40 SIP Application Filter Editor (Options Default Tab)

- 3 In the **Name** and **Description** fields, enter a name for the application filter and a brief description. The name is required, but the description is optional.

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the [“Global Application Filters” \(p. 5-127\)](#) section for an explanation of global application filters.

4 The **Options** tab contains these fields:

- **Maximum Lengths - Entire Message**
Sets the maximum length of the entire SIP message. The default is blank, which implies no limit. If set to a value, it must be greater than the **Maximum Lengths - Message Header Part** and **Maximum Lengths - Message Body Parts** settings, if specified.
- **Maximum Lengths - Message Header Part**
Sets the maximum combined length of all headers in a SIP message. The default is blank, which implies no limit.
- **Maximum Lengths - Message Body Parts**
Sets the maximum length of message body parts. The default is blank, which implies no limit.
- **Maximum Lengths - Header Line**
Sets the maximum length (default = 500) of a SIP header line such as:
Via: SIP/2.0/UDP bobspc.lucent.com:5060;branch=z9hG4bKnashds7;
received=192.0.2.4
or
Route: <sip:alice@lucent.com>,<sip:carol@lucent.com>,
<sip:bob@lucent.com>
The length is the length of the entire logical header *line*, including all its continuation lines (lines that begin with a blank or a tab).
- **Maximum Lengths - Request/Response Line**
Sets the maximum length (default = 500) of a request/response line such as:
INVITE sip:bob@lucent.com SIP/2.0
SIP/2.0 200 OK
- **Maximum Lengths - Keyword**
Sets the maximum length of keywords, such as header names, parameter names, etc. The default value is 200.
- **Configure Inside and Outside of Zone Identically**
If checked (the default), configures the SIP application identically both inside and outside of the zone. If not checked, the **Maximum Forward Count** and **Dynamic Ports** fields are presented for both inside and outside the zone (see [Figure 5-41, “SIP Application Filter Editor \(Options Tab\)- both zones” \(p. 5-95\)](#)).
- **Inside and Outside of Zone - Maximum Forward Count**
Can be set from 1 to 255. The default is 70.
- **Inside and Outside of Zone - Dynamic Ports**
Sets the range that dynamic ports may take. The default is 1024-65535.

Figure 5-41 SIP Application Filter Editor (Options Tab)- both zones

The screenshot shows the 'Options' tab of the SIP Application Filter Editor. The window title is 'Application Filter Editor - /system/Policies/Application Filters/sipDefault'. The 'Name' field is 'sipDefault'. The 'Type' is 'SIP'. The 'Options' tab is selected, showing the following configuration:

- Maximum Lengths:**
 - Entire Message: []
 - Message Header Part: []
 - Message Body Parts: []
 - Header Line: 500
 - Request/Response Line: 500
 - Keyword: 200
- SIP Audit:**
 - Session Exception
 - Drop re-connect attempts after a violation (for this many seconds): []
 - Media Transport: UDP TCP
 - Media Maximum Streams in Session: 10
 - Session Media VPN: Like rule
 - Signaling-Only Calls: Timeout (sec) 3600 Never Timeout
 - Calls on Hold Timeout (sec): 3600
- Configure Inside and Outside of Zone Identically
- Inside of Zone:**
 - Maximum Forward Count: 70
 - Dynamic Ports: 1024-65535
- Outside of Zone:**
 - Maximum Forward Count: 70
 - Dynamic Ports: 1024-65535

The following explains these fields:

- **SIP Audit - Session**

If this checkbox is not selected, no SIP-specific audit messages will be generated. If the option is selected, audit records for SIP will be generated depending on the values in the session audit field on the zone rule. If the rule audit is set to **None**, then no SIP audits will be generated. If it is set to **Basic**, then an audit for each successful SIP dialog will be generated. If it is set to **Detailed**, then an audit record will be created for each SIP message.

- **SIP Audit - Exception**

If this checkbox is selected and the zone rule for Audit Exception field is set to **Basic** or **Detailed**, audit records will be generated for errors encountered while filtering SIP messages.

- **Drop re-connect attempts after a violation (for this many seconds)**

The number of seconds for which packets with the current IP source address will be discarded after a violation of this filter. Care should be taken when using this option because IP source addresses do not always reflect the true source of the packet. For example, they may be deliberately "spoofed" or may have been inserted by an upstream NAT device.

- **Media Transport**

Two media transport options are available, UDP and TCP. UDP is checked as the default.

- **Media Maximum Streams in Session**

A number from 0 to 9999. The default is 10. This option allows five RTP streams, each paired with an RTCP stream.

- **Session Media VPN**

The following options are available: **Like rule** and **No**. The default is **Like rule**, which treats media streams the same as the SIP streams with respect to VPN tunneling. Select **No** if the media should not be sent through a tunnel, even when the SIP messages are.

This option applies when the media address is the same as the SIP User Agent address. In cases where the addresses are different and the address is taken from the Names and Other Addresses table, the value of Session Media VPN in that table determines whether that media will be tunneled.

- **Signaling-Only Calls** (Configurable Timeout setting)

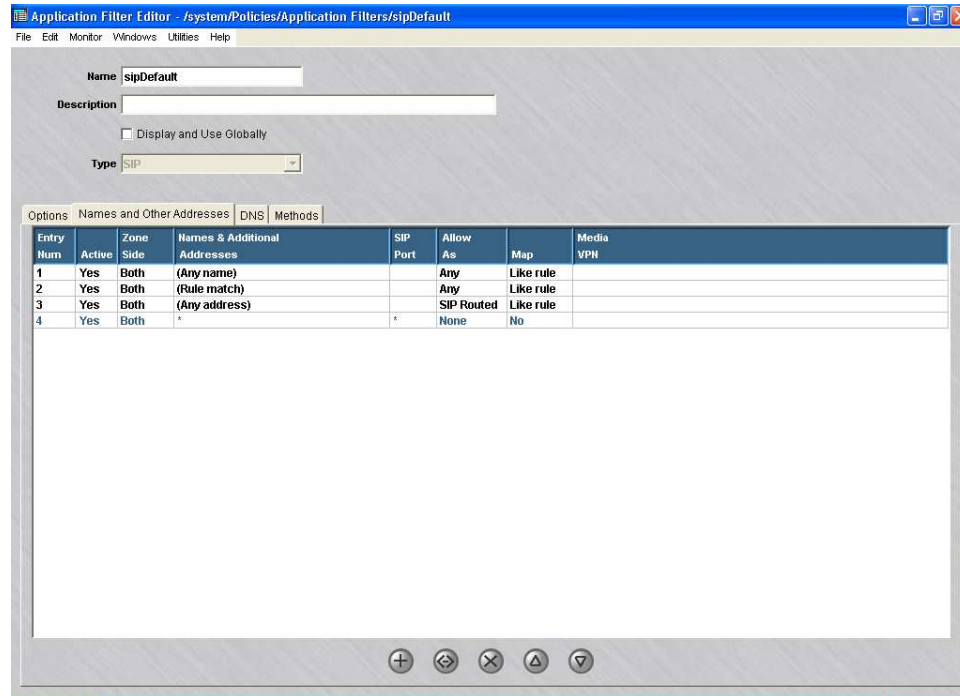
In many configurations, SIP control messages (the signaling) and the media flow through the same zone. In such cases, media activity is sufficient to prevent the SIP dialog from timing out. In signaling-only configurations, the media does not flow through the zone and the SIP dialog will time out based on the value in this field. The default value is 3600 (seconds). Valid values are from 0 (a value of 0 means time out immediately) to 99999.

If the **Never Timeout** option is selected, the call will not time out and must be ended with a BYE. When the **Never Timeout** option is selected, the **Signaling-Only Calls** timer field is disabled and greyed out. The **Never Timeout** option should not be used unless it is guaranteed that a BYE will be sent through the zone, by, for example, an IP Multimedia Subsystem (IMS) component.

- **Calls-on Hold** (Configurable Timeout setting)

Some calls on hold may have RTP traffic that will keep them from timing out, but some do not. This field provides a means to keep the calls on hold from timing out too soon even if there is no media traffic while on hold. Calls on hold will time out based on the value in this field. The default value is 3600 (seconds). Valid values are from 0 (a value of 0 means time out immediately) to 99999.

-
- 5 The **Names and Other Addresses** tab (see [Figure 5-42, “SIP Application Filter Editor \(Names and Other Addresses Tab\)”](#) (p. 5-97)) is a table used to control the names and addresses that can be used inside and outside of the zone and their purpose. A name or address can be used as a SIP, and possibly Media endpoint, whose IP address will be seen by the firewall in a transport layer stream; as a Media-only address for the media stream; or as otherwise mentioned in the SIP Via, Route and Record-Route headers, but not seen by the firewall at the transport layer (SIP application routing).


Figure 5-42 SIP Application Filter Editor (Names and Other Addresses Tab)

The initial default table allows any name or address to be mentioned in a SIP message, but all transport layer streams crossing the firewall must be between the initial source and destination endpoints. In many cases, however, one or both of the original endpoints is a proxy, not a user agent. A proxy may help setup a call and then drop out of the dialog, allowing the user agents to send SIP messages directly to each other. In such a case, the user agents must be explicitly allowed, as must any other proxies through which the dialog is routed.

Additionally, this table can be used to ensure that certain domain names only originate from inside the zone.

To create new names and addresses, right-click in the tab folder and select **New** or click on the button with the + sign. The SIP editor will appear (see [Figure 5-43, “SIP Editor \(Names Entry Screen\)”](#) (p. 5-98)). The **Entry Number** is automatically sequentially assigned by the system beginning with 1, giving the position of entry within the table. The number is used to identify the entry in audit logs.

Click on **OK** to save the entry.

Figure 5-43 SIP Editor (Names Entry Screen)

The following entry fields are available:

- **Active**
If checked (default), the entry is enabled.
- **Zone Side**
Options are **Inside**, **Outside**, or **Both**. The default is **Both**.
- **Name/Address**
The entry may be an IPv4 address, a range of addresses, a subnet, an LVF host group, a Fully Qualified Domain Name (FQDN), a partially qualified domain name, beginning with a star-dot (*.), the wildcard (*), meaning any name or address, or any of the options (**Rule match**), (**Any address**), or (**Any name**). The default for new entries is (**Rule match**). The option (**Any name**) will match any Fully Qualified Domain Name, such as host.example.com.
If this table is left to its initial defaults, then the only addresses that are permitted are the two that initiated the dialog. The initial table is suitable when it is known that the SIP proxy will not drop out or there is no proxy between the user agent and the firewall.
The (**Rule match**) option is less restrictive, allowing any agent or proxy that matches the rule to participate in the dialog. When this option is selected, proxies can drop out and be replaced by any user agent that matches the rule that is controlled by this filter. The inside zone addresses are compared to the source address of the rule for rules with direction out of the zone and to the destination side for rules whose direction is into the zone. Addresses may be single addresses, a range or a subnet (i.e., 10.1.1.1; 10.1.1.1-10.1.1.10; 10.1.1.0/24).
This table is also used to control the domains and host names that are allowed in

SIP messages coming from the zone and those entering the zone from elsewhere. It can be used to prevent Domain Name spoofing and to allow or deny calls to or from certain Domains. Domain Names may be fully qualified host names or partially qualified where the left part of the name is replaced with an asterisk (i.e., server.example.com; *.example.com; *.com; **(Any name)**).

- **Port**

Options are a port number, range of port numbers, the wildcard (*), or blank. The default is blank, which is equivalent to the default port (5060 for SIP). Normally, this field may be left blank unless a non-standard port number is used.

- **Allow As**

Options from the dropdown menu are: **Any, Media, SIP Routed, or None**, which is the default. This option determines whether the name or address is allowed for a given usage. A name or address allowed as "SIP routed" can be mentioned in a Via, Route or Record-route header, but not in any other context. This implies that the name/address was used or will be used in some transport leg that does not traverse the zone. A name or address designated as "Any" can be used as a SIP routed address, or as a SIP or media endpoint of a transport stream that does traverse the zone. In general, hosts should be specified by both name and address, since either may appear in the SIP message. The SIP Awareness feature does not do reverse DNS resolutions to map address into names. It does normal name to address DNS resolution only when forced to map a media name to an address. The principal reason for this is to minimize the dependency of security policy on DNS. Additionally, avoiding the resolution, where possible, provides better performance.

- **Map**

This field has three options: **Like rule, Yes** or **No**. The default is **Like rule**. The **Yes** option forces Network Address Translation to take effect even if that was not specified in the original rule. The VBA of the Brick replaces the zone addresses in the SIP message so that the internal network configuration is not revealed outside of the zone. The **Yes** option has no additional effect if the original rule specified Network Address Translation for the zone. The **No** option will prevent Network Address Translation from being performed on the zone, irregardless of what was specified in the rule.

- **Media VPN**

Options from the dropdown menu are: **Like rule** and **No**. The default is blank. This field is not relevant if "Pass" is "No" or if usage is "Other." If **No** is selected, then the media is not VPN tunneled regardless of whether the SIP stream is. If **Like rule** is selected, the media is tunneled if the original rule specifies VPN processing.

-
- 6 The **DNS** tab has selections for a DNS Application Filter to be applied and for specifying the DNS server (see [Figure 5-44, "SIP Application Filter Editor \(DNS Default Tab\)"](#) (p. 5-100)). An IP address or Host Group can be used. If two entries are in the Host Group, the first one is treated as a primary DNS and the second as a

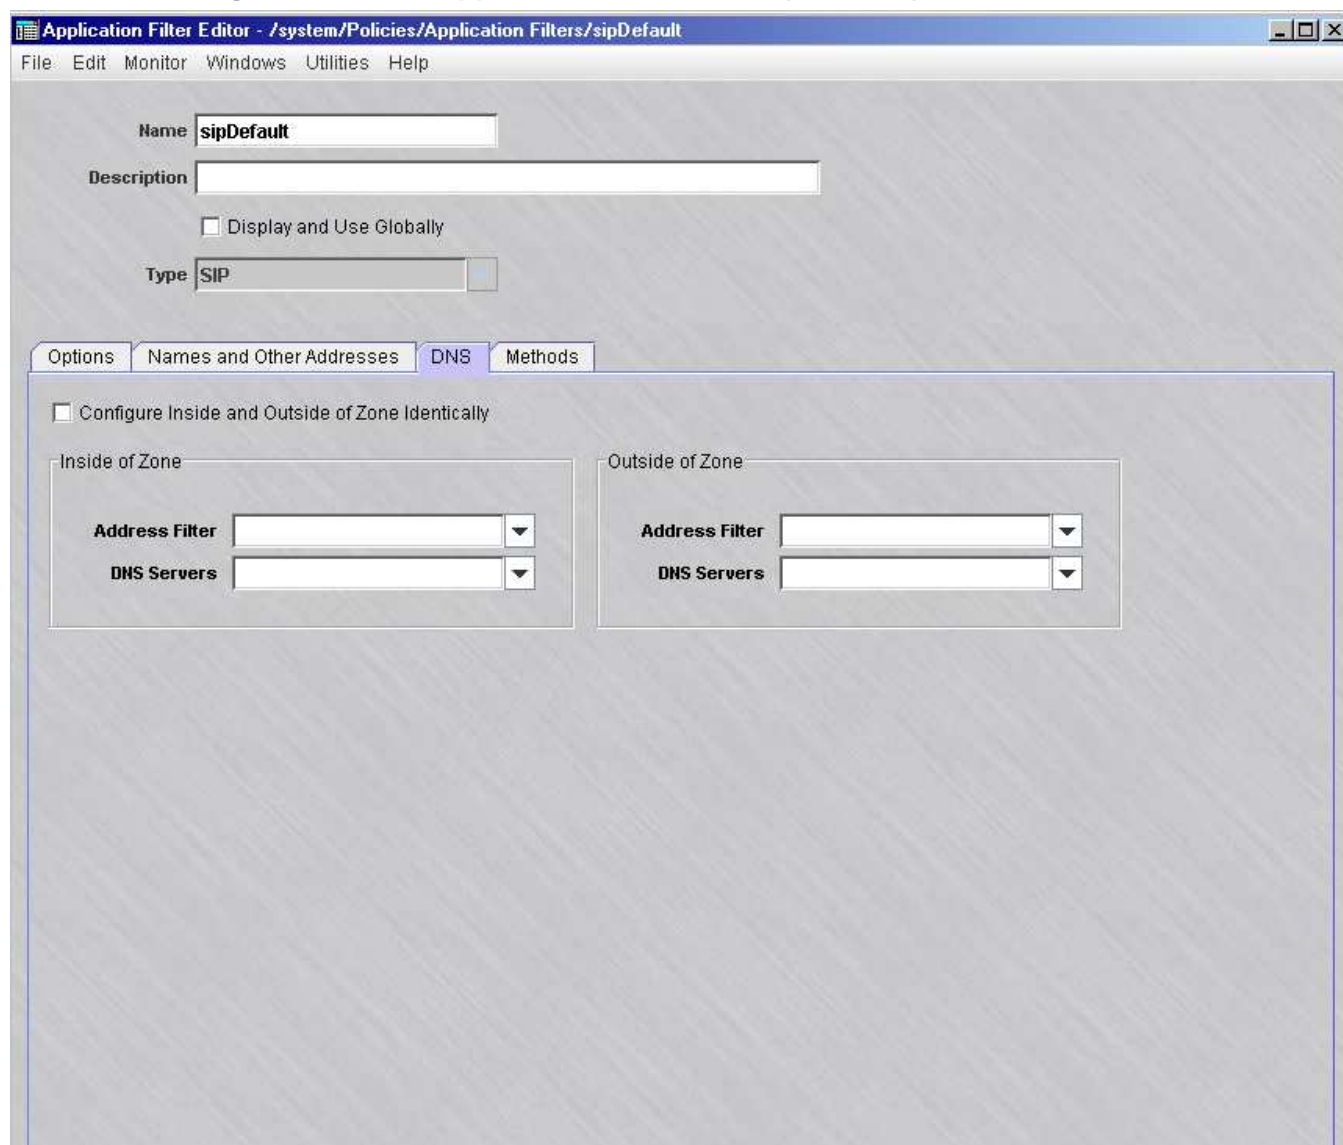
secondary DNS. If there are more than two entries, they are ignored. Separate DNS servers can be configured for Inside of Zone and for Outside of Zone (see [Figure 5-45, “SIP Application Filter Editor \(DNS Tab\)-both zones”](#) (p. 5-101)).

DNS name resolution queries and responses must be filtered through a DNS application filter. Queries are sent from the VBA of the zone to the DNS servers. If there is a rule defined in the zone for such queries with a DNS application filter, that rule and filter will be used. If not, then a DNS filter must be specified in the address filter drop-down window. In this case, a dynamic rule will be automatically created to call out the queries to pass using the DNS filter.

Figure 5-44 SIP Application Filter Editor (DNS Default Tab)

The screenshot shows the 'SIP Application Filter Editor' window with the 'DNS' tab selected. The 'Name' field is set to 'sipDefault'. The 'Type' is set to 'SIP'. The 'Display and Use Globally' checkbox is unchecked. The 'DNS' tab is active, showing the 'Configure Inside and Outside of Zone Identically' checkbox checked. Below this, the 'Inside and Outside of Zone' section contains two dropdown menus: 'Address Filter' and 'DNS Servers'.

Figure 5-45 SIP Application Filter Editor (DNS Tab)-both zones



-
- 7 The **Methods** tab contains these options (see [Figure 5-46, “SIP Application Filter Editor \(Methods Default Tab\)”](#) (p. 5-102)):
- **Configure Inside and Outside of Zone Identically**
If checked (default), configures the SIP application identically both inside and outside of the zone.

Figure 5-46 SIP Application Filter Editor (Methods Default Tab)

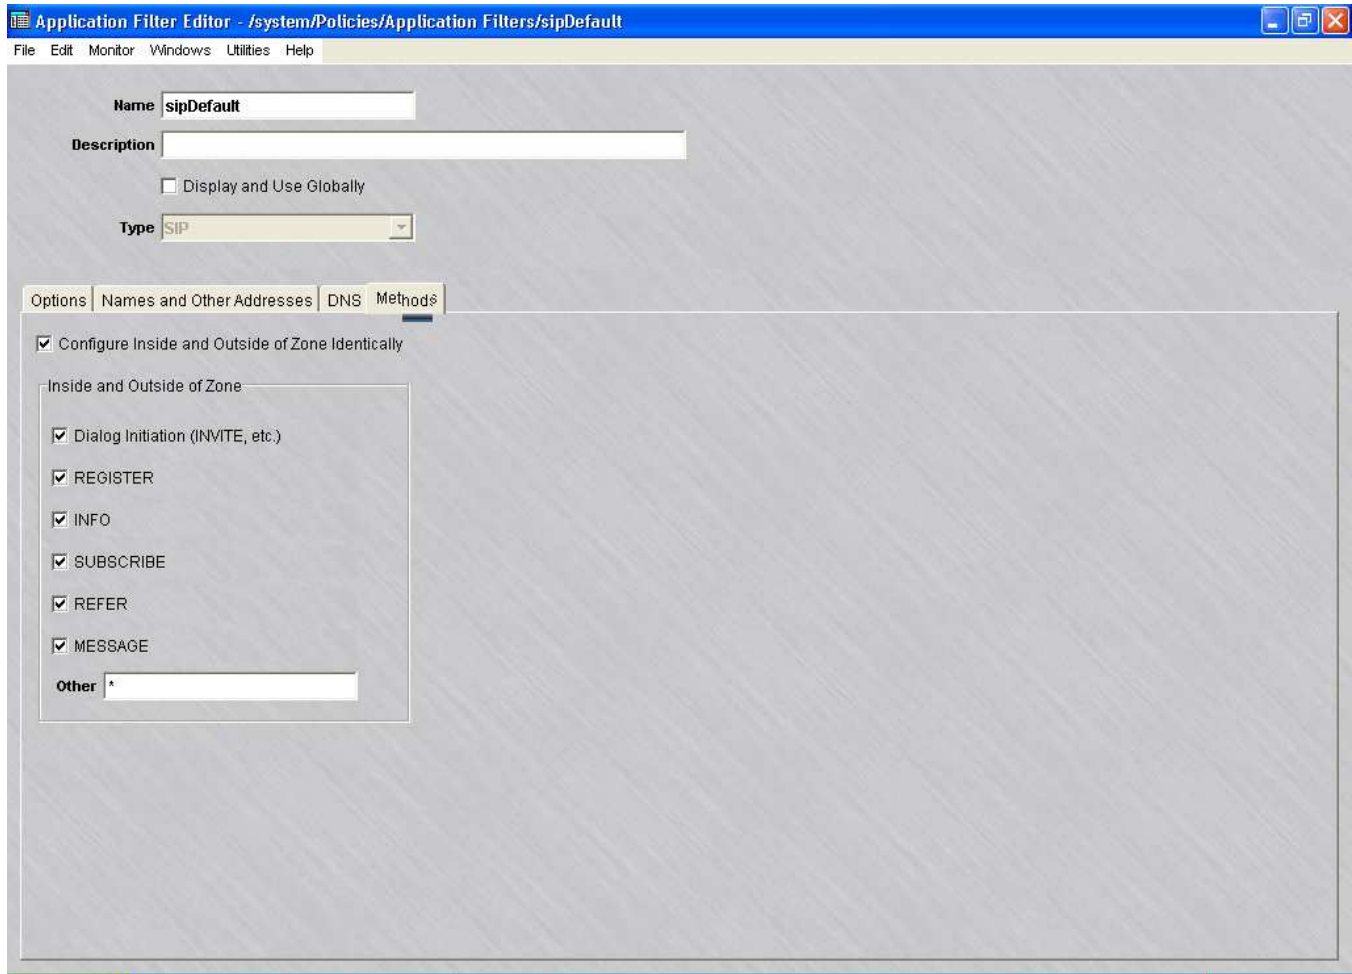
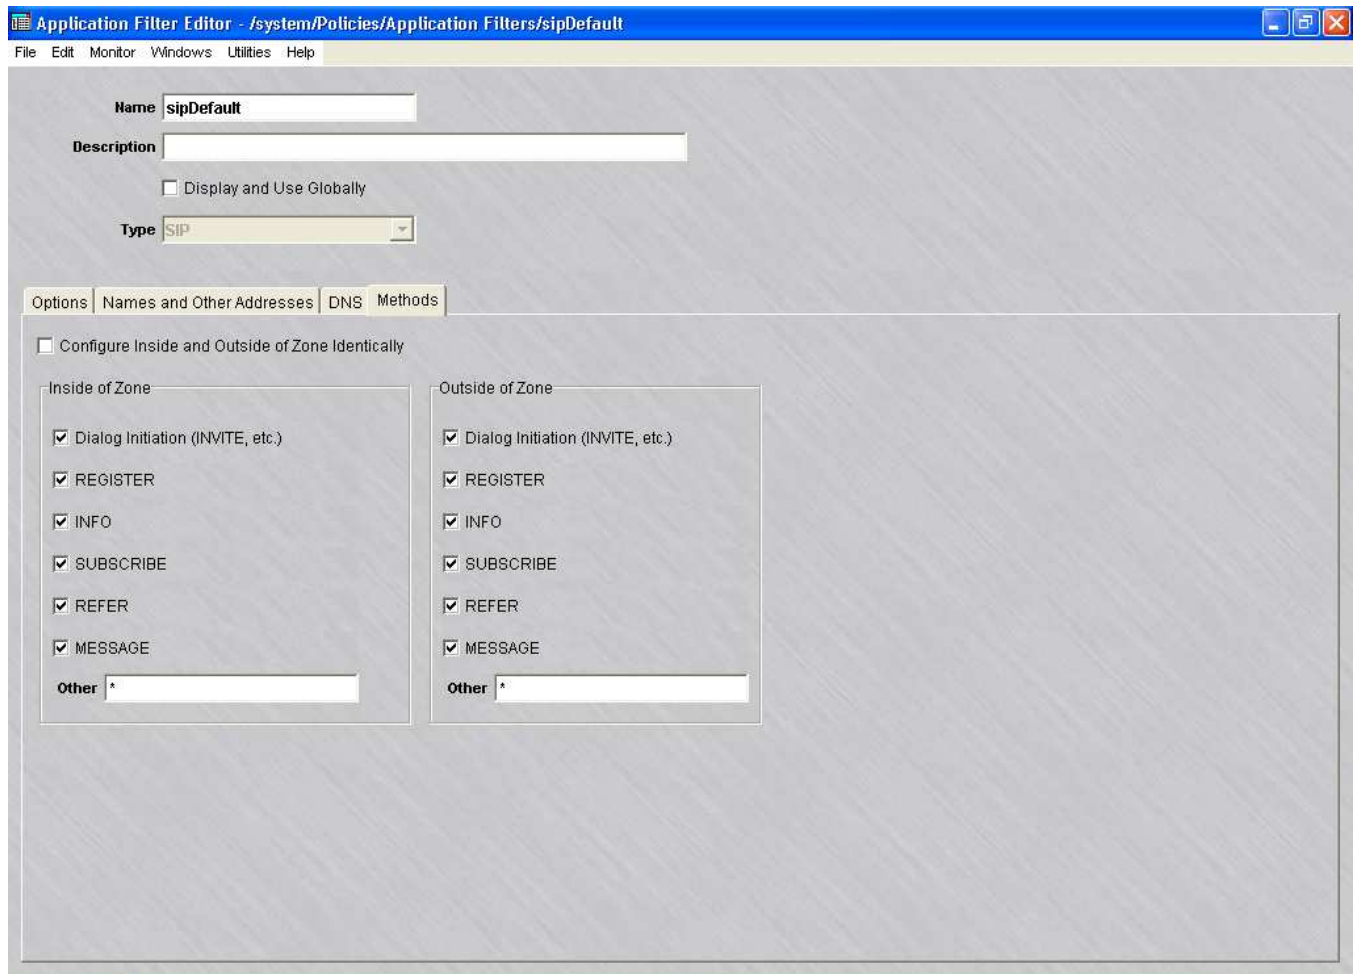


Figure 5-47 SIP Application Filter Editor (Methods Tab) - both zones



The following explains these fields:

- **Dialog Initiation (INVITE, etc.)**
This method allows a user to initiate a call.
- **REGISTER**
This method allows users to register their SIP IDs with their registrar server. Unselecting this option on one side of the zone prevents users from registering with the registrar server on the other side.
- **INFO**
This method is used to send control information in the middle of a call.
- **SUBSCRIBE**
Inside zone SUBSCRIBE implies outside zone NOTIFY capability and outside zone SUBSCRIBE implies inside zone NOTIFY.
- **REFER**
Inside zone REFER also implies outside zone NOTIFY.

- **MESSAGE**
This method is used in Instant Messaging services.
 - **Other**
The other field is writable, if selected, and can contain a comma-delimited list of methods, a blank or a wildcard (*). The default is selected and wild-card.
-

- 8 From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.

END OF STEPS



SMTP Application Filters

Overview

E-mail is one of the oldest programs on the Internet and the very same protocol that launched e-mail, SMTP, is still in use today with only minor revisions. SMTP was not designed to be very secure, and consequently, is one of the most hacked applications, accounting for about 19% of all vulnerabilities on the Internet. SMTP can be exploited with various mail relay attacks, Multipurpose Internet Mail Extension (MIME) attacks, buffer overflow attacks, address spoofing attacks, and covert channel attacks. The SMTP application filter prevents all of these types of attacks.

Configuring an SMTP Application Filter

An SMTP application filter called *smtpDefault* is created automatically in the system group for a newly installed SMS and, once installed, is created automatically for newly created groups. Similarly, an *smtpDefault* filter is created when an SMS upgrade is done for each of your existing groups. However, *smtpDefault* is not automatically assigned to any service group after installation or upgrade.

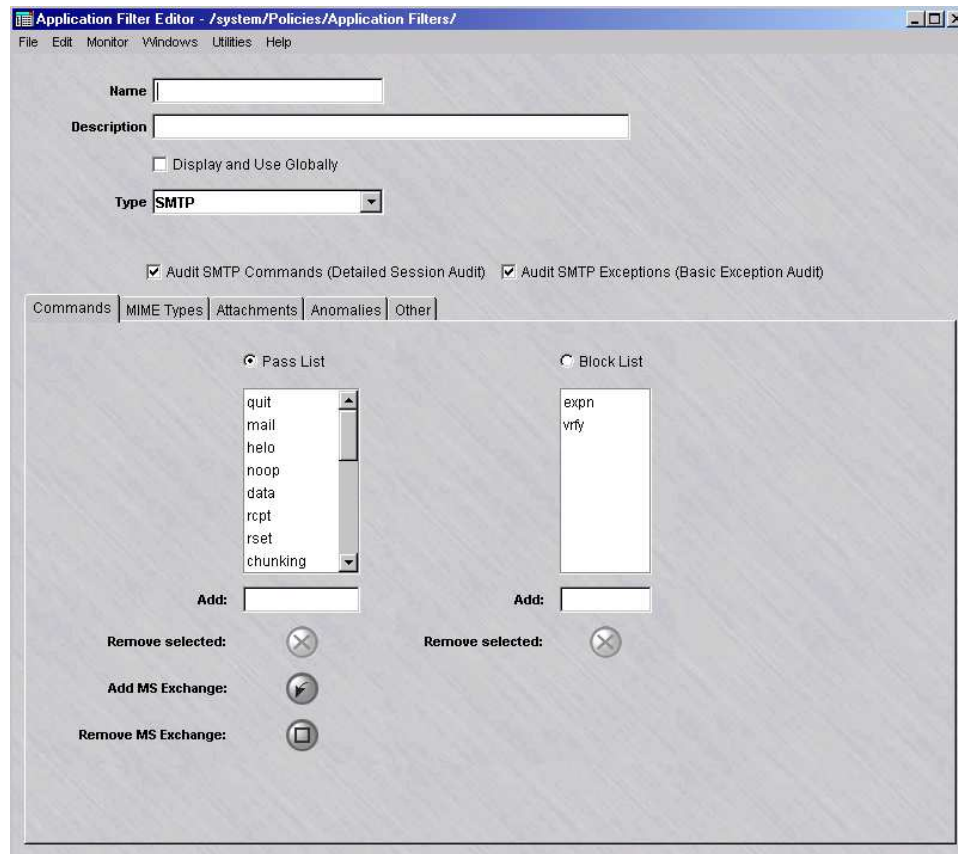
You may decide to just reconfigure the *smtpDefault* application filter or create a new one.

To configure an SMTP application filter

To configure an SMTP application filter, follow the steps below:

- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.
 - 2 In the right hand column, right-click the mouse and select **New Application Filter** from the pop-up menu. After the Application Filter Editor appears, click the Type drop-down menu and select **SMTP**.
-

The Application Filter Editor is displayed with the parameters for setting up an SMTP application filter. [“To configure an SMTP application filter” \(p. 5-105\)](#) shows a sample of the screen.

Figure 5-48 SMTP Application Filter Editor (Commands Tab)

- 3 In the **Name** and **Description** fields, enter a name for the application filter and a brief description. The **Description** field is optional.

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the [“Global Application Filters”](#) (p. 5-127) section for an explanation of global application filters.

- 4 Click the **Audit SMTP Commands (Detailed Session Audit)** checkbox to monitor and log FTP commands issued when an SMTP connection is established.

Click the **Audit SMTP Exceptions (Basic Exception Audit)** checkbox to monitor and log application filter violations and the SMTP traffic.

5 The Commands tab contains these fields:

- **Pass List / Block List**

Command filtering is accomplished in one of two separate modes: pass or block. If you select the **Pass List** radio button, only those commands in the pass list are allowed. If you select the **Block List** radio button, commands in the block list are prohibited.

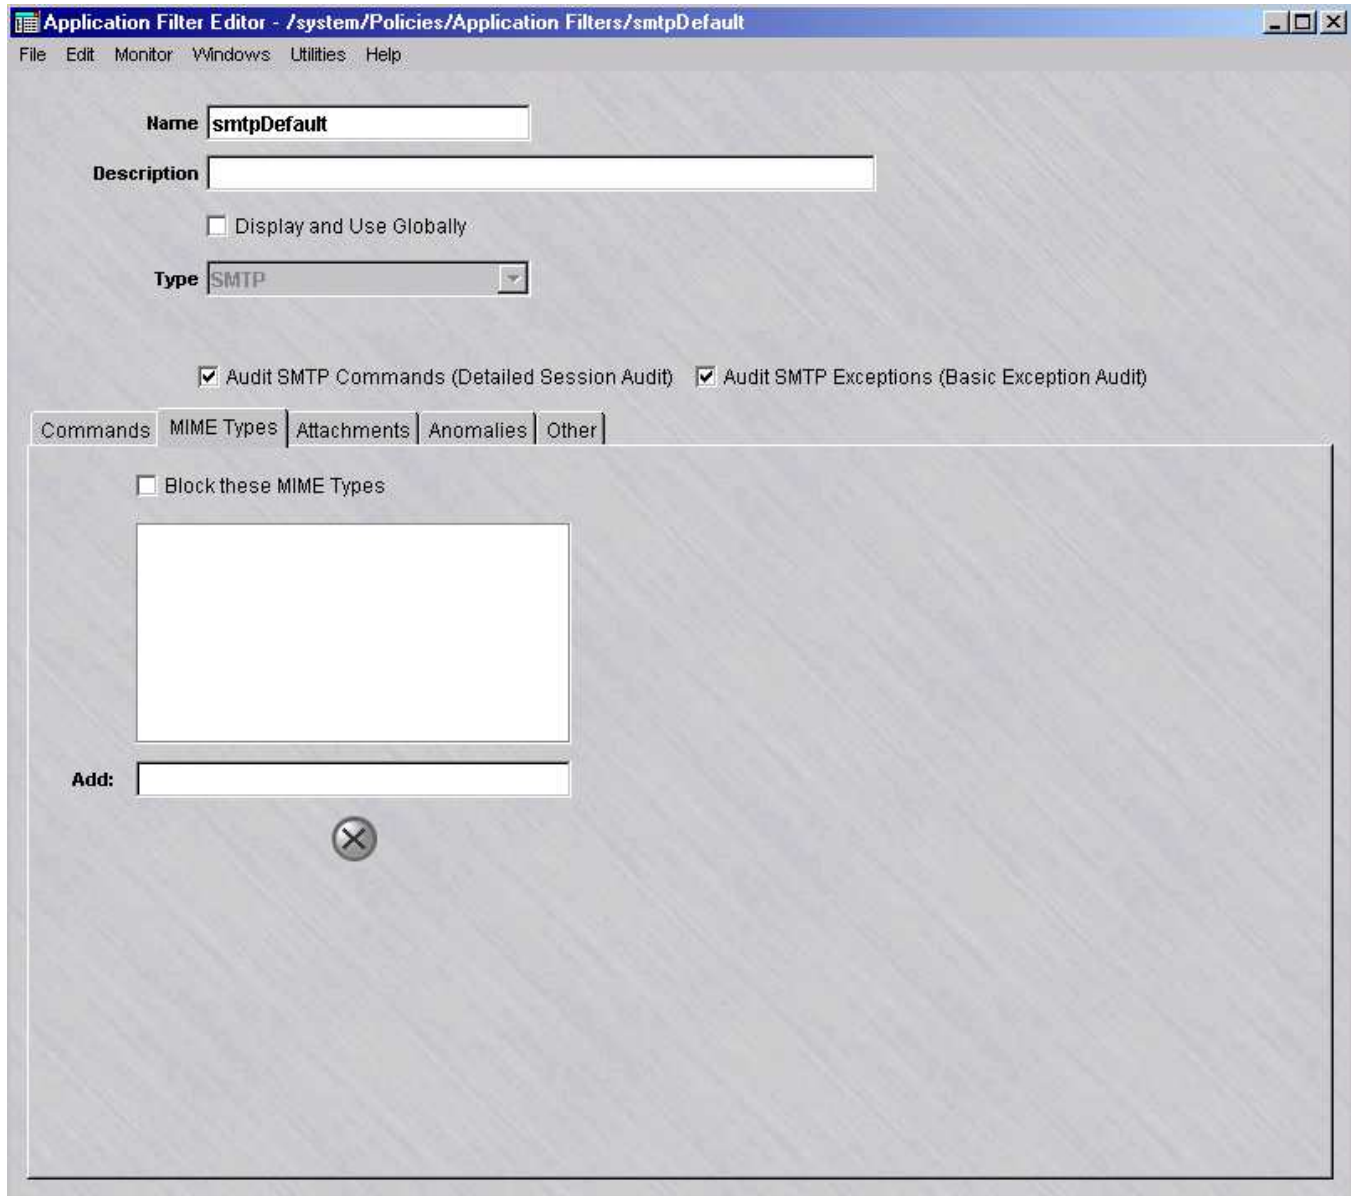
Commands can be added to either list by typing the command into the **Add** field under the respective list and pressing the Enter key. Only one list can be active at a time for updating.

A command in either list can be deleted by highlighting it and then clicking the **Remove selected** button, or by right-clicking on the command and choosing **Delete**. You can select multiple entries for deletion by clicking the left mouse button while pressing the Ctrl key, or by clicking the left mouse button while pressing the Shift key to select a range of commands in the list.

The **Add MS Exchange** button allows you to add SMTP commands used by the Microsoft Exchange Server. The **Remove MS Exchange** button allows you to delete commands used by the Microsoft Exchange Server. You can selectively delete some of these commands or use the **Remove MS Exchange** button to delete all of them (or the remaining commands).

6 The MIME Types tab contains these fields ([Figure 5-49, “SMTP Application Filter Editor \(MIME Types Tab\)”](#) (p. 5-108) shows a sample screen):

Figure 5-49 SMTP Application Filter Editor (MIME Types Tab)



This tab allows you to configure a list of MIME types to disallow by typing each in the following format:

<top level media type>/<subtype>

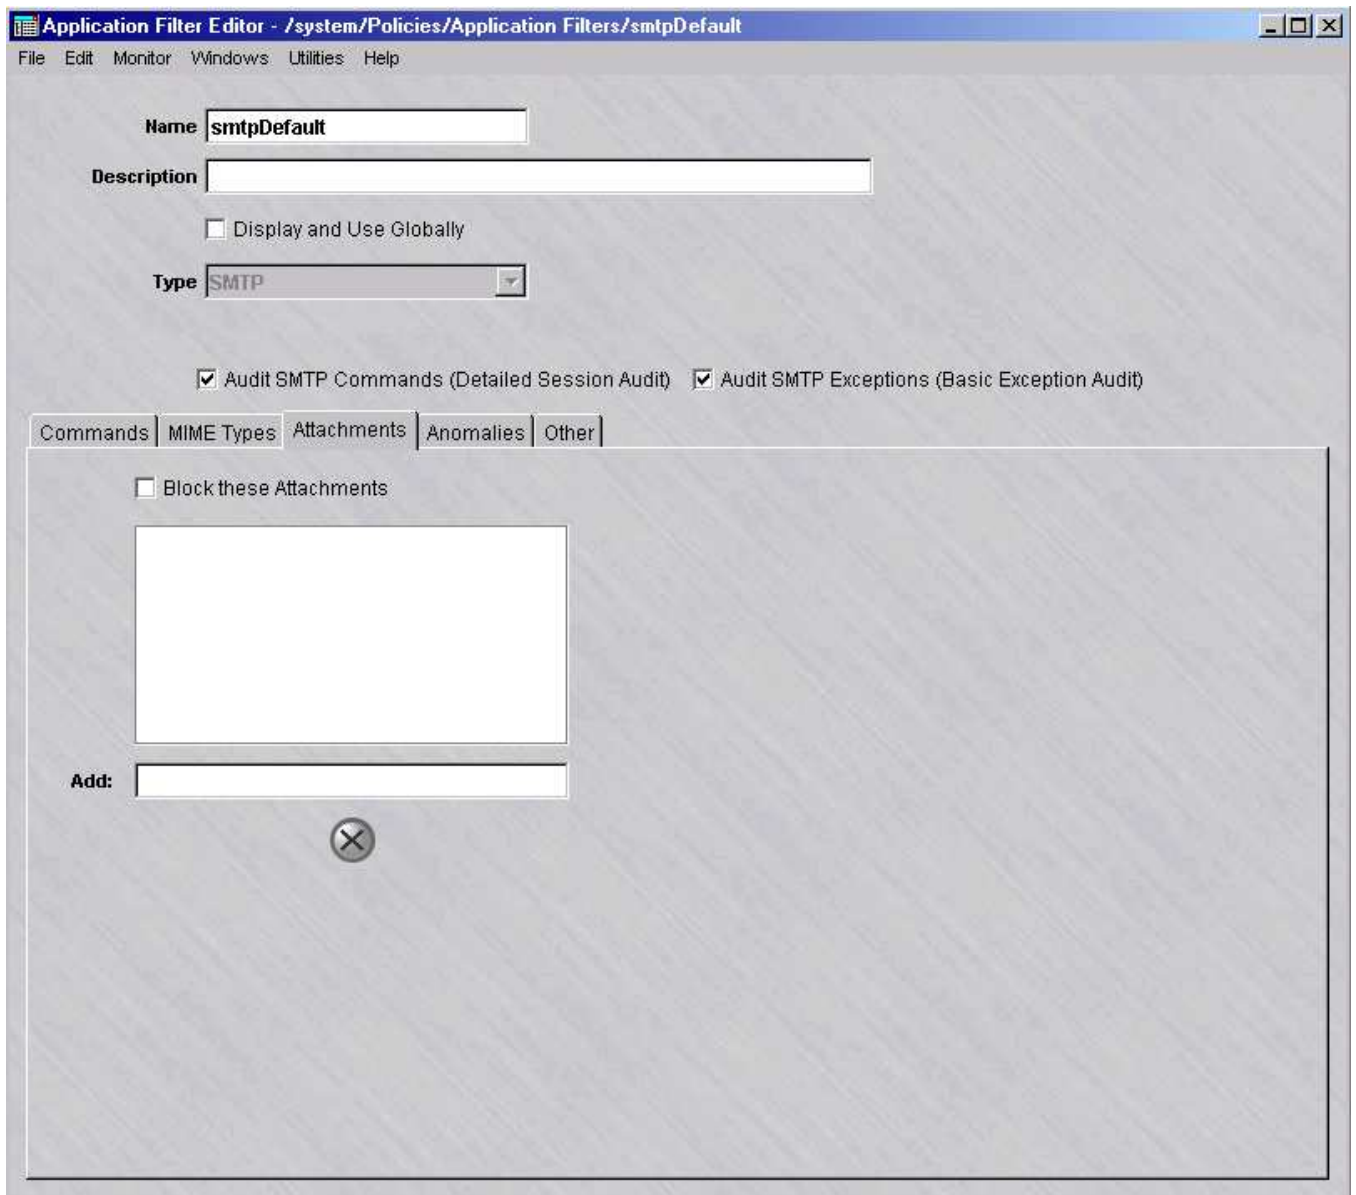
If the incoming mail contains a MIME type found in this list, then the entire MIME object (content) is removed and a message to this effect is inserted in the mail. An error message is sent back to the mail sender.

The checkbox can enable and disable the checking while remembering the list values.

MIME types can be added and deleted in the same way as commands on the Commands tab.

- 7 The Attachments tab contains these fields (Figure 5-50, “SMTP Application Editor (Attachment Tab)” (p. 5-109) shows a sample screen):

Figure 5-50 SMTP Application Editor (Attachment Tab)



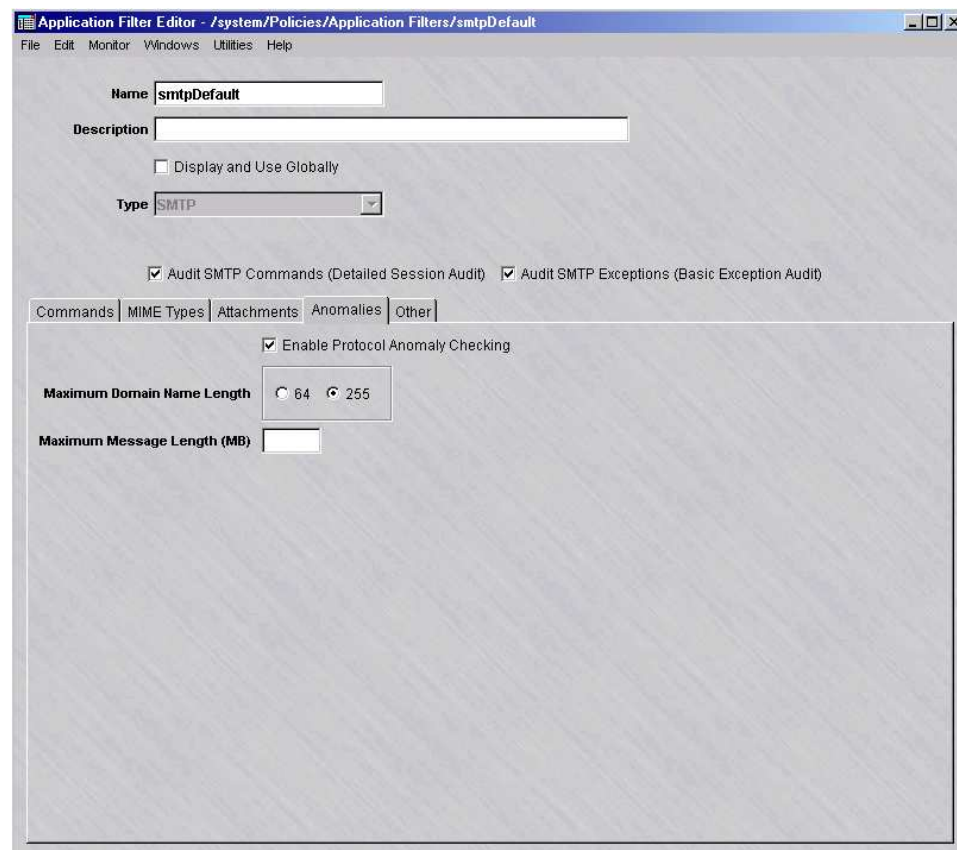
This tab allows you to configure a list of attachments to be blocked. The attachment name may contain the wildcard character, such as *.exe'. The length of the attachment file name cannot be greater than 255 characters, and it cannot contain non-printable characters or the following special characters (\/:..).

The checkbox can enable and disable the checking while remembering the list values.

Attachment names can be added and deleted in the same way as commands on the Commands tab. If the mail contains an attachment with a name that matches an entry in this list, then the attachment is removed and a message to this effect is inserted in the mail. An error code is sent back to the mail sender.

- 8 The Anomalies tab contains these fields (Figure 5-51, “SMTP Application Filter Editor (Anomalies Tab)” (p. 5-110) shows a sample screen):

Figure 5-51 SMTP Application Filter Editor (Anomalies Tab)



The maximum length of domain names in the e-mail address can be set to either **64** (RFC 821) or **255** (RFC 2821) characters. Some older server implementations may not accept more than 64 characters. The default is **255**.

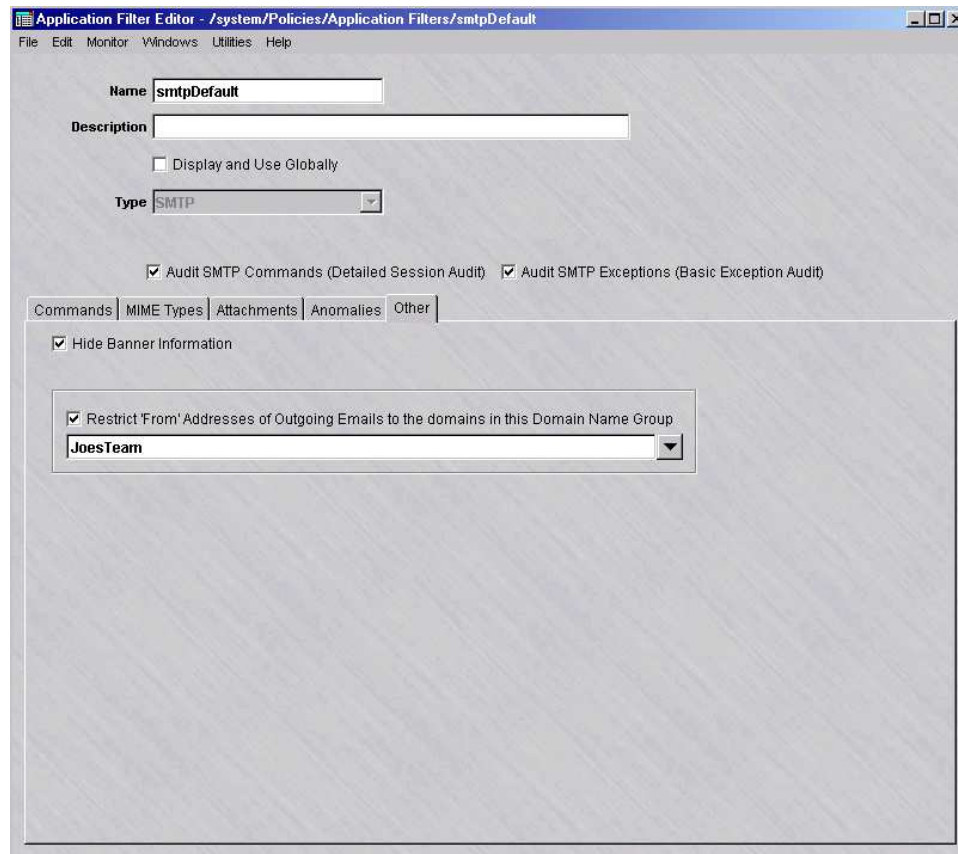
To thwart mail bomb attacks, a maximum message size may be specified. The allowed range is 0-9999. Setting the value to zero (0) means that the limitation is 64k. *If left blank, no limit is enforced.*

Other anomaly checks include:

- The total SMTP command line, including parameters and CR-LF, should not exceed 1000 characters.
- The encoded portion of the mail message must not violate the encoding character set.
- File names must be no larger than 255 characters and only printable characters are allowed and do not contain these special characters: "/", "\", ":", and "..".
- The boundary in MIME characters are checked to be no larger than 70 characters and that they contain characters defined by RFC 2046.
The checkbox can enable and disable the anomaly checking while remembering the configured settings.

-
- 9** The Other tab contains these fields ([Figure 5-52, “SMTP Application Filter Editor \(Other Tab\)”](#) (p. 5-112) shows a sample screen):

Figure 5-52 SMTP Application Filter Editor (Other Tab)



The following explains these fields:

- **Hide Banner Information**

The banner message returned by the server upon a successful client connection usually shows the kind of mail server being used, which could be exploited for e-mail attacks. If the **Hide Banner Information** checkbox is checked, the banner characters are replaced by asterisks, except for the three-digit response code and terminating <CR><LF>.

- **Restrict 'From' Addresses of Outgoing Emails to the domains in this Domain Name Group**

It may be desirable to prevent employees from using their corporate e-mail network to carry out personal correspondence and business using forged names. This feature can stipulate that any e-mail leaving an organization must have an authorized domain name.

The **Restrict 'From' Addresses of Outgoing Emails to the domains in this Domain Name Group** checkbox can enable and disable this checking while remembering the Domain Name Group.

-
- 10 From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.

END OF STEPS

To add the filter to a service group

To add this application filter to a service group, follow the steps below:

-
- 1 In the Navigator window, open the appropriate Service Group folder and double-click on the service group or create a new one.
 - 2 Edit the tcp protocol entry, specifying the destination port for incoming messages using this application filter.
 - 3 Click the pulldown for the application filter, select **SMTP** from the pull-down menu, and click **OK**.
 - 4 Open the File menu and select one of the **Save** options..

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information on Brick Zone Rulesets, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”. “Overview” \(p. 1-1\)](#). For additional information on service groups, see [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”. Chapter 4, “Service Groups”](#).

END OF STEPS



SQL*Net Application Filter

Overview

Oracle Corporation's SQL*Net is remote data access software that enables both client-server and server-server communications across any network. With SQL*Net, databases and their applications can reside on different computers and communicate as peer applications.

SQL*Net enables clients and servers to connect to each other, send data such as SQL statements and data responses, initiate interrupts from client or server, and disconnect when the session is complete.

When a client or server makes a connection request, SQL*Net receives the request and, if more than one machine is involved, passes the request to its underlying layer, the transparent network substrate (TNS) to be transmitted over the appropriate communications protocol to the appropriate server. On the server, SQL*Net receives the request from TNS and passes it to the database as a network message with one or more parameters (via an SQL statement).

When this application filter is applied to a service group and used within a zone ruleset, the Brick looks at the appropriate string(s) and determines the exact port (pinhole) for relaying messages between the client and server.

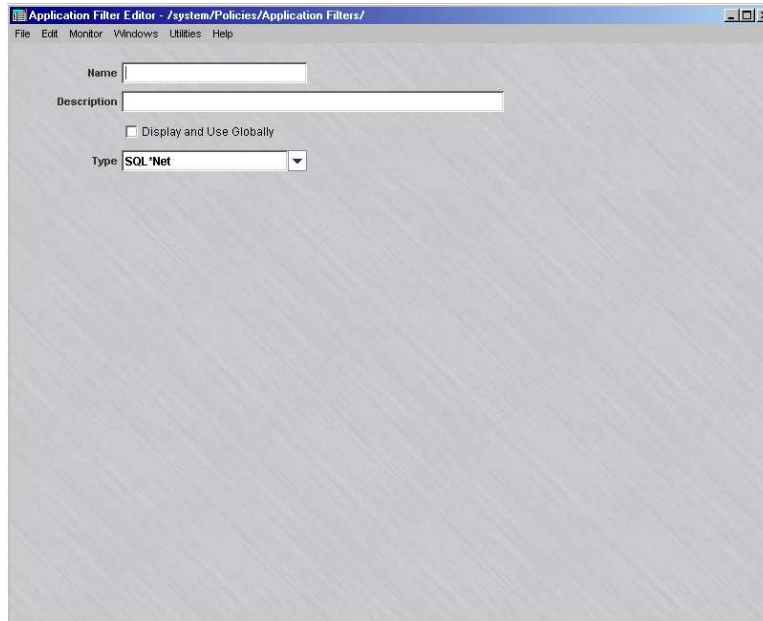
To configure an SQL*Net application filter

Complete the following steps to configure an SQL*Net application filter.

- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.
- 2 In the right hand column, right-click the mouse and select **New Application Filter** from the pop-up menu. After the Application Filter Editor appears, click the Type drop-down menu and select **SQL*Net**.

Result The Application Filter Editor is displayed with the parameters for setting up an SQL*Net application filter (Figure 5-53, “SQL*Net Application Filter Editor” (p. 5-115)).

Figure 5-53 SQL*Net Application Filter Editor



-
- 3 In the **Name** and **Description** fields, enter a name for the application filter and a brief description. The **Description** field is optional.

Important! If you are an SMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the “[Global Application Filters](#)” (p. 5-127) section for an explanation of global application filters.

-
- 4 From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.

END OF STEPS

To add the filter to a service group

To add this application filter to a service group, follow the steps below:

- 1 In the Navigator window, open the appropriate Service Group folder and double-click on the service group or create a new one.

Result The Service Group Editor window is displayed.

- 2 Right-click and select **New** from the pop-up menu if you are creating a new service group.

Result The Service Editor window is displayed.

- 3 In the **Protocol** field, enter **50**.

- 4 Click the pulldown for the application filter, select the name of the application filter from the pull-down menu, and click **OK**.

- 5 Open the File menu and select one of the **Save** options.

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information on Brick Zone Rulesets, refer to [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). For additional information on service groups, refer to [Chapter 4, “Service Groups”](#).

END OF STEPS



TFTP Application Filter

Overview

The Trivial File Transfer Protocol (TFTP) is a simple Internet-based file transfer protocol that has been used for a long time. TFTP is commonly employed to initially configure devices or load new versions of operating system code by transferring files from a server to another device. The basic TFTP Uniform Resource Indicator (URI) scheme specifies a host, a filename, and, optionally, the data mode for the transfer (which is octet mode if not specified). When a TFTP URI is specified, the file is transmitted via TFTP to a specified server using the optionally specified mode.

Despite its long-standing usage, TFTP has several disadvantages from a security standpoint in handling file transfers. File size information is not available prior to retrieval, so it is difficult to determine the integrity of the file that was transferred. TFTP has no inherent integrity check to determine if what was sent was received. It has no mechanism for access control within the protocol, thereby leaving any file transfer vulnerable to an intrusive attack en route. It is not recommended for very large files or where memory and CPU resources are limited.

These factors make it advisable to use a TFTP application filter within a service group and Brick zone rule to monitor and verify the source and destination of each TFTP session.

TFTP application filter and NOE traffic filter option

An NOE traffic filtering option can be activated within a TFTP application filter for checking IP phone voice and maintenance traffic in a VoIP application environment. The combined TFTP/NOE application filtering mechanism inspects and validates the various types of data, voice, and maintenance traffic that is exchanged between the various network elements in a VoIP call network (IP Touch phones, call servers, Media Gateway units) . In this context, the TFTP application filter also inspects and validates TFTP sessions during downloads of configuration files and software updates from the call server(s) to IP Touch phones.

When the NOE traffic filtering option is enabled within a TFTP application filter, the Brick creates dynamic rules in the zone ruleset for inspecting RTP and NOE signalling traffic between the IP phones and other network elements (call servers, MGW) in a VoIP network. The Brick also creates a dynamic rule for any telnet debug sessions to IP phones.

The TFTP application filter, with the NOE traffic filter option enabled and selected NOE application filter, is applied to a service group (at UDP destination port 69), which is, in turn, assigned to a rule within the zone ruleset of the Brick that is protecting the call server(s) in the VoIP call network.

A set of pre-defined TFTP application filter templates is provided with the SMS application when it is installed to make the configuration of the Brick as an Application Layer Gateway (ALG) to handle VoIP/NOE traffic easier. For more information about these templates, refer to the section [“TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-123).

Additional information about the Application Layer Gateway (ALG)/NOE feature is provided in the section *Deployment of a Brick device as an application layer gateway (ALG) for VoIP/NOE phone communications* in the *SMS Administration Guide*.

Before you begin

Before you begin this task, if you are configuring a TFTP application filter and using the NOE Traffic filter option within a TFTP application filter, you must first create and configure an NOE application filter. This NOE application filter can then be selected and called within the TFTP application filter in a Brick zone rule to establish the parameters for inspecting Real Time Protocol (RTP) sessions between VoIP network elements (such as call/data traffic between an IP Touch phone and another phone or MGW unit or MGW-to-CS IPlink protocol signalling traffic). Refer to the [“NOE Application Filter”](#) (p. 5-62) section in this chapter for details about creating and configuring an NOE application filter.

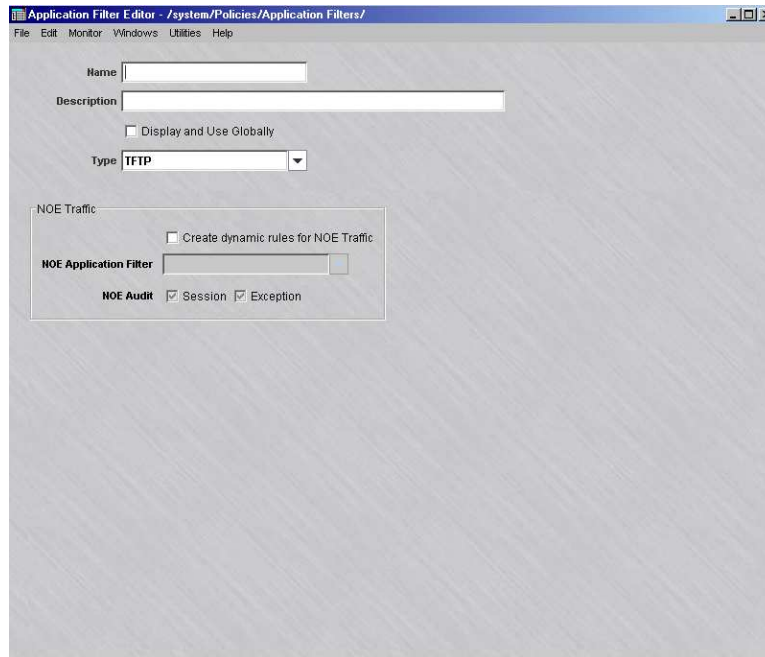
To configure a TFTP application filter

Complete the following steps to configure a TFTP application filter.

- 1 From the SMS Navigator, open the Policies folder in the desired group and click **Application Filters**.
- 2 In the right-hand column, right-click the mouse and select **New Application Filter** from the pop-up menu.
Result The Application Filter Editor is displayed.
- 3 Click the down arrow next to the **Type** field to display a drop-down menu and select **TFTP**.

Result The TFTP Application Filter Editor window is displayed (Figure 5-54, “TFTP Application Filter Editor” (p. 5-119)).

Figure 5-54 TFTP Application Filter Editor



-
- 4 In the **Name** and **Description** fields, enter a name for the application filter and a textual description, respectively. **Name** is a required field. The **Description** field is optional.

A checkbox labeled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the “[Global Application Filters](#)” (p. 5-127) section for an explanation of global application filters.

If you are using the TFTP application with the NOE Traffic filter option, go to [Step 5](#).

Otherwise, go to [Step 9](#).

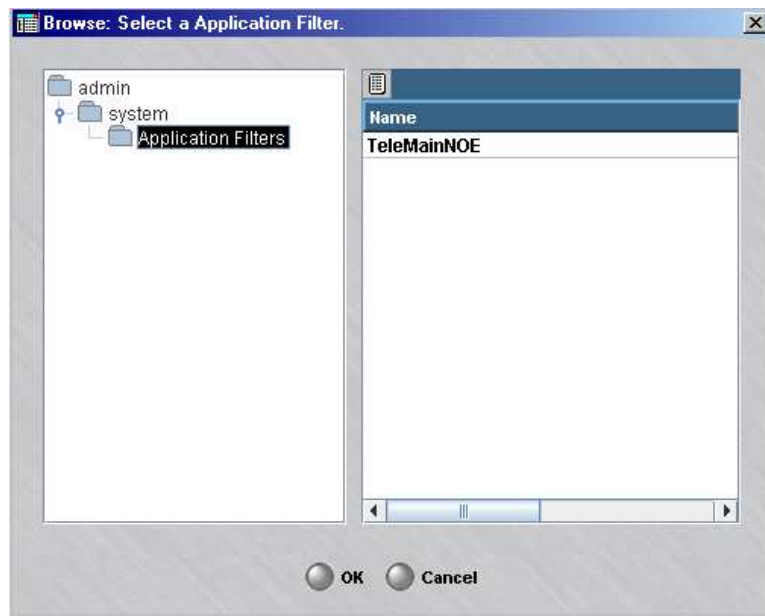
-
- 5 Click the **Enable creation of dynamic rules for NOE traffic** checkbox (checkbox is checked) to enable the NOE Traffic filter option.

Result The **NOE Application Filter** field and NOE Audit fields are activated.

- 6 Select an NOE application filter by clicking the down arrow next to the **NOE Application Filter** field and select **BROWSE** from the drop-down menu.

Result The Select a Application Filter window is displayed (Figure 5-55, “Select a Application Filter Window” (p. 5-120) shows a sample window).

Figure 5-55 Select a Application Filter Window



- 7 Select the desired NOE application filter from the Name column (expand the group folders in the explorer tree in the left column if necessary to display the application filter) and click **OK**.

Result The selected NOE application filter appears in the field below the NOE Traffic filter option checkbox on the TFTP Application Filter Editor.

- 8 The **NOE Audit** portion of the window has the following options:
 - **Session**— If this option is enabled (checkbox is checked) and the **Session Audit** parameter in a rule is set to **Basic** or **Detailed** (this parameter is set for a rule on the Brick Zone Rule Editor), the Session Log will record and provide details for all TFTP sessions that use this filter. If this option is disabled (checkbox is unchecked), the Session Log does not provide any TFTP session details. This option is enabled by default.
 - **Exception**— If this option is enabled (checkbox is checked) and the **Exception Audit** parameter in a rule is set to a value other than **None** (this parameter is set for a rule on the Brick Zone Rule Editor), the Session Log will record and provide details on blocked TFTP sessions that use this filter. This option is enabled by default.

 - 9 From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.
-

END OF STEPS

To add the filter to a service group

Complete the following steps to add this application filter to the **TFTP_App** Service Group or another existing service group. For complete details about creating a service group, refer to [Chapter 4, “Service Groups”](#) in the *SMS Policy Guide*.

- 1 In the Navigator window, open the appropriate Service Group folder and double-click the **TFTP_App** service group.

 - 2 Highlight and double click the entry for udp port 69.

 - 3 Click the pulldown for the application filter, select your TFTP application filter, and click **OK**.

 - 4 Open the File menu and select one of the **Save** options.
-

Important! For this filter to be active on the Brick, the service group must be included as part of a rule in a Brick zone ruleset. For additional information about Brick Zone Rulesets, refer to [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#). For additional information about service groups, refer to [Chapter 4, “Service Groups”](#).

END OF STEPS



TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic

Overview

A set of pre-defined TFTP application filter templates are provided with the SMS application when it is installed to make it easier to configure the Brick to handle VoIP/NOE traffic. These TFTP application filter templates are automatically pre-selected in the appropriate Brick zone ruleset. The name of the application filter template indicates where the Brick is being deployed to serve as an ALG in the VoIP network and whether the NOE application filter called within the TFTP application filter is pre-configured for Brick NATing between the TFTP server (MGW device) behind the Brick and call server(s) for IPlink protocol signalling messages. For example, the pre-defined TFTP application filter **TFTP_with_NOE_NAT_Branch** is pre-configured to filter VoIP call and data traffic to/from IP phone devices located behind the Brick zone at a Branch Office site and for the Brick to perform NAT for MGW-CS IPlink protocol message communications.

Related information

Review the following sections in this Guide for related information on use of the pre-defined templates provided with the SMS application for VoIP traffic:

- [“Brick Zone Ruleset Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 1-66)
- [“Host Group Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 2-25)
- [“Service Group Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 4-27)
- [“NOE Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-68)
- [Appendix E, “Configuring the Brick for VoIP/NOE Traffic Using Pre-Defined SMS Templates”](#)

Downloading the NOE TFTP application filter templates for VoIP/NOE traffic

After the SMS application is installed, the NOE templates can be downloaded by right-clicking on a group folder (such as the *system* folder initially created when the SMS is installed) and choosing **Create NOE Template** from the pop-up menu.

When the NOE template set is created, the pre-defined TFTP application filter templates are downloaded to a subfolder named **NOE** under the **Application Filters** folder in the SMS Navigator (refer to [Figure 5-56, “Pre-Defined TFTP Application Filter Templates \(noe Subfolder in Navigator\)”](#) (p. 5-124)).

Figure 5-56 Pre-Defined TFTP Application Filter Templates (noe Subfolder in Navigator)

The screenshot shows the Alcatel-Lucent Security Management Server Navigator window. The left pane displays a tree view of the system hierarchy, with the 'NOE' subfolder under 'Application Filters' selected. The right pane shows a table of 19 pre-defined TFTP Application Filter Templates.

Name	Admin	Last Modified	Description
NOE_HTTP	setup	2009-04-29 16:17:41	
NOE_NAT_Branch	setup	2009-04-29 16:17:40	
NOE_NAT_Remote	setup	2009-04-29 16:17:40	
NOE_Standard_PCS	setup	2009-04-29 16:17:40	
NOE_Standard_inside_Branch	setup	2009-04-29 16:17:40	
NOE_Standard_inside_HQ	setup	2009-04-29 16:17:41	
NOE_Standard_inside_Regional	setup	2009-04-29 16:17:41	
NOE_Standard_inside_Remote	setup	2009-04-29 16:17:41	
NOE_Standard_outside_HQ	setup	2009-04-29 16:17:41	
NOE_Standard_outside_Regional	setup	2009-04-29 16:17:41	
TFTP_with_NOE_NAT_Branch	setup	2009-04-29 16:17:41	
TFTP_with_NOE_NAT_Remote	setup	2009-04-29 16:17:41	
TFTP_with_NOE_PCS	setup	2009-04-29 16:17:41	
TFTP_with_NOE_inside_Branch	setup	2009-04-29 16:17:41	
TFTP_with_NOE_inside_HQ	setup	2009-04-29 16:17:41	
TFTP_with_NOE_inside_Regional	setup	2009-04-29 16:17:41	
TFTP_with_NOE_inside_Remote	setup	2009-04-29 16:17:41	
TFTP_with_NOE_outside_HQ	setup	2009-04-29 16:17:41	
TFTP_with_NOE_outside_Regional	setup	2009-04-29 16:17:41	

To remove the NOE zone ruleset templates and all component templates from the Policies folder, right-click on the group folder and choose **Delete NOE Template**. A confirmation dialog box is displayed, asking if you are sure that you want to delete the NOE policy template (which includes all of the zone ruleset templates and component templates) under the selected group folder. Choose **Yes**. The NOE templates are removed from the Policies folder under the selected group. If, however, one of the zone ruleset templates is in use on one or more Bricks, an error message is generated and the templates are not removed.

Summary description of Pre-Defined TFTP Application Filter Templates for VoIP/NOE Traffic

The following sections describe the default settings within each of the pre-defined TFTP application filter templates (refer to [Figure 5-54](#), “TFTP Application Filter Editor” (p. 5-119) for a sample of the TFTP Application Filter Editor window).

TFTP_with_NOE_NAT_Branch application filter

The **TFTP_with_NOE_Branch** application filter has the standard TFTP default settings with other fields configured as follows:

- The **Create dynamic rules for NOE Traffic** field is enabled (checkbox is checked).
- Both **NOE Audit** options are selected.
- In the **NOE Application Filter** field, the pre-defined NOE application filter **NOE_NAT_Branch** is already selected.

TFTP_with_NOE_NAT_Remote

The **TFTP_with_NOE_NAT_Remote** application filter has the standard TFTP default settings with other fields configured as follows:

- The **Create dynamic rules for NOE Traffic** field is enabled (checkbox is checked).
- Both **NOE Audit** options are selected.
- In the **NOE Application Filter** field, the pre-defined NOE application filter **NOE_NAT_Remote** is already selected.

TFTP_with_NOE_inside_Branch

The **TFTP_with_NOE_inside_Branch** application filter has the standard TFTP default settings with other fields configured as follows:

- The **Create dynamic rules for NOE Traffic** field is enabled (checkbox is checked).
- Both **NOE Audit** options are selected.
- In the **NOE Application Filter** field, the pre-defined NOE application filter **NOE_Standard_inside_Branch** is already selected.

TFTP_with_NOE_inside_HQ

The **TFTP_with_NOE_inside_HQ** application filter has the standard TFTP default settings with other fields configured as follows:

- The **Create dynamic rules for NOE Traffic** field is enabled (checkbox is checked).
- Both **NOE Audit** options are selected.
- In the **NOE Application Filter** field, the pre-defined NOE application filter **NOE_Standard_inside_HQ** is already selected.

TFTP_with_NOE_inside_Regional

The **TFTP_with_NOE_inside_Regional** application filter has the standard TFTP default settings with other fields configured as follows:

- The **Create dynamic rules for NOE Traffic** field is enabled (checkbox is checked).
- Both **NOE Audit** options are selected.
- In the **NOE Application Filter** field, the pre-defined NOE application filter **NOE_Standard_inside_Regional** is already selected.

TFTP_with_NOE_inside_Remote

The **TFTP_with_NOE_inside_Remote** application filter has the standard TFTP default settings with other fields configured as follows:

- The **Create dynamic rules for NOE Traffic** field is enabled (checkbox is checked).
- Both **NOE Audit** options are selected.
- In the **NOE Application Filter** field, the pre-defined NOE application filter **filterNOE_Standard_inside_Remote** is already selected.

TFTP_with_NOE_outside_HQ

The **TFTP_with_NOE_outside_HQ** application filter has the standard TFTP default settings with other fields configured as follows:

- The **Create dynamic rules for NOE Traffic** field is enabled (checkbox is checked).
- Both **NOE Audit** options are selected.
- In the **NOE Application Filter** field, the pre-defined NOE application filter **filterNOE_Standard_outside_HQ** is already selected.

TFTP_with_NOE_outside_Regional

The **TFTP_with_NOE_outside_Regional** application filter has the standard TFTP default settings with other fields configured as follows:

- The **Create dynamic rules for NOE Traffic** field is enabled (checkbox is checked).
- Both **NOE Audit** options are selected.
- In the **NOE Application Filter** field, the pre-defined NOE application filter **filterNOE_Standard_outside_Regional** is already selected.

TFTP_with_NOE_PCS

The **TFTP_with_NOE_PCS** application filter has the standard TFTP default settings with other fields configured as follows:

- The **Create dynamic rules for NOE Traffic** field is enabled (checkbox is checked).
- Both **NOE Audit** options are selected.
- In the **NOE Application Filter** field, the pre-defined NOE application filter **filterNOE_Standard_PCS** is already selected.



Global Application Filters

Definition

A global application filter is an application filter that is created in one group, but can be seen and used in every other group. Only SMS Administrators can create global application filters.

Create a Global Application Filter

If you are an SMS Administrator, you create a global application filter by clicking the **Display and Use Globally** checkbox on the Application Filter Editor. You can do this when you create the application filter, or you can do this after the application filter has been created by editing the application filter.

When creating a global application filter, make sure the name you give the application filter is unique across all groups. If you attempt to give a global application filter a name that is in use elsewhere, the application filter will not be created, and you will get an error message indicating the save failed because there is an object in another group with the same name.

View a Global Application Filter

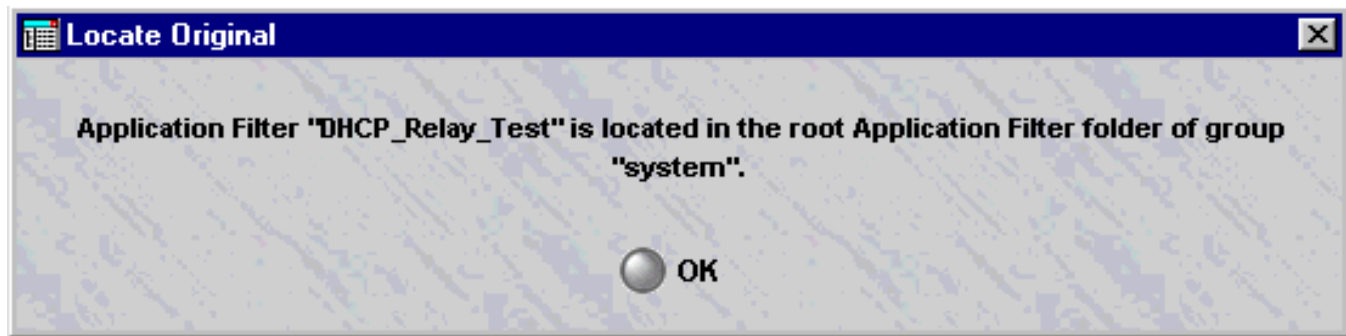
Global application filters appear in the Navigator window, just like a standard application filter. When you view the application filters that have been created in a specific group, all the application filters *created in that group* appear the same, regardless of whether or not they are global.

However, any global application filters created in a *different group* can be identified by a globe icon that appears to the left of the entry, as shown below.

Name	Admin	Last Modified	Description
 DHCP_Relay_Test	art	2002-06-03 23:07:18	Dynamic Host Configuration Protocol Relay

To determine the group in which the global application filter was originally created, right-click the application filter in the Navigator window and select **Original Location** from the pop-up menu. A window similar to the one shown in [Figure 2-4, “Locate Original Window”](#) (p. 2-9) will appear.

Figure 5-57 Locate Original Window (Application Filters)



Removing the Global Status of an Application Filter

Just as it is possible to make a non-global application filter global by clicking the **Display and Use Globally** checkbox on the Application Filter Editor, it is possible to remove the global status of an application filter by unchecking this checkbox. However, this can only be done if the application filter is *not* in use globally. If the application filter is in use in any group *other than the one in which it was created*, you cannot remove its global status.

It is also possible to delete a global application filter — but only the original source of the application filter. You cannot delete a global application filter from any folder in which it appears except the folder in which it was originally created. No application filter — global or standard — can be deleted if it is in use anywhere.

Similarly, only the original source of a global application filter can be moved.

Permissions

Permissions over global application filters are based on the group in which the application filter was originally created. If an administrator has FULL policy permissions for that group, then that administrator has FULL permissions over all global application filters created in that group.

If an administrator has FULL permission over a global application filter, the administrator can edit the application filter in any of the groups in which it appears. An administrator can create a copy of a global application filter as long as the administrator has FULL permissions over the *destination* group.



To Copy an Application Filter

Overview

Use this procedure to copy an application filter. You can copy an application filter to a different folder in the same group, or to a folder in another group, provided that you have *Full* permission for that folder. When you edit an application filter (using the **Save As** option to save the edited version with a new name) or copy an application filter, you have the option of copying components from the original application filter (host groups, other application filters), specifying how those components will be named and where they will be created and stored.

Task

Complete the following steps to copy an application filter:

- 1 With the application filters displayed in the Navigator window, right-click the application filter to be copied and select **Copy** from the pop-up menu.

Result A Copy window is displayed (Figure 5-58, “Copy Window” (p. 5-129)).

Figure 5-58 Copy Window



- 2 Complete the following fields, as applicable:
 - **Name**— Enter a new name for the application filter copy (Example: **TFTP_2**). This field is required. The name entered must be different from the name of the original application filter by at least one character if it is in the same group.
 - **Folder**— The copy of the application filter will be created and stored in the **Application Filters** folder under the *system* group, by default. To specify a different folder, click the **Browse** button, which brings up a Select a Application Filter Folder window, and select the folder where the application filter copy will be created and stored.

- **Copy components when copying within a Group**— Click this checkbox (place a check in it) to copy the components from the original application filter. If this checkbox is unchecked, and the copied application filter is in the same group as the original one, the copied application filter uses the same components as the original. If the application filter is copied across groups, the components are copied regardless of this field setting.
- **Component Suffix**— If the **Copy components when copying within a Group** option is checked, enter a text string that will be appended to the name of each copied component. If this field is left blank, the name of each copied component will be spelled the same way as the original one. Name collisions are possible, regardless of whether or not an entry is made in this field. If this occurs, the content of the conflicting component(s) is examined. If the content is the same, the duplicated host group uses the existing component(s) and the new copy(ies) is not made. If the content is different, a copy of the component(s) is created and stored under a unique name by appending its database ID to the name.
Note: the Component Suffix used must be unique and cannot be reused across zone rulesets.
- **Component Folder**— This field is used to specify a common folder name for the copied components. The folder name is relative to the top-level folder of each type of component. If this field is left blank, the copied components are stored in the top-level folder of that component. This folder is created as needed (if it does not already exist). A path string can be specified; / or \ delimiters can be used to delineate folder names.

3 Click **OK** to copy the application filter and dismiss the Copy window.

Result The system returns to the Navigator window. The copy of the application filter appears in the Contents panel.

The application filter (and components, if specified) is copied to the chosen folder(s).

END OF STEPS



6 Network Address Translation

Overview

Purpose

This chapter explains how to set up Network Address Translation (NAT). NAT is a feature that enables the Alcatel-Lucent *VPN Firewall Brick*® Security Appliance to map the source or destination addresses of inbound and outbound sessions to other addresses.

Contents

What is Network Address Translation?	6-2
To Set Up Source Address Mapping	6-4
To Set Up Destination Address Mapping	6-8
To Set Up Destination Port Mapping	6-12
Dynamic NAT	6-13
To Perform Source Address Mapping with a Router	6-16
To Perform Source Address Mapping without a Router	6-18
To Perform Destination Address Mapping without a Router	6-20
Other Examples of a Brick Responding to ARPs	6-22



What is Network Address Translation?

Definition: NAT

The purpose of NAT is to map the local IP addresses used in one network to different IP addresses known by another network.

For example, for outbound traffic an enterprise might map its local IP addresses to one or more registered IP addresses — and then unmap the registered IP addresses on incoming packets back to the local IP addresses.

Important! Since there is only one Local Map Address pool, and Client Tunnels typically assign private addresses from the pool, and Dynamic NAT typically assign public addresses from the pool, it is typically not possible to use Client Tunnels and Dynamic NAT in the same zone.

Types of Network Address Translation

The Brick supports three types of mapping:

- Source address mapping
- Destination address mapping
- Destination port mapping

Reasons to Use Network Address Translation

As an Administrator, there are a number of reasons for you to consider making use of the Brick NAT feature:

- *Connect unregistered addresses*
If your network contains unregistered Internet addresses, you can use the NAT feature to map these unregistered IP addresses to other, registered addresses. This allows these hosts to communicate with servers on the Internet.
- *Minimize the number of IP addresses required*
Using the NAT feature, you can limit the number of IP addresses required by your organization. You can, for example, use a limited pool of IP addresses, or even one IP address, to represent a network to the outside world.
- *Tighten security*
NAT enables you to increase security by concealing sensitive IP addresses from the Internet. Even if these addresses are registered, you can map them to other servers to mask their true identities.
- *Perform load balancing*
NAT allows you to relieve overworked servers by distributing the load across multiple machines. You can, for example, provide a single URL to individuals seeking access to a corporate website, and then map the address to a pool of servers to share the load.

ICMP Messages

A Brick supports many-to-one NAT for ICMP *pings* that are sent simultaneously from multiple hosts to a host protected behind the Brick, customers that use ICMP *pings* for network health checks. The Brick uses the ID field in an ICMP message header as a unique identifier to associate requests and responses with the originating host. A Brick supports a minimum of 100 simultaneous ICMP NAT sessions per zone.



To Set Up Source Address Mapping

Overview

Source address mapping is used when you want to mask the identity of your local hosts when they communicate with the Internet, for example, because the local addresses are unregistered or for reasons of security. Therefore, source address mapping is usually used for outbound sessions.

The first thing you have to do is create a security rule to allow outbound traffic through the Brick. The source in this rule would be a host group containing the local IP addresses. If there is only one local host, you can enter its IP address instead of a host group. Then, you have to provide the addresses to be mapped to the source address(es) and indicate the address selection type.

Task

To perform source address mapping, follow the steps below:

- 1 With the Navigator window displayed, open the appropriate Group folder, and then open the Policies folder.
- 2 Click **Brick Zone Rulesets** to display all existing Brick zone rulesets.
- 3 Double-click the ruleset you want. The Brick Zone Ruleset Editor will appear.
- 4 Right-click in the Rule Viewer and select **New** from the pop-up menu. The Brick Zone Rule Editor will appear (see [Figure 1-3, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-17) in [“Overview”](#) (p. 1-1)).
- 5 Create an outbound rule. [Figure 6-1, “Web Access Rule”](#) (p. 6-4) shows an example of an outbound rule, in this case a rule to allow hosts in the host group *sales_hosts* access to the Web.

Figure 6-1 Web Access Rule

Rule Number	Active	Direction	Source	Destination	Service	Action	Drop Action	Audit	Description
1000	Yes	← Out	sales_hosts	*	http	Pass		Yes	sales department

-
- 6 Click **Address Translation** to display the Address Translation tab. The **Source Address Mapping** box is displayed (see [Figure 6-2, “Source Address Mapping Box”](#) (p. 6-5)).

Figure 6-2 Source Address Mapping Box



-
- 7 In the **Address or Host Group** field, enter the address that you want to map the source address(es) in the rule to. In the example in [Figure 6-1, “Web Access Rule”](#) (p. 6-4), the source addresses are the addresses in the host group *brickLocalAddresses*.

There are three ways to do this, and these are described below. If there *is* a router deployed between the Brick and the hosts in the zone, you have to use either the second or third method. If there *is no* router, or if the Brick is the next hop, you must use the first.

Do the following:

- *Enter the keyword “Virtual Brick Address”*
If there is no router between the Brick and the internal hosts, the Brick must respond to ARP requests. For this reason, you have to select the keyword *Virtual Brick Address* from the **Address or Host Group** field drop-down list. The Brick responds when an outside router ARPs for a VBA.

The VBA is entered when a Brick zone ruleset is assigned to an interface on a Brick (it can also be the endpoint of a LAN-LAN or client tunnel). See [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#) for instructions on assigning a Brick zone ruleset to a Brick interface.

This is an example of “Many to One NAT”.

- *Enter an IP address*

If there is a router positioned to respond to ARP requests, you can map the source addresses to a single public IP address. Type the IP address directly into the **Address or Host Group** field.

- *Enter a host group*

If there is a router positioned to respond to ARP requests, you can also map the source addresses to multiple public IP addresses. To do this, you must first create a host group that contains these addresses. Then, select **Browse** from the drop-down list and enter the host group. See [Chapter 2, “Host Groups”](#) for instructions on how to create a host group.

-
- 8** In the **Type** box, indicate the selection type. The selection type determines how the Brick selects the addresses it is mapping the source address to.

If you are mapping a single source address to another IP address, the selection type does not matter, and you can leave the default in place. However, if you are using host groups to map multiple source addresses to multiple other IP addresses, the selection type is important. The following explains both selection methods:

- **Pool**

Pool is the default and is used in the great majority of cases. In pool selection, the Brick takes addresses from the pool you have created as they are needed, and returns the addresses to the pool when the session is over.

If the number of requests for addresses exceeds the number of addresses in the pool, the Brick continues to distribute IP address in a “round robin” fashion. In other words, if there are three addresses in the pool, the first three users receive addresses A, B and C in order. Assuming no addresses are returned to the pool, the fourth user would also use address A; the fifth user would use address B, and so forth. A different source port would be assigned when another user is using a previously allocated address.

- **Direct**

Direct selection is typically used when a number of addresses are being mapped to an equal number of addresses. However, it can also be used to map many private IP addresses to a few public IP addresses; Port Address Translation (PAT) is performed as needed to preserve uniqueness.

When mapping a large group of private IP addresses to a smaller group of public IP addresses, Direct NAT divides the group of private IP addresses into smaller groups of addresses, and establishes a 1:1 mapping between a private IP address in

one of those groups to a public IP address in a session. For example, if there are 200 private IP addresses mapped to 10 public IP addresses, the 200 private IP addresses will be broken down into 20 groups of 10 IP addresses each. Each of the private IP addresses in the groups will have a 1:1 mapping with the 10 public IP addresses.

- **Local**

This option is only used in certain cases with Alcatel-Lucent IPSec Client users that need the "Local Presence" feature. Please refer to the *Local Presence* appendix for more details.

- **Dynamic**

This option is used when the Dynamic NAT feature is applied to source IP address mapping. When Dynamic NAT address mapping is used, all sessions that originate from the same private source IP address of a client host are mapped to an IP address from a per-zone pool of public IP addresses by its associated Brick. This dynamically assigned public IP address is allocated for the duration of the host connection to the service provider network. The mapped IP address is released back to the dynamic IP address pool after all forward sessions have stopped and the HD_timeout has elapsed. The HD_timeout is set to the rule Session Timeout value plus the value assigned to the **Dynamic NAT Release Delay** field on the Options tab of the Brick Editor. The **Dynamic NAT Release Delay** is the period of session inactivity that determines when the assigned IP address is released back to the pool of IP addresses by the Brick. When all sessions associated with the source (host) are terminated, and after the HD_timeout has elapsed, all sessions that are using the same mapped IP address to send data back to the host are deleted.

Two host groups are required for the **Dynamic** option:

- A host group containing a pool of addresses assigned in the **Local Map Address** field on the Policy Assignment tab of the Brick Editor
- A blank host group assigned in the **Address or Host Group** field on the **Source Address Mapping** box.

For more details, refer to the "Dynamic NAT" (p. 6-13) section.

END OF STEPS



To Set Up Destination Address Mapping

When to use

Destination address mapping is used when you want to mask the identity of your local hosts when they are the destination of traffic originating on the Internet. You may be doing this for security reasons, to minimize the number of registered IP addresses you require, or to perform load balancing.

Whatever the reason, the first thing you have to do is create a security rule to allow inbound traffic through the Brick. The destination address in that rule is not the actual addresses of your local hosts, but the address(es) you are mapping the actual addresses to.

It is important to understand this crucial difference between source and destination mapping.

- In *source* address mapping, you enter the actual addresses of your local hosts in the **Source** field of the rule under the Basic tab, and then enter the mapping address in the Address Translation tab.
- In *destination* address mapping, you do the opposite, you enter the actual address of the local hosts in the Address Translation tab, and the mapping address in the **Destination** field of the rule under the Basic tab.

Task

To perform destination address mapping, follow the steps below:

- 1 With the Navigator window displayed, open the appropriate group folder, and then open the Policies folder.
-

- 2 Click **Brick Zone Rulesets** to display all existing zones (rulesets). Right-click the ruleset you want and select **Edit** from the pop-up menu.

Result The Brick Zone Ruleset Editor is displayed .

- 3 Right-click an existing rule and select **New** from the pop-up menu. The Brick Zone Rule Editor will appear. See [Figure 1-3, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-17) in [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#).

- 4 Create an inbound rule. [Figure 6-3, “Local Access Rule” \(p. 6-9\)](#) shows an example of an inbound rule; in this case, a rule to allow Internet hosts to access local servers in the host group *sales_servers*.

Figure 6-3 Local Access Rule

Rule Number	Active	Direction	Source	Destination	Service	VLAN ID	Action
65535	Yes	↔ Both	*	*	*	*	● Drop

- 5 Click **Address Translation** to display the Address Translation tab. The **Destination Address Mapping** box appears below the **Source Address Mapping** box. It is shown in [Figure 6-4, “Destination Address Mapping Box” \(p. 6-9\)](#).

Figure 6-4 Destination Address Mapping Box

- 6 In the **Address or Host Group** field, enter the actual IP addresses of the local hosts. These are the addresses that the destination addresses in the rule will be mapped to. In the example in [Figure 6-3, “Local Access Rule” \(p. 6-9\)](#), the destination addresses are the addresses in the host group *sales_servers*.

Do the following:

- *Enter an IP address*
If you are mapping a single local host to a single destination address, enter the local host’s address directly into the **Address or Host Group** field.
- *Enter a host group*
If there is a router positioned to respond to ARP requests, you can also map the source addresses to multiple IP addresses. To do this, you must first create a host group that contains these addresses. Then, select **Browse** from the drop-down list and enter the host group. See [Chapter 2, “Host Groups”](#) for instructions on how to create a host group.

-
- 7 In the **Type** box, indicate the selection type. The selection type determines how the Brick selects the addresses it is mapping the source address to.

If you are mapping a single destination address to another IP address, the selection type does not matter, and you can leave the default in place. However, if you are using host groups to map multiple destination addresses to multiple other IP addresses, the selection type is important. The following explains both selection methods:

- **Pool**

Pool is the default and is used in the great majority of cases. In pool selection, the Brick takes addresses from the pool you have created as they are needed, and returns the addresses to the pool when the session is over.

If the number of requests for addresses exceeds the number of addresses in the pool, the Brick continues to distribute IP address in a "round robin" fashion. In other words, if there are three addresses in the pool, the first three users receive addresses A, B and C in order. Assuming no addresses are returned to the pool, the fourth user would also use address A; the fifth user would use address B, and so forth. A different source port would be assigned when another user is using a previously allocated address.

- **Direct**

In direct selection, the Brick maps the addresses on a one-to-one basis.

For example, if you are mapping ten addresses in Host Group A to ten addresses in Host Group B, the first address in A will be mapped to the first address in B, the second address in A to the second address in B, and so forth. If you were mapping ten addresses to 15, the last five would go unused.

Direct selection is typically used when a number of addresses are being mapped to an equal number of addresses. The purpose is to ensure that each address is mapped to a unique address that is shared by no other machine.

- **Local**

This option is only used in certain cases with Alcatel-Lucent IPsec Client users that need the "Local Presence" feature. Please refer to [Appendix A, "Local Presence"](#) in this Guide for more details.

- **Dynamic**

This option is used when the Dynamic NAT feature is applied to destination IP address mapping. When Dynamic NAT address mapping is used, and an incoming data packet is being sent back to the originator (source) host, the Brick looks up the public IP address being used to see if it is currently mapped to the private IP address of a host (source) through the Dynamic NAT feature. If no reverse mapping is found for the destination address of the packet, then the packet is dropped.

Two host groups are required for the **Dynamic** option:

- A host group containing a pool of addresses assigned in the **Local Map Address** field on the Policy Assignment tab of the Brick Editor
- A blank host group assigned in the **Address or Host Group** field on the **Source Address Mapping** box.

For more details, refer to the [“Dynamic NAT” \(p. 6-13\)](#) section.

END OF STEPS



To Set Up Destination Port Mapping

When to use

To perform destination port mapping, you have to enter the ports you want to map the destination port to. The ports in this entry will be used as the destination ports for all traffic that matches the security rule.

Task

- 1 With the Navigator window displayed, open the appropriate group folder, and then open the Policies folder.

- 2 Click **Brick Zone Rulesets** to display all existing zones (rulesets). Right-click the ruleset you want and select **Edit** from the pop-up menu. The Brick Zone Ruleset Editor is displayed.

- 3 Right-click an existing rule and select **New** from the pop-up menu. The Brick Zone Rule Editor will appear. See [Figure 1-3, “Brick Zone Rule Editor \(Basic Tab\)” \(p. 1-17\)](#) in [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#).

- 4 Create an inbound rule (see [Figure 6-3, “Local Access Rule” \(p. 6-9\)](#)), and then click **Address Translation** to display the Address Translation tab.

- 5 In the **Destination Port Mapping** field, enter the number of the port you want to map the destination port to.

END OF STEPS



Dynamic NAT

Definition

Dynamic NAT is a variation of network address translation, in which the original (usually private) IP address of a client that is connecting to a service provider network is dynamically mapped to another (usually public) IP address by its supporting Brick from a pool of IP addresses, on a per-zone basis. Unlike Pool NAT, all sessions associated with a single client - and only that client - are mapped to the same IP address.

The initial assignment of the IP address is done when the first outbound session is established from the client and a properly configured source NAT rule is triggered. Subsequent sessions may go through the Brick in either direction using source NAT for outbound or destination NAT for inbound sessions. If an inbound packet matches the rule, but does not match any current mapping, then it is dropped.

The mapping is deleted after all the outbound sessions from the same source are terminated and no new outbound sessions are created for a period of time, called the **Dynamic NAT Release Delay**, has elapsed. Any remaining inbound sessions will also be deleted at that time.

The Brick assigns IP addresses from the Dynamic NAT pool to clients until the pool of IP addresses is exhausted. When it runs out of available IP addresses, the Brick issues an error code and drops the client session.

An alarm trigger can be configured in the SMS managing the Brick to generate an alarm when the Dynamic NAT pool of IP addresses is exhausted. For details about configuring alarm triggers, refer to the *SMS Report, Alarms, and Logs* guide.

Important! Dynamic NAT and Alcatel-Lucent IPSec Clients are mutually exclusive. While it is possible to enable both in the policy, they will not work together correctly.

To set up dynamic NAT

Setting up Dynamic NAT is a multi-step process that requires the creation of a pool of IP addresses within a host group that can be assigned to a zone ruleset, and, in turn, assigned to the managing Brick, which can be configured on a zone basis for Dynamic NAT source or destination IP address mapping.

The following is the recommended procedure to follow:

- 1 Create a "real" host group that contains the list/range of IP addresses to be used as the dynamic NAT IP address pool. To do this, select **Host Groups > New Host Group**. The Host Group Editor is displayed (see [Figure 2-1, "Host Group Editor" \(p. 2-4\)](#)).

Enter a Name and Description (if desired) for the Host Group (for example, **JRealHostGroup**). Click the plus (+) button, or right-click and select **New**. The Host Group Entry window is displayed. Enter a list or range of IP addresses in the **Address or Range** field, and click **OK**. Go to the File menu and select **Save and Close** to save the host group.

- 2 Create a "blank" host group. To do this, select **Host Groups > New Host Group**. The Host Group Editor is displayed. Enter a Name and Description (if desired) for this Host Group. Do not enter any IP addresses for this group. Go to the File menu and *select* **Save and Close** to save the host group.

- 3 Select the **Bricks** folder, then select the Brick to be used for Dynamic NAT. The Brick Editor is displayed. Click on the Options tab. In the Miscellaneous Options tab section, change the value of the **Dynamic NAT Release Delay (secs)** field, if desired. This field sets the amount of time to wait before releasing a dynamically assigned IP address back to the Dynamic NAT pool after all sessions have terminated. The default value is **60** (seconds). A value of zero (**0**) causes the mapped IP address (and any remaining inbound sessions) to be released immediately after the last outbound session has terminated.

- 4 Click on the Policy Assignment tab (see [Figure 1-13, "Brick Policy Assignment Editor" \(p. 1-47\)](#)), and select the interface with the zone ruleset to be used for Dynamic NAT.
Result The Brick Policy Assignment Editor is displayed.

- 5 If a Virtual Brick Address is not assigned to the interface, enter a "dummy" IP address in the **Tunnel Endpoint/Virtual Brick Address** field to open the **Local Map Address** field.

- 6 Click the down arrow next to the **Local Map Addresses** field to display a drop-down list, and select the "real" Host Group that contains the list of IP addresses to be used for the Dynamic NAT pool. Click **OK**. Go to the File menu and select **Save and Apply** to save and apply these changes to the Brick.

- 7 Select the **Brick Zone Rulesets** folder and select the Zone Ruleset that was applied on the Brick interface for Dynamic NAT in step 4. The Brick Zone Ruleset Editor is displayed.

-
- 8 Double-click on a rule in the Zone Ruleset, or right-click and select **Edit**. The Brick Zone Rule Editor is displayed (see [Figure 1-3, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-17)).
 - 9 Click on the Address Translation tab. Depending on whether you want the Brick to perform **Source IP Address Mapping** or **Destination Address Mapping** using Dynamic NAT, click the down arrow next to the respective field to display a drop-down list and select the “blank” Host Group created in step 2. Select **Dynamic** as the Type.

END OF STEPS




To Perform Source Address Mapping with a Router

When to use

In any NAT scenario, the administrator must determine which device will answer the ARP requests. If there is a router deployed between the hosts in the zone and the Brick, you can map source (or destination) IP addresses to any address that the router can route to. Since the Brick is optimally used as a bridge-level device, this is a typical configuration.

Scenario

In the diagram below, client 10.1.1.1 is attempting to establish a telnet session to server 123.1.1.1. In order to hide this private source IP address, the Administrator will create a rule (see below) mapping 10.1.1.1 to 100.1.1.1, a public IP address that Router 2 can route to. In this example, we assume that the Brick must be



used as a bridge and all of its ports are in the 100.1.1.0/24 subnet. In addition, the router ports connected to the Brick are also set to 100.1.1.0/24 subnet. Finally, all devices are assumed to have a netmask of /24.

In this scenario, the following happens when the first packet from client 10.1.1.1 reaches the Brick:

- The Brick receives the packet and maps the private source IP address (10.1.1.1) to the public IP address entered by the Administrator (100.1.1.1).
- The Brick forwards the packet to Router 2, which sends it to the destination IP address (123.1.1.1).
- The return packet has a source IP address of 123.1.1.1 and a destination IP address of 100.1.1.1.
- When the return packet reaches Router 2, the router ARPs for the gateway to 100.1.1.1 (Router 1).

- The Brick passes the ARP through to Router 1, which responds to the ARP.
- Router 2 then sends the packet to the Brick, which maps the destination address in the packet (100.1.1.1) to the original private source address (10.1.1.1).
- The Brick then sends the packet to 10.1.1.1 to complete the telnet session.

Rule

The following is the rule the Administrator created for the above scenario:

Field	Purpose
Direction	Out of zone
Source	10.1.1.1
Destination	*
Service	telnet
Action	Pass
Source Address Mapping	100.1.1.1



To Perform Source Address Mapping without a Router

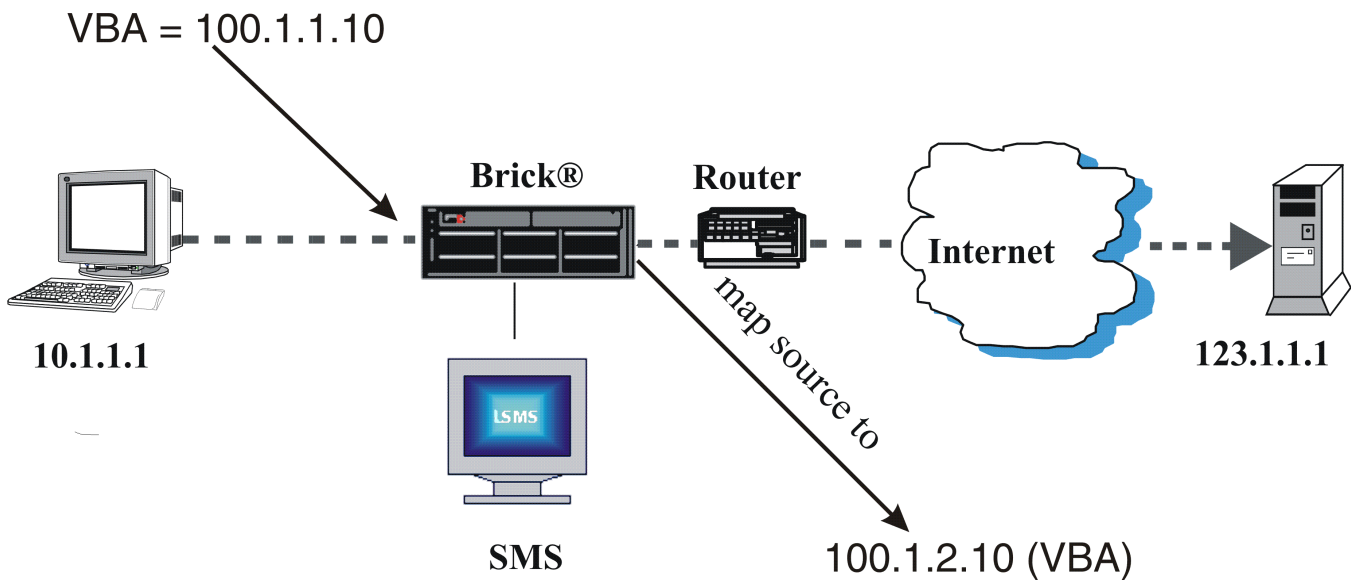
When to use

If there is no router between the hosts in your zone and the Brick, you can map source IP addresses to the VBA assigned to the Brick interface, so the Brick can respond to an ARP request.

If you do not do this, there will be no device to respond when the router on the Internet side of the Brick receives the return packet and ARPs for the public source IP address (the packet's destination IP address).

The diagram below shows the same scenario as the preceding section, except that this time there is no router between host 10.1.1.1 and the Brick. Since there is no router, the Administrator has to map the private source IP address to the VBA.

The Administrator does not actually have to enter the VBA. All the Administrator has to do is select the keyword **Virtual Brick Address** from the drop-down list in the field.



In the above scenario, the following happens when the first packet from client 10.1.1.1 reaches the Brick:

- The Brick receives the packet and maps the source IP address (10.1.1.1) to the VBA (100.1.1.10).
- The Brick forwards the packet to the router, which sends it to its destination IP address (123.1.1.1).
- The return packet has a source IP address of 123.1.1.1 and a destination IP address of 100.1.1.10 (VBA).
- When the return packet reaches the router, the router ARPs for 100.1.1.10 (VBA).

- The Brick (rather than a router) responds to the ARP with the MAC address of the interface card.
- The router sends the packet to the Brick, which maps 100.1.1.10 (VBA) to the private source address (10.1.1.1).
- The Brick then sends the packet to 10.1.1.1 to complete the telnet session.

Rule

The following is the rule the Administrator created for the above scenario:

Field	Purpose
Direction	Out of zone
Source	10.1.1.1
Destination	*
Service	telnet
Action	Pass
Source Address Mapping	Virtual Brick Address




To Perform Destination Address Mapping without a Router

When to use

If there is no router between the hosts in a zone and the Brick, you can map destination IP addresses to the VBA of the interface.

In the diagram below, the address of a web server is 10.1.1.3. Since this address is not legal, the Administrator tells all users who want to connect to the web server that its address is 100.1.1.10, which is really the VBA of the Brick interface to which the web server is connected.



In the above scenario, the following happens when the first packet destined for the web server reaches the router:

- The router ARPs for 100.1.1.10 (the destination address, and the VBA).
- The Brick recognizes the VBA and responds to the ARP with the MAC address of the interface card.
- The router sends the packet to the Brick, which maps the packet to 10.1.1.3, the web server's private address.
- The Brick then sends the packet to the server, completing the session.

Rule

The following is the rule the Administrator created for the above scenario:

Field	Purpose
Direction	In to zone
Source	*
Destination	Virtual Brick Address
Service	webServices
Action	Pass
DestAddrMap	10.1.1.3



Other Examples of a Brick Responding to ARPs

Overview

As noted elsewhere in this chapter, a key consideration for NAT configuration is determining whether a router or the Brick will be responding to ARP requests from the network.

In the illustrations provided, we have demonstrated that NAT may be configured so that either a nearby router will respond to an ARP request that may involve the Brick -OR- a Virtual Brick Address (VBA) on the Brick will respond to an ARP request.

However, it is also possible for the Brick to perform NAT and respond to ARPs on addresses other than the VBA. This can be done by utilizing a field on the Brick under the Policy Assignment tab called "Local Map Addresses".

One to One NAT with Local Map Addresses

In this example, we will assume that we have five internal hosts with private IP addresses that we would like to NAT to five public IP addresses on a one-to-one basis. We'll create "Hostgroup A" for the five internal hosts and "Hostgroup B" for the group of five routable IP addresses. It is important that both host groups have an identical number of entries.

After creating the host groups, the following steps are needed to implement this NAT scenario:

-
- 1 In the Brick zone ruleset, create a rule using source NAT with these parameters:

Direction = OUT

Source = Hostgroup A

Destination = <your_value>

Service = <your_value>

And under the "Address Translation" tab:

Source Mapping = Hostgroup B

Type = Direct

- 2 Next, create a rule using destination NAT with these parameters:

Direction = IN

Source = <your_value>

Destination = Hostgroup B

Service = <your_value>

And under the "Address Translation" tab:

Destination Mapping = Hostgroup A

Type = Direct

Do a "Save and Apply" to download these new rules to the Brick.

For more information on updating Brick zone rulesets, refer to [Chapter 1, "Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets"](#) in this Guide.

-
- 3 Open the Brick editor for the desired Brick and proceed to the Policy Assignment tab. Highlight and double click the port where your ruleset has been assigned. In the Brick Policy Assignment Editor, click on the pulldown for Local Map Addresses and select Hostgroup B. Click OK to accept this entry. Do a "Save and Apply" to download this change to the Brick.

For more information on editing a Brick, refer to the *Configuring Alcatel-Lucent VPN Firewall Brick® Security Appliance Ports* in the *SMS Administration Guide*.

The Brick is now ready to respond to ARP requests to the public addresses listed in Hostgroup B and to NAT packets as needed to the internal hosts in Hostgroup A.

END OF STEPS



7 Dependency Masks

Overview

Purpose

This chapter explains how to set up, use, and maintain dependency masks.

Contents

What is a Dependency Mask?	7-2
To Set Up a Dependency Mask	7-3
Maintaining Dependency Masks	7-8
Example: RealAudio Session	7-15



What is a Dependency Mask?

Definition

A dependency mask is a tool that allows an Administrator to set up a dependency between a particular rule in a Alcatel-Lucent *VPN Firewall Brick*[™] Security Appliance zone ruleset and a specific session in the session cache.

This means that even if a packet matches the rule, and the rule is a pass rule, that packet will still not be permitted to pass through the Brick until the Brick verifies that a certain session, identified in the dependency mask, already exists in the session cache.

Dependency masks are especially useful for handling multimedia applications, because multimedia applications frequently require two sessions, one to send control information, and a second to send the actual data.



To Set Up a Dependency Mask

When to use

To set up a dependency mask, you have to display the Dependency Masks Editor and enter the information requested. The purpose of this information is to define the session that the Brick will look for in the session cache. Once this has been done, you have to associate the dependency mask with a specific rule in a Brick zone ruleset.

To display the dependency masks editor

The easiest way to display the Dependency Masks Editor is to:

- 1 Open the folder of the group that will contain this dependency mask.
- 2 Open the Policies folder.
- 3 Right-click the Dependency Masks folder and select **New Dependency Mask** from the pop-up menu. The Dependency Masks Editor will appear. It is shown in [Figure 7-1, “Dependency Masks Editor”](#) (p. 7-3).

Figure 7-1 Dependency Masks Editor



END OF STEPS

To create a dependency mask

The only fields on the Dependency Masks Editor that are required are the **Name** of the dependency mask, and *one* of the following: **Source IP Addr. or Group**, **Destination IP Addr. or Group**, or **Service or Group**. The more information you enter, the more specifically you will define the session that the Brick is looking for. For example, if you enter TCP in the **Service or Group** field, and nothing else, the rule associated with the dependency mask will execute if the Brick finds *any* TCP session. However, if you also enter a source and destination IP address, the rule will only execute if the Brick finds a TCP session with those particular IP addresses as the source and destination.

To create a dependency mask, follow the steps below.

-
- 1 In the **Name** field, enter a unique name to identify this dependency mask. The name can contain up to 80 characters. It can consist of lower case letters, numbers and certain special characters.

-
- 2 In the **Description** field, you can enter an optional description of the dependency mask. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers, and certain special characters.

This is the description that will appear on the Navigator window when you view existing dependency masks. See the section below entitled [“To view dependency masks”](#) (p. 7-8).

-
- 3 In the **Source IP Addr. or Group** field, enter the source address of the session that must be found in the session cache. There are a number of ways to enter these addresses:

Do the following:

- *Type a specific IP address directly into the field*
This address must be the source address found in the header of the packet in the session cache.
- *Select a host group from the drop-down list*
If you have created a host group containing the source addresses, display the drop-down list and select **Browse**. Then, use the browse window that appears to select the host group. This is the way to enter more than one source address.
- *Select the keyword SOURCE from the drop-down list*
This means the source address in the header of the packet in the session cache must match the *source* IP address of the packets invoking the rule.

- *Select the keyword DESTINATION from the drop-down list*
This means the source address in the header of the packet in the session cache must match the *destination* IP address of the packets invoking the rule.
- *Select an asterisk from the drop-down list*
This means the header of the packet in the session cache can have any address as its source address.

-
- 4 In the **Destination IP Addr. or Group** field, enter the destination address of the session that must be found in the session cache. There are a number of ways to enter this address:

There are a number of ways to enter this address:

- *Type a specific IP address directly into the field*
This address must be the destination address in the header of the packet in the session cache.
- *Select a host group from the drop-down list*
If you have created a host group containing the destination addresses, display the drop-down list and select **Browse**. Then, use the browse window that appears to select the host group. This is the way to enter more than one destination address.
- *Select the keyword SOURCE from the drop-down list*
This means the destination address in the header of the packet in the session cache must match the *source* IP address of the packets invoking the rule.
- *Select the keyword DESTINATION from the drop-down list*
This means the destination address in the header of the packet in the session cache must match the *destination* IP address of the packets invoking the rule.
- *Select an asterisk from the drop-down list*
This means the header of the packet in the session cache can have any address as its destination address.

-
- 5 In the *Service or Group* field, enter the protocol, or protocol and port(s), that must be found in the header of the packet in the session cache.

There are a number of ways to enter this information:

- *Enter the protocol and port(s), using one of the following formats:*
 - Protocol name or number (for example, TCP = 6)
 - Protocol number/destination port (for example, TCP/80)

- Protocol number/destination port/source port (for example, UDP/520/520)
- For ICMP messages, the format is protocol/type/code (for example, 1/type/code).
- *Select a service group from the drop-down list*
Display the drop-down list and select **Browse**. Use the Browse window that appears to select the service group. It can be a service group you created, or one of the service groups provided with the SMS.
- *Select an asterisk from the drop-down list*
This means there are no restrictions on the protocol and ports.

6 In the **Action** box, indicate whether the Brick should look for passed sessions, dropped sessions, or any session (pass or drop). The session cache keeps both types of sessions. The default is *Drop*, but you can change this by clicking the **Pass** or **Any** button.

7 In the **Alarm Code** field, enter an alarm code. An entry in this field is optional.

Some rules have an alarm code associated with them. When a packet arriving at the Brick matches the rule, the alarm associated with the code is triggered, and the appropriate Administrator notified (if configured).

You have the option of entering an alarm code in the dependency mask. If you do, the Brick will search the session cache *only* for sessions matching that alarm code and will not match sessions without the same code.

The efficiency of dependency mask matching is significantly better if an alarm code is included as part of the matching criteria. Whenever possible, associate some unique alarm code with the sessions that the dependency mask should match and include that alarm code in the dependency mask.

8 In the **Hit Count** field, enter the minimum number of occurrences of the session that must be found in the session cache.

The default is 1, but you can change this to any number between 2 - 65,535.

9 Display the File menu and select one of the **Save** options.

END OF STEPS

To associate the dependency mask with a rule

Once you have created a dependency mask, you have to associate it with a rule. This rule will operate like any other rule, except it will not pass or drop a session that it matches until it first searches the session cache for the session described in the dependency mask.

To associate a dependency mask with a rule, you first have to create the rule, using the standard procedure explained in [“To Create a Security Rule”](#) (p. 1-16). Then, once you have created the rule, follow the steps below:

- 1 Display the rule in the Brick Zone Ruleset Editor.
.....
- 2 Click **Advanced** to display the Advanced tab.
.....
- 3 In the **Dependency Mask** field, display the drop-down list and select **Browse**. A Browse window will appear.
.....
- 4 Select the folder you want to put this dependency mask in. It must be the folder labeled “Dependency Masks” or a subfolder under that folder.
.....
- 5 Click **OK** to return to the Brick Ruleset Editor.
.....
- 6 Display the File menu and select one of the **Save** options.

.....
E N D O F S T E P S
.....



Maintaining Dependency Masks

When to use

Once a dependency mask has been created and saved, it can be viewed, modified, copied, or moved. If the dependency mask is no longer needed it can be deleted.

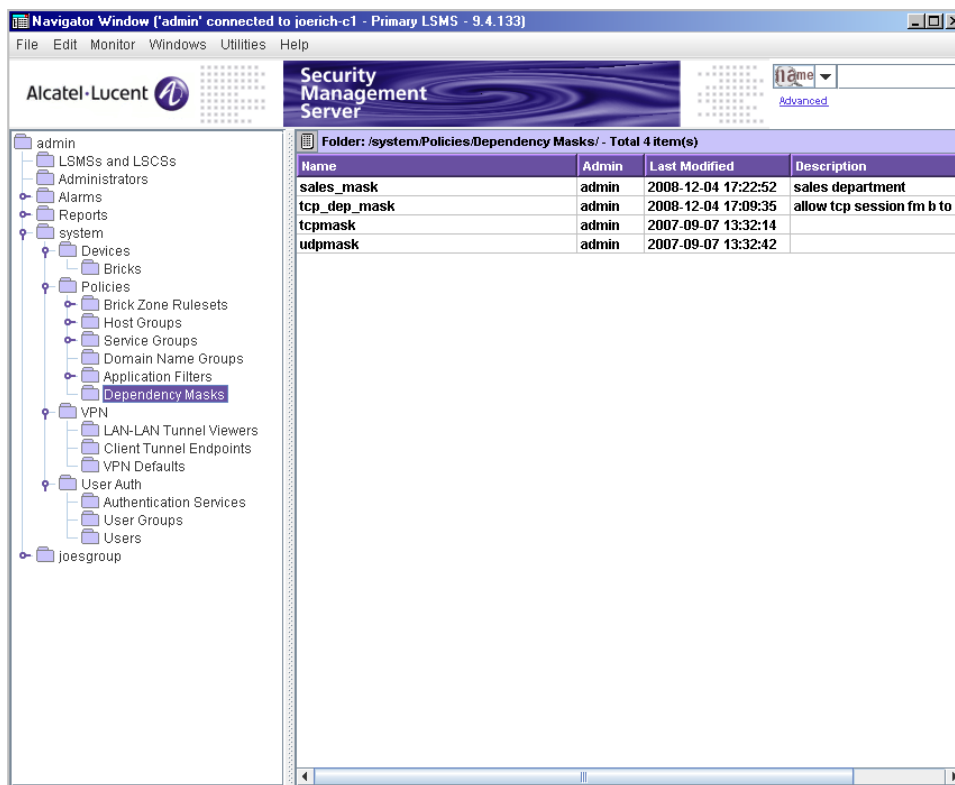
To view dependency masks

You must view the dependency masks before you can edit, move, copy, or delete a specific dependency mask. To view all the dependency masks that have been created to date in a particular group, follow the steps below:

- 1 Open the appropriate group folder, and then open the Policies folder.
- 2 Click the Dependency Masks folder. All existing dependency masks will be displayed in the Navigator window (see [Figure 7-2, “Navigator Window \(View Dependency Masks\)”](#) (p. 7-9)).

For each dependency mask, the Navigator window shows the Administrator who created the mask, the date and time the mask was created, and a brief description, if one was entered when the mask was created.

Figure 7-2 Navigator Window (View Dependency Masks)



END OF STEPS


To find entities that use a dependency mask

The SMS application allows you to perform a search for all entities (such as a zone ruleset) that are using the selected dependency mask.

- 1 With the dependency masks displayed in the Navigator window, right-click the desired host group and select **Find Entities Using this Dependency Mask** from the pop-up menu.

Result An Entities Found Using this Dependency Mask window is displayed (Figure 7-3, “Entities Found Using this Host Group Window” (p. 7-10)) showing the results of the search.

Figure 7-3 Entities Found Using this Host Group Window



The display results on the window show the zone ruleset that uses or contains the selected dependency mask, the group and subfolder in which the entity is found, the name of the dependency mask being used, and details about its use in the ruleset (the display lists the specific rules in which the dependency mask is used).

To view additional details about the entity, right-click on the entry in the table and select **View** from the pop-up menu.

To edit the entity, right-click on the entry in the table and select **Edit** from the pop-up menu.

- 2 Click the Close (X) button at the top of the window to close the window.

END OF STEPS

To modify a dependency mask

You can modify any field in a dependency mask except the **Name** field. To modify a dependency mask, follow the steps below:

- 1 With the dependency masks displayed in the Navigator window, right-click the dependency mask you want to modify and select **Edit** from the pop-up menu.

The Dependency Masks Editor will appear, with the **Name** field greyed-out and the other fields active, as shown in [Figure 7-4, “Dependency Masks Editor \(Edit Mode\)”](#) (p. 7-11).

Figure 7-4 Dependency Masks Editor (Edit Mode)

-
- 2 Change any information in the active fields.
-
- 3 Display the File menu and select one of the **Save** options.

END OF STEPS

To copy a dependency mask

You can copy a dependency mask to a different folder in the same group, or to a folder in another group, provided you have *Full* permission for that folder. Optionally, you can also copy components from the original dependency mask (host group, service group) and specify how those components will be named and where they will be created and stored.

Complete the following steps to copy a dependency mask:

- 1 With the dependency masks displayed in the Navigator window, right-click the dependency mask you want to copy and select **Copy** from the pop-up menu.

Result A Copy window is displayed (Figure 7-5, “Copy Window” (p. 7-12)).

Figure 7-5 Copy Window



- 2 Complete the following fields, as applicable:

- **Name**— Enter a new name for the dependency mask copy (Example: **sales_mask_2**). This field is required. The name entered must be different from the name of the original dependency mask by at least one character if it is in the same group.
- **Folder**— The copy of the dependency mask will be created and stored in the **Dependency Masks** folder under the *system* group, by default. To specify a different folder, click the **Browse** button, which brings up a Select a Dependency Mask Folder window, and select the folder where the dependency mask copy will be created and stored.
- **Copy components when copying within a Group**— Click this checkbox (place a check in it) to copy the components from the original dependency mask. If this checkbox is unchecked, and the copied dependency mask is in the same group as the original one, the copied dependency mask uses the same components as the original. If the dependency mask is copied across groups, the components are copied regardless of this field setting.
- **Component Suffix**— If the **Copy components when copying within a Group** option is checked, enter a text string that will be appended to the name of each copied component (host group, service group). If this field is left blank, the name of each copied component will be spelled the same way as the original one. Name collisions are possible, regardless of whether or not an entry is made in this field.

If this occurs, the content of the conflicting component(s) is examined. If the content is the same, the duplicated dependency mask uses the existing component(s) and the new copy(ies) is not made. If the content is different, a copy of the component(s) is created and stored under a unique name by appending its database ID to the name.

- **Component Folder**— This field is used to specify a common folder name for the copied components. The folder name is relative to the top-level folder of each type of component. If this field is left blank, the copied components are stored in the top-level folder of that component. This folder is created as needed (if it does not already exist). A path string can be specified; / or \ delimiters can be used to delineate folder names.

-
- 3 Click **OK** to copy the dependency mask and dismiss the Copy window.

Result The system returns to the Navigator window. The copy of the dependency mask appears in the Contents panel.

The dependency mask (and components, if specified) is copied to the chosen folder(s).

END OF STEPS

Move a Dependency Mask

You can move a dependency mask to a different folder in the same group, or to a folder in another group, provided you have *Full* permission for that folder. However, you can only move it to the folder labeled "Dependency Masks" or to a subfolder under that folder. If you try to move it to another folder, you will get an error message. Before moving a dependency mask, make sure that it is not currently used in a rule. If you attempt to move a dependency mask used in a rule, you will get an error message.

Complete the following steps to move a dependency mask:

-
- 1 With the dependency masks displayed in the Navigator window, right-click the dependency mask you want to move and select **Move** from the pop-up menu. A Browse window will appear.
 - 2 Select the folder you want to move this dependency mask to. The dependency mask will be moved to the folder you selected, and you will be returned to the Navigator window.

END OF STEPS

Delete a Dependency Mask

Before deleting a dependency mask, make sure that it is not currently used in a rule. If you attempt to delete a dependency mask used in a rule, you will get an error message.

Complete the following steps to delete a dependency mask:

- 1 With the dependency masks displayed in the Navigator window, right-click the dependency mask you want to delete and select **Delete** from the pop-up menu. A Confirmation window similar to the one shown in [Figure 3-9, “Confirmation Window \(Domain Name Groups\)”](#) (p. 3-17) is displayed.

Figure 7-6 Confirmation Window (Dependency Masks)



- 2 Click **OK** to delete the dependency mask and dismiss the Confirmation window. The dependency mask will be removed from the Navigator window.

END OF STEPS



Example: RealAudio Session


Overview

RealAudio is a multimedia application that allows users to connect to a RealAudio Server and receive audio or video packets from that machine. A RealAudio session is really two sessions, a TCP session to pass control data and a UDP session to pass the audio/video stream. [Figure 7-7, “RealAudio Sessions” \(p. 7-15\)](#) on the next page illustrates this scenario.

To permit these sessions through the Brick without compromising security, an Administrator would have to create two rules, one to pass the TCP control session and the second to pass the audio/video traffic, and a dependency mask. The purpose of the dependency mask is to ensure that the Brick does not pass the audio/video traffic without first looking in the session cache and verifying the existence of the TCP control session.

The alternative to using a dependency mask would be to create one rule passing the audio/video traffic and leave the port that Host A will be listening on open. From a security standpoint, leaving a port open like this is obviously not desirable.

Figure 7-7 RealAudio Sessions



First Rule

To allow Host A to connect to the RealAudio server on the Internet, you first have to create a rule permitting a TCP control channel to be established from Host A to the RealAudio server.

The following shows this rule:

Field	Explanation
Direction	Out Of zone
Source	IP address of Host A
Destination	IP address of RealAudio server
Service	realaudio_control (service group)
Action	Pass
Alarm Code	1234

The service group *realaudio_control* is included with the SMS application. The protocol is TCP and the destination port is 7070, which is the standard RealAudio port.

Alarm code 1234 is a numeric tag that is associated with all sessions created using this rule.

Dependency Mask

Before creating the second rule, you have to create the dependency mask. The reason you have to create the dependency mask first is that you will need to select it from a drop-down list when creating the rule.

In this example, the alarm code in the dependency mask is matching the alarm code in the first rule.

The following shows the dependency mask:

Field	Explanation
Name	realaudio
Source IP	DESTINATION (keyword)
Destination IP	SOURCE (keyword)
Service	realaudio_control (service group)
Action	Pass
Hit Count	1

Field	Explanation
Alarm Code	1234

The alarm code is 1234 to match the alarm code in the sessions created by the first rule. Including the alarm code in the matching criteria significantly improves the efficiency of the session lookup and helps ensure that only the intended sessions are matched.

The name given to this dependency mask is *realaudio*. The source and destination IP addresses are the keywords *DESTINATION* and *SOURCE*, which can be selected from the drop-down lists in these fields. This means:

- The source IP of the dependency mask is the destination IP address in the packet that matches rule #1 (the IP address of the RealAudio server).
- The destination IP of the dependency mask is source IP address in the packet that matches rule #1 (the IP address of Host A).

The service is *realaudio_control* because this is the service group in rule #1.

The hit count is one because only one instance of this session has to be found in the session cache.

Second Rule

Once you have created the dependency mask, you can create the second rule. This rule has to permit a UDP session back from the RealAudio server. This session will transmit the audio/video stream.

This rule will also have to have the dependency mask associated with it. The purpose of the dependency mask is to verify that the TCP control channel enabled by the first rule does in fact exist.

The following is such a rule:

Field	Explanation
Direction	In To Zone
Source	IP address of RealAudio server
Destination	IP address of Host A
Service	realaudio_data (service group)
Action	Pass
Dependency Mask	realaudio

The service group *realaudio_control* is included with the SMS application. The protocol is UDP and the destination ports are 6970-7170, which are the standard RealAudio ports.

How It Works

The following explains how the Brick uses the two rules and the dependency mask:

1. Host A initiates the TCP control session by sending a packet to the RealAudio server via port 7070.
2. The Brick intercepts the packet and applies its rules to the packet.
3. The packet matches the first rule, and is passed through the Brick to the RealAudio server. An entry is created in the session cache.
4. Host A requests an audio stream via this TCP control channel to the RealAudio server.
5. The RealAudio server sends the UDP packets containing the audio signals to Host A on a destination port in the range 6970-7170.
6. The UDP packet arrives at the Brick, and the Brick runs through the rules.
7. The packet matches the second rule, which contains the dependency mask. The Brick checks the dependency mask and then examines the session cache, looking for a session with the following parameters:
 - Source = IP address of Host A
 - Destination = IP address of RealAudio server
 - Service = *realaudio_control*.
 - Alarm code=1234
8. Since such a session does exist, the Brick allows the packet through. All subsequent packets with the same attributes as the first packet (direction, source and destination addresses, source and destination ports, protocol) are also allowed through.

Important! As soon as the source (Host A) closes the RealAudio session, no UDP packets will be allowed into ports 6970-7170 because there will no longer be a session cache entry for the TCP session.

□

8 User Authentication

Overview

Purpose

This chapter explains how to set up user authentication. The purpose of user authentication is to allow users to:

- access resources on another side of an Alcatel-Lucent *VPN Firewall Brick™* Security Appliance or set of Brick devices
OR
- establish a tunnel to a Brick device

on the basis of their individual identities, not the IP addresses of the hosts they are using.

Contents

What is User Authentication?	8-2
How Authentication Works	8-4
Before Setting Up User Authentication	8-8
To Set Up Local Password Authentication	8-11
To Set Up RADIUS and SecurID Authentication	8-21
To Set Up VPN Certificate Authentication	8-32



What is User Authentication?

Overview

Since many users today are mobile and do not have fixed IP addresses, it is important to have the ability to authenticate users on the basis of their identity, rather the address of the computer they are using.

User authentication is a means of controlling access to a network on the basis of the *identities of the individuals requesting access*. It enables a Brick to verify the credentials of potential users before the users are allowed to access protected resources or establish a tunnel.

Types of Users

You can set up user authentication to verify the credentials of two types of users:

- *Application users*
The term "application user" refers to individuals attempting to access an application or service (such as HTTP, FTP or telnet) through a Brick.
For any of a number of reasons, including address assignment by DHCP, and multi-user service, these individuals' IP addresses may not be known ahead of time. Therefore, it is necessary to verify their personal credentials before allowing them access to the hosts protected by the Brick.
- *Client users*
The term "client user" refers to individuals attempting to establish a tunnel, (also referred to as a virtual private network, or VPN), to a Brick.
Client users have to be running the Alcatel-Lucent IPSec Client application on their computers to establish the tunnel. The application will ask for their credentials, and not allow the tunnel to be enabled until their user ID and password, or other authentication credentials, have been verified.
A copy of the Alcatel-Lucent IPSec Client application is provided on a separate CD-ROM that comes with the SMS. For instructions on installing and using the client application, see the *Alcatel-Lucent IPSec Client Installation and User Guide*.

Types of Authentication

The SMS supports four types of authentication:

- **Local Password**

The simplest way to perform authentication is to set up your own user database on the SMS. This database will have an account for each user, and this account will contain the user's ID and password. When the user attempts to login, the SMS will verify the ID and password based on information in the database (see [Figure 8-1, "User Editor" \(p. 8-12\)](#)).

You can use Local Password to authenticate both application and client users.

- **RADIUS**

RADIUS (Remote Authentication Dial-In User Service) is a protocol that permits the remote authentication of users. Various vendors make RADIUS server software. If you already have a database of user accounts on a RADIUS host, you can use that database to authenticate users. If you want to create user accounts on the SMS, and have these users authenticated by RADIUS, you can do that, too. In either case, the SMS will contact the RADIUS server after the user has requested authentication and supplied an ID and password (see [Figure 8-1, "User Editor" \(p. 8-12\)](#)).

You can use RADIUS to authenticate both application and client users.

- **SecurID**

SecurID is a token-based authentication package. The token generates a random code that changes every 60 seconds and is unique to each user. Users enter their user ID and passcode (PIN plus token code), and a SecurID ACE/Server verifies their identity before granting them access.

If you already have a database of user accounts on an ACE/Server, you can use that database to authenticate users. If you want to create user accounts on the SMS, and have these users authenticated by SecurID, you can do that, too. In either case, the SMS will contact the ACE/Server after the user has requested authentication and supplied an ID and password (see [Figure 8-1, "User Editor" \(p. 8-12\)](#)).

You can use SecurID to authenticate both application and client users.

- **VPN Certificate**

A VPN certificate is an electronic means of establishing a user's credentials. It is issued by a certification authority (CA), and it contains a public/private key pair that is used to verify the user's credentials.

□

How Authentication Works

Overview

Users seeking authentication connect to the VBA or tunnel endpoint of the Brick protecting the hosts they seek to access. This sets up a connection with the SMS, initiating the authentication process.

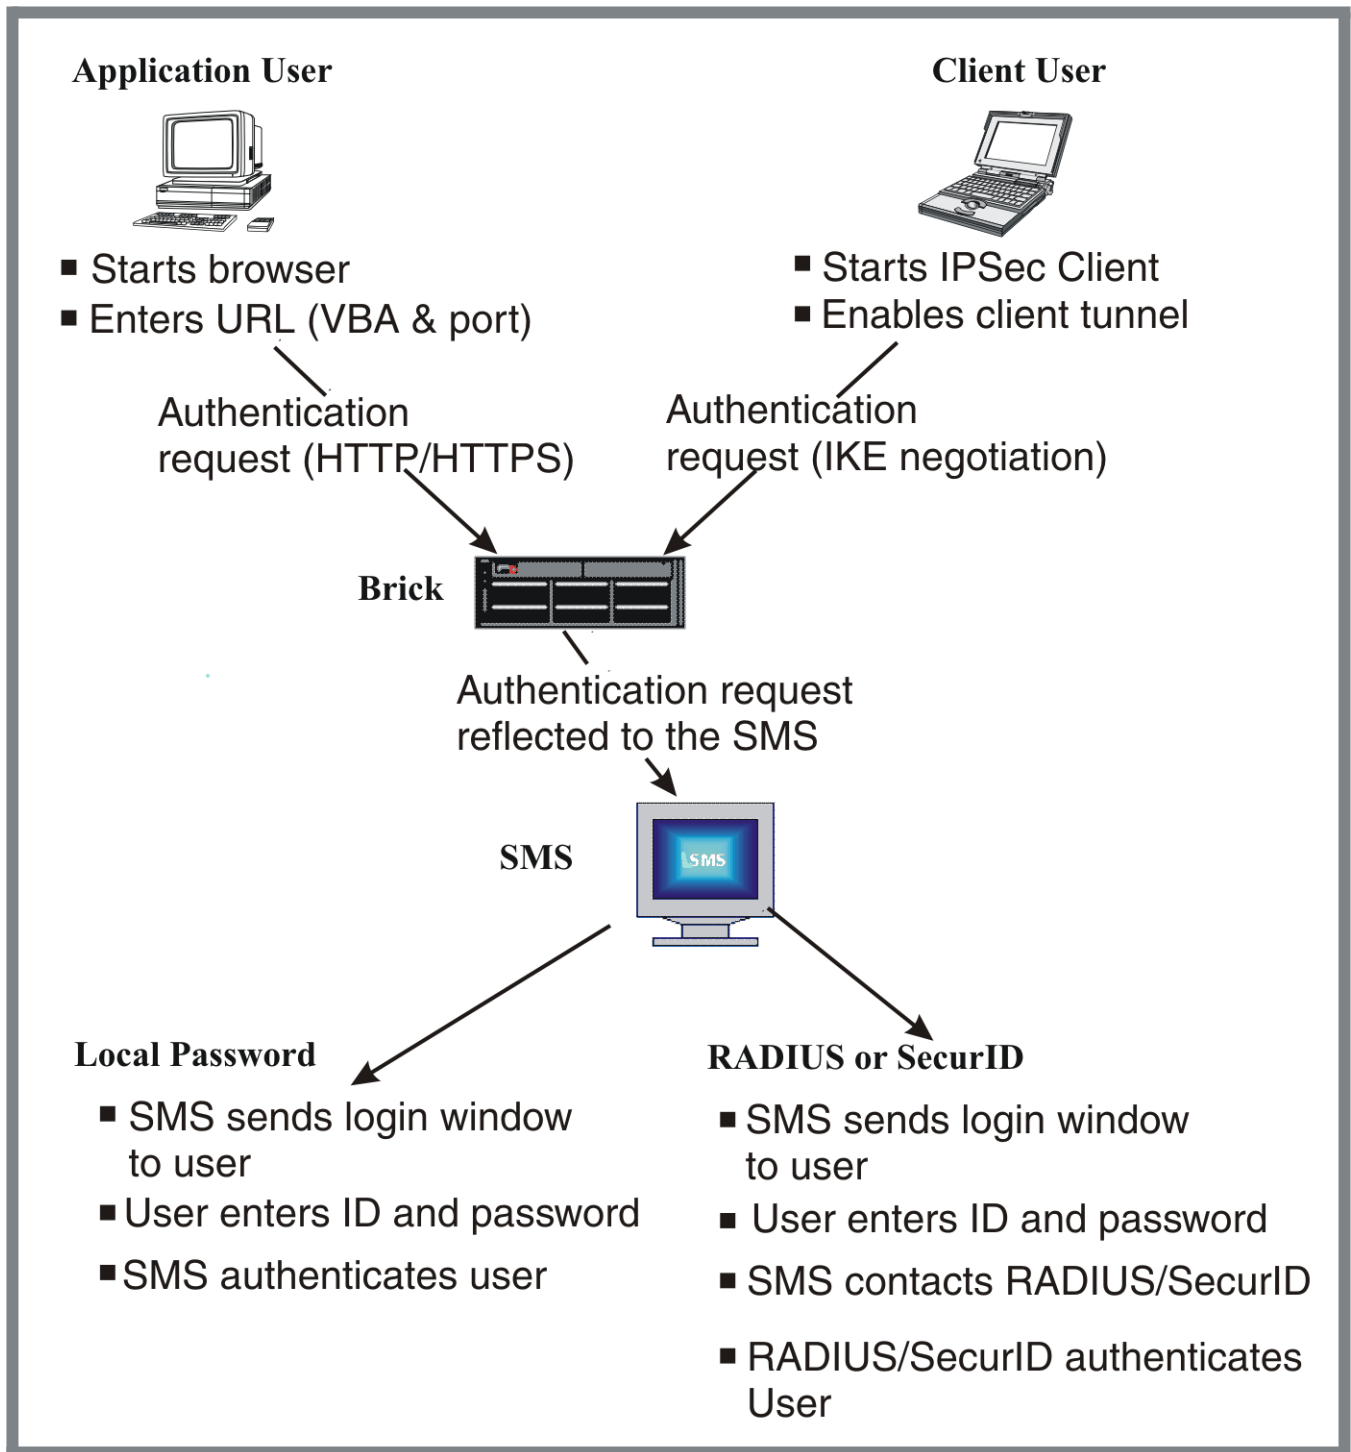
Login Procedure

The login procedure is slightly different for application users than for client users.

- *Application users*
Application users seeking access to network resources protected by a Brick initiate the authentication process by opening a web browser and entering the URL that has been provided by their Security Administrators. This is the URL of the SMS authentication server, and it consists of the virtual Brick address (VBA) of the zone and the authentication server port.
Once the application user enters the URL and establishes a connection with the authentication server, a login window is sent back to the user, who enters the ID and password.
- *Client users*
Client users initiate the authentication process by starting the Alcatel-Lucent IPSec Client application and entering their ID, password, and the tunnel endpoint address in the initial client window. The tunnel endpoint address is used to make the connection to the authentication server.

Authentication Process

The diagram below illustrates the user authentication process.



Authentication Server

The SMS application includes an authentication server that processes incoming authentication requests from application users only.

During installation of the SMS application, you were given a choice of configuring the authentication server as either an HTTP server or an HTTPS server. The advantage of HTTPS is that it ensures the confidentiality of the transmitted credentials. This is extremely important when using a fixed password, but less important when using a changing password like SecurID.

If the authentication server is configured as HTTP, the port number is usually set to 80; if it is HTTPS, the port is usually 443. To change the application user authentication type and port, open the Configuration Assistant and edit the User Authentication parameter. Refer to the *Using the Configuration Assistant* chapter in the *SMS Administration Guide* for instructions.

If the authentication server is HTTPS, you will need a digital certificate. See [Chapter 9, “Digital Certificates”](#) for instructions on how to obtain and install one.

Proxy Table

When a zone is assigned to a port on a Brick, an entry is automatically placed in the Proxy Table for user authentication. Each entry has a **Destination Port** field and a **Proxy Port** field:

- **Destination port**
The destination port is the port assigned to the authentication server. It is usually 80 or 443.
- **Proxy port**
The proxy port is the port used by the SMS to listen for authentication requests (see [Figure 8-1, “User Editor”](#) (p. 8-12)). If the authentication server is HTTP, the proxy port is set to 9010. If it is HTTPS, the proxy port is set to 9011.

Important! If you change the authentication server type *after* one or more zones have been assigned to an interface, you will have to manually edit the entries in the Proxy Table for those zones and change both the destination port and the proxy port.

In addition, proxy rules are automatically added to a zone’s ruleset when the first user authentication rule is first created. A user authentication rule is any rule with a user group as the source. If the destination port was changed after this rule was created, you will have to change the service in the rule accordingly.

Authentication Timeout

When a user is authenticated, the authentication is for a specific period of time, known as the *timeout period*. The timeout period determines how long the user will have to begin new sessions. Once the timeout period expires, the user will have to be re-authenticated to begin new sessions.

The timeout period can range from one minute to 48 hours. If the user database is maintained locally on the SMS, you can assign each user a different timeout period. If you are using an external database, you can assign one timeout period that pertains to all users.

It is important to understand that the session timeout is not an idle timeout. User traffic does not affect it. When a user times out, this only affects the user's ability to begin new sessions. Existing sessions are not affected, primarily because they continue to exist in the Brick's session cache.

The amount of time existing sessions remain in the session cache is determined by the session timeout value in the rule that allowed the user's session through the Brick. For additional information about the session timeout, see [Chapter 1, "Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets"](#)

Important! Although there is no requirement to leave the browser window open, for convenience a small Java applet is downloaded to the web browser. This applet is essentially a countdown timer that indicates the time before the user has to re-authenticate.



Before Setting Up User Authentication

Prerequisites

Before you set up user authentication, there are several actions you have to perform.

First, if you are using either RADIUS or SecurID to perform user authentication, you must have a rule in the same zone allowing the SMS to communicate with the RADIUS or SecurID ACE/Server. You also must add several entries to the *services* file on the SMS.

In addition, if you are using SecurID, you have to copy a file known as a *resource encrypted file* from the ACE/Server to a specific directory on the SMS.

Add a RADIUS or SecurID Rule

If you are using RADIUS or SecurID to authenticate users, and the RADIUS or SecurID ACE/Server is located on the other side of the Brick from the SMS, you have to add a rule to the NOC Gateway or Administrative Zone to permit the SMS to communicate with the server.

To create this rule, follow the steps below:

-
- 1 With the Navigator window displayed, open the appropriate Group and Policies folders, and click the Bricks Zone Rulesets folder to display all existing rulesets.
 - 2 Double-click **nocgwzone** or **administrativezone**. The Brick Zone Ruleset Editor will appear, with the ruleset you selected displayed.
 - 3 Right-click in the **Rules** panel and select **New** from the pop-up menu. The Brick Zone Rule Editor will appear.
 - 4 In the **Direction** field, select **Out Of Zone** from the drop-down list.
The direction is out of the zone because the SMS is sending an authentication request to the RADIUS or ACE/Server, which is external to the zone.
 - 5 In the **Source** field, enter the IP address of the SMS.
 - 6 In the **Destination** field, enter the IP address of the RADIUS or ACE/Server.

-
- 7 In the **Service or Group** field, display the drop-down list and select **Browse**. Use the Browse window that appears to select either **RADIUS1** or **RADIUS2** if you are using RADIUS, or *securId* if you are using SecurID.

The **RADIUS1** (UDP/1645-1646/*), **RADIUS2** (UDP/1812-1813/*), and *securID* (UDP/5500/*) service groups are provided with the SMS application specifically for this rule.

-
- 8 In the **Action** field, select **Pass** from the drop-down list.
 - 9 Click **OK** to dismiss the Brick Zone Rule Editor and return to the Brick Zone Ruleset Editor.

-
- 10 Display the File menu and select one of the **Save** options.

END OF STEPS

Copy the Resource Encrypted File

Note the following about the resource encrypted file:

-
- 1 The ACE/Server has a resource encrypted file that must be copied to a specific location on the SMS. This file contains information about the ACE/Server, such as the IP address and port of the server, and information to encrypt communications between the SMS and the ACE/Server.

This file is usually found on the ACE/Server at

<ACE_SRVR_ROOT>/data/sdconf.rec

and must be copied to the following directory on the SMS:

<SMS_ROOT>/groups/<group_name>/securID

If you have multiple ACE/Servers, the files can be renamed for each server. Multiple ACE/Servers are only supported on the *Windows*[®] platform, not *Solaris*[®].

END OF STEPS

Update the Services File

If you will be using an external SecurID or RADIUS server for user authentication, you must add a number of entries to the *services* file. The *services* file is an ASCII file that is found in the directory: c:\Winnt\system32\drivers\etc on Windows systems, and /etc on Solaris systems.

Complete the following steps to add the appropriate entries to the *services* file:

- 1 Using a standard text editor such as Notepad or vi, open the *services* file.

- 2 Add the services shown in the box below to the file.

securid	5500/UDP
securidprop	5510/TCP
radius	1812/UDP
radacct	1813/UDP
sdxauthd	5540/UDP
sdreport	5540/TCP
sdlog	5520/TCP
sdserv	5530/TCP

The port assignments for radius and radacct could also be 1645/UDP and 1646/UDP.

- 3 Save the changes, close the file, and close the text editor. Be sure to save the file under its original name in the same directory, with no extensions.

- 4 Reboot the computer.

END OF STEPS



To Set Up Local Password Authentication

When to use

The simplest and quickest way to set up user authentication is to create a database of users on the SMS itself. This method of authentication is referred to as **Local Password**, because the ID and password that the users enter to establish their credentials reside in the SMS database, and are verified by the SMS.

If you choose to set up **Local Password** authentication, you do not have to create an authentication service, one has been provided with the SMS application. All you have to do is create an account for each user who will be authenticated. If you choose, you can also create one or more user groups to help manage these accounts.

Once this has been done, you can begin creating rules to permit authenticated users to access resources protected by a Brick, or to enable client users to establish a tunnel to a Brick or router.

Strong password (SOX compliance) restrictions for local password authentication

When a new password is set for a user that is authenticated using Local Password authentication, or an existing local password is changed, if the Strong Password option is enabled via the Configuration Assistant, password restrictions that comply with Sarbanes-Oxley (SOX) requirements can be applied for local password creation and authentication. One or more password restrictions can be set via the Strong Passwords parameters screen of the Configuration Assistant so that a new or modified password:

- Must be a minimum of eight characters, or the **Minimum Password Length** set for the **Local Password** Authentication Service, whichever is greater
- Must contain at least one alpha character and one non-alpha character (0-9, special characters, no restrictions)
- Cannot contain three or more repeated alphanumeric characters in a row
- Cannot contain three or more consecutive, ascending or descending, alphanumeric characters in a row
- Not contain the User Account name or its mirror (reverse character format)
- Not be one of the previous three passwords most recently used
- Must contain at least one non-alpha character not in the first or last position

For more information about setting the strong password (SOX compliance) options via the Configuration Assistant, refer to the *Using the Configuration Assistant* chapter in the *SMS Administration Guide*.

To create a user account

To create a user account, you have to display the User Editor and enter the information requested. To do this, follow the steps below:

- 1 With the Navigator window displayed, open the appropriate Group and User Authentication folders.
- 2 Right-click the Users folder and select **New User** from the pop-up menu. The User Editor will appear. It is shown in [Figure 8-1, “User Editor”](#) (p. 8-12).

Figure 8-1 User Editor

The screenshot shows a window titled "User Editor - /system/User Authentication/Users/". The window has a menu bar with "File", "Edit", "Monitor", "Windows", "Utilities", and "Help". The main area contains several form fields and a checkbox:

- User ID**: A text input field.
- Enable User**: A checked checkbox.
- Full Name**: A text input field.
- Description**: A large text input field.
- Email**: A text input field.
- Telephone**: A text input field.
- Page Info**: A text input field.
- Authentication Service**: A dropdown menu.
- Authentication Timeout (mins)**: A text input field with the value "60".
- Disable User After**: A text input field with the value "0".
- Failed Logins**: A section with two text input fields. The first is "Active From" with the value "06-16-2000 11:15:37". The second is "to" with the value "06-16-2001 11:15:37".
- Allowed Source IP Addresses**: A text input field with a small icon on the left.

- 3 In the **User ID** field, enter an ID to identify this user. The ID can contain up to 40 characters (letters and numbers). The ID must be unique within this group. This is the ID the user will have to enter to log in and be authenticated.

-
- 4 In the **Full Name** and **Description** fields, enter the user's complete name and a description, such as the user's title, function, or organization. Both of these fields are optional.

The description is useful because it appears on the Navigator window when you click the Users folder to display existing users. It can help you identify specific users by function, department, or user group.

- 5 In the **E-mail**, **Telephone**, and **Pager Info** fields, enter the user's e-mail address, telephone number and pager number. These fields are also optional.
-

- 6 In the **Authentication Service** field, select **Local Password** from the drop-down list. The authentication service determines how the user will be authenticated.

As soon as you enter the authentication service, the **Additional Options for Local Passwords** panel will appear at the bottom of the User Editor. This panel is shown in [Figure 8-2, "Password Panel" \(p. 8-13\)](#).

Figure 8-2 Password Panel

Additional Options For Local Passwords

Password
Verify Password
 Require User To Change Password On Next Login
Type One-Time Continuous
Expiration Every

-
- 7 In the **Authentication Timeout** field, specify the timeout period (in minutes). This is the length of time after authentication that the user will be permitted to begin new sessions through the Brick. Existing sessions will not be affected.

The default is 60 minutes, but you can change this to any number from 1 to 2880 (48 hours).

Important! For client users, the timeout period is the smaller of two values: the **Authentication Timeout** and the **VPN ISAKMP Proposal Security Association Lifetime**.

The **SA Lifetime** is set in the VPN Client Defaults Editor when you create the client tunnel. This is explained in [“Overview” \(p. 11-1\)](#).

- 8** In the **Disable User After** field, specify the number of consecutive failed logins permitted before the user is disabled. The default is 0, which means that failed logins will be ignored.

Important! If a user is disabled, you have to enable the user before that user can log in. This is done by clicking the **Enable User** box at the top of the User Editor.

- 9** In the **Active From... To** field, indicate how long this account will be active. The default is one year from the current date.
-

- 10** In the **Allowed Source IP Addresses** field, enter the appropriate IP addresses. Only users coming from these IP addresses will be able to be authenticated.

There are several ways to enter this information:

- Leave the default asterisk in place. This means the user can connect from any IP address.
 - Enter a specific IP address, or several addresses separated by commas.
 - Enter a range of IP addresses, using the format 10.1.1.1-10.1.1.10.
 - Enter an IP address and subnet mask, for example, 10.10.10.10/24.
-

- 11** Enter password information for this user in the **Additional Options for Local Passwords** panel (see [Figure 8-2, “Password Panel” \(p. 8-13\)](#)).

Do the following:

1. In the **Password** and **Verify Password** fields, enter this user’s password. These fields are case-sensitive, so make sure the capitalization is consistent.
2. If you want the user to change the password on the next login, click the **Require User To Change Password On Next Login** checkbox. The next time this user logs in, the user will be prompted to choose a different password.
3. In the **Type** field, specify the type of password. The following are the options:
 - **One-time** = when the password expires, the user is disabled
 - **Continuous** = when the password expires, the user is prompted to enter another one. This is the default.
4. In the **Expiration Every** field, indicate how often this password will expire. Once a year is the default.

To change the default, do the following:

- Enter the *frequency* in the left field (1 is the default) and
- Enter the *time period* in the right field. It can be logins, days, months or years (years is the default).

If the frequency is set to 0, the password will never expire, regardless of the time period.

-
- 12** Display the File menu and select one of the **Save** options.

END OF STEPS

To copy a user account

When you edit a user account (using the **Save As** option to save the edited version with a new name) or copy a user account, you have the option of copying components from the original user account, specifying how those components will be named and where they will be created and stored.

Complete the following steps to copy a user account:

-
- 1** With the user accounts displayed in the Navigator window, right-click the user account to be copied and select **Copy** from the pop-up menu.

Result A Copy window is displayed ([Figure 8-3, “Copy Window” \(p. 8-15\)](#)).

Figure 8-3 Copy Window



2 Complete the following fields, as applicable:

- **Name**— Enter a new name for the user account copy (Example: **Jean Smith**). This field is required. The name entered must be different from the name of the original user account by at least one character if it is in the same group.
- **Folder**— The copy of the user account will be created and stored in the **Users** folder under the *system* group, by default. To specify a different folder, click the **Browse** button, which brings up a Select a User Folder window, and select the folder where the user account copy will be created and stored.
- **Copy components when copying within a Group**— Click this checkbox (place a check in it) to copy the components from the original user account. If this checkbox is unchecked, and the copied user account is in the same group as the original one, the copied user account uses the same components as the original. If the user account is copied across groups, the components are copied regardless of this field setting.
- **Component Suffix**— If the **Copy components when copying within a Group** option is checked, enter a text string that will be appended to the name of each copied component. If this field is left blank, the name of each copied component will be spelled the same way as the original one. Name collisions are possible, regardless of whether or not an entry is made in this field. If this occurs, the content of the conflicting component(s) is examined. If the content is the same, the duplicated user account uses the existing component(s) and the new copy(ies) is not made. If the content is different, a copy of the component(s) is created and stored under a unique name by appending its database ID to the name.
- **Component Folder**— This field is used to specify a common folder name for the copied components. The folder name is relative to the top-level folder of each type of component. If this field is left blank, the copied components are stored in the top-level folder of that component. This folder is created as needed (if it does not already exist). A path string can be specified; / or \ delimiters can be used to delineate folder names.

3 Click **OK** to copy the user account and dismiss the Copy window.

Result The system returns to the Navigator window. The copy of the user account appears in the Contents panel.

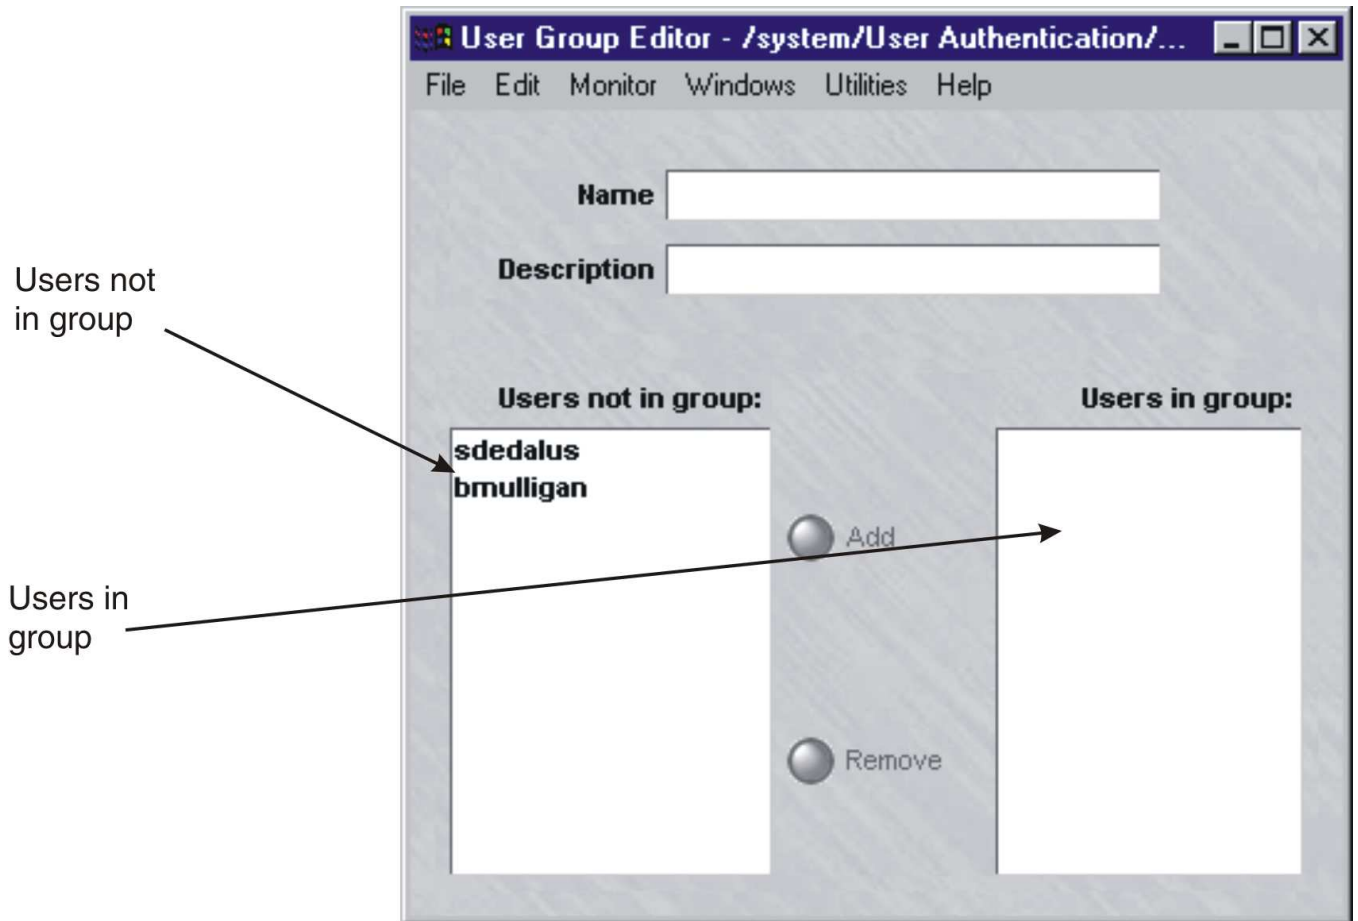
END OF STEPS

To create a user group

To create a user group, you have to display the User Group Editor and indicate the user accounts that will be included in the user group. To do this, follow these steps:

- 1 With the Navigator window displayed, and the appropriate Group and User Authentication folders open, right-click the User Groups folder and select **New User Group** from the pop-up menu. The User Group Editor will appear. It is shown in Figure 8-4, “User Group Editor” (p. 8-17).

Figure 8-4 User Group Editor



- 2 In the **Name** field, enter a unique name to identify this user group. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers, and certain special characters.

-
- 3 In the **Description** field, you can enter an optional description of the user group. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers, and certain special characters.

The description is useful because it appears on the Navigator window when you click the User Groups folder to display existing user groups. It can help you identify user groups by function, department, or other easily recognized characteristic.

- 4 To add a user to the group, select the user in the **Users not in group** panel and click **Add**. This moves the user to the **Users in group** panel. Repeat for each user you are adding to this group.
-

- 5 When you have finished populating the user group, display the File menu and select one of the **Save** options.
-

END OF STEPS

To copy a user group


When you edit a user group (using the **Save As** option to save the edited version with a new name) or copy a user group, you have the option of copying components from the original user group, specifying how those components will be named and where they will be created and stored.

Complete the following steps to copy a user group:

- 1 With the user groups displayed in the Navigator window, right-click the user group to be copied and select **Copy** from the pop-up menu.

Result A Copy window is displayed ([Figure 8-5, “Copy Window”](#) (p. 8-18)).

Figure 8-5 Copy Window



-
- 2 Complete the following fields, as applicable:
- **Name**— Enter a new name for the user group copy (Example: **salesgroup_2**). This field is required. The name entered must be different from the name of the original user group by at least one character if it is in the same group.
 - **Folder**— The copy of the user group will be created and stored in the **User Groups** folder under the *system* group, by default. To specify a different folder, click the **Browse** button, which brings up a Select a User Group Folder window, and select the folder where the user group copy will be created and stored.
 - **Copy components when copying within a Group**— Click this checkbox (place a check in it) to copy the components from the original user group. If this checkbox is unchecked, and the copied user group is in the same group as the original one, the copied user group uses the same components as the original. If the user group is copied across groups, the components are copied regardless of this field setting.
 - **Component Suffix**— If the **Copy components when copying within a Group** option is checked, enter a text string that will be appended to the name of each copied component. If this field is left blank, the name of each copied component will be spelled the same way as the original one. Name collisions are possible, regardless of whether or not an entry is made in this field. If this occurs, the content of the conflicting component(s) is examined. If the content is the same, the duplicated user group uses the existing component(s) and the new copy(ies) is not made. If the content is different, a copy of the component(s) is created and stored under a unique name by appending its database ID to the name.
 - **Component Folder**— This field is used to specify a common folder name for the copied components. The folder name is relative to the top-level folder of each type of component. If this field is left blank, the copied components are stored in the top-level folder of that component. This folder is created as needed (if it does not already exist). A path string can be specified; / or \ delimiters can be used to delineate folder names.

-
- 3 Click **OK** to copy the user group and dismiss the Copy window.

Result The system returns to the Navigator window. The copy of the user group appears in the Contents panel.

END OF STEPS

SMS-defined user groups

When you create the rule to authenticate the users whose accounts you just created, you identify these accounts by means of a user group. The SMS automatically creates four user groups, which are described in [Table 8-1, “SMS-defined user groups”](#) (p. 8-20). These user groups are automatically populated with each user account you

create. You can see these user groups when you create the rule. They will appear in the drop-down list in the **Source** and **Destination** fields after you click the **User** radio button. See “[Overview](#)” (p. 1-1).

Table 8-1 SMS-defined user groups

User group	Description
All_Users	Contains the IP addresses of any VPN Client user or Firewall Authentication user that has been successfully been authenticated by connecting to the zone TEP/VBA.
Active_VPN_Users	Contains the IP addresses of any VPN Client user that has successfully been authenticated by connecting to the zone TEP/VBA. If a “local presence” IP address has been assigned to the client, this is the address that is added to this user group.
Active_VPN_UserTEPs	Similar to Active_VPN_Users except the client’s actual IP address is added to the group rather than the “local presence” IP address.
Default_Auth_Users	<p>If all of your users will be authenticated by means of a Local Password, you can consider using this user group as the source or destination in the security rules you create.</p> <p>However, if you are using different methods of authentication for different groups of users, or if you have a large body of users spanning different geographic locations, organizations or roles, you may want to consider creating your own user groups. Assigning users to specific groups will make it easier for you to manage these users.</p>



To Set Up RADIUS and SecurID Authentication

When to use

If you will be using RADIUS or SecurID to authenticate application or client users, you first have to create a RADIUS or SecurID authentication service. Then, you have to decide whether your user database will reside on the SMS host as well as the external servers, or only on a RADIUS server or SecurID ACE/Server.

You may choose to have some users configured on the SMS, and the remainder on a given external server.

- **SMS Host**

If the database will reside on the SMS host, you have to create an account for each user, and you have the option of creating user groups to help manage these users. The procedure for creating a user account is the same as described for local password authentication (“[When to use](#)” (p. 8-11)). Be sure to enter the RADIUS or SecurID authentication service you just created in the **Authentication Service** field.

The procedure for creating a user group is the same as described for local password authentication (“[To create a user group](#)” (p. 8-17)).

- **RADIUS Host or ACE/Server**

If the database will reside on the RADIUS host or ACE/Server, you do not have to create user accounts or user groups. You can use the SMS-created **All_Users** group in the authentication rule.

However, you do have to enter the authentication service you just created when configuring the port on the Brick, or when configuring any client tunnel endpoints.

Important! Each user can have his or her own authentication service if he or she is configured in the SMS database. However, to avoid replicating your entire external database on the SMS, you can choose default authentication services for each zone for both application and client user authentication. If the user is not in the internal database, the SMS will try the default authentication service for that zone.

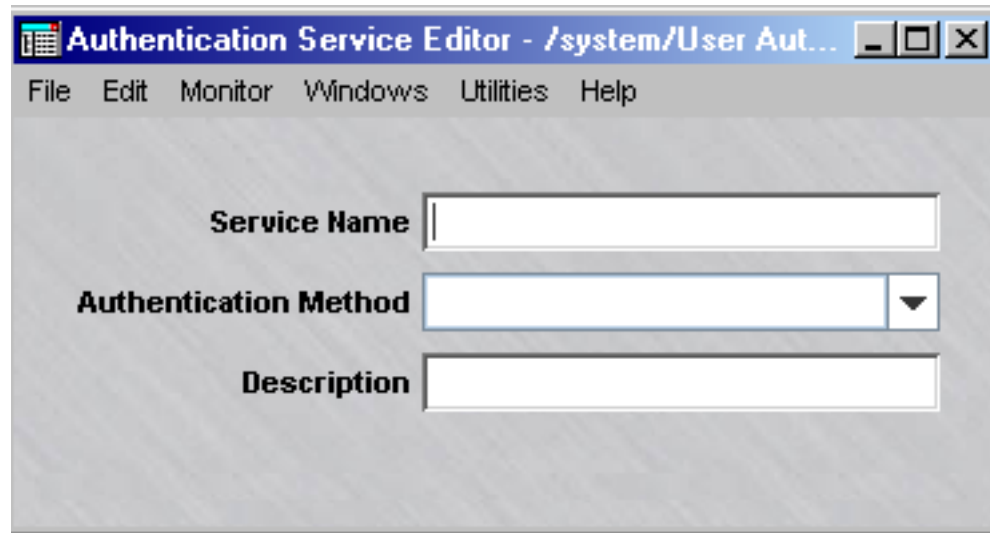
To create a RADIUS Authentication Service

To create a RADIUS authentication service, you have to display the Authentication Service Editor and enter the information requested. To do this, follow the steps below:

-
- 1 With the Navigator window displayed, and the appropriate Group and User Authentication folders open, right-click the Authentication Services folder and select **New Auth Service** from the pop-up menu.

Result The Authentication Service Editor is displayed (Figure 8-6, “Authentication Service Editor” (p. 8-22)).

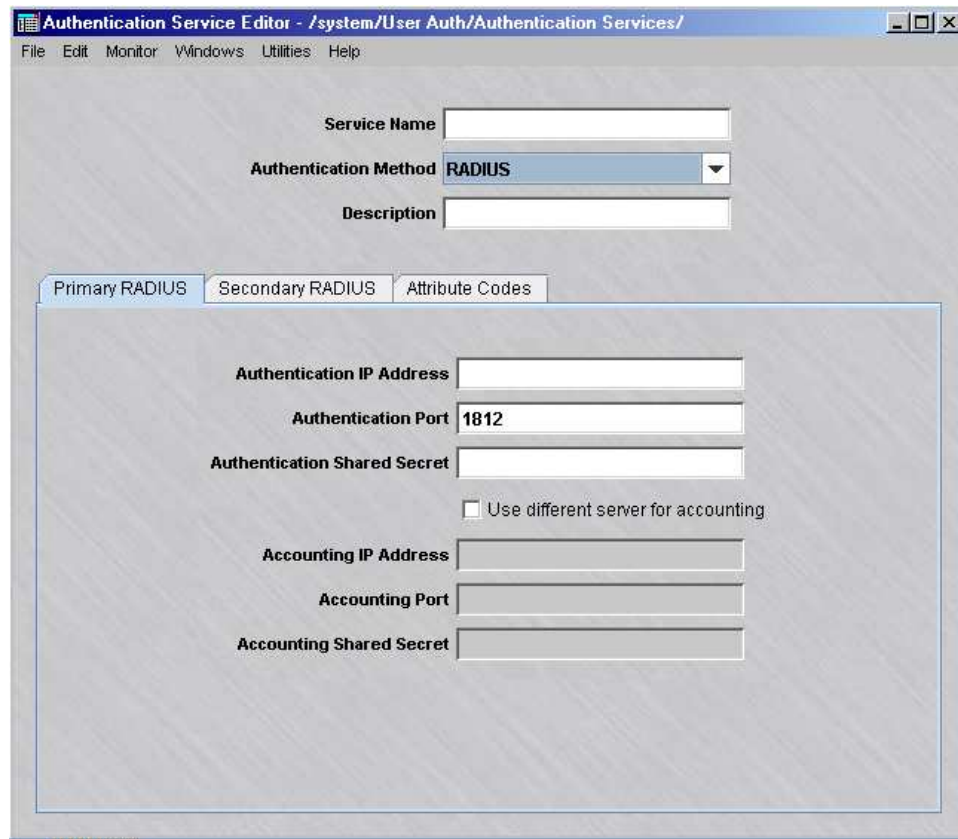
Figure 8-6 Authentication Service Editor



- 2 In the **Service Name** field, enter a unique name to identify this authentication service. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers, and certain special characters.
- 3 In the **Description** field, you can enter an optional description of this authentication service. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers, and certain special characters.
The description is useful because it appears on the Navigator window when you click the Authentication Services folder to display existing authentication services. It can help you identify specific services.
- 4 In the **Authentication Method** field, select **RADIUS** from the drop-down list. The Authentication Service Editor will expand to include three tabs (**Primary RADIUS**, **Secondary RADIUS**, and **Attribute Codes**).

Result The Primary RADIUS tab is displayed **RADIUS Options** boxes, is the tab that is displayed first ([Figure 8-7, “Authentication Service Editor \(Primary RADIUS Tab\)”](#) (p. 8-23)).

Figure 8-7 Authentication Service Editor (Primary RADIUS Tab)



- 5 Enter the following information about the Primary RADIUS server:

Field	Explanation
Authentication IP Address	The IP address of the RADIUS authentication server
Authentication Port	The number of the port on which the RADIUS authentication service is listening. The default is 1812, but you can change this if necessary.

Field	Explanation
Authentication Shared Secret	The Authentication Shared Secret is the string used to encrypt communication between the SMS and the RADIUS authentication server. The same value must be entered on both servers. The value is case-sensitive.
Accounting Shared Secret	The Accounting Shared Secret is the string used to encrypt communication between the SMS and the RADIUS accounting server. The same value must be entered on both servers. The value is case-sensitive.
Use different server for accounting	Check this box if you need to specify a different IP address, port, or shared secret for the accounting server
Accounting IP Address	The IP address of the RADIUS accounting server
Accounting Port	The number of the port on which the RADIUS accounting service is listening. The default is 1813, but you can change this if necessary.

-
- 6 If your environment includes a secondary RADIUS server, you can enter the same information for this server in the **Secondary RADIUS** tab. If the primary RADIUS server cannot be reached, the SMS will try to reach the secondary RADIUS server.
-
- 7 If you want the SMS to use certain information contained in the RADIUS protocol messages that the RADIUS server sends to the SMS when authenticating users, click **Attribute Codes** to display the Attribute Codes tab ([Figure 8-8, “Authentication Service Editor \(Attribute Codes Tab\)”](#) (p. 8-25)).

Figure 8-8 Authentication Service Editor (Attribute Codes Tab)

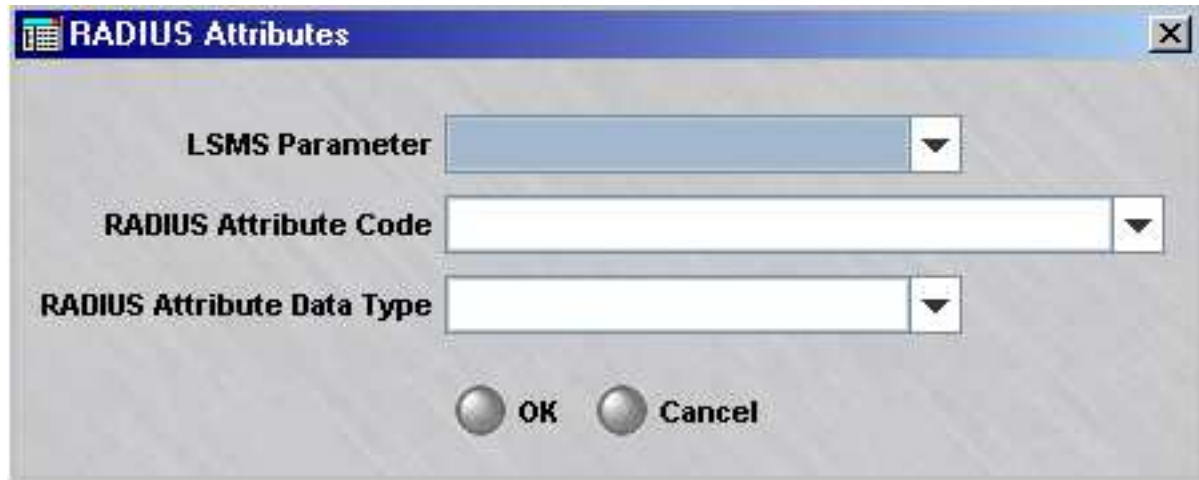
The screenshot shows the 'Authentication Service Editor' window with the 'Attribute Codes' tab selected. The window title is 'Authentication Service Editor - /system/User Auth/Authentication Services/'. The menu bar includes 'File', 'Edit', 'Monitor', 'Windows', 'Utilities', and 'Help'. The main area contains three input fields: 'Service Name', 'Authentication Method' (set to 'RADIUS'), and 'Description'. Below these are three tabs: 'Primary RADIUS', 'Secondary RADIUS', and 'Attribute Codes'. The 'Attribute Codes' tab contains a table with the following columns: 'LSMS Parameter', 'RADIUS Attribute Code', 'RADIUS Attribute Description', and 'RADIUS Attribute Data Type'. The table is currently empty. At the bottom, there is a 'Log Level' section with three radio buttons: 'Do not Log Attributes' (selected), 'Log Configured Attributes', and 'Log All Attributes'.

LSMS Parameter	RADIUS Attribute Code	RADIUS Attribute Description	RADIUS Attribute Data Type
----------------	-----------------------	------------------------------	----------------------------

-
- 8 Right-click in the Attribute Codes panel and select **NEW** from the pop-up menu.

Result The RADIUS Attributes window is displayed (Figure 8-9, “RADIUS Attributes Window” (p. 8-26)).

Figure 8-9 RADIUS Attributes Window



- 9 To configure the SMS to use an attribute returned by RADIUS, do the following:

In the **LSMS Parameter** field, select the SMS parameter that corresponds to the RADIUS attribute code you want to configure on the RADIUS server. Then, select the appropriate **RADIUS Attribute Code** and **RADIUS Attribute Data Type**.

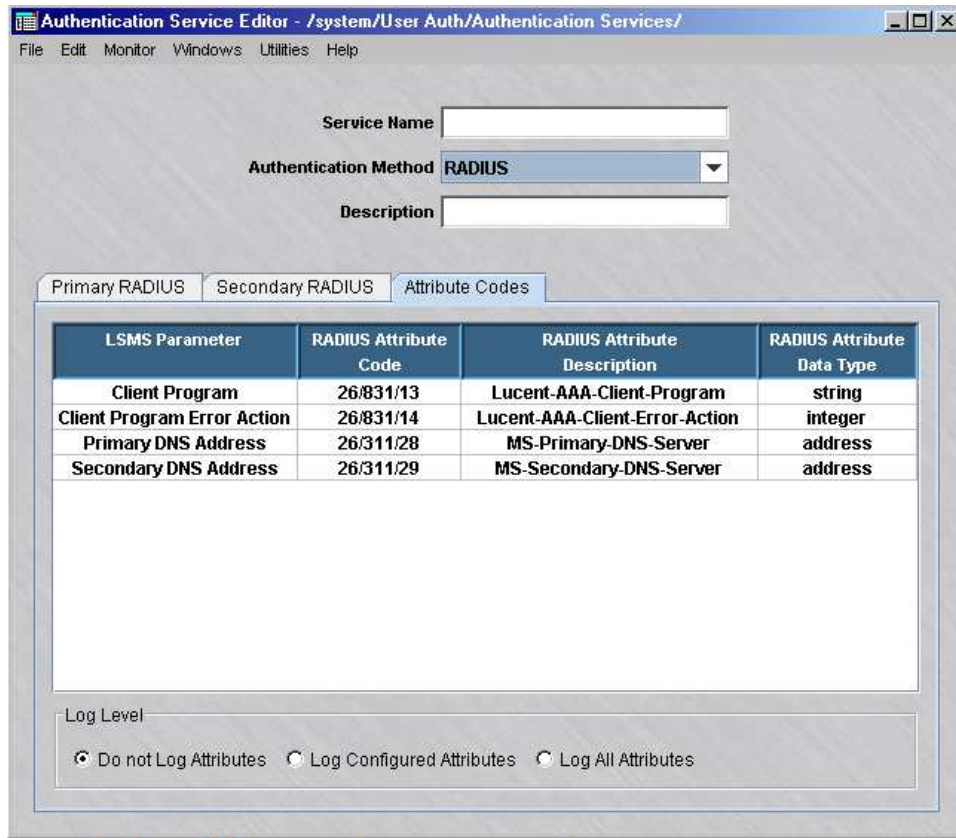
If the RADIUS attribute that you want to assign is not provided in the drop-down list, you may enter an attribute code in the range 0-255. Attribute 26 is defined as the Vendor Specific Attribute (VSA), and should be entered in the format `26/<vendorID>//<vendorType>`

Result The selected SMS parameter, RADIUS attribute code and RADIUS data type is displayed on the Attribute Codes tab.

To select additional SMS parameters and RADIUS response attributes, right-click in the Attribute Codes panel and select **New**. The RADIUS Attributes window is displayed (Figure 8-9, “RADIUS Attributes Window” (p. 8-26)). Add as many RADIUS response attributes as needed.

Figure 8-10, “Attribute Codes tab (Configured SMS Parameters With RADIUS Attributes)” (p. 8-27) shows a sample of the Attributes Code tab with some configured SMS parameters. For each SMS parameter, the tab panel shows the RADIUS attribute assigned, a description of the RADIUS attribute, and the RADIUS code type.

Figure 8-10 Attribute Codes tab (Configured SMS Parameters With RADIUS Attributes)



[Appendix D, “RADIUS Attributes”](#) provides a description of SMS parameters and the RADIUS attributes that can be assigned for user authentication.

- 10 To debug the RADIUS interface, select one of the following Log Level options:

To	Select
Turn off logging of attributes (this option should be chosen for normal operation; only use one of the logging options for debugging the RADIUS interface)	Do not Log Attributes
Log the configured attributes that are returned from RADIUS in the User Authentication log	Log Configured Attributes
Log all response attributes from RADIUS in the User Authentication log	Log All Attributes

-
- 11 Display the **File** menu and select one of the **Save** options. The service you just created will now appear in the **Contents** panel of the Navigator window

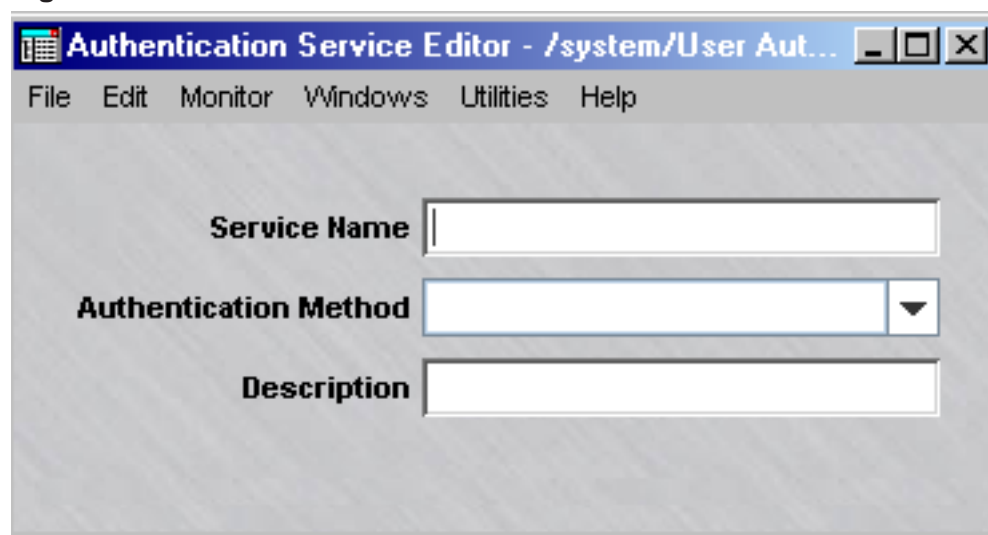
END OF STEPS

To create a SecurID authentication service

To create a SecurID authentication service, you have to display the Authentication Service Editor and enter the information requested. To do this, follow the steps below:

-
- 1 With the Navigator window displayed, and the appropriate Group and User Authentication folders open, right-click the Authentication Services folder and select **New Auth Service** from the pop-up menu. The Authentication Service Editor will appear. It is shown in [Figure 8-11, “Authentication Service Editor”](#) (p. 8-28).

Figure 8-11 Authentication Service Editor

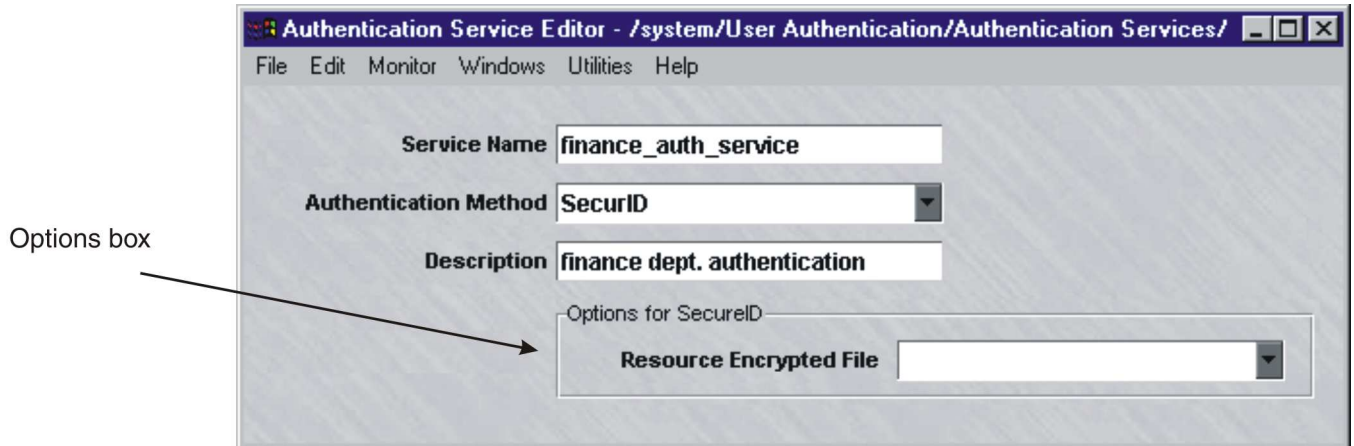


-
- 2 In the **Service Name** field, enter a unique name to identify this authentication service. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers, and certain special characters.
 - 3 In the **Description** field, you can enter an optional description of this authentication service. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers, and certain special characters.

The description is useful because it appears on the Navigator window when you click the Authentication Services folder to display existing authentication services. It can help you identify specific services.

- 4 In the **Authentication Method** field, select **SecurID** from the drop-down list. An **Options for SecurID** box will appear in the Authentication Service Editor. This box is shown in [Figure 8-12, “SecurID Options”](#) (p. 8-29).

Figure 8-12 SecurID Options



- 5 In the **Resource Encrypted File** field, select the resource encrypted file from the drop-down list.

This file is generated on the ACE/Server. It contains information about the ACE/Server, such as the IP address and port of the server, and information to encrypt communications between the SMS and the ACE/Server. This file is usually found on the ACE/Server at

```
<ACE_SRVR_ROOT>/data/sdconf.rec
```

and must be copied to the following directory on the SMS:

```
<LSMS_ROOT>/group/<group_name>/securID
```

After it is copied here, it will be visible in the drop-down list. If you have multiple ACE/Servers, the files can be renamed for each server. Multiple ACE/Servers are only supported on the Windows platform, not Solaris.

- 6 Display the File menu and select one of the **Save** options. The service you just created will now appear in the Navigator window.

END OF STEPS

Enter the Authentication Service

If the user accounts reside on the SMS host even though these users are externally authenticated, the only time you have to enter the authentication service is when you create the user accounts. However, if the user database resides solely on a RADIUS host or ACE/Server, you have to associate the authentication service with the appropriate Brick port. This can be done when assigning a Brick zone ruleset or router main ruleset to the interface, or when setting up a client tunnel endpoint. The following explains.

1 *When configuring the Brick port*

1. Open the appropriate Group and Devices folders, right-click the Bricks folder, and select **New Brick** from the pop-up menu. If the Brick has already been configured, double-click it instead. The Brick Editor is displayed.
2. If you are configuring a Brick port, click **Policy Assignment** to display the Policy Assignment tab. If you are configuring a router **Interfaces**, click Interfaces to display the Interfaces tab
3. Double-click an existing port, or right-click and select **New** from the pop-up menu. The Brick Policy Assignment Editor or Router Interface Editor will appear.
4. In the **Authentication Service** field, display the drop-down list and select the authentication service.
5. In the **Authentication Timeout** field, enter the timeout period in minutes. The default is 480 minutes (eight hours).
The timeout can be from one minute to 5 years (2628000 minutes). The timeout period will apply to all users authenticated by this authentication service.
6. In the **Allowed Source IP or Range** field, enter the appropriate IP addresses. Only users coming from these IP addresses will be able to be authenticated.
There are several ways to enter this information:
 - Leave the default asterisk in place. This means users can connect from any IP address.
 - Enter a specific IP address, or several addresses separated by commas.
 - Enter a range of IP addresses, using the format 10.1.1.1-10.1.1.10.
 - Enter an IP address and subnet mask, for example, 10.10.10.10/24.

2 *When setting up the client tunnel endpoint*

1. Open the appropriate Group and VPN folders, right-click the Client Tunnel Endpoints folder, and select **New Client Tunnel Endpoint** from the pop-up menu. If the client tunnel endpoint has already been configured, double-click it instead.
2. In the **Authentication Service** field, display the drop-down list and select the authentication service.
3. In the **Authentication Timeout** field, enter the timeout period in minutes. The default is 480 minutes (eight hours).
The timeout can be from one minute to 5 years (2628000 minutes). The timeout period will apply to all users authenticated by this authentication service.

Important! If you enter the authentication service when assigning a zone (ruleset) to a port, the authentication service will only apply to that zone on that particular port.

If the zone is assigned to another port, the authentication service will not automatically apply to traffic on that port. Similarly, if another zone is assigned to that same port, the authentication service will not apply to that zone.

If you enter the authentication service when setting up a client tunnel endpoint, the service will only apply to the zone whose VBA is the client tunnel endpoint.

END OF STEPS



To Set Up VPN Certificate Authentication

When to use

You can use a digital certificate to authenticate VPN client users. Digital certificate authentication is not available for application users (the Firewall authentication feature). Digital certificate authentication requires a Certificate Authority (CA) certificate, Certificate Revocation List (CRL), and a VPN certificate to be installed on the SMS host. The procedures for creating Certificate Signing Requests (CSRs) and importing certificates are described in [Chapter 9, “Digital Certificates”](#).

Digital certificate authentication for VPN client users can be configured to use either the VPN Certificate authentication service or the RADIUS authentication service. For both options, the Brick will verify the validity of the client certificate against the Certificate Authority (CA) certificate and the Certificate Revocation List (CRL). RADIUS authentication can only be used with IKEv2 clients.

When the VPN Certificate authentication service is used, the Brick sends the client certificate Subject Name attributes to the SMS, and the SMS verifies that client certificate attributes match the required attributes specified in the authentication service.

When the RADIUS authentication service is used, the SMS forwards the authentication request to RADIUS, using the certificate email address attribute as the UserID and the client certificate Subject Name as the password. Then, RADIUS, checks the validity of the userID and password.

For regular RADIUS user authentications, the SMS sends Authenticate Only (8) in the Service=Type attribute of RADIUS authentication and accounting messages. To differentiate digital certificate authentication, the SMS sends Authorize Only (17) in the Service=Type attribute to inform RADIUS that the Brick has validated the client certificate.

Decide which authentication service you want to use to authenticate the VPN clients, and assign it using the Client Tunnel Endpoint Editor. The procedure for creating a RADIUS authentication service is described in the section [“To Set Up RADIUS and SecurID Authentication”](#) (p. 8-21).

If a RADIUS authentication service is used, the following additional items must also be configured on the Client Tunnel Endpoint Editor:

- On the Endpoint tab, IKEv2 must be enabled
- On the IKEv2 Gateway tab, X.509 Certificate must be selected
- On the Remote Client ID tab, check the box to Allow IKEv2 clients to authenticate with X.509 Certificates and select Email Address for the User ID.

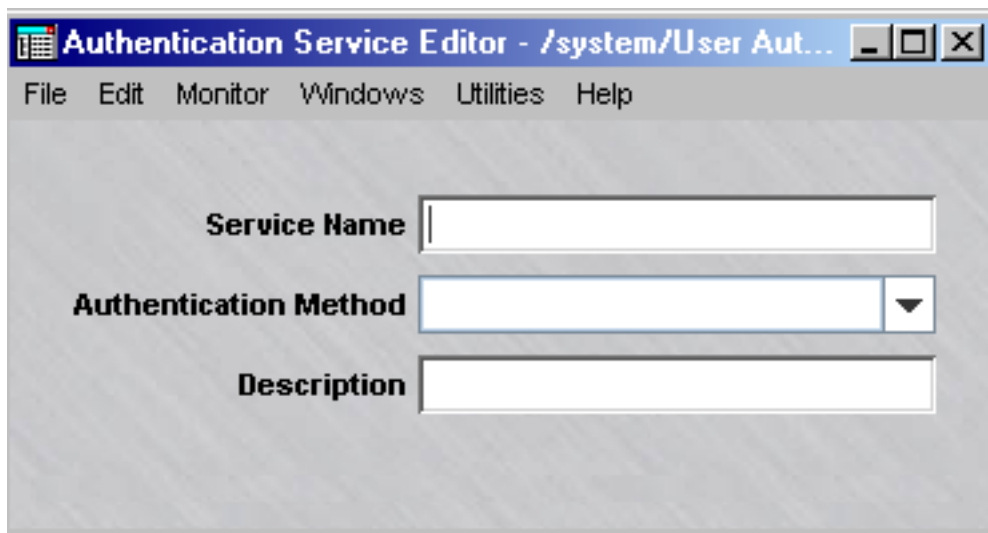
The procedure for creating a VPN Certificate authentication service is described below.

Task

Complete the following steps to set up a VPN Certificate Authentication service:

- 1 With the Navigator window displayed, and the appropriate Group and User Authentication folders open, right-click the Authentication Services folder and select **New Auth Service** from the pop-up menu. The Authentication Service Editor will appear. It is shown in “[When to use](#)” (p. 8-11).

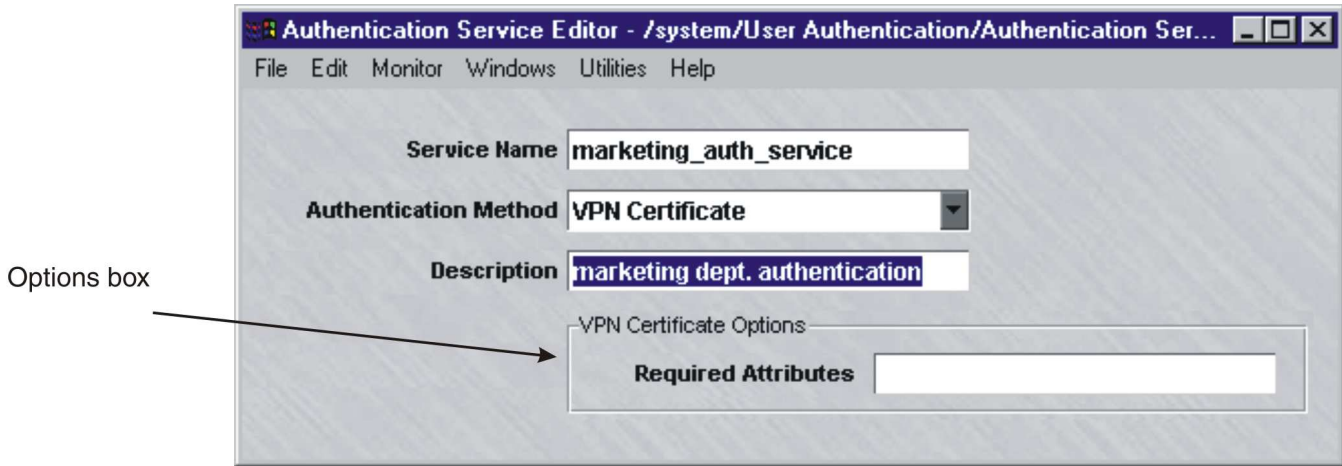
Figure 8-13 Authentication Service Editor



- 2 In the **Service Name** field, enter a unique name to identify this authentication service. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers, and certain special characters.
- 3 In the **Description** field, you can enter an optional description of this authentication service. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers, and certain special characters.

The description is useful because it appears on the Navigator window when you click the Authentication Services folder to display existing authentication services. It can help you identify specific services.
- 4 In the **Authentication Method** field, select **VPN Certificate** from the drop-down list. A **VPN Certificate Options** box will appear in the Authentication Service Editor, as shown in [Figure 8-14, “VPN Certificate Options”](#) (p. 8-34).

Figure 8-14 VPN Certificate Options




- 5 In the **Required Attributes** field, enter the same attributes found in the user’s X.509 digital certificate. If any of the required attributes are not present in the user’s certificate, the user will be denied access to the system.

This is a free-form field and you can enter any attributes specific to your Public Key Infrastructure (PKI). Separate individual attributes with a comma. The format of each attribute is shown below:

attribute=value, attribute=value, ...

Typical attributes are *O* (the user’s organization), *OU* (the user’s organizational unit), and *C* (the user’s country). [Figure 8-15, “Attributes” \(p. 8-34\)](#) shows an example of attributes entered in the **Required Attributes** field.

Figure 8-15 Attributes



- 6 Display the File menu and select one of the **Save** options. The service you just created will now appear in the **Contents** panel of the Navigator window.

END OF STEPS

To copy an authentication service

When you edit an authentication service (using the **Save As** option to save the edited version with a new name) or copy an authentication service, you have the option of copying components from the original authentication service, specifying how those components will be named and where they will be created and stored.

Complete the following steps to copy an authentication service:

- 1 With the authentication services displayed in the Navigator window, right-click the authentication service to be copied and select **Copy** from the pop-up menu.

Result A Copy window is displayed (Figure 8-16, “Copy Window” (p. 8-35)).

Figure 8-16 Copy Window



- 2 Complete the following fields, as applicable:
 - **Name**— Enter a new name for the authentication service copy (Example: **RADIUS_local**). This field is required. The name entered must be different from the name of the original authentication service by at least one character if it is in the same group.
 - **Folder**— The copy of the authentication service will be created and stored in the **Authentication Services** folder under the *system* group, by default. To specify a different folder, click the **Browse** button, which brings up a Select a Auth Service Folder window, and select the folder where the authentication service copy will be created and stored.
 - **Copy components when copying within a Group**— Click this checkbox (place a check in it) to copy the components from the original authentication service. If this checkbox is unchecked, and the copied authentication service is in the same group as the original one, the copied authentication service uses the same components as the original. If the authentication service is copied across groups, the components are copied regardless of this field setting.

- **Component Suffix**— If the **Copy components when copying within a Group** option is checked, enter a text string that will be appended to the name of each copied component. If this field is left blank, the name of each copied component will be spelled the same way as the original one. Name collisions are possible, regardless of whether or not an entry is made in this field. If this occurs, the content of the conflicting component(s) is examined. If the content is the same, the duplicated authentication service uses the existing component(s) and the new copy(ies) is not made. If the content is different, a copy of the component(s) is created and stored under a unique name by appending its database ID to the name.
 - **Component Folder**— This field is used to specify a common folder name for the copied components. The folder name is relative to the top-level folder of each type of component. If this field is left blank, the copied components are stored in the top-level folder of that component. This folder is created as needed (if it does not already exist). A path string can be specified; / or \ delimiters can be used to delineate folder names.
-

- 3 Click **OK** to copy the authentication service and dismiss the Copy window.

Result The system returns to the Navigator window. The copy of the authentication service appears in the Contents panel.

.....
E N D O F S T E P S
.....



9 Digital Certificates

Overview

Purpose

This chapter explains how to obtain and install digital certificates. Digital certificates are used:

- To configure a secure (HTTPS) SMS web server
- To configure a secure (HTTPS) User Authentication web server
- To authenticate Client and LAN-LAN VPN Tunnels

Contents

What is a Digital Certificate?	9-2
What is the Certificate Manager?	9-3
To Start the Certificate Manager	9-5
To Obtain and Install a Server Certificate	9-7
To Obtain and Import a VPN Certificate	9-13
Group Assignment for VPN Certificates	9-21
Pending Certificate Signing Requests	9-25
Server Versus VPN Certificate Signing Request	9-26



What is a Digital Certificate?

Definition

A digital certificate is a means of establishing a user's credentials. It is issued by a Certificate Authority (CA), and it contains information about the user, such as:

- The user's name, company, and location.
- The user's public key, which is used for encryption.
- The digital signature of the CA that is used to verify that the certificate is real.

Digital certificates can be obtained from Certificate Authorities such as Entrust, Verisign, and Thawte.

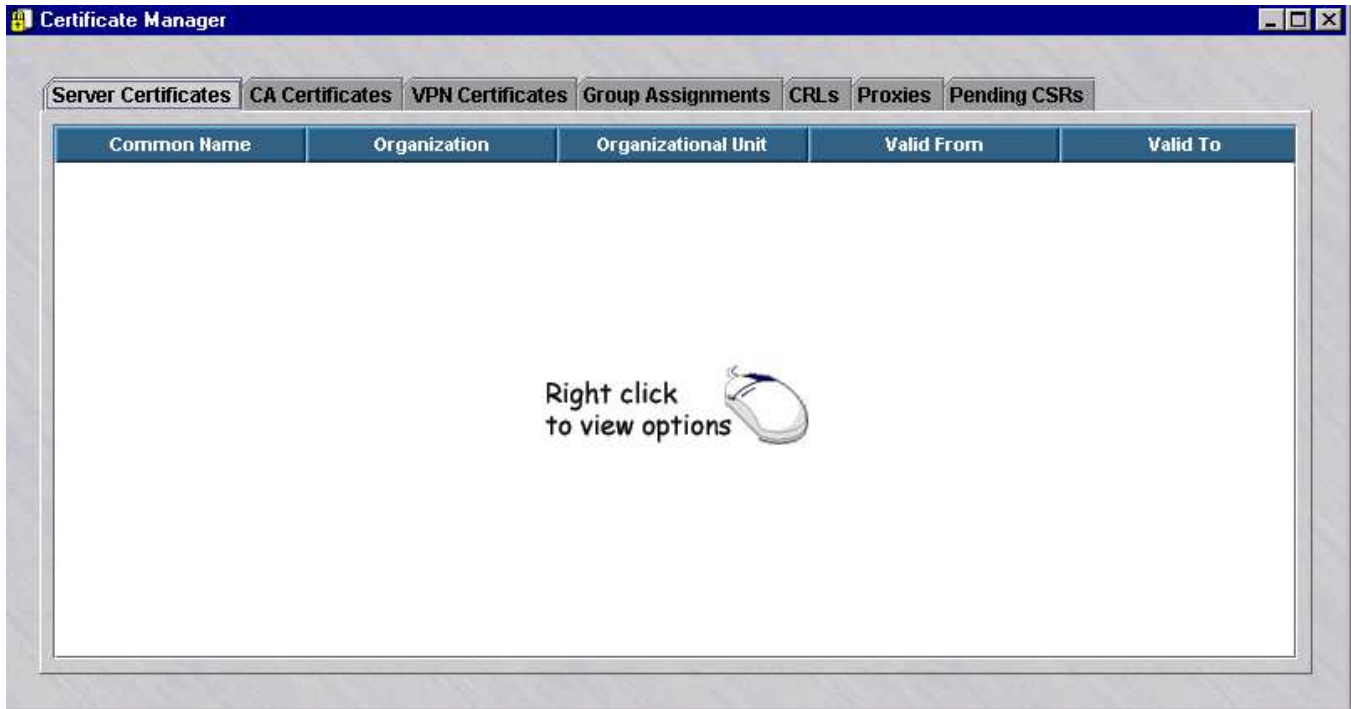


What is the Certificate Manager?

Definition

Certificate Manager is the name given to the SMS utility that is used to obtain and import digital certificates. An example is shown in [Figure 9-1, “Certificate Manager”](#) (p. 9-3)

Figure 9-1 Certificate Manager



The Certificate Manager contains seven tabs. These tabs are explained below.

Tab	Explanation
Server Certificates	Click this tab to manage certificates for the SMS web server and the User Authentication Web server when they are configured for HTTPS.
CA Certificates	Click this tab to manage the Certificate Authority certificates. You must install the CA certificate(s) of the Certificate Authorities that issue your Server and VPN certificates.
VPN Certificates	Click this tab to manage certificates that are used for VPN Client and LAN-LAN tunnels.
Group Assignments	Click this tab to specify which SMS group(s) can use each VPN certificate.

Tab	Explanation
CRLs	Click the tab to manage Certificate Revocation Lists (CRLs). These lists are provided by the Certificate Authority to identify certificates that are no longer valid.
Proxies	CRLs can be updated periodically or on demand from a URL that you specify on the CRLs tab. If there is an HTTP proxy server between the SMS and the web site specified in the CRL Update URL, click this tab to configure the proxy server and port.
Pending CSRs	Click this tab to view or delete Certificate Signing Requests (CSRs) for Server Certificates and VPN Certificates that are still pending (CSR has been generated, but the certificate has not been imported into the SMS yet).

All digital certificates that have been obtained and installed will be shown in the Certificates Panel.



To Start the Certificate Manager

Overview

Use this task to access the Certificate Manager from the SMS Navigator or Remote Navigator.

To access the Certificate Manager from the Navigator or Remote Navigator

- 1 Log into the SMS Navigator or Remote Navigator as described in the *To Log On and Off the SMS Server or Compute Server* section in the *SMS Administration Guide*.

Result The SMS Navigator or Remote Navigator is displayed, depending on whether you logged into the SMS from the local SMS host or logged into the SMS from a remote host.

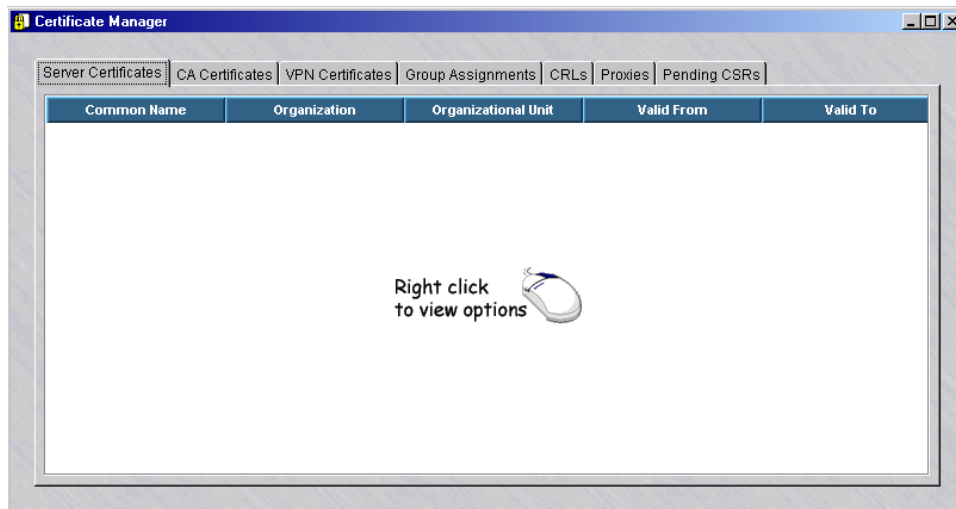
- 2 Select **Utilities** from the menu bar.

Result A drop-down menu is displayed.

- 3 Select **Certificate Manager** from the drop-down menu.

Result The Certificate Manager window is displayed ([Figure 9-2, “Certificate Manager Window”](#) (p. 9-5)).

Figure 9-2 Certificate Manager Window



END OF STEPS



To Obtain and Install a Server Certificate

Overview

Server certificates are used when the SMS web server or User Authentication web server is configured for HTTPS.

To obtain and import a server certificate you have to:

- Create a Certificate Signing Request (CSR)
- Submit the CSR to the Certificate Authority (CA) to obtain the server certificate
- Import the server certificate on the SMS
- Obtain and import the Certificate Authority's CA certificate(s) on the SMS

To create the Certificate Signing Request

To create a CSR, follow the steps below. The procedure is the same regardless of the CA that you are using.

-
- 1 Start the Certificate Manager.

 - 2 With the Server Certificates tab displayed, right-click and select New CSR from the menu.

 - 3 Fill in the fields on the Certificate Request form.

Request Options:

Field	Description	Optional/Required
Key Length	The length of the public and private keys determines how strong the encryption will be. Longer key lengths provide stronger security but slower performance.	Required
Key Algorithm	Two algorithms are supported: RSA (Rivest, Shamir, and Adleman) and DSA (Digital Signature Algorithm).	Required

The following fields define the Subject Name of your certificate:

Field	Description	Optional/Required
Common Name	The fully qualified domain name used for DNS lookup of your web server. <ul style="list-style-type: none"> For the SMS web server, this must resolve to the IP address of the SMS host. For the User Authentication web server, this must resolve to the IP address of the VBA (Virtual Brick Address) assigned to the zone where users will be authenticated. 	Required
Organization	The legal name of your company.	Optional
Organizational Unit	The name of your division or department.	Optional
Locality	The city where you are located.	Optional
State/Province Name	The full name of the state or province where you are located.	Optional
Country	The two-letter abbreviation of the country where you are located.	Required

The following fields can be used to optionally define Subject Alternate Name attributes in your certificate. These are normally used in VPN certificates rather than Server certificates:

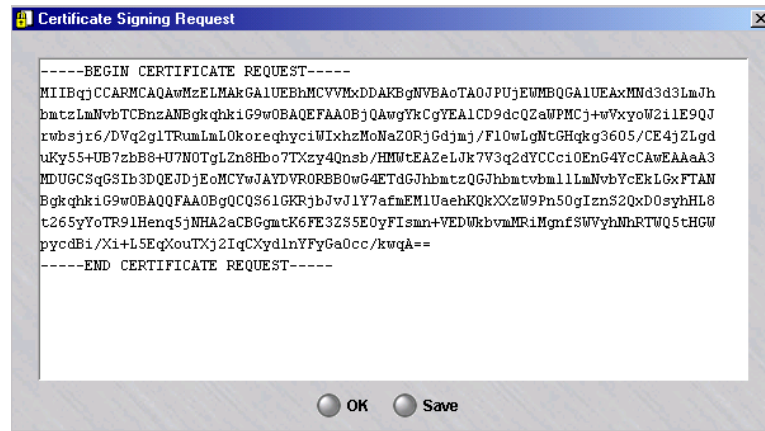
Field	Description	Optional/Required
Email	Your email address	Optional
Domain Name	The fully qualified domain name of your web server.	Optional
IP Address	The IP address of your web server.	Optional

-
- When you have finished entering this information, click **OK**.

Result A private/public key pair will now be generated. The public key will be wrapped in the CSR and displayed in a pop-up Certificate Signing Request window.

[Figure 9-3, “Certificate Signing Request Window”](#) (p. 9-9) shows a sample of the Certificate Signing Request pop-up window.

Figure 9-3 Certificate Signing Request Window



The CSR can be copied and pasted to a CA website in a browser (refer to the [“Obtain the Server Certificate”](#) (p. 9-9) procedure), or saved to a specified folder and filename by clicking the **Save** button to display a Select File pop-up window and specifying a folder and filename for the CSR file, then clicking OK to save the file.

END OF STEPS

Obtain the Server Certificate

Once the CSR has been created, you have to present it to a Certificate Authority to obtain your certificate. The CA will check the information in the CSR, verify your credentials, and then issue a certificate. The entire process should take from two to five business days. The specific procedure to follow to obtain the certificate varies according to the CA. The following is a general procedure you can use as a guide. Check with your CA for the specific steps they require to issue a digital certificate.

To present the CSR to a Certificate Authority, follow the steps below:

- 1 Open a browser and access your CA’s website. Follow the instructions on the website to obtain a digital certificate.

-
- 2 The instructions will take you to a page that will request that you enter the CSR in the space provided.

To do this:

1. From Step 4 of the [“To create the Certificate Signing Request”](#) (p. 9-7) task, leave the Certificate Signing Request window open with the CSR file, or open the saved CSR file in a separate window using a text editor such as Notepad or vi.
2. Select the text in the file, copy it, and paste it into the space provided on the web site.
Be sure to copy the entire text, including the lines that say
“-----BEGIN CERTIFICATE REQUEST-----” and
“-----END CERTIFICATE REQUEST-----”.
3. Follow the instructions on the web site for generating a digital certificate. You may be prompted to select a format for your certificate. You can choose Base64 (PEM), binary, or a PKCS7 format that also packages the CA certificate along with your sever certificate. See [“Import the Server Certificate”](#) (p. 9-10) for more information on supported certificate formats.
4. In two to five business days you should receive an email containing your certificate or instructions on how to pick up your certificate from the Certificate Authority’s web site. Save the certificate in a file. You are now ready to import the certificate. The procedure is described in the next section.

END OF STEPS

Import the Server Certificate

Once the CA issues your server certificate, you have to import it on the SMS. The following explains how to do this:

-
- 1 Start the Certificate Manager.
 - 2 With the Server Certificates tab displayed, right-click and select **Import** from the menu.
 - 3 Select the format of the certificate file in the File Format field. If the file is in ASCII format with the line “-----BEGIN CERTIFICATE-----” at the top of the file, and “-----END CERTIFICATE-----” at the bottom of the file, select PEM. This is sometimes called Base64 format. Other formats that are supported include Binary, PKCS7 PEM (usually used when the server certificate and CA certificate are sent

together in the same file), PKCS7 Binary, or PKCS12. PKCS12 is a password-protected format and is usually used when both the private key and certificate are packaged together.

-
- 4 Click the **Browse** button and select the folder and file where you saved the certificate file that you received from the certificate authority.

-
- 5 Click **OK** to import the certificate.

-
- 6 If you imported your server certificate in a PKCS7 format that included the CA certificate, the CA certificate should be imported automatically along with the server certificate. Click on the CA Certificates tab to see if the CA certificate is listed there.

If the CA certificate for the certificate authority that issued your server certificate has not been imported into the SMS, follow the procedure in the next section. *Server certificates will not work unless the CA certificate of the CA that issued the certificate is imported into the SMS.*

END OF STEPS

Obtain and Import the CA Certificate

The CA certificate for the Certificate Authority that issues your Server or VPN certificate must also be imported on the SMS. Your Server or VPN certificate will not work until you import the CA certificate of the Certificate Authority that issued the certificate. If the CA you are using has issued you more than one digital certificate, you only need to import the CA certificate once on the SMS. If the CA that issued your server certificate is an Intermediate Certificate Authority, you must import the Intermediate CA certificate, and any other Intermediate CA certificates in the certificate chain, as well as the Root CA certificate. CA certificates for Verisign and Thawte are provided with the SMS software and are already imported on the SMS. If you are using another CA, obtain their CA certificate and save it in a text file.

To import a CA certificate, follow the steps below:

-
- 1 Start the Certificate Manager.

-
- 2 Click on the CA Certificates tab.

.....
3 Right-click and select **Import** from the menu.

.....
4 Select the format of the certificate file in the **File Format** field.

.....
5 Click the Browse button and select the folder and file where you saved the CA certificate file.

.....
6 Click **OK** to import the certificate.

.....
E N D O F S T E P S
.....



To Obtain and Import a VPN Certificate

Overview

You can use X.509 digital certificates to authenticate Alcatel-Lucent IPsec Client users who are attempting to establish a VPN tunnel between their computer and a Brick. X.509 certificates can also be used to authenticate LAN-LAN tunnels.

To obtain and import a VPN certificate you have to:

- Create a Certificate Signing Request (CSR)
- Submit the CSR to the Certificate Authority (CA) or your PKI (Public Key Infrastructure) Administrator, to obtain the VPN certificate
- Import the VPN certificate on the SMS
- Obtain and import the Certificate Authority's CA certificate(s) on the SMS

Create a Certificate Signing Request

To create a CSR, follow the steps below. The procedure is the same regardless of the CA that you are using.

-
- 1 Start the Certificate Manager.

 - 2 Click on the VPN Certificates tab.

 - 3 Right-click and select New CSR from the menu.

 - 4 Fill in the fields on the Certificate Request form.

Request Options::

Field	Description	Optional/Required
File Name	Click on the Browse button and select the folder and file name where the new CSR will be created.	Required
Key Length	The length of the public and private keys determines how strong the encryption will be. Longer key lengths provide stronger security but slower performance.	Required

Field	Description	Optional/Required
Key Algorithm	Two algorithms are supported: RSA (Rivest, Shamir, and Adleman) and DSA (Digital Signature Algorithm).	Required

The following fields define the Subject Name of your certificate::

Field	Description	Optional/Required
Common Name	The fully qualified domain name used for DNS lookup of your tunnel endpoint. This must resolve to the IP address of the VBA (Virtual Brick Address) assigned to the Tunnel Endpoint on the Brick Policy Assignment tab.	Required
Organization	The legal name of your company.	Optional
Organizational Unit	The name of your division or department.	Optional
Locality	The city where you are located.	Optional
State/Province Name	The full name of the state or province where you are located.	Optional
Country	The two-letter abbreviation of the country where you are located.	Required

The following fields can be used to optionally define Subject Alternate Name attributes in your certificate::

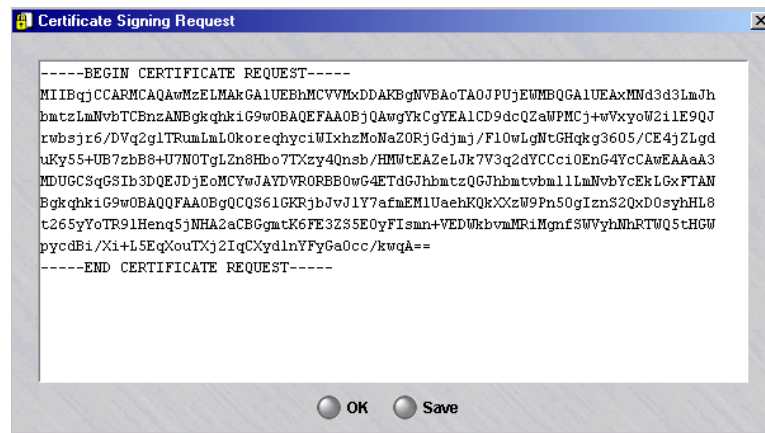
Field	Description	Optional/Required
Email	Your email address	Optional
Domain Name	The fully qualified domain name of your tunnel endpoint	Optional
IP Address	The IP address of your tunnel endpoint.	Optional

-
- 5** When you have finished entering this information, click **OK**.

Result A private/public key pair will now be generated. The public key will be wrapped in the CSR and displayed in a pop-up Certificate Signing Request window.

Figure 9-4, “Certificate Signing Request Window” (p. 9-15) shows a sample of the Certificate Signing Request pop-up window.

Figure 9-4 Certificate Signing Request Window



The CSR can be copied and pasted to a CA website in a browser (refer to the “Obtain the Server Certificate” (p. 9-9) procedure), or saved to a specified folder and filename by clicking the **Save** button to display a Select File pop-up window and specifying a folder and filename for the CSR file, then clicking OK to save the file.

END OF STEPS

Obtain the VPN Certificate

Once the CSR has been created and saved in a file, you have to present it to a Certificate Authority or your PKI (Public Key Infrastructure) Administrator to obtain your certificate. The CA will check the information in the CSR, verify your credentials, and then issue a certificate. The entire process should take from two to five business days. The specific procedure to follow to obtain the certificate varies according to the CA. The following is a general procedure you can use as a guide. Check with your CA for the specific steps they require to issue a digital certificate.

To present the CSR to a Certificate Authority, follow the steps below:

- 1 Open a browser and access your CA’s website. Follow the instructions on the website to obtain a digital certificate.

-
- 2** The instructions will take you to a page that will request that you enter the CSR in the space provided.

To do this:

1. From Step 4 of the [“To create the Certificate Signing Request”](#) (p. 9-7) task, leave the Certificate Signing Request window open with the CSR file, or open the saved CSR file in a separate window using a text editor such as Notepad or vi.
2. Select the text in the file, copy it, and paste it into the space provided on the web site.
Be sure to copy the entire text, including the lines that say
“-----BEGIN CERTIFICATE REQUEST-----” and
“-----END CERTIFICATE REQUEST-----”.
3. Follow the instructions on the web site for generating a digital certificate. You may be prompted to select a format for your certificate. You can choose Base64 (PEM), binary, or a PKCS7 format that also packages the CA certificate along with your sever certificate. See [“Import the Server Certificate”](#) (p. 9-10) for more information on supported certificate formats.
4. In two to five business days you should receive an email containing your certificate or instructions on how to pick up your certificate from the Certificate Authority’s web site. Save the certificate in a file. You are now ready to import the certificate. The procedure is described in the next section.

-
- 3** The instructions will take you to a page that will request that you enter the CSR in the space provided. To do this:

To do this:

1. In a separate window, open the CSR file in a text editor such as Notepad or vi.
2. Select the text in the file, copy it, and paste it into the space provided on the web site.
Be sure to copy the entire text, including the lines that say
“-----BEGIN CERTIFICATE REQUEST-----” and
“-----END CERTIFICATE REQUEST-----”.

3. Follow the instructions on the web site for generating a digital certificate. You may be prompted to select a format for your certificate. You can choose Base64 (PEM), binary, or a PKCS7 format that also packages the CA certificate along with your VPN certificate. See [Import the VPN Certificate](#) for more information on supported certificate formats.
4. In two to five business days you should receive an email containing your certificate or instructions on how to pick up your certificate from the Certificate Authority's web site. Save the certificate in a file. You are now ready to import the certificate. The procedure is described in the next section.

.....
E N D O F S T E P S
.....

Import the VPN Certificate

Once the CA issues your VPN certificate, you have to import it on the SMS. The following explains how to do this:

- 1 Start the Certificate Manager.
.....
- 2 Click on the VPN Certificates tab.
.....
- 3 Right-click and select **Import** from the menu.
.....
- 4 Select the format of the certificate file in the File Format field. If the file is in ASCII format with the line "-----BEGIN CERTIFICATE-----" at the top of the file, and "-----END CERTIFICATE-----" at the bottom of the file, select PEM. This is sometimes called Base64 format. Other formats that are supported include Binary, PKCS7 PEM (usually used when the server certificate and CA certificate are sent together in the same file), PKCS7 Binary, or PKCS12. PKCS12 is a password-protected format and is usually used when both the private key and certificate are packaged together.

Important! The Entrust EPF file format is no longer supported by the SMS. If you have VPN certificates in Entrust EPF format, follow the procedure in the [“How to Convert an Entrust EPF File to PKCS12”](#) (p. 9-19) section.

- 5 Click the Browse button and select the folder and file where you saved the certificate file that you received from the certificate authority.

-
- 6 Click **OK** to import the certificate.

The Group Assignment window is displayed.

- 7 Select the Group where the VPN Certificate will be used and click **OK**. If the certificate will be used in more than one group, click on the Group Assignments tab and create additional entries.
-

- 8 If you imported your VPN certificate in a PKCS7 format that included the CA certificate, the CA certificate should be imported automatically along with the VPN certificate. Click on the CA Certificates tab to see if the CA certificate is listed there.

If the CA certificate for the certificate authority that issued your VPN certificate has not been imported into the SMS, follow the procedure in the next section. *VPN certificates will not work unless the CA certificate of the CA that issued the certificate is imported into the SMS.*

END OF STEPS

Obtain and Import the CA Certificate

The CA certificate for the Certificate Authority that issues your Server or VPN certificate must also be imported on the SMS. Your Server or VPN certificate will not work until you import the CA certificate of the Certificate Authority that issued the certificate. If the CA you are using has issued you more than one digital certificate, you only need to import the CA certificate once on the SMS. If the CA that issued your server certificate is an Intermediate Certificate Authority, you must import the Intermediate CA certificate, and any other Intermediate CA certificates in the certificate chain, as well as the Root CA certificate. CA certificates for Verisign and Thawte are provided with the SMS software and are already imported on the SMS. If you are using another CA, obtain their CA certificate and save it in a text file.

To import a CA certificate, follow the steps below:

- 1 Start the Certificate Manager.

- 2 Click on the CA Certificates tab.

- 3 Right-click and select **Import** from the menu.

- 4 Select the format of the certificate file in the File Format field.
- 5 Click the Browse button and select the folder and file where you saved the CA certificate file.
- 6 Click **OK** to import the certificate.

END OF STEPS

How to Convert an Entrust EPF File to PKCS12

If you have Entrust EPF file certificates, use the utility provided on the SMS CD to convert the files to PKCS12 format.

Copy the directory *Lucent-EPF-Conversion* and its contents from the SMS CD to a folder on a Windows machine where you will do the conversion. This directory contains the files and directories needed to perform the following conversion procedure:

- 1 To get the certificate from the Entrust Server, on your Entrust Server, create a user as a regular user.
- 2 Go to the Certificate Info and click **Export**→**Apply**→**OK**→**OK**
- 3 To get the certificate from the Alcatel-Lucent IPsec Client, click on the **File** menu and select **Get Entrust Certificate**.
- 4 Type in the necessary information and save the certificate on a floppy disk (certificate.epf and certificate.ini, where certificate is any name to identify the certificate).
- 5 Copy these files to the directory *Lucent-EPF-Conversion/Data* that was created from the SMS CD prior to beginning the conversion procedure.
- 6 Edit the *certificate.ini* file to remove the FIPS information at the end of the file.

.....
7 Double-click on the executable file *LuCerCov.exe* to launch the conversion utility.
.....

8 On the conversion screen, select the radio button **Entrust EPF to PKCS#12**.
.....

9 Use the search button (>>) next to the drop-down window to locate the certificate.epf file.
.....

10 Enter the password for the certificate in both the password and confirm password windows and click **Next**.
.....

A new screen opens to allow you to enter the PKCS#12 file name.
.....

11 Enter the file name with a .p12 extension.
.....

12 Enter the password in both the password and confirm password windows and click **OK**.
A window confirming that you have successfully converted the certificate is displayed.
.....

13 Click **OK** and you have completed the conversion process.
.....

END OF STEPS
.....



Group Assignment for VPN Certificates

Overview

VPN certificates must be assigned to an SMS Group before they can be used. Your VPN certificate will not be listed as a choice in the VPN Certificate field on the Brick Editor Policy Assignment tab until you assign the certificate to a group. The group should be the group of the Zone that is assigned to your tunnel endpoint on the Brick Editor Policy Assignment tab.

To assign a VPN Certificate to a group:

- 1 Start the Certificate Manager.
.....
- 2 Click on the Group Assignments tab.
.....
- 3 Right-click and select **New** from the menu.
.....
- 4 Select a group from the choices in the **Group** field.
.....
- 5 Select your VPN certificate from the list of Common Names in the Common Name field.
.....
- 6 Click **OK** to finish the assignment.

END OF STEPS
.....

Certificate Revocation Lists

Complete the following steps to configure Certificate Revocation Lists:

- 1 Certificate Authorities provide Certificate Revocation Lists (CRLs) to identify certificates that are no longer valid. To prevent users from establishing VPN tunnels using invalid certificates, you should configure the CRL for your Certificate Authority. The CRL can usually be downloaded from the CA's website, and they usually provide a URL that can be used to automatically update the CRL.

To configure a CRL for your Certificate Authority you should:

- Obtain and Import the current CRL on the SMS
- Configure the URL for automatic CRL updates
- Configure HTTP Proxies needed to access the CRL Update URL

END OF STEPS

Obtain and Import the CRL

To obtain and import the CRL for your Certificate Authority:

- 1 Open a browser and go to the CA web site.
.....
- 2 Download a copy of the current CRL to the SMS. The file can be in Base64 (PEM) or binary format.
.....
- 3 Start the Certificate Manager.
.....
- 4 Click on the CRLs tab.
.....
- 5 Right-click and select **Import** from the menu.
.....
- 6 Select the format of the certificate file in the **File Format** field.
.....
- 7 Click the Browse button and select the folder and file where you saved the CRL file.
.....
- 8 Click **OK** to import the CRL.

END OF STEPS

Configure a CRL for Automatic Updates

Certificate Authorities update their Certificate Revocation List periodically to include additional certificates that have become invalid. You can manually update CRLs by going to your CA's website and downloading the latest CRL, but it is more convenient to have the SMS automatically retrieve the updated CRL for you. When CRL Update

is configured, the SMS will contact the CA's web site once a day, download the new CRL file, and download the CRL file to any Bricks that use VPN certificates from that Certificate Authority.

To configure a CRL for automatic updates:

- 1 Start the Certificate Manager.
- 2 Click on the CRLs tab.
- 3 Right-click on the row containing the CRL you want to automatically update and select Edit from the menu.
- 4 In the Update URL field, enter the URL that should be used to update the CRL. You should be able to find information on what URL to use for CRL updates on your Certificate Authority's web site.
- 5 Click **OK**.

Once a CRL has been configured with an Update URL, you can update the CRL on demand by right-clicking on that CRL and selecting Update Now from the menu.

END OF STEPS

Configure Proxies for CRL Updates

If you have configured Automatic CRL Updates, you may have to configure a proxy to allow the SMS to contact the Certificate Authority web site. Each SMS (primary, secondary, and compute server) will perform CRL updates once a day for the Bricks that are homed to that SMS. If there is an HTTP Proxy server between the SMS and the Certificate Authority web site, you must configure the proxy in the Certificate Manager.

To configure a proxy in the Certificate Manager:

- 1 Start the Certificate Manager.
- 2 Click on the Proxies tab.

.....
3 Right-click on the row containing the SMS you want to configure and select **Edit** from the menu.

.....
4 In the **Proxy Server** field, enter the domain name or IP Address of the HTTP proxy server.

.....
5 Enter the port number in the **Proxy Port** field.

.....
6 Click **OK**.

.....
E N D O F S T E P S
.....



Pending Certificate Signing Requests

Overview

Whenever you generate a CSR on the Server Certificates or VPN Certificates tab, an entry is created on the Pending Certificates tab. This allows you to view any CSRs that are currently pending (CSR has been generated, but the certificate has not been imported on the SMS yet). Entries are automatically removed from this table when you import the certificate.

If for some reason you have decided not to import the certificate for a CSR (for example, if the information in the CSR was wrong), you can delete the CSR and its associated private key from the SMS database by deleting the CSR on the Pending CSRs tab.

Delete a Pending CSR

You should only delete a CSR from the Pending CSRs tab if you are sure you will not be importing the certificate for that CSR. Deleting the CSR from the Pending CSRs tab deletes the private key information for that certificate. Once this is done, you will not be able to import a certificate obtained from a Certificate Authority for that CSR.

To delete a pending CSR:

-
- 1 Start the Certificate Manager.

 - 2 Click on **Pending CSRs**.

 - 3 Right-click on the row containing the CSR you want to delete and select **Delete** from the menu.

 - 4 Click **OK** in the delete confirmation dialog box.

END OF STEPS



Server Versus VPN Certificate Signing Request

Certificate signing request

If you are obtaining a Server Certificate, you should create the CSR on the Server Certificates tab. If you are obtaining a VPN Certificate, you should create the CSR on the VPN Certificates tab. If you create the CSR on the wrong tab by mistake, all is not lost. Obtain your certificate and import it using the procedures described previously.

After the certificate is imported (on the wrong tab), right-click on the certificate and select Move to VPN (or Move to Server) to move your certificate to the other tab.



10 LAN-LAN Tunnels

Overview

Purpose

This chapter explains how to set up a LAN-LAN tunnel between two devices so that hosts behind both devices will be able to communicate securely with one another over a public network such as the Internet. It also explains how to set up a LAN-LAN tunnel that uses UDP encapsulation instead of pure IPSec as the transport mode.

Contents

What is a LAN-LAN Tunnel?	10-2
To Set LAN-LAN Tunnel Defaults	10-6
To Set Up a LAN-LAN Tunnel	10-13
To Set Up a LAN-LAN Tunnel with UDP Encapsulation	10-25
Redundant LAN-LAN Tunnels	10-27
Maintaining LAN-LAN Tunnels	10-28
To Set Up Service Level Agreements	10-32



What is a LAN-LAN Tunnel?

Definition

A LAN-LAN tunnel is a Virtual Private Network (VPN), or encrypted path, through the Internet. The devices at either end of the path are called *tunnel endpoints*.

Types of Endpoints

Brick devices and unmanaged devices can all serve as the endpoint of a LAN-LAN tunnel. The drawing in [Figure 10-1, “Types of LAN-LAN Tunnel Endpoints”](#) (p. 10-2) illustrates the different types of endpoints.

Figure 10-1 Types of LAN-LAN Tunnel Endpoints



The following briefly describes what is required to make each type of device a tunnel endpoint:

- Brick*

For a Brick to serve as a tunnel endpoint, it must have a tunnel endpoint address associated with each port that will terminate a tunnel. Tunnel endpoint addresses are entered when a Brick zone ruleset is assigned to a port on the Brick. Refer to the *Configuring Alcatel-Lucent VPN Firewall Brick™ Security Appliance Ports* chapter in the *SMS Administration Guide*.

A Brick port can have multiple tunnel endpoint addresses, one for each ruleset assigned to it, and each endpoint address can terminate multiple tunnels.
- Unmanaged Device*

An unmanaged device is any device that is not controlled by your SMS. It could be a Brick that is managed by another SMS, another Group Administrator, another vendor router or firewall, or a non-Alcatel-Lucent IPSec Client application.

Tunnel modes

A LAN-LAN tunnel can be set up in two different modes:

- Tunnel
- Transport

In Tunnel mode, a LAN-LAN tunnel can be set up to allow secure, encrypted communications between two LAN-connected devices (such as two laptops) on the same VPN.

In Transport mode, which is mainly used in end-to-end IPsec deployments, the endpoint devices of the LAN-LAN tunnel are two hosts at the A and Z end of the network. The hosts, or endpoints, are hidden behind each tunnel device (Brick) when Transport mode is used. Each Brick device performs NAT to deliver encrypted communications to its destination endpoint (host). Each Brick's NAT (public) address is part of the key exchange information when the tunnel is set up between the endpoints. A tunnel can be set up in Transport mode to secure, for example, SIP signaling between an IMS client and P-CSCF server or IP-encapsulated GRE tunneling between routers.

Tunnel Keying

Regardless of the type of device at each endpoint, you have to configure both endpoints so that they can exchange keys and communicate. SMS supports both automated key exchange and manual key assignment:

The following explains the differences between automated and manual key exchange:

- *Automated Key Exchange*

If you choose automated key exchange, the encryption and authentication types and keys are automatically negotiated by the devices at both endpoints. For greater security, the keys are automatically updated periodically.

The mechanism that automatically negotiates the information that goes into the Security Associations (SAs) is the Internet Key Exchange (IKE) process. Default parameters for automatic key exchange are provided with the SMS application. You can use those defaults, or you can change them to suit your needs.
- *Manual Key Assignment*

If the key is to be assigned manually, an Administrator has to create manually create SAs, containing the encryption/authentication types and keys, at both endpoints. This must be done precisely, because the encryption/authentication types and keys have to be exactly the same, or the tunnel will not work.

The SAs that are manually created are used indefinitely by the tunnel endpoints, until new SAs are manually created and downloaded to them.

Internet key exchange version 2 (IKEv2)

The SMS and managed Bricks, as of Release 9.0, implement the Internet Key Exchange, Version 2 protocol (IKEv2) as defined in Internet Engineering Task Force (IETF) draft-ietf-ipsec-ikev2-17, including support for NAT Traversal and Extensible Authentication Protocol (EAP), in addition to IKEv1 protocol.

LAN-LAN tunnels can be configured with Manual Keys, or automatic key exchange using IKEv1 or IKEv2. Preshared Key and X.509 Certificate authentication are supported.

IKE hardware acceleration

Some IKE (v1 and v2) operations are now executed with hardware acceleration on the following Brick models:

- Model 50 Brick
- Model 150 Brick
- Model 700 Basic and 700 VPN Bricks
- Model 1100 and 1100A Bricks
- Model 1200 and 1200 HS Bricks

IKE hardware acceleration greatly increases the number of tunnels per second that can be established on these devices as well as reducing the overall load on Brick devices.

LAN-LAN tunnel defaults

Before you set up a LAN-LAN tunnel, you should examine the LAN-LAN tunnel defaults provided with the SMS application to determine whether or not any of them need to be modified.

The SMS application comes configured with certain session and policy (IKE SA/IPSec) parameters already set. These default settings apply only to automatic key exchange, and they make it easy for you to configure Bricks to serve as LAN-LAN tunnel endpoints.

You can use these default settings as they are, or you can change them to suit your own specific needs. Moreover, the settings are group-specific, which means if you have created groups other than the *System* group, you can leave the defaults in place in one group, but change them in another, or change them differently in both groups.

For this reason, it is recommended that before you attempt to set up a tunnel, you take a look at the default parameters that apply to the group you are using and determine whether to keep them or change them.

Important! It is best to be cautious when changing the defaults. If you have active tunnels based on the defaults — and you change the defaults — the active tunnels may be adversely affected, particularly if one endpoint is a non-Alcatel-Lucent device.



To Set LAN-LAN Tunnel Defaults

When to use

Use this task to configure or modify the default parameter settings for LAN-LAN tunnels.

Task

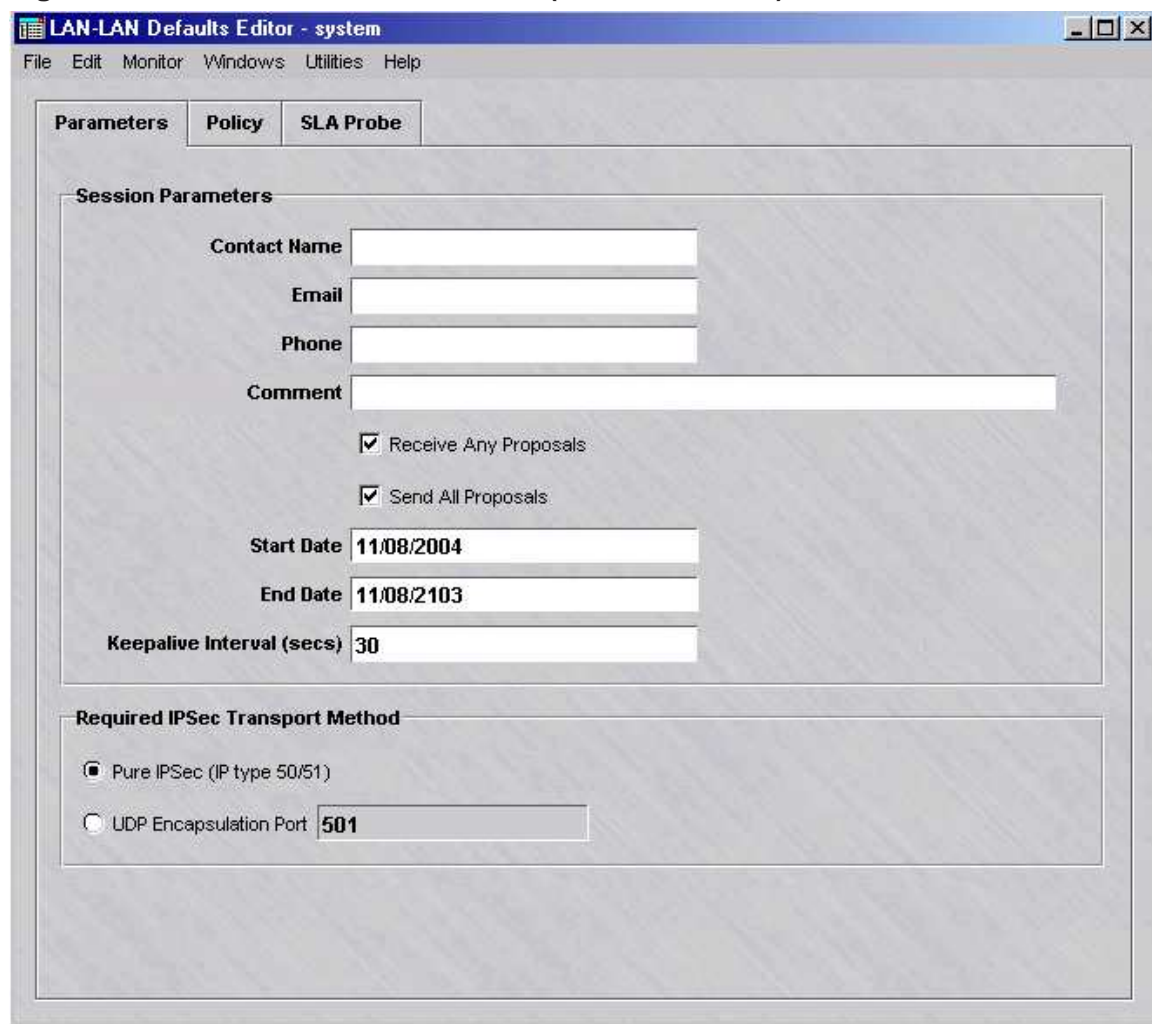
Complete the following step to display and/or modify defaults for LAN-LAN tunnels:

- 1 With the Navigator window displayed, open the appropriate group folder, and then open the VPN folder.

- 2 Click **VPN Defaults** to display entries for the two defaults editors in the Contents panel.

- 3 Double-click **LAN-LAN Defaults**. The LAN-LAN Defaults Editor is displayed, with the Parameters tab displayed. This tab is shown in [Figure 10-2, “LAN-LAN Defaults Editor \(Parameters Tab\)”](#) (p. 10-7).

Figure 10-2 LAN-LAN Defaults Editor (Parameters Tab)



- 4 Several of the fields are populated with default values, and others are left blank. You can change any of the defaults and fill in any of the blanks. The table below explains.

Field	Explanation
Contact Name	If one of the endpoints is administered by another individual, you can enter information about that person in these fields. This data is purely informational, and not used in tunnel negotiations.
E-mail	
Phone	
Comment	

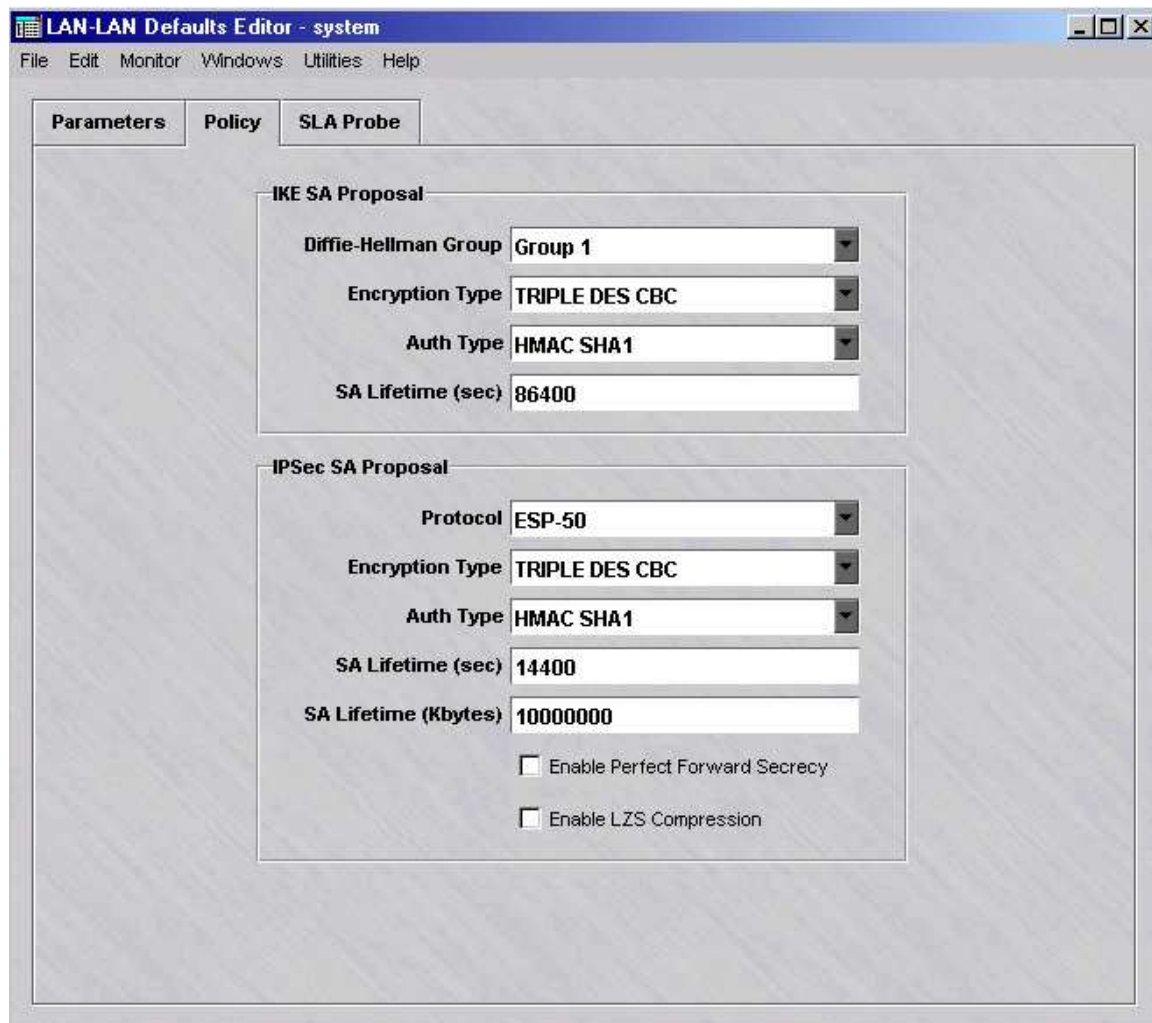
Field	Explanation
Receive Any Proposals <i>and</i> Send All Proposals	These two boxes are checked by default. The purpose is to enhance interoperability. When these boxes are checked, the IKE SA parameters and IPSec parameters will be negotiated at a possibly lower security level than you specified, to allow devices configured differently to still serve as one tunnel endpoint. However, if you want to ensure that only devices sharing your IKE SA parameters and IPSec parameters can serve as a tunnel endpoint, uncheck either or both of the checkboxes.
Start Date <i>and</i> End Date	These dates determine the time period during which this tunnel is operable. The default is a 99-year time period, beginning with the current date. You can change this.
Keepalive Interval (secs)	The Brick sends heartbeat/keepalive messages through the tunnel at regular intervals. This field allows you to set the interval (in seconds). The default is 30 seconds. Enter a value of zero (0) to disable keepalive/heartbeat messages. The heartbeat is used to determine the status of the tunnel that is displayed in the LAN-LAN Tunnel Viewer and the Status Overview portion of the Status Monitor.
IPSecTransport Method	These radio buttons determine the transport method. The default is pure IPSec (IP type 50/51), but you can change this to UDP Encapsulated (IP type 17). See “To Set Up a LAN-LAN Tunnel with UDP Encapsulation” (p. 10-25) for a more complete discussion of UDP encapsulation.

-
- 5 Once you have completed the fields on the Parameters tab, click **Policy** to display the Policy tab.

Result The Policy tab is displayed ([Figure 10-3, “LAN-LAN Defaults Editor \(Policy Tab\)” \(p. 10-9\)](#))

The Policy tab contains the IKE SA Proposal and IPSec Proposal defaults.

Figure 10-3 LAN-LAN Defaults Editor (Policy Tab)



- 6 The fields in the **IKE SA Proposal** box are populated with default values. You can change any of these values. The following explains:

Field	Explanation
Diffie-Hellman Group	Higher numbered groups provide more security than lower numbered groups, but also slow down re-keying and tunnel establishment. Generally, Group 1 should not be used.

Field	Explanation
Encryption Type	The default is TRIPLE DES CBC . The other encryption method choices that can be selected are DES CBC, AES CBC 128, AES CBC 192, AES CBC 256 (which is the strongest encryption method available).
Auth Type	The default is HMAC SHA1 , but you can change this to HMAC MD5. <i>Note: MD5 signed certificates are no longer considered secure.</i>
SA Lifetime (secs)	This Security Association has a lifetime specified in seconds. The default is 86400 seconds (24 hours). You can change the lifetime in seconds to any value between 120-172,860 seconds (between 2 minutes and 48 hours).

- 7 The fields in the IPsec SA Proposal box are populated with default values. You can change any of these values. The following explains:

Field	Explanation
Protocol	The default is ESP-50, but this can be changed to AH-51. ESP-50 provides both encryption and authentication for every packet, while AH-51 only provides authentication.
Encryption Type	The default is TRIPLE DES CBC. Other available choices are AES CBC 128, AES CBC 192, AES CBC 256, NULL, and DES CBC. DES CBC is not recommended and is not available as a choice when FIPS is enabled.
Auth Type	The default is HMAC SHA1, but you can change this to HMAC MD5. <i>Note: MD5 signed certificates are no longer considered secure.</i>

Field	Explanation
SA Lifetime (secs)	This Security Association has a lifetime specified in seconds. The default is 14400 seconds (6 hours). You can change the lifetime in seconds to any value between 120-172,860 seconds (between 2 minutes and 48 hours).
SA Lifetime (Kbytes) ¹	This Security Association has a lifetime specified in kilobytes. The default is 10,000,000 kilobytes. To change this, enter a number between 61000 and 10,000,000 kilobytes.
Enable Perfect Forward Secrecy	By default, this checkbox is unchecked. If it is checked, a new Diffie-Hillman key exchange will take place at every rekey interval, thereby increasing rekey time and the load on the SMS. This can improve security, but it can also affect performance.
Enable LZS Compression	This checkbox applies only to LAN-LAN tunnels between a Brick with an encryption accelerator cards and any device that supports LZS compression. By default it is unchecked. If it is checked, traffic through the tunnel will be compressed. The advantage of compression is that it means less data has to be sent over the wires, which may help conserve bandwidth and speed up transmission. The disadvantage is that it requires extra processing (to compress and decompress) on either end of the tunnel — which has performance implications at the tunnel endpoints
1. The Security Association will expire after the <i>first</i> of the two lifetimes is reached. However, no session will be permitted to timeout before one minute, even if one of the above two lifetimes is reached first.	

-
- 8 Once you have completed both tabs, display the File menu and select **Save and Apply**.
.....
 - 9 When the changes have been applied, a message box will appear. Click **OK** to dismiss this box. The new defaults are now in effect. If any tunnels are currently using these defaults, the changes will apply to them immediately.

You may be asked to download policies to devices if any are out of date.

END OF STEPS



To Set Up a LAN-LAN Tunnel

When to use

Use this task to set up a LAN-LAN tunnel.

Before you begin

Before you begin this task, be aware that there are two methods of accessing the LAN-LAN Tunnel Editor: from the VPN folder or from the Device folder. Use Method 1 to display the LAN-LAN Tunnel Editor from the VPN folder. Use Method 2 to display the LAN-LAN Tunnel Editor from the Devices folder.

Make sure that at least one of the tunnel endpoints being selected for the LAN-LAN Tunnel has a VPN Rule

Task, Method 1: from the VPN folder

Complete the following steps to display the LAN-LAN Tunnel Editor from the VPN folder.

- 1 With the Navigator window displayed, open the appropriate group folder, and then open the **VPN** folder.
- 2 Right-click **LAN-LAN Tunnel Viewers** and select **New LAN Tunnel** from the pop-up menu.

Result The LAN-LAN Tunnel Editor is displayed.

END OF STEPS

Task, Method 2: from the Devices folder

Complete the following steps to display the LAN-LAN Tunnel Editor from the Devices folder.

- 1 With the Navigator window displayed, open the appropriate group folder, and then open the **Devices** folder.
- 2 Click the **Bricks** folder to display all configured Bricks in the Contents panel.
- 3 Double-click the Brick that will be serving as the tunnel endpoint.

Result The Brick Editor is displayed.

- 4 Click **Policy Assignment** to display the Policy Assignment tab.
-

- 5 Right-click the port or interface serving as the tunnel endpoint and select **LAN-LAN VPN** from the pop-up menu.

Result If you opt for this method, the Tunnel tab of the LAN-LAN Tunnel Editor is displayed and the **Device** and **Tunnel Endpoint** fields for Endpoint 1 will be populated on the Endpoint 1 Tab of the Editor.

END OF STEPS

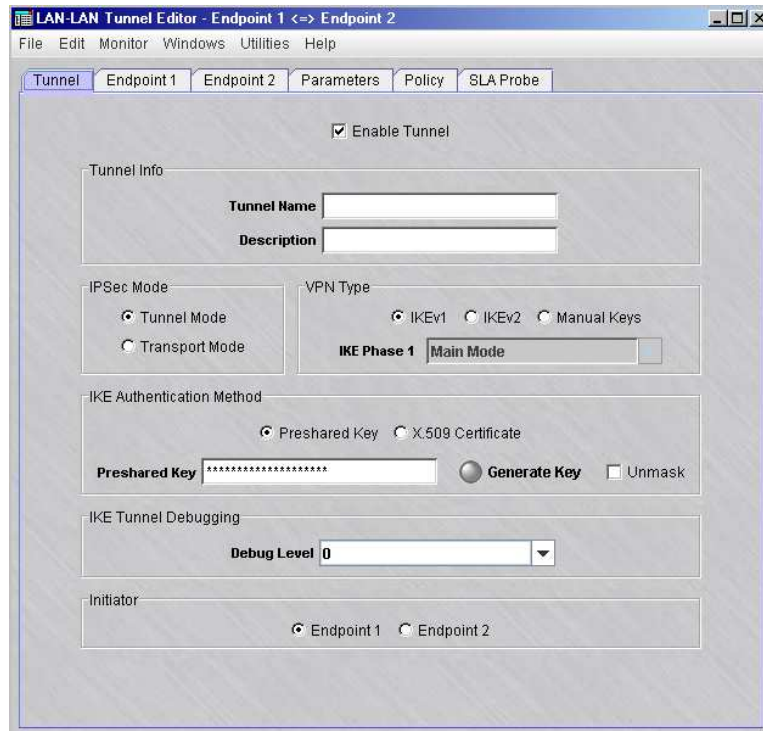
Task: complete the Tunnel Tab

Complete the following steps to complete the Tunnel tab. The purpose of the Tunnel tab is to configure general information about the tunnel.

- 1 Display the LAN-LAN Tunnel Editor using [“Task, Method 1: from the VPN folder”](#) (p. 10-13) or [“Task, Method 2: from the Devices folder”](#) (p. 10-13).

Result The LAN-LAN Tunnel Editor is displayed, showing the Tunnel tab (Figure 10-4, “LAN-LAN Tunnel Editor (Tunnel Tab)” (p. 10-15)).

Figure 10-4 LAN-LAN Tunnel Editor (Tunnel Tab)



-
- 2 The **Enable Tunnel** checkbox is checked by default, meaning that the tunnel will be activated as soon as you save and apply the tunnel. If you want the tunnel saved but not activated at this point, uncheck this box.

 - 3 In the **Tunnel Name** field, enter a name for the tunnel.
This information is used for debugging purposes, to identify the debug message logged in the VPN log with the TEP and tunnel.

 - 4 In the **Description** field, enter a text description of the tunnel. This field is optional.

 - 5 In the **IPSec Mode** box, choose the tunnel mode that will be used in setting up the tunnel. The choices are **Tunnel** or **Transport**. The default mode is **Tunnel**.

If **Transport** mode is chosen, only the IKEv2 key is allowed. This method will be pre-selected in the **VPN Type** box and all choices will be greyed out (skip [Step 6](#)).

- 6 In the **VPN Type** box, choose the VPN type of key exchange that will be used in setting up the tunnel.

If...	Then...
IKEv1 is chosen	Key exchange is handled automatically using the IKEv1 method, with Security Associations (SAs) automatically created and downloaded to each TEP, and updated periodically. (This is the default). If IKEv1 is chosen, go to Step 7 .
IKEv2 is chosen	Key exchange is handled automatically using the IKEv2 method, with SAs automatically created and downloaded to each TEP, and updated periodically.
Manual Keys is chosen	You have to manually enter the SA information for both tunnel endpoints. The SA information will not change until it is manually updated.

7

If...	Then...
IKEv1 was chosen as the VPN Type in Step 6	The IKE Phase 1 field is displayed. This is a read-only field that indicates whether Main Mode or Aggressive Mode is used for Phase 1 IKE. Main Mode is used when both tunnel endpoints are identified by IP address. Aggressive Mode is used if there is a mobile tunnel endpoint.

If...	Then...
<p>IKEv1 was chosen as the VPN Type in Step 6</p>	<p>The Initiator field is displayed.</p> <p>Select Endpoint 1 or Endpoint 2 to be the TEP that initiates the tunnel. Endpoint 1 is the default. It is recommended that Endpoint 1 be used as the initiator.</p> <p>If a TEP does not have a fixed IP address, it must be the initiator, to allow SMS to automatically select the address.</p>

-
- In the **IKE Authentication Method** box, select the authentication method used in setting up the tunnel between the two endpoints, if the VPN Type is **IKEv1** or **IKEv2**. The choices are: **Preshared Key** or **X509 Certificate**.

If...	Then...
<p>Preshared Key is chosen</p>	<p>A default preshared key is randomly generated by the system and hidden by asterisks (*) in the Preshared Key field.</p> <p>To generate a new preshared key, click the Generate Key button. The system generates a random 20-character key.</p> <p>To view the key, click the Unmask checkbox. The preshared key is displayed in the Preshared Key field.</p>

-
- Click the down arrow next to the **Debug Level** field to select a value to enable tunnel debugging.

If the Debug Level option is enabled, the Brick sends debug messages to the VPN log, identified by zone, tunnel name, and remote TEP. The available choices are **0** (debug off) to **3** (most verbose).

.....

END OF STEPS

.....

Task: complete the Endpoint 1 tab

Complete the following steps to complete the Endpoint 1 tab. This tab is used to select Endpoint 1 of the tunnel, if it has not already been selected, the IKE authentication method, and the identification method for the TEP.

- 1 Click the **Endpoint 1** tab of the LAN-LAN Tunnel Editor.

Result The Endpoint 1 tab is displayed.

- 2 Click the down-arrow next to the **Brick** field to display and select one of the Bricks in the list or select **Browse** to choose a Brick from another group.

If “[Task, Method 2: from the Devices folder](#)” (p. 10-13) was used to display the LAN-LAN Tunnel Editor, this field is already populated.

- 3 Click the down arrow next to the **Tunnel Endpoint** field and select one of the tunnel endpoints that has been configured. The associated Brick zone ruleset is displayed in parentheses next to the tunnel endpoint.

If “[Task, Method 2: from the Devices folder](#)” (p. 10-13) was used to display the LAN-LAN Tunnel Editor, this field is already populated.

If the tunnel mode chosen in the **IPSec Mode** box on the Tunnel tab is **Transport** (refer to the “[Task: complete the Tunnel Tab](#)” (p. 10-14) section), skip **Step 4**. Hosts behind the tunnel (Brick) device are hidden when Transport mode is used.

- 4 In the **Hosts Behind Tunnel** field, enter the hosts behind this endpoint that will be allowed to access the tunnel. The hosts that were entered in the same field on the Brick Policy Assignment Editor. Refer to the *Configuring Alcatel-Lucent VPN Firewall Brick® Security Appliance Ports* chapter in the *SMS Administration Guide* for more details.

5	If...	Then...
	<p>Preshared Key was chosen as the Authentication Method on the Tunnel tab</p>	<p>Click the down arrow next to the ID Type and choose the ID Type. The choices are: IP Address, Domain Name, or Email Address (for IKEv2 only). If IP Address is chosen, enter an IP address in the ID field. If the TEP has a fixed IP address, the ID is set to the VBA. If IP Address is chosen in the ID Type field, the ID is pre-set to Virtual Brick Address.</p> <p>If Domain Name or Email Address is chosen, enter the appropriate value in the ID field.</p> <p>For IKEv1 mobile endpoints (dhcp or pppoe), the IKE ID Type must be Domain Name. For IKEv2 mobile endpoints, all ID Type choices are allowed. For IKEv1 endpoints where the Endpoint Type is Name, the ID Type must be Domain Name.</p>
	<p>X509 Certificate was chosen as the Authentication Method on the Tunnel tab</p>	<p>Click the down arrow next to the ID Type field and select an ID type to be used for the tunnel endpoints. The choices are: Distinguished Name(the default), IP Address, Domain Name(IKEv2 only), and Email Address (IKEv2 only).</p> <p>The Certificate box is a read-only display that shows the attributes of the certificate assigned to the TEP on the Brick Policy Assignment tab.</p>

END OF STEPS

Task: complete the Endpoint 2 tab

Complete the following steps to complete the Endpoint 2 tab.

- 1 Click the **Endpoint 2** tab of the LAN-LAN Tunnel Editor.

Result The Endpoint 2 tab is displayed.

- 2 In the **Endpoint Type** field, choose **Brick** if both endpoints are devices in your group, or in two groups for which you have *Full* policy and VPN privileges. Click the **Browse** button and select the Brick.

Choose **IP Address** if the endpoints are in two groups, but you only have full privileges over one of the groups. You can also use the IP address if one of the endpoints is not managed by the SMS, such as a TEP on a non-Alcatel-Lucent device. Both you and the Administrator of the other group have to define Endpoint 1 as the device in your group, and then specify an IP address and a list of hosts behind the tunnel as Endpoint 2.

Choose **Name** if the TEP is a non-Alcatel-Lucent device or a Brick that is not managed by you (a Brick in a group where you do not have Full Policy and VPN privileges, or a Brick managed by another SMS). In this case, you select the host group that contains the hosts behind the tunnel or an asterisk (*) to represent all hosts.

3

If...	Then...
<p>Brick was chosen in Step 2</p>	<p>Click the down-arrow next to the Brick field and select a Brick from the list or select Browse to choose a Brick from another group.</p> <p>Click the down arrow next to the Tunnel Endpoint field and select the TEP.</p> <p><i>For Tunnel Mode only:</i>click the down arrow next to the Hosts Behind Tunnel field and enter an IP address or host group of the host(s) behind the tunnel.</p>
<p>IP Address was chosen in Step 2</p>	<p><i>For Tunnel Mode only:</i>enter the IP address of the device, click the down arrow next to the Hosts Behind Tunnel field to enter the IP address or host group of the host(s) behind the tunnel.</p>
<p>If Name was chosen in Step 2</p>	<p><i>For Tunnel Mode only:</i> click the down arrow next to the Hosts Behind Tunnel field to enter the IP address or host group of the host(s) behind the tunnel.</p>

4	If...	Then...
	<p>Preshared Key was chosen as the Authentication Method on the Tunnel tab</p>	<p>Click the down arrow next to the ID Type and choose the ID Type. The choices are: IP Address, Domain Name, or Email Address (for IKEv2 only). If IP Address is chosen, enter an IP address in the ID field. If the TEP has a fixed IP address, the ID is set to the VBA. If IP Address is chosen in the ID Type field, the ID is pre-set to Virtual Brick Address.</p> <p>If Domain Name or Email Address is chosen, enter the appropriate value in the ID field.</p>
	<p>X.509 Certificate was chosen as the Authentication Method on the Tunnel tab</p>	<p>For IKEv1 negotiations, click the down arrow next to the ID Type and choose the ID Type. The choices are: Distinguished Name (the default), IP Address, Domain Name (for IKEv2 only), or Email Address (for IKEv2 only).</p> <p>If the Endpoint Type is Brick, the Certificate box displays the VPN Certificate assigned to the TEP on the Brick Policy Assignment tab.</p> <p>If the Endpoint Type is IP Address or Name, and the ID Type is Distinguished Name, fields are displayed for entering the attributes that will be used to authenticate the tunnel against the certificate selected in the Certificate Authority field.</p> <p>If IP Address is chosen, enter an IP address in the ID field. If the TEP has a fixed IP address, the ID is set to the VBA.</p> <p>If Domain Name or Email Address is chosen, enter the appropriate value in the ID field.</p> <p>If the Endpoint Type is Brick, the Certificate Authority field is read-only and is automatically set.</p>

END OF STEPS

Task, complete the Parameters tab

Complete the following steps to complete the Parameters tab.

.....

- 1 Click **Parameters** to display the Parameters tab.

Result The Parameters tab is displayed.

.....

- 2 Click the **Use Group Defaults** checkbox to use the LAN-LAN tunnel defaults that are defined for tunnel endpoints in this group on the LAN-LAN Defaults Editor.
-

- 3 Change or enter any of the default values on the Parameters tab.

Refer to the field definitions in [“To Set LAN-LAN Tunnel Defaults”](#) (p. 10-6).

END OF STEPS

.....

Task: complete the Policy tab

Complete the following steps to complete the Policy tab.

.....

- 1 Click **Policy** to display the Policy tab.

Result The Policy tab is displayed.

.....

- 2 Click the **Use Group Defaults** to use the LAN-LAN tunnel defaults that are defined for tunnel endpoints in this group on the LAN-LAN Defaults Editor.
-

- 3 Change or enter any of the default values on the Policy tab.

Refer to the field definitions in [“To Set LAN-LAN Tunnel Defaults”](#) (p. 10-6).

END OF STEPS

.....

Task: complete the SLA Probe tab

- 1 Click the **SLA Probe** tab to display the SLA Probe tab.
-

Result The SLA Probe tab is displayed.

- 2 Refer to [“To Set Up Service Level Agreements” \(p. 10-32\)](#) for instructions on how to set up Service Level Agreements and configure the fields for the SLA Probe tab for a tunnel.

.....
E N D O F S T E P S
.....

Save and Apply the Tunnel

Once you have completed all tabs, you are ready to save and apply the tunnel. It is necessary to do this to make the tunnel operational. To save and apply the tunnel, follow the steps below:

- 1 From any tab in the LAN-LAN Tunnel Editor, display the File menu and select **Save and Apply**.

Result A dialog box is displayed, advising you that the **Save and Apply** action will recompile and apply the policy to the tunnel endpoint, and asks if you want to continue.

- 2 Click **OK** to proceed with recompiling and applying the policy to the LAN-LAN tunnel.

Result The LAN-LAN tunnel is created.

If any errors occur during processing, an error message dialog box is displayed with information about the error.

Sometimes a manually-keyed tunnel does not work properly if used right away, or if used immediately after rebooting one of the tunnel endpoints.

This is not necessarily a problem. It may simply be that one of the endpoints has not had an opportunity to “learn” about the other. The sequence of events is usually something like the following scenario:

- A Brick that has been freshly rebooted (or had its MAC table refreshed) is presented with a packet needing VPN tunneling.
- The Brick tunnels the packet, giving it a destination address (a tunnel endpoint) equal to some other Brick Virtual Brick Address (VBA).

- Because the Brick has not yet discovered what MAC address is associated with that VBA, it does not yet know where to send the packet. It issues an ARP to try to find that MAC address.
- Meanwhile (before the ARP has been answered), the Brick does not know where to send the packet. It creates an administrative event log entry saying it was unable to find a route for the packet and drops the packet.

Usually, all that is necessary to get the tunnel working is to wait until the ARP reply has been received and try again. TCP should automatically retransmit packets, so TCP users should not "see" the above behavior as a problem; administrators are most likely to perceive that a problem exists when using ping to verify that a tunnel is up— in which case the first few echo requests may fail. Subsequent attempts should nonetheless succeed.

END OF STEPS



To Set Up a LAN-LAN Tunnel with UDP Encapsulation

Overview

In addition to setting up LAN-LAN tunnels using pure IPsec as the transport mode, administrators have the option of configuring the Bricks in a LAN-LAN tunnel to transmit and receive IPsec packets that have been encapsulated inside UDP segments.

UDP Encapsulation

When UDP encapsulation is employed, the original IP packet, including the entire IP header, is included as UDP payload data. The Brick can exchange these packets bi-directionally with other Bricks, and with hosts running the Alcatel-Lucent IPsec Client. The Brick will encapsulate the packets when the Security Association specifies, and decapsulate the packets upon receipt from the network.

This feature is designed to allow VPNs to work in many different network topologies. For example, a teleworker at home could use it with a WAN connection (DSL/cable) with a single IP address, using a Brick (such as a Model 20) for tunnel termination between the hosts on the network and the WAN device.

NAT traversal

NAT Traversal is the standard for performing UDP Encapsulation between IKEv2 tunnel endpoints. NAT Traversal for IKEv2 works with Alcatel-Lucent IPsec Client software deployed by other vendors.

If the **IKEv2** checkbox is chosen as the **VPN Type** on the LAN-LAN Tunnel Editor, NAT Traversal is performed automatically, based on the IP addresses of the packets during the IKE negotiation process.

If NAT Traversal is enabled, the SMS automatically adds the NAT Traversal ports to the **UDP_Encapsulation_Ports** service group. This service group is only visible on the Brick, not via the SMS GUI. Port 4500 is added to the UDP Encapsulation Ports list (which is also used for negotiating IKEv1 tunnels) and is used for negotiating tunnels using IKEv2.

Set Up UDP Encapsulation

To set up UDP encapsulation, follow the steps below:

-
- 1 Set up the tunnel endpoints, just as you would any LAN-LAN tunnel. Enter all the information required in the Main tab to define both endpoints (refer to the [“Task: complete the Tunnel Tab”](#) (p. 10-14) for instructions).

-
- 2 Click **Parameters** to display the Parameters tab, and then click the checkbox labeled **Use LAN-LAN Default Parameters**. This allows you to use the information you are about to enter in place of the default parameters.

-
- 3 In the **Required IPSec Transport Method** box, click the radio button labeled **UDP Encapsulation Port** on the Parameters tab of the LAN-LAN Tunnel Editor.

Important! If you want all your LAN-LAN tunnels to use UDP encapsulation, you should change the default setting by clicking the **UDP Encapsulation Port** radio button on the LAN-LAN Defaults Editor instead of the LAN-LAN Tunnel Editor.

This way, every LAN-LAN tunnel you set up will automatically use UDP encapsulation, and you will not have to perform this procedure each time you set up a tunnel.

-
- 4 For UDP Encapsulation, the initiator destination port is 501.

For NAT Traversal, the initiator destination port is 4500.

However, when you clicked either **IKEv1** or **IKEv2**checkbox, a keyword called **UDP-Encapsulation-Ports** was added to the group in which the zone resides.

This keyword is automatically modified to include all the UDP destination ports specified in both the Client Tunnel Endpoint Editor and the LAN-LAN Tunnel Editor for all such entities in the zone.

END OF STEPS



Redundant LAN-LAN Tunnels

Redundant LAN-LAN tunnels defined

This feature, which is not intended for use with manual tunnels, provides the ability to provision a pair of LAN-LAN tunnels to improve network reliability. An example of the application of this feature would be to support a hub and spoke type of network where the hub is replicated at a different location so that, if the primary site is unavailable, the secondary site can take over the traffic.

To provision this capability in the SMS, simply create two tunnels from the same Tunnel EndPoint on a Brick to two different Remote EndPoints and provision the host access lists in both tunnels to be the same. Upon saving the second tunnel, you will get a warning message that this tunnel is redundant, and you must click **save anyway**.

From this point on, it will route traffic to the tunnel with the lowest SLA probe ID (if more than one applicable tunnel is up or if no tunnel is up) or to whichever tunnel is up.

Important! A tunnel is up if it has received a heartbeat or any other traffic within two times of the heartbeat interval.



Maintaining LAN-LAN Tunnels

Overview

Once a tunnel has been configured, you can view, modify and delete existing tunnels. You can also enable and disable the tunnels.

To view existing tunnels

The SMS provides these six viewers that you can use to view existing LAN-LAN tunnels. The table below indicates what each viewer shows.

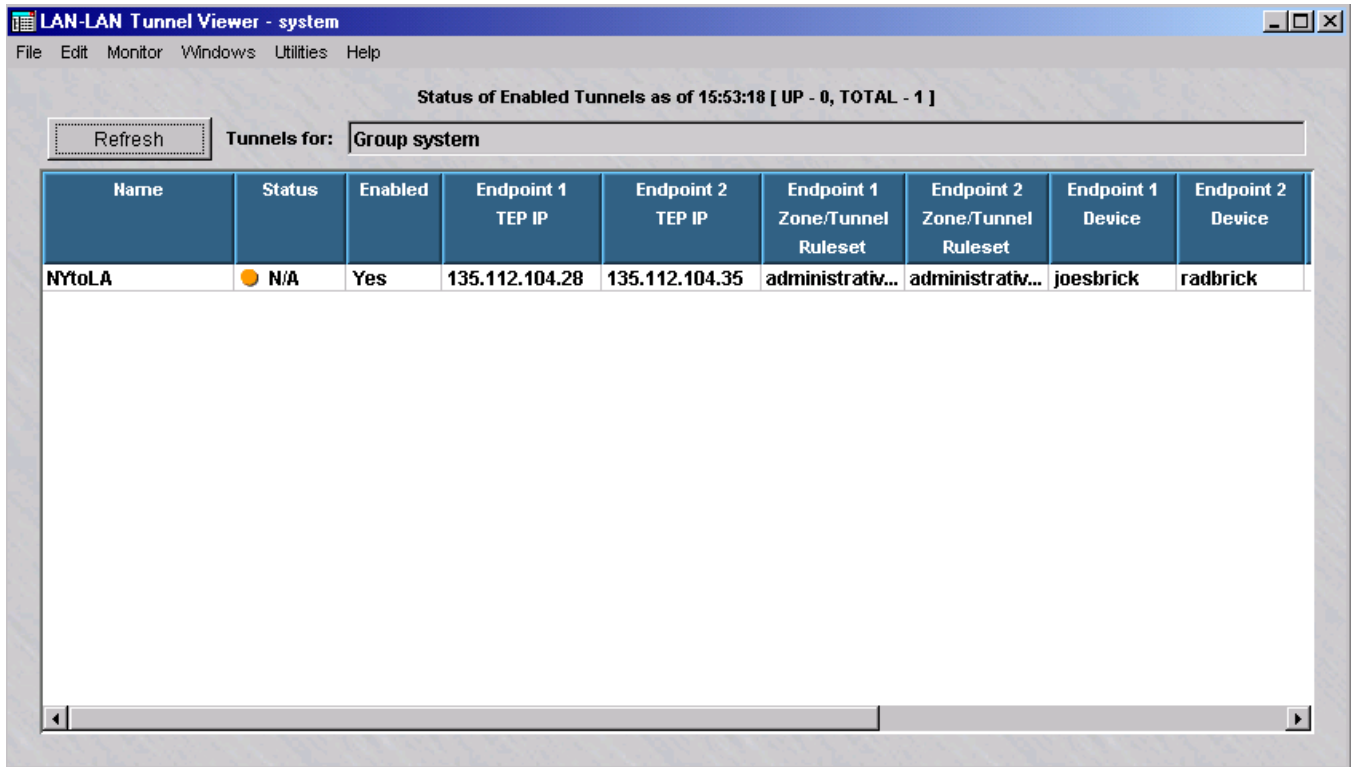
Viewer	Shows . . .
All Brick Tunnels	All tunnels terminating on Bricks in the current group
All Group Tunnels	All tunnels terminating on devices (Bricks) in the current group
Selected Device Tunnels	All tunnels terminating on the device you specify
Selected Folder Tunnels	All tunnels terminating on devices in the folder you specify
System Group Tunnels	All tunnels terminating in the System group with policies in the current group

To access one of these viewers, click the LAN-LAN Tunnel Viewers folder in the desired group to display the viewers in the Navigator window, and then double-click the viewer you want.

The viewer will appear. If you chose Selected Folder Tunnels or Selected Device Tunnels, a browse window will appear first to enable you to select the folder or device.

[Figure 10-5, “All Group Tunnels Viewer” \(p. 10-29\)](#) shows an example of the All Group Tunnels Viewer with one tunnel displayed.

Figure 10-5 All Group Tunnels Viewer



The **Status** column at the far left will show either *Up*, *Down*, *Manual* or *NA*. A status of *NA* means that the SMS is unable to contact the Brick at either end of the tunnel for status. The status is determined by means of periodic “heartbeats” exchanged by both Bricks. The frequency of the heartbeat is determined by the keepalive interval, which is set by default to 30 seconds, but can be changed by the administrator (refer to the description of the **Keepalive Interval (secs)** field in the section “[To Set LAN-LAN Tunnel Defaults](#)” (p. 10-6)).

To refresh the status, click the **Refresh** button.

To modify a tunnel

Complete the following steps to modify the configuration of an existing tunnel:

- 1 With a tunnel viewer open, right-click the tunnel you want to modify and select **Edit** from the pop-up menu. The LAN-LAN Tunnel Editor will appear, with the Tunnel tab displayed.

-
- 2 Make any necessary changes to the information in any of the tabs, including changing any default values set on the **SLA Probe** tab for the tunnel (for details about Service Level Agreement (SLA) probes, refer to the section [“To Set Up Service Level Agreements”](#) (p. 10-32)).
-

- 3 Display the File menu and select **Save and Apply**.

END OF STEPS

To delete a tunnel

Complete the following steps to delete the configuration of one or more existing tunnels:

-
- 1 With a tunnel viewer open, select the tunnel you want to delete.

To select multiple tunnels in sequence in the viewer, press the Shift or Ctrl key while clicking on each tunnel. To selectively choose multiple tunnels in the viewer (for example, the first and third tunnel in the listing), press the Shift or Ctrl key while clicking on each tunnel.
 - 2 Right-click and select **Delete** from the pop-up menu.
A confirmation window is displayed
 - 3 Click **OK** to delete the tunnel(s) and dismiss the Confirmation window.
The tunnel(s) are removed from the LAN-LAN Tunnel Viewer.

END OF STEPS

To enable/disable a tunnel

To enable a tunnel, right click the tunnel in a viewer and select **Enable** from the pop-up menu. To disable the tunnel, right-click the tunnel and select **Disable**. A confirmation window is displayed, asking if you really want to enable or disable the tunnel. Click **OK** to enable or disable the tunnel.

To select multiple tunnels in sequence in the viewer, press the Shift or Ctrl key while clicking on each tunnel. To selectively choose multiple tunnels in the viewer (for example, the first and third tunnel in the listing), press the Shift or Ctrl key while clicking on each tunnel.

If you disable a previously enabled tunnel, the entry in the **Enabled** column will change from **Yes** to **No**. However, the entry in the **Status** column will not change to **Down** until you refresh the tunnel viewer by right-clicking the tunnel in the viewer and selecting **Refresh Tunnels**.

Tunnel status is not automatically updated in a viewer, so it is a good idea to refresh the viewer periodically to ensure that the status of the tunnels shown is up-to-the-minute.



To Set Up Service Level Agreements

Service level agreements (SLAs) defined

Service Level Agreements (SLAs) are generally established between Service Providers and their customers. They outline the types of service to be furnished by the provider and the expected levels of reliability and bandwidth to be maintained.

Generally, SLAs have clauses that specify actionable conditions on the part of the customer should the service provider not live up to the conditions set therein. As such, Service Providers must keep records of the actual measured service available over time, to prevent such actions on the part of their customers.

SLA probes on the Brick can test the following metrics for VPN tunnels:

- Round-trip latency measurements
- Packet loss measurements

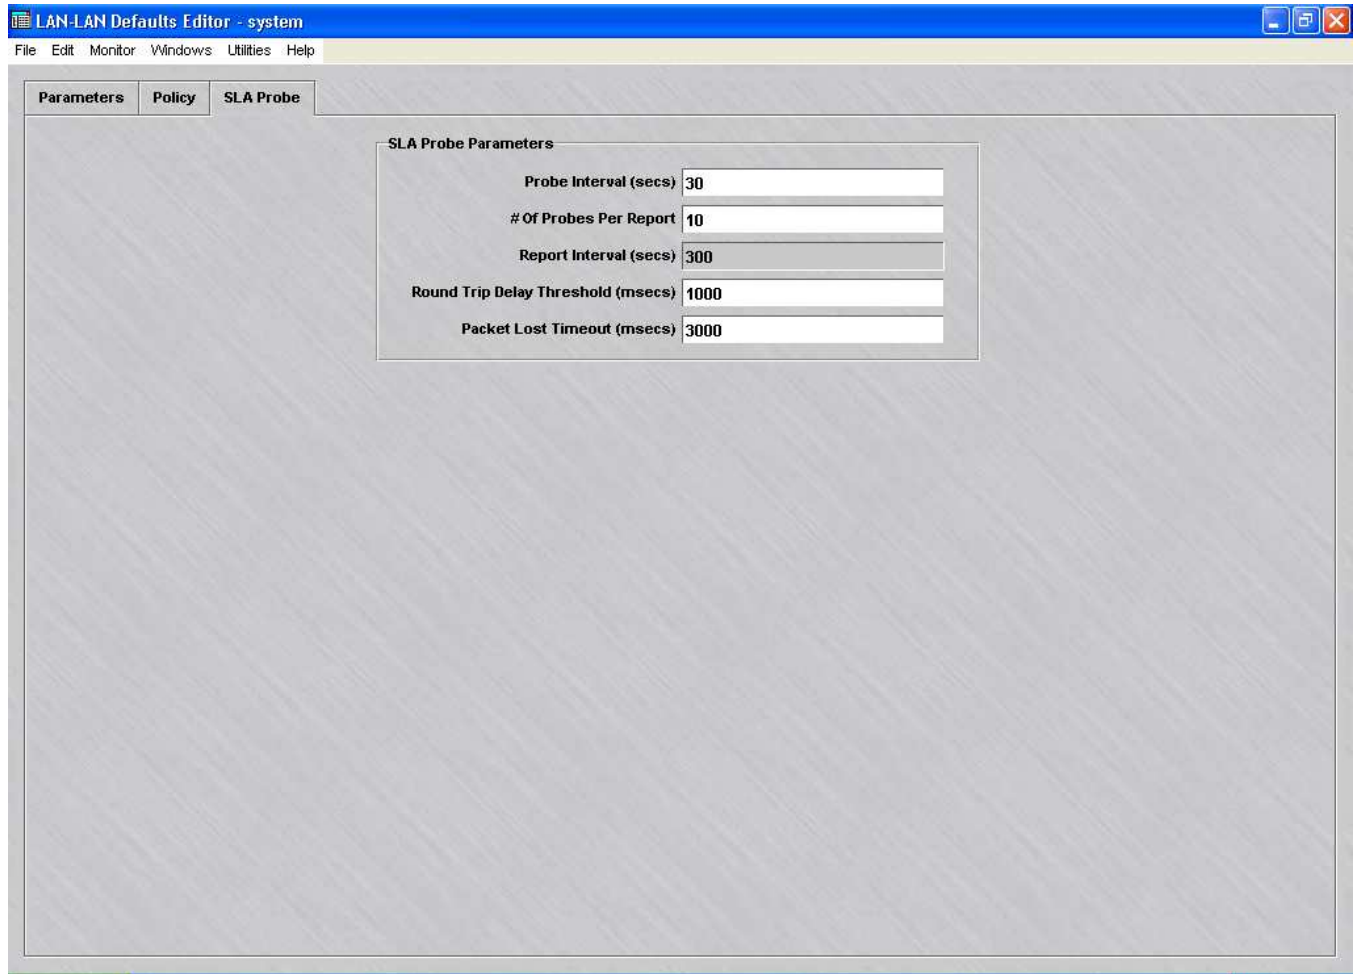
Any or all of these may be in place for a given customer of a Service Provider. SLAs are particularly well suited to work with the QoS/bandwidth management feature. This feature creates limits and guarantees on session and packet throughput, and can report when those limits and guarantees are not met.

Set the Default SLA Parameters

The SMS application comes with certain SLA default parameters set for each LAN-LAN tunnel. These parameters are defined and configurable on the **SLA Probe** tab of the LAN Tunnel Defaults Editor (refer to [Figure 10-6, “Default SLA Parameters”](#) (p. 10-33) for a sample of this tab).

SLA Probe defaults can be set for each Group. For each LAN-LAN tunnel, you can choose to apply the LAN-LAN Tunnel Group Defaults, or set unique values for the tunnel.

Figure 10-6 Default SLA Parameters



You can leave these defaults in place, or you can change the probe interval (default = 30 seconds), the report interval (10 probes), the round trip delay threshold (1000 milliseconds) or the packet lost timeout period (3000 milliseconds).

Set up SLA Probes for a Specific Tunnel

When setting up a LAN to LAN tunnel, you can use the default SLA parameters, or you can change the parameters for this particular tunnel.

Complete the following steps to set up SLA Probes for a specific tunnel.

- 1 Display the LAN-LAN Tunnel Editor for the tunnel (either when you are creating the tunnel, or after the tunnel has been set up) and click the **SLA Probe** tab. The display shown in will appear.

Result The SLA Probes tab is displayed (Figure 10-7, “LAN-LAN Tunnel Editor (SLA Probes tab)” (p. 10-34)).

Figure 10-7 LAN-LAN Tunnel Editor (SLA Probes tab)

The screenshot shows the 'LAN-LAN Tunnel Editor - Endpoint 1 <=> Endpoint 2' window with the 'SLA Probe' tab selected. The interface includes a menu bar (File, Edit, Monitor, Windows, Utilities, Help) and a tabbed interface with 'Endpoints', 'Parameters', 'Authentication', 'Policy', and 'SLA Probe' tabs. The 'SLA Probe' tab contains the following elements:

- An unchecked checkbox labeled 'Enable SLA Probe'.
- A section titled 'SLA Probe IP Addresses' containing two text input fields: 'Probe Source IP Address' and 'Probe Destination IP Address'.
- A section titled 'SLA Probe Parameters' containing:
 - A checked checkbox labeled 'Use LAN-LAN Default SLA Probe Parameters'.
 - Five text input fields with the following values: 'Probe Interval (secs)' (30), '# Of Probes Per Report' (10), 'Report Interval (secs)' (300), 'Round Trip Delay Threshold (msecs)' (1000), and 'Packet Lost Timeout (msecs)' (3000).

- 2 Click the **Enable SLA Probe** checkbox to set up the SLA probe feature for this tunnel.
- 3 Enter the probe source and destination IP addresses. The destination field has a drop-down to select TEP (Tunnel End Point) as the destination address.

-
- 4 If you want to use the default SLA parameters, leave the **Use LAN-LAN Default SLA Probe Parameters** checkbox checked.

If you want to change one or more of the default parameters for this tunnel, uncheck the checkbox and change the appropriate parameter(s). Unless you change the defaults, these parameters will only apply to this tunnel.

END OF STEPS

View SLA Probe Data

To view SLA probe data for a given tunnel, you have to display the LAN-LAN tunnel in one of the LAN-LAN tunnel viewers (see [“To view existing tunnels”](#) (p. 10-28) above). Once this is done, display either the round trip delay statistics for this particular tunnel or the round trip delay statistics for all the tunnels in the group. To do this, either:

- Right-click the tunnel and select **Show Tunnel Round Trip Delay Statistics** or **Show Group Round Trip Delay Statistics** from the pop-up menu, or
- Click the **Show Tunnel Round Trip Delay Statistics** or **Show Group Round Trip Delay Statistics** button at the bottom of the view (see [Figure 10-5, “All Group Tunnels Viewer”](#) (p. 10-29)).

[Figure 10-8, “Tunnel Round Trip Delay Statistics”](#) (p. 10-36) shows a typical tunnel round trip delay statistics display, and [Figure 10-9, “Group Round Trip Delay Statistics”](#) (p. 10-37) shows a group round trip delay statistics display.

Figure 10-8 Tunnel Round Trip Delay Statistics

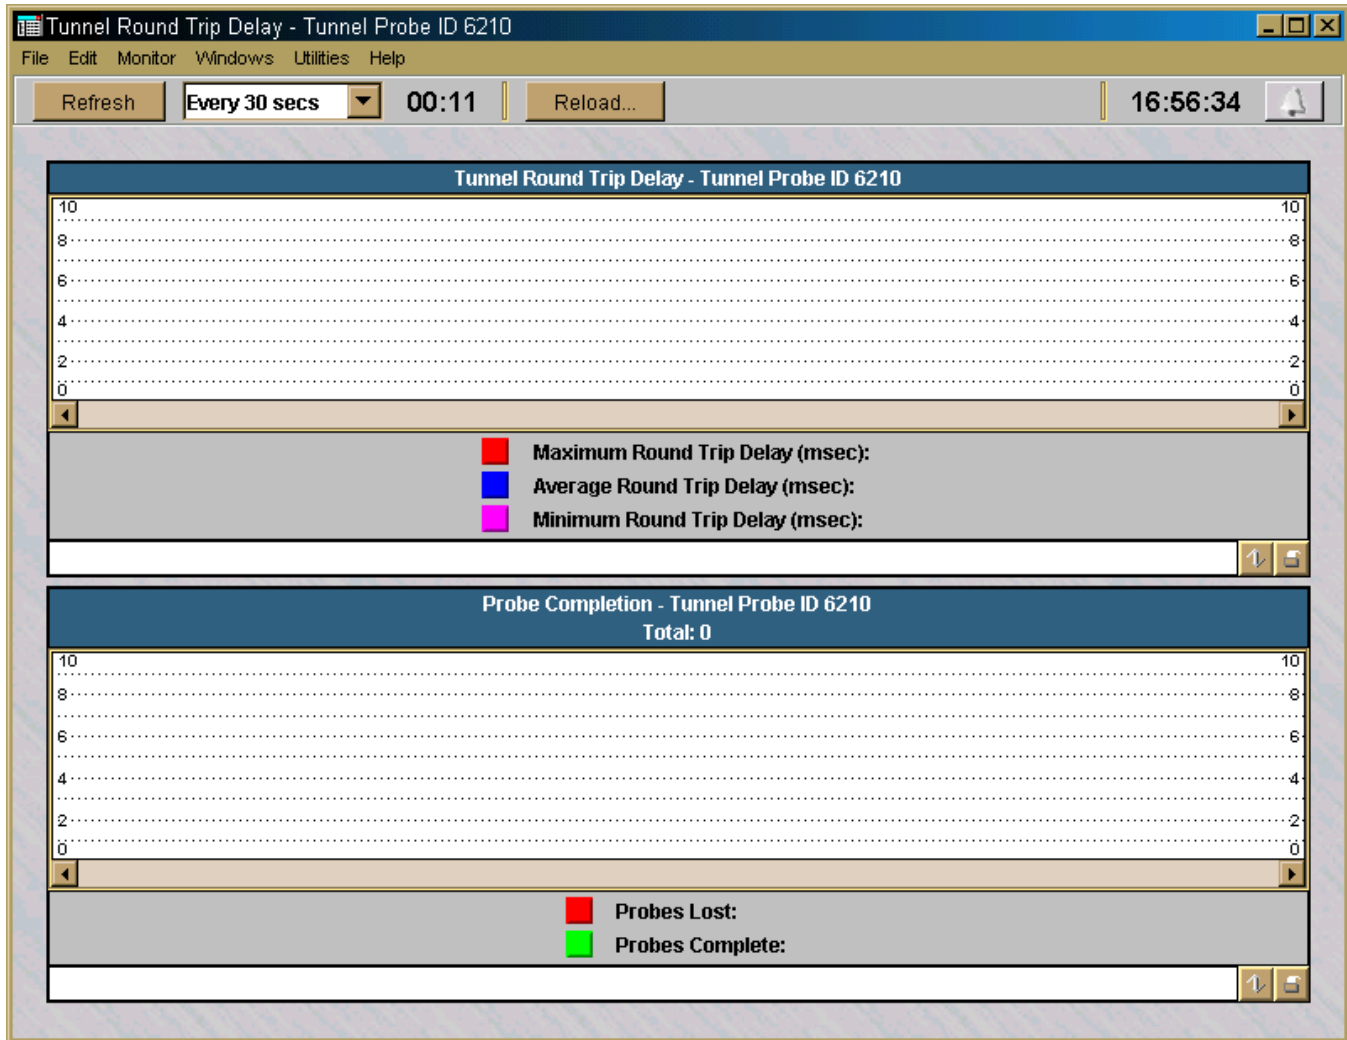
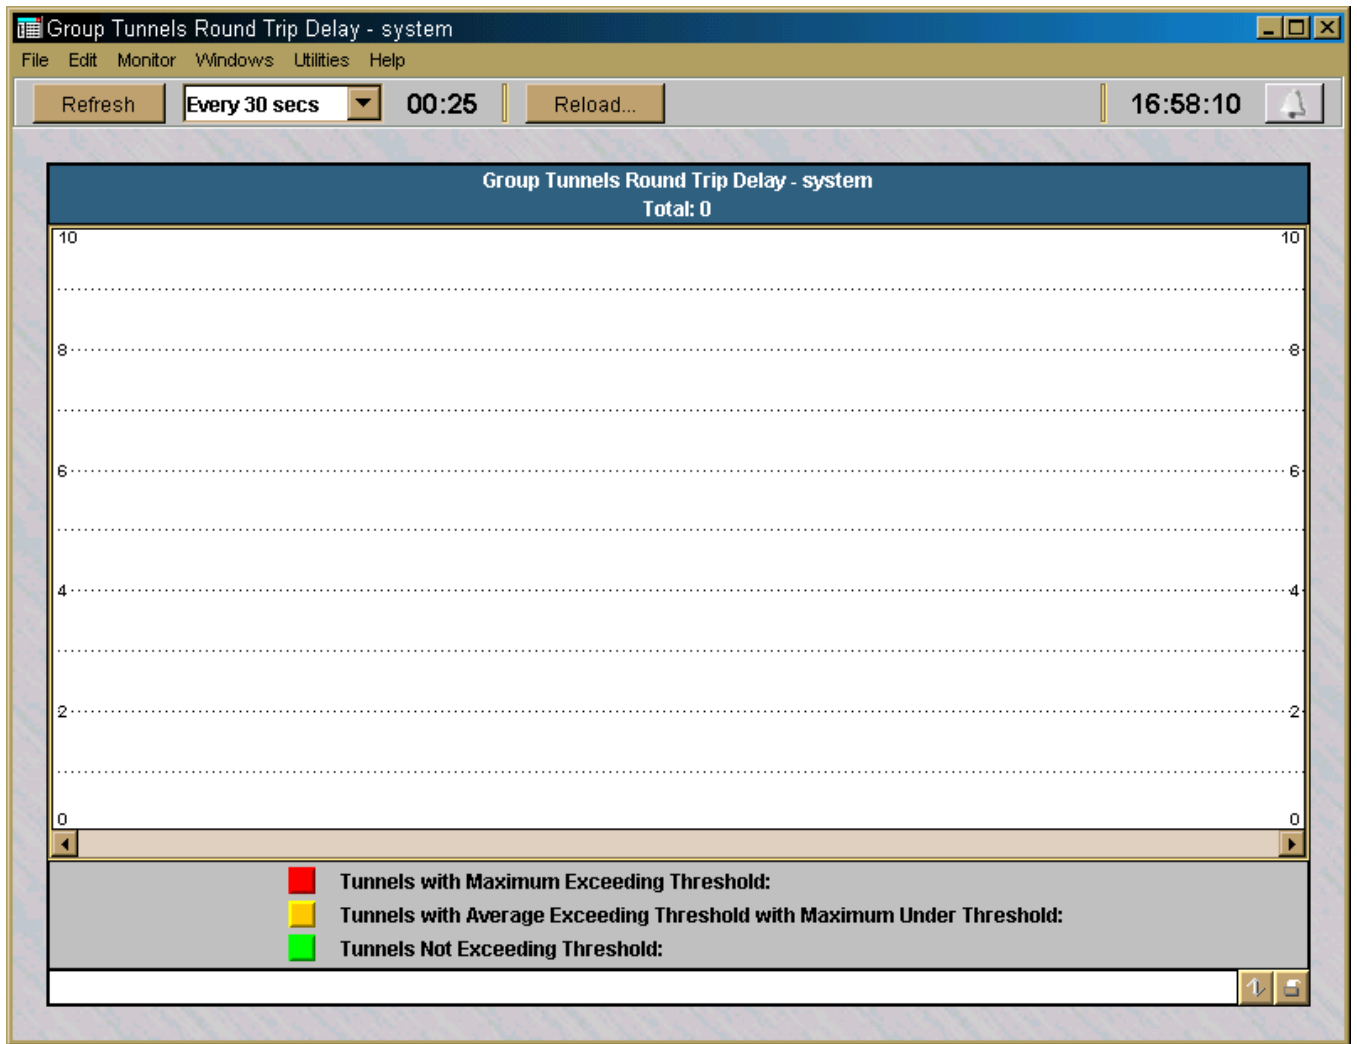


Figure 10-9 Group Round Trip Delay Statistics



□

11 Client Tunnel Endpoints

Overview

Purpose

This chapter explains how to set up a client tunnel endpoint so that users of the Alcatel-Lucent IPsec Client application (or a compatible IPsec Client application) can establish a tunnel between their hosts and an Alcatel-Lucent *VPN Firewall Brick*TM Security Appliance. The chapter also explains how to encapsulate IPsec packets within UDP packets in the tunnel.

This tunnel will allow the client user to communicate securely with another computer over the public Internet.

Contents

What is a Client Tunnel?	11-2
To Set the Client Tunnel Defaults	11-8
To Set Up a Client Tunnel Endpoint	11-24
To Set Up a Client Tunnel with UDP Encapsulation	11-33
What to Do Next	11-35
Maintaining Client Tunnel Endpoints	11-40
To Create a Message for Alcatel-Lucent IPsec Client Users	11-46



What is a Client Tunnel?

Definition

A client tunnel is a Virtual Private Network (VPN) or encrypted path through the Internet. The endpoints of the tunnel are a host running the Alcatel-Lucent IPsec Client and a port on a Brick.

Tunnel endpoints


The client user endpoint is represented by the IP address of the host running the Alcatel-Lucent IPsec Client. This can be either a fixed IP address or an IP address temporarily assigned by an Internet Service Provider (ISP).

The other endpoint is a port on a Brick, represented by the tunnel endpoint address (Virtual Brick Address) assigned to that port on the Policy Assignment tab of the Brick Editor. Refer to the *Configuring Alcatel-Lucent VPN Firewall Brick® Security Appliance Ports* chapter in the *SMS Administration Guide*.

Scenarios

The diagram in [Figure 11-1, “Client Tunnels” \(p. 11-3\)](#) illustrates the following scenario.

Figure 11-1 Client Tunnels



In Figure 11-1, “Client Tunnels” (p. 11-3), the endpoints of the tunnel are a mobile user laptop computer running the Alcatel-Lucent IPsec Client and a Brick device. The mobile user accesses the Internet by connecting to a local ISP and uses the tunnel to communicate securely with servers on the corporate LAN.

The encrypted traffic travels from the clients to the Brick device, where it is decrypted and sent in the clear to the designated host(s) behind the Brick device. The return traffic is then encrypted back to the clients.

Internet key exchange version 2 (IKEv2)

The SMS and managed Brick devices, as of Release 9.0, implement the Internet Key Exchange, Version 2 protocol (IKEv2) as defined in Internet Engineering Task Force (IETF) draft-ietf-ipsec-ikev2-17, including support for NAT Traversal and Extensible Authentication Protocol (EAP), in addition to IKEv1 protocol.

Client Tunnel Endpoints can now be configured to accept connections from both IKEv1 and IKEv2 IPsec Clients. Key and X.509 Certificate authentication are supported.

Multiple Certificate Authorities (CAs) for client tunnel endpoint access

Some customers have more than one Certificate Authority (CA) signing certificates that connect to a client tunnel endpoint. There may also be a separate CA used by the customer for the Brick certificates. The SMS allows you to select and assign multiple CAs to the same client tunnel endpoint for multiple groups of IPSec clients when certificates for those clients or Bricks are administered by separate organizations that each have their own CA.

The CA(s) chosen for a client TEP are included in the compiled policy downloaded to the Brick. If a CA is an intermediate CA, the CAs in the chain, up to and including the root CA, are included automatically.

Multiple SA proposals

The Multiple SA proposals feature allows you to create two separate IPSec Security Association (SA) proposals, with different security parameters, for Alcatel-Lucent IPSec Client tunnel traffic passing through a Brick tunnel endpoint (port). The SA proposal that is applied is based on the IP address of the host (client tunnel endpoint) and TCP/UDP port(s) of the incoming/outgoing data packet. This feature can be used, for example, to apply one SA proposal, with little or no encryption (and thereby, allowing quicker transit), to less confidential data from one group of IP addresses and UDP ports, and apply another SA proposal, with strict encryption, to more secure traffic from another group of IP addresses and UDP ports.

This feature can only be used if IKEv2 is enabled for the Brick tunnel endpoint.

When both SA proposals apply to a particular packet, the Brick chooses the narrower of the two SAs when encrypting the packet.

A Brick that does not currently support this feature can be deployed in a failover pair with a Brick that supports this feature, as long as a single SA proposal is provisioned and the protocol and ports fields are configured with an asterisk (*) to accept any value.

If the **Receive Any Proposals** checkbox on the Client Defaults Editor or Client Tunnel Endpoint Editor is checked, and the tunnel encryption policy sent by the Alcatel-Lucent IPSec Client host does not match either SA proposal, the Brick uses the first set of traffic selectors and IPSec SA Proposal 1 for encryption and handling.

IKE hardware acceleration

Some IKE (v1 and v2) operations are now executed with hardware acceleration on the following Brick models:

- Model 50 Brick
- Model 150 Brick
- Model 700 Basic and 700 VPN Bricks

- Model 1100 and 1100A Bricks
- Model 1200 and 1200 HS Bricks

IKE hardware acceleration greatly increases the number of tunnels per second that can be established on these devices as well as reducing the overall load on Brick devices.

Packet data gateway (PDG) accounting

The Packet Data Gateway (PDG) accounting feature is an optional feature that can be enabled for a client tunnel endpoint. When this feature is enabled, the SMS and Brick that supports the client tunnel endpoint provide interim accounting data updates, at configurable intervals, for all traffic that passes through the client tunnel endpoint in an active client tunnel. The collection interval for this accounting data can be based on time (minutes), traffic volume (number of bytes), or both. This PDG accounting function is a key requirement of the IP Multimedia Core Network subsystem (IMS) solution, to gather traffic data whenever a user with a Dual Mode Handset (DMH) establishes a secure connection with a Brick device, using IKEv2 EAP to access Packet Switched services over the converged network such as Multimedia Messaging Service (MMS), Wireless Access Protocol (WAP), Internet access, Voice Over Internet Protocol (VoIP), and Peer-to-Peer applications.

DMH users are authenticated and given permission to access the converged network and its suite of services via a RADIUS server and its authentication and access control protocols.

The PDG Accounting feature is an optional feature that must be purchased and installed using an installation key via the New Feature Setup utility. If the PDG Accounting feature has not been installed, all the fields on the PDG Accounting tab are greyed out. For details about the New Feature Setup utility, refer to the *SMS Administration Guide*.

PDG lawful intercept

The PDG Lawful Intercept feature, which is related to the PDG Accounting feature, allows a law enforcement agency to monitor an encrypted client tunnel session by forwarding an unencrypted copy of the data traffic to a designated IP address and port.

If the Lawful Intercept feature is enabled, the SMS informs the AAA server during the user authentication process. If law enforcement has provisioned monitoring for that user, the AAA server will include a lawful intercept request in the authentication response, which specifies an IP address and port where intercepted data should be forwarded.

To turn lawful intercept on or off during an active session, law enforcement provisions a request, which is forwarded to the AAA server.

The next time the AAA server receives a PDG Accounting update message from the SMS for that client session, the AAA will include a lawful intercept request in the accounting response. The intercept request will either include the forwarding IP address and port to turn the intercept on, or include an indicator to turn the intercept off. The SMS forwards this information to the Brick device in the PDG accounting response.

The PDG Lawful Intercept feature is an optional feature that is enabled for a client tunnel endpoint via a checkbox on the PDG Accounting tab of the Client Tunnel Endpoint Editor. The PDG Accounting feature must be purchased and installed first using an installation key before the Lawful Intercept feature can be installed. The PDG Lawful Intercept feature must be purchased and installed separately using a installation key via the New Feature Setup utility to activate the checkbox on the PDG Accounting tab for enabling or disabling the feature.

For details about the New Feature Setup utility, refer to the *SMS Administration Guide*.

Client tunnel defaults

Before you set up a client tunnel, you should examine the client tunnel defaults provided with the SMS application to determine whether or not any of them need to be modified.

The SMS application comes configured with certain tunnel parameters already set. These default settings enable a Brick device and a client to negotiate IPsec tunnel parameters and exchange encryption and authentication keys.

You can use these default settings as they are, or you can change them to suit your own specific needs. Moreover, the settings are group-specific. If you have created groups other than the *system* group, you can leave the defaults in place in one group, but change them in another — or change them differently in both groups.

For this reason, it is recommended that before you attempt to set up a tunnel, you take a look at the default parameters that apply to the group you are using and determine whether to keep them or change them.

Host groups and client tunnels

Before setting up a client tunnel, it is also recommended that you create a host group to identify the hosts behind the tunnel that the client users will be permitted to access. Client users can access all of the hosts in a zone assigned to a Brick device port. If you want to restrict access to specific hosts in a zone, you need to create a host group containing the appropriate IP addresses.

Authentication service

You may also need to create an authentication service. If you are using Local Password to authenticate client users, you do not have to create an authentication service, because one is provided with the SMS application. See [Chapter 8, “User Authentication”](#) for a more detailed discussion of user authentication.

If you are using RADIUS, SecurID, or a VPN certificate, an authentication service is required to set up the tunnel.



To Set the Client Tunnel Defaults

When to use

Use this task to configure default settings for client tunnels.

Task

Complete the following steps to set client tunnel defaults.

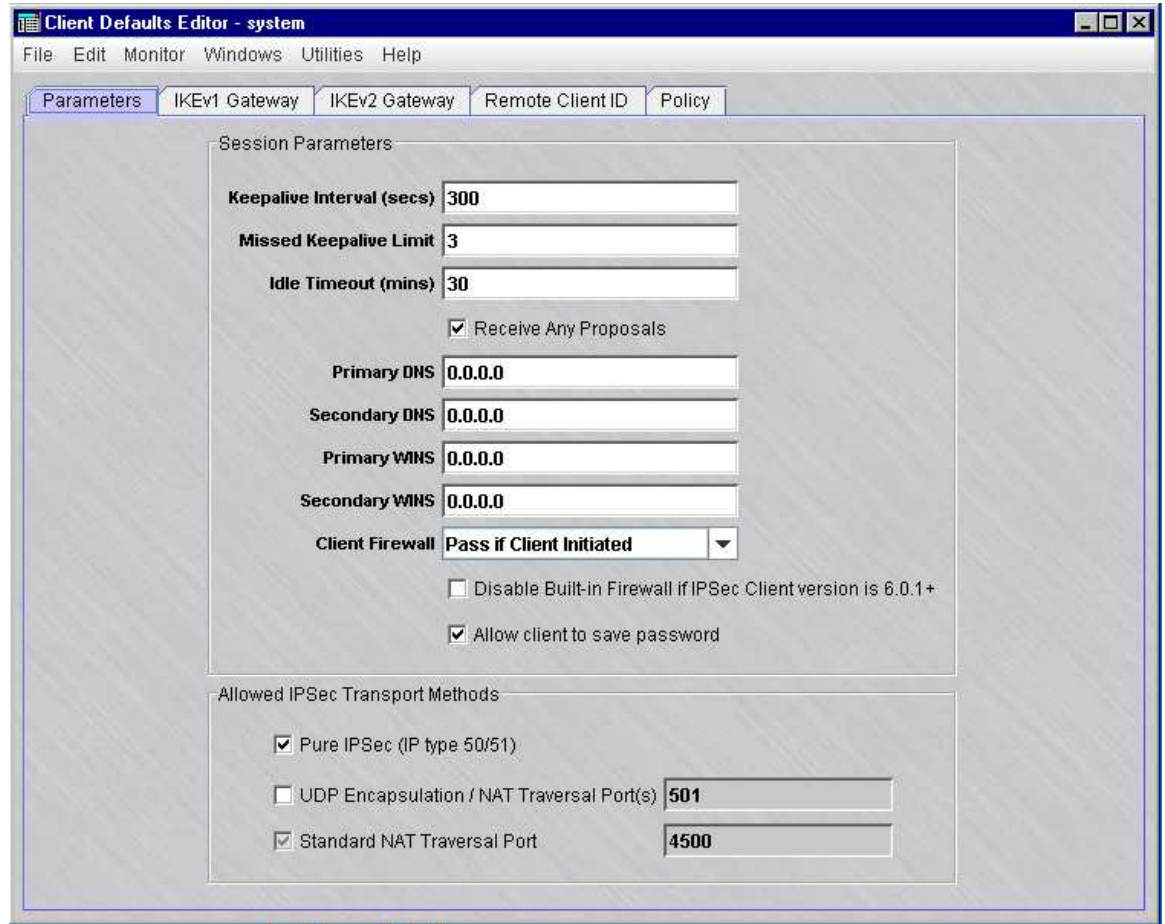
-
- 1 With the Navigator window displayed, open the appropriate group folder, and then open the **VPN** folder.

 - 2 Click **VPN Defaults** to display entries for the two defaults editors in the Contents panel.

 - 3 Double-click **Client Defaults**.

Result The Client Defaults Editor screen is displayed. This screen consists of multiple tabs. The Parameters tab is initially displayed. This display is shown in [Figure 11-2, “Client Defaults Editor \(Parameters Tab\)”](#) (p. 11-9).

Figure 11-2 Client Defaults Editor (Parameters Tab)



- 4 Enter or change the default value(s) of any of the following fields in the Parameters tab:

Field	Explanation
Keepalive Interval (secs)	<p>The Alcatel-Lucent IPsec Client is programmed to send out a keepalive/heartbeat message at regularly scheduled intervals to verify that it still has network connectivity to the other end of the tunnel.</p> <p>The Brick device that receives the message generates a message back to the client.</p> <p>The default is 300 (seconds). The valid range is 0-172800.</p> <p>An entry of zero (0) disables keepalive/heartbeat messages. <i>Note: If the Keepalive Interval is set to a value less than 30 seconds, it can adversely impact Brick performance as the number of IPsec clients increases.</i></p>
Missed Keepalive Limit	<p>The number of missed heartbeat messages that the Brick and IPsec Client will allow before determining that the other end is not responding and the tunnel is taken down.</p> <p>The default is 3 (missed heartbeat messages). The valid range is 3-999.</p>
Idle Timeout (mins)	<p>A client tunnel will time out after a specified period of time if there is no activity in the tunnel in either direction.</p> <p>The default is 30 minutes. The valid range is 1-2880 minutes.</p>
Receive Any Proposals	<p>If this box is unchecked, the client parameter and policy settings must match the TEP settings defined in the Editor exactly for the tunnel to come up. If this checkbox is checked, the parameter and policy settings defined in the Editor are the preferred settings, but the Brick accepts any combination of settings that the client proposes, provided the Brick supports those options. The box is checked, by default.</p>

Field	Explanation
Primary/Secondary WINS/DNS	If your environment includes DNS and/or WINS servers, you can enter the IP addresses of these servers in the appropriate fields. The default in each field is 0.0.0.0.
Client Firewall	<p>This field determines whether packets to and from the client that are <i>not</i> going through the tunnel will be passed or dropped. There are three options:</p> <ul style="list-style-type: none"> • Pass All (allows the packets through the Brick) • Drop All (allows <i>no</i> packets through the Brick) • Pass if Client Initiated (sessions started on the Client will be allowed out, but no <i>new</i> sessions allowed in; supported by v3.1 of the Client and above) <p>If you leave the default asterisk in the Hosts Behind Tunnel field when setting up a client tunnel, all traffic will automatically go through the tunnel. In this case, it does not matter what you enter in this field.</p> <p>The value of this field automatically overrides the firewall setting in the Alcatel-Lucent IPsec Client, if the two settings are different.</p>
Disable Built-in Firewall if IPsec Client version is 6.0.1+	Placing a check in this checkbox will disable the built-in firewall. This feature only works if using Alcatel-Lucent IPsec Client V6.0.1 or higher. Note: Disabling the firewall also removes the "Allow Multi-session Protocol" feature from the client's GUI.
Allow Client to Save Password	<p>The Allow Client to Save Password checkbox is a convenience for the Alcatel-Lucent IPsec Client user.</p> <p>By default, this box is checked. This means client users have the option of saving their password the first time they enable a tunnel, so that they do not have to enter it again each time they enable the tunnel.</p> <p>If you uncheck this box, client users will not have this option, and will have to enter their password each time they enable a tunnel.</p>

Field	Explanation
Allowed IPSec Transport Methods	<p>These checkboxes determine the transport method. The choices are:</p> <ul style="list-style-type: none"> • Pure IPSec (IP Type 50/51) • UDP Encapsulation Port(s)(for IKEv1 TEPs only) • NAT Traversal Port (for IKEv2 TEPs only)¹ <p>By default, both the Pure IPSec (IP Type 50/51) and NAT Traversal Port (for IKEv2 TEPs only) are checked. More than one method can be selected. Refer to “To Set Up a Client Tunnel with UDP Encapsulation” (p. 11-33) for a more detailed discussion of UDP encapsulation.</p> <p>If NAT Traversal Port or UDP Encapsulation Port(s) is selected, SMS automatically adds the NAT Traversal ports to the UDP_Encapsulation_Ports service group. This service group is only visible via the Brick console.</p> <p>If NAT Traversal Port or UDP Encapsulation Port(s) is selected, SMS automatically inserts Rules 190, 191, 192, and 193 into the zone assigned to the Client Tunnel TEP ruleset. These rules allow VPN traffic to pass into and out of the Zone on the specified UDP/NAT-T ports.</p>

Notes:

1. The NAT Traversal Port checkbox is a read-only field (cannot be modified).

-
- 5 Click **IKEv1 Gateway** or **IKEv2 Gateway** to configure the default authentication settings for client tunnels using the IKEv1 or IKEv2 key negotiation method, respectively.

Result The IKEv1 Gateway or IKEv2 Gateway tab is displayed, depending on the tab selected.

-
- 6 Choose the default client TEP authentication method.

The choices are:

- **Preshared Key.** This is the default for IKEv1 and IKEv2 key negotiation. If the Unlicensed Mobile Access (UMA) feature is enabled, the associated **Preshared Key** and **Group ID** selection fields for this choice are *not* displayed on the IKEv1 Gateway tab. If you choose the **Preshared Key** option, go to [Step 7](#).
- **Get Preshared Key from RADIUS.** This checkbox is only displayed on the IKEv1 Gateway tab if the UMA feature is enabled. To enable this authentication method, click the checkbox to place a check in the box. To disable this method, click the checkbox again to remove the check.
- **X.509 Certificate.** This choice is only displayed on the IKEv2 Gateway tab. If you choose the **X.509 Certificate** option, go to [Step 7](#).

7	If	Then
	<p>Preshared Key was selected on the IKEv1 Gateway tab in Step 6</p>	<p>Use the default Preshared Key and Group ID assigned by the system, or change these field values.</p> <p>Click the Generate Key button to generate a random 20-character Preshared Key. A default group key is generated randomly by the system. For security reasons, the key does not appear on screen, but is masked by a series of asterisk.</p> <p>If you are not using digital certificates to authenticate client users, this key is used by all client users in the group when they set up a tunnel from their PCs to the TEP.</p> <p>To change the key, click the Generate Key key, or manually enter a key in the Preshared Key field. the new key must be 8-20 characters. Valid characters include a-z, A-Z, 0-9, and the following special characters: : ; + ? " () < > ^ % \$ # &. Click the Unmask checkbox to view the hidden key.</p> <p>Enter a Group ID in the Group ID field. The Group ID is only used by non-Alcatel-Lucent IPsec Client programs. It is used in conjunction with the Preshared Key for the first phase of the IKE negotiation process. If you are using the Alcatel-Lucent IPsec Client, you can ignore this field.</p> <p>The default value for this field is gatewaygroupID. If you are using a non-Lucent client, keep the default value or change it. The group ID entered must be used when configuring the non-Lucent client.</p>

If	Then
<p>Preshared Key was selected on the IKEv2 Gateway tab in Step 6</p>	<p>Use the default Preshared Key assigned by the system, or change this field value.</p> <p>Click the Generate Key button to generate a random 20-character Preshared Key. To change the key, click the Generate Key key, or manually enter a key in the Preshared Key field. the new key must be 8-20 characters. Valid characters include a-z, A-Z, 0-9, and the following special characters: : ; + ? " () < > ^ % \$ # &.</p> <p>Click the Unmask checkbox to view the hidden key.</p> <p>Click the down arrow next to the ID Type field to display a list of choices and select the Identity that the Brick uses for the TEP during IKEv2 negotiations. The choices are: IP Address(default), Email Address, and Domain Name.</p> <p>If IP Address is chosen as the ID Type, click the down arrow next to the ID field and select Virtual Brick Address (the default) or enter the IP address of the TEP.</p> <p>If Domain Name is chosen as the ID Type, enter a Fully Qualified Domain Name (example: salesdept.mycompany.com) in the ID field that is associated with the IP address assigned to the TEP.</p> <p>If Email Address is chosen as the ID Type, enter the e-mail address of the client.</p>
<p>X.509 Certificate was selected on the IKEv2 Gateway tab in Step 6</p>	<p>Click the down arrow next to the ID Type field to select an Identity used by the Brick during IKEv2 negotiations. The choices are: Distinguished Name (the default), IP Address, Email Address, and Domain Name.</p>

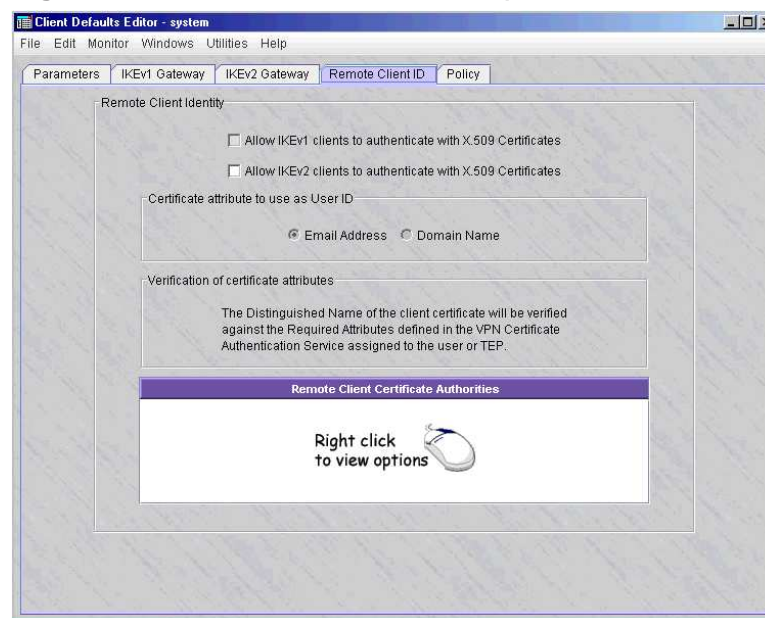
-
- 8 If you are on the IKEv2Gateway tab, choose a method for the Brick device to obtain the IKEv2 Client EAP Identity. If **Send EAP Identity Request** is chosen (the default), the Brick device initiates a EAP ID Request message to the Client to obtain the Client ID. If **Use Client IKE ID** is chosen, the Brick device uses the Client ID obtained in a message response from the Client during IKE negotiation.

Important! For UMA applications, the **Use Client IKE ID** method should be chosen.

- 9 Once you have completed the fields on the IKEv1 Gateway or IKEv2 Gateway tab, click **Remote Client ID**.

Result The Remote Client ID tab is displayed ([Figure 11-3, “Client Defaults Editor \(Remote Client ID Tab\)”](#) (p. 11-16)).

Figure 11-3 Client Defaults Editor (Remote Client ID Tab)



-
- 10** Use the system defaults or change the values of the following fields, as needed:
- **Allow IKEv1 clients to authenticate with X.509 Certificates**— this checkbox is not editable and is automatically checked if a VPN Certificate is assigned to the TEP.
 - **Allow IKEv2 clients to authenticate with X.509 Certificates**— click this checkbox to allow TEPs using IKEv2 key negotiation to respond to VPN Certificate authentication requests. This checkbox is unchecked, by default.
 - **Certificate attribute to use as User ID**— choose **Email Address** or **Domain Name** as the field of the certificate to be used by the Brick as the User ID for authentication. The default value is **Email Address**.
 - **Remote Client Certificate Authorities** (Selection of Multiple Certificate Authorities for a Client Tunnel Endpoint)— this option allows you to select and assign multiple Certificate Authorities (CAs) for a Client TEP. This option allows customers to configure a single Client TEP to provide access for multiple groups of IPSec Clients when certificates for those groups, or the Bricks, are administered by separate organizations that each have their own CA.
To use this option, right-click below the **Remote Client Certificate Authorities** label and choose **New** from the pop-up menu. The Remote Client CA Editor is displayed. Select a CA from the list and click **OK**. The selected CA is added to the table of allowed CAs on the Remote Client ID tab. To remove a CA from the list, right-click on it in the table below the **Remote Client Certificate Authorities** label and choose **Delete** from the pop-up menu. A confirmation dialog box is displayed, asking if you are sure that you want to delete the selected CA from the list. Choose **Yes**. The CA is removed from the list of allowed CAs.

The Distinguished Name of the client certificate will be verified against the Required Attributes defined in the VPN Certificate Authentication Service assigned to the User or TEP.
-
- 11** Click **Policy** to display the Policy tab.

Result The Policy Tab is displayed (Figure 11-4, “Client Defaults Editor (Policy Tab)” (p. 11-18)).

Figure 11-4 Client Defaults Editor (Policy Tab)

The Policy tab contains the IKE SA Proposal and IPsec SA Proposal defaults for IPsec SA Proposal 1 and IPsec SA Proposal 2 (if the Multiple SA Proposals feature is used).

The IKE SA Proposal is used to authenticate IPsec peers and negotiate IKE SAs to set up a secure channel for negotiating IPsec SAs in Phase 2.

IPsec SA Proposal 1 and IPsec SA Proposal 2 (if it is enabled) are used to negotiate IPsec SA parameters (such as Encryption and Auth Type) and to set up matching IPsec SAs in the peers.

- 12 The fields in the **IPSec SA Proposal 1** box are populated with default values. You can change any of these values. The following explains:

Field	Explanation
Protocol	The default is ESP-50 , but this can be changed to AH-51 . ESP-50 provides both encryption and authentication for every packet, while AH-51 only provides authentication.
Encryption Type	The options are TRIPLE DES CBC (default), DES CBC , AES CBC 128 , AES CBC 192 , AES CBC 256 , and NULL (Null Encryption).
Auth Type	The choices are: HMAC SHA1 (the default), AES XCBC (valid only if the Protocol selection is ESP-50 and for the IKEv2 client TEP setting), ANY NOT NULL (valid only if the Protocol selection is ESP-50), and HMAC MD5 . <i>Note: MD5 signed certificates are no longer considered secure.</i> If ANY NOT NULL is chosen, the Brick device negotiates the Auth Type in the following order: <ul style="list-style-type: none"> • For the IKEv1 client TEP setting: 1) HMAC SHA1, 2) HMAC MD5. • For the IKEv2 client TEP setting: 1) HMAC SHA1, 2) HMAC MD5, 3) AES XCBC.
Hosts Behind Tunnel	For IPSec SA Proposal 1, this is a read-only field (not editable) and contains the same value as the Hosts Behind Tunnel field on the Endpoint tab.
Protocol	The choices are: *, icmp , sctp , tcp , udp , or an integer in the range 1-255.
Local Ports	The choices are *, an integer in the range 0-65535, or a low-high range of integers from 0-65535.
Remote Ports	The choices are *, an integer in the range 0-65535, or a low-high range of integers from 0-65535.

- 13 The fields in the **IPSec SA Parameters** box are populated with default values. You can change any of these values. The following explains:

Field	Explanation
Enable Perfect Forward Secrecy	By default, this checkbox is unchecked. If it is checked, a second Diffie-Hellman key exchange will take place during processing. This can improve security, but it also can impact re-keying performance. This parameter setting applies to IPSec SA Proposal 1 and IPSec SA Proposal 2 (if enabled).
Enable LZS compression	<p>The compression feature only applies when the tunnel endpoint is a Brick or another device that supports LZS compression.</p> <p>If the device is a Brick, it must be equipped with an encryption accelerator card (either a Model 201, 300, 700, 1200 Min, 1000 or 1200 Max). For Brick devices with an encryption card, this feature should be enabled.</p> <p>This feature is supported in v3.1 and above of the Alcatel-Lucent IPSec Client.</p> <p>This parameter setting applies to IPSec SA Proposal 1 and IPSec SA Proposal 2 (if enabled).</p>
SA Lifetime (secs and Kbytes) (IKEv1 and IKEv2 client TEP settings)	<p>The Security Association has a lifetime specified in seconds and kilobytes. The defaults are 14400 seconds (4 hours) and 5,000,000 kilobytes (approximately T1 speed for 8 hours).</p> <p>You can change the lifetime in seconds to any value between 120 -157,680,000 seconds (between 2 minutes and 5 years).</p> <p>You can change the lifetime in kilobytes to any value between 1000 - 10,000,000 kilobytes.</p> <p>You can set the value of this field to 0 to disable SA expiration.</p> <p>The Security Association will expire after the <i>first</i> of the above two lifetimes is reached. The session will then have to re-key.</p>

To use the Multiple SA Proposals feature, go to [Step 14](#). Otherwise, go to “[Task: save and apply the tunnel defaults](#)” (p. 11-21).

- 14** Click the **Enable IPsec SA Proposal 2** checkbox to enable a second IPsec SA proposal, if you intend to use the Multiple SA Proposals feature. This feature is disabled, by default.

Important! The Multiple SA Proposals feature can only be enabled if IKEv2 is enabled for the Brick tunnel endpoint.

Result If the **Enable IPsec SA Proposal 2** checkbox is checked (enabled), the **IPsec Protocol, Encryption Type, Auth Type**, and second set of traffic selector fields are activated for input.

When defining IPsec SA Proposal 2, the value(s) of at least one of the following fields must be different than IPsec SA Proposal 1:

- **Protocol**
- **Encryption**
- **Auth Type**

END OF STEPS

Task: save and apply the tunnel defaults

Once you have defined the client tunnel defaults, you must save and apply them to any Client TEPs that are using Group Defaults, and to compile and download the new/updated policies to the affected Bricks. Complete the following steps:

- 1** Once you have completed all tabs on the Client Tunnel Defaults Editor, display the File menu and select **Save and Apply**.

Result The Apply Group window is displayed (Figure 11-5, “Apply Group Window” (p. 11-22))

Figure 11-5 Apply Group Window



-
- 2 Click **OK** to save and apply the client tunnel defaults to the affected Brick device(s).

Result The new policies are compiled and downloaded to the affected Brick device(s).

END OF STEPS

Create a Host Group of Hosts Behind the Tunnel

When setting up the tunnel, you will be asked to identify the hosts behind the tunnel. These are the hosts that client users will be permitted to access through the tunnel.

The default is asterisk. If you leave the default in place, all outbound traffic from the client will automatically go through the tunnel. The alternative is to create a host group with the IP addresses of specific hosts on the other side of the tunnel. In this case, all traffic to these hosts will automatically go through the tunnel.

You should create the host group before setting up the tunnel, so that it will appear on the drop-down list for you to select. For instructions on creating a host group, see [Chapter 2, “Host Groups”](#).

If you create a host group, traffic to hosts not in the host group will *not* go through the tunnel. In this case, you need to consider whether or not to permit this traffic through the Brick or router. The **Client Firewall** field on the Parameters tab of the Client Defaults Editor (see [Step 4](#) above) determines whether or not to pass this traffic out of the client.

Create an Authentication Service

The purpose of an authentication service is to determine how client users will be authenticated. You do not need to create an authentication service if clients will be authenticated by means of Local Password.

However, if the users will be authenticated by means of one of the external authentication methods (RADIUS, SecurID, or VPN Certificate), you must create an authentication service. This needs to be done before setting up the tunnel, so that the authentication service will appear on the drop-down list for you to select.

The procedure for creating an authentication service is described in [Chapter 8, “User Authentication”](#). When setting up the tunnel, be sure to make this the default authentication service. The procedure for obtaining and installing a digital certificate is explained in [Chapter 9, “Digital Certificates”](#).

□

To Set Up a Client Tunnel Endpoint

When to use

Use this task to set up a client tunnel endpoint (TEP).

Before you begin

Before you begin this task, be aware that there are two methods of accessing the Client Tunnel Endpoint Editor: from the VPN folder or from the Device folder. Use Method 1 to display the Client Tunnel Endpoint Editor from the VPN folder. Use Method 2 to display the Client Tunnel Endpoint Editor from the Devices folder.

The information requested on the Endpoint tab of the Client Tunnel Endpoint Editor is required. The fields on the Parameters, IKEv1 Gateway, IKEv2 Gateway, Remote Client ID, or Policy tab only need to be modified if you want to change one or more of the default settings.

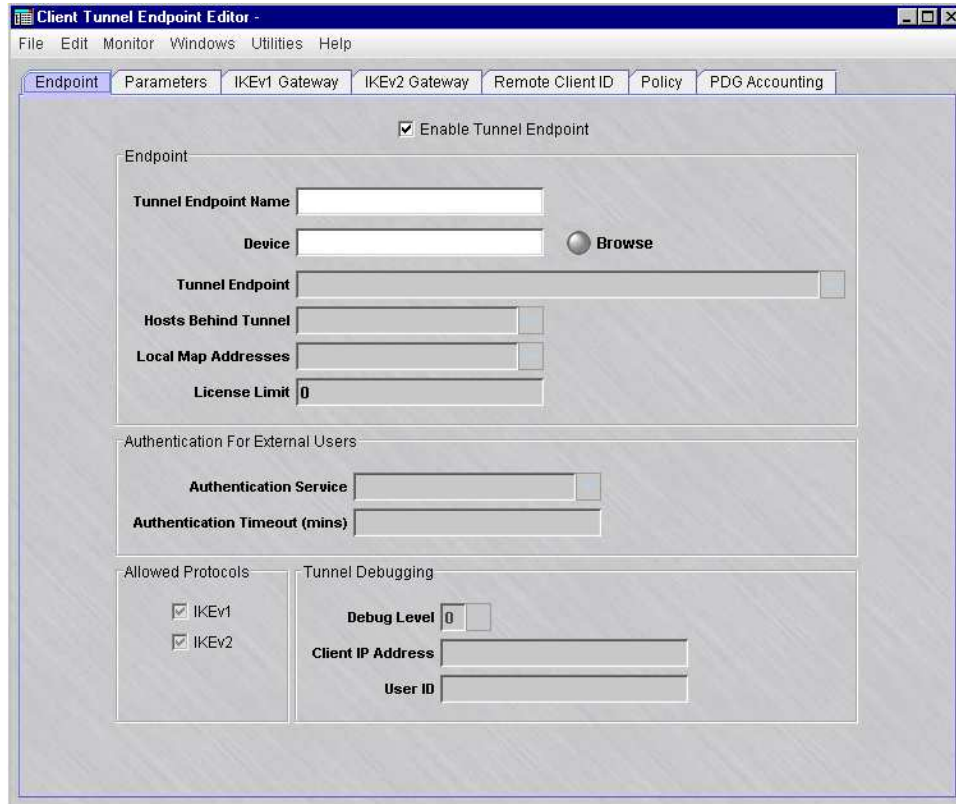
Task, Method 1: from the VPN folder

Complete the following steps to display the Client Tunnel Endpoint Editor from the VPN folder.

-
- 1 With the Navigator window displayed, open the appropriate group folder, and then open the **VPN** folder.
-
- 2 Right-click **Client Tunnel Endpoints** and select **New Client Tunnel Endpoints** from the pop-up menu.

Result If you opt for this method, the Main tab of the Client Tunnel Endpoint Editor is displayed with the **Tunnel Endpoint Name** and **Device** fields active and unpopulated, and all the other fields greyed-out, as shown in Figure 11-6, “Client Tunnel Endpoint Editor (Endpoint Tab)” (p. 11-25).

Figure 11-6 Client Tunnel Endpoint Editor (Endpoint Tab)



- 3 In the **Tunnel Endpoint Name**, enter a unique name for the client tunnel endpoint.
- 4 Click the **Browse** button next to the **Device** field to select the Brick of the client tunnel endpoint, from the Bricks folder.

Result The greyed-out fields on the Endpoint tab are activated.

END OF STEPS

Task, Method 2: from the Devices folder

Complete the following steps to display the Client Tunnel Endpoint Editor from the Devices folder.

- 1 With the Navigator window displayed, open the appropriate group folder, and then open the **Devices** folder.

- 2 Click the **Bricks** folder to display all configured Bricks in the Contents panel.

- 3 Double-click the Brick device that will be serving as the tunnel endpoint.

Result The Brick Editor is displayed.

- 4 Click **Policy Assignment** to display the Policy Assignment tab.

- 5 Right-click the port or interface serving as the tunnel endpoint and select **Client VPN** from the pop-up menu.

Result If you opt for this method, the Main tab of the Client Tunnel Endpoint Editor is displayed (refer to [Figure 11-6, “Client Tunnel Endpoint Editor \(Endpoint Tab\)”](#) (p. 11-25)) with the **Tunnel Endpoint Name** and **Device** fields already populated with the device you selected and its existing endpoint address. These fields are greyed-out, so they cannot be modified. The other fields are activated and can be modified.

END OF STEPS

Task: complete the Endpoint tab

Complete the following steps to complete the Endpoint tab. The fields on the Endpoint tab identify and configure the device serving as the tunnel endpoint.

- 1 Display the Client Tunnel Endpoint Editor using [“Task, Method 1: from the VPN folder”](#) (p. 11-24) or [“Task, Method 2: from the Devices folder”](#) (p. 11-26).

Result The Client Tunnel Endpoint Editor is displayed, showing the Endpoint tab.

If [“Task, Method 1: from the VPN folder”](#) (p. 11-24) is used to display the Client Tunnel Endpoint Editor, the **Device** field is already populated with the selected Brick.

If “[Task, Method 2: from the Devices folder](#)” (p. 11-26) is used to display the Client Tunnel Endpoint Editor, the **Device** and **Tunnel Endpoint** fields are already populated.

If the zone on the Brick you selected was configured to provide Alcatel-Lucent IPsec Client users with a presence on the local LAN, the appropriate address(es) will be automatically entered in the **Local Map Addresses** field. See [Appendix A, “Local Presence”](#) for a discussion of local presence.

-
- 2 The **Enable Endpoint** checkbox is checked, by default. If you do not want this TEP enabled at this point, click the checkbox to disable it.

 - 3 In the **Tunnel Endpoint Name** field, enter a name for this TEP.

 - 4 In the **Tunnel Endpoint** field, select the endpoint address from the drop-down list.

The endpoint address of a Brick is created when a Brick zone ruleset is assigned to a port. If a Brick has multiple Brick zone rulesets assigned to one port, each zone will have a different tunnel endpoint address. As a result, you may see more than one entry in the drop-down list. Be sure to select the entry that corresponds to the specific zone you want (the name is included after each endpoint address in the list).

 - 5 In the **Hosts Behind Tunnel** field, the default is asterisk. This means all outbound traffic from the client will automatically go through the tunnel.

The alternative is to enter a host group containing specific IP addresses. If you created such a host group, display the drop-down list, select **Browse**, and select the host group from the Browse window that appears.

If you enter a host group, make sure the **Client Firewall** field on the Parameters tab is set properly. This field determines whether traffic to other hosts — traffic that does not go through the tunnel — will be passed or dropped by the Brick device.

 - 6 If you want to provide client users with a presence on the local LAN, the local presence has to be entered using the Brick Policy Assignment Editor. Refer to the *Configuring Alcatel-Lucent VPN Firewall Brick® Security Appliance Ports* in the *SMS Administration Guide*. Usually this is done when the Brick is initially configured, but if it has not been done, and you require local presence, you can edit the Brick configuration and add the local presence information now. The local presence address must be part of the zone IP address range.

Important! Once the local presence has been set up using the Brick Policy Assignment Editor, the local presence information appears automatically in the Client Tunnel Endpoint Editor when this window is displayed. The information will be greyed-out, so it can only be changed or deleted from the Brick Policy Assignment Editor.

- 7** In the **Authentication for External Users** box, select a pre-configured authentication service from the drop-down list in the **Authentication Service** field.

The only authentication service provided with the SMS application is Local Password. If you are using RADIUS, SecurID or VPN Certificate, you have to create the authentication service yourself. Refer to [Chapter 8, “User Authentication”](#).

By default, the authentication will time out after 480 minutes (8 hours). You can change this by entering a new time in the **Authentication Timeout** field. The time can range from 1 to 2628000 minutes (5 years).

- 8** In the **Allowed Protocols** box, choose at least one key negotiation method: **IKEv1** or **IKEv2**. For new client TEPs, both methods are chosen, by default.

Depending on the method(s) chosen, the fields on the IKEv1 Gateway and/or IKEv2 Gateway tabs are greyed out or activated.

- 9** In the **Tunnel Debugging** box, you have the option of enabling tunnel debugging and the **Debug Level**. When tunnel debugging is enabled, the Brick writes debug messages to the VPN log, identified by zone, tunnel endpoint name, remote TEP (client IP), and User ID. The available debug levels are **0** (debugging off), **1**, **2**, and **3** (most verbose).

In the **Client IP Address** field, enter the IP address of the remote client machine to debug client tunnel connections for a specific endpoint (IP address).

In the **User ID** field, enter a User ID to debug client tunnel connections for a specific User ID. Valid values are upper/lowercase letters, 0-9, spaces, and special characters .?~!@\$%^&_-+=*<>. An asterisk (*) can be used as a wildcard character to specify a partial User ID. The **User ID** field is blank by default.

END OF STEPS

Task: complete the Parameters, IKEv1 Gateway, IKEv2 Gateway, Remote Client ID, and Policy tabs

-
- 1 The **Parameters, IKEv1 Gateway, IKEv2 Gateway, Remote Client ID, and Policy** tabs appear with fields already populated with their default settings. If the defaults are acceptable for this tunnel, regardless of whether they are the original defaults or modified defaults, you can ignore these tabs and save and apply the tunnel. Refer to the [“Task: save and apply the tunnel”](#) (p. 11-31).

However, if you need to change one or more of the defaults for this tunnel, click the appropriate tab, and uncheck the **Use Group Defaults** checkbox at the top. Then, make any changes as necessary.

Refer to the procedure [“To Set the Client Tunnel Defaults”](#) (p. 11-8) for an explanation of each tab field on the Client Defaults Editor window, which are the same fields shown on the Client Tunnel Endpoint Editor window.

END OF STEPS

Task: to complete the PDG Accounting tab (optional)

Important! The PDG Accounting and Lawful Intercept features are optional features that must be purchased and installed using separate installation keys via the New Feature Setup utility. If the PDG Lawful Intercept feature has not been installed, all fields on the PDG Accounting tab are greyed out.

For details about the New Feature Setup utility, refer to the *SMS Administration Guide*.

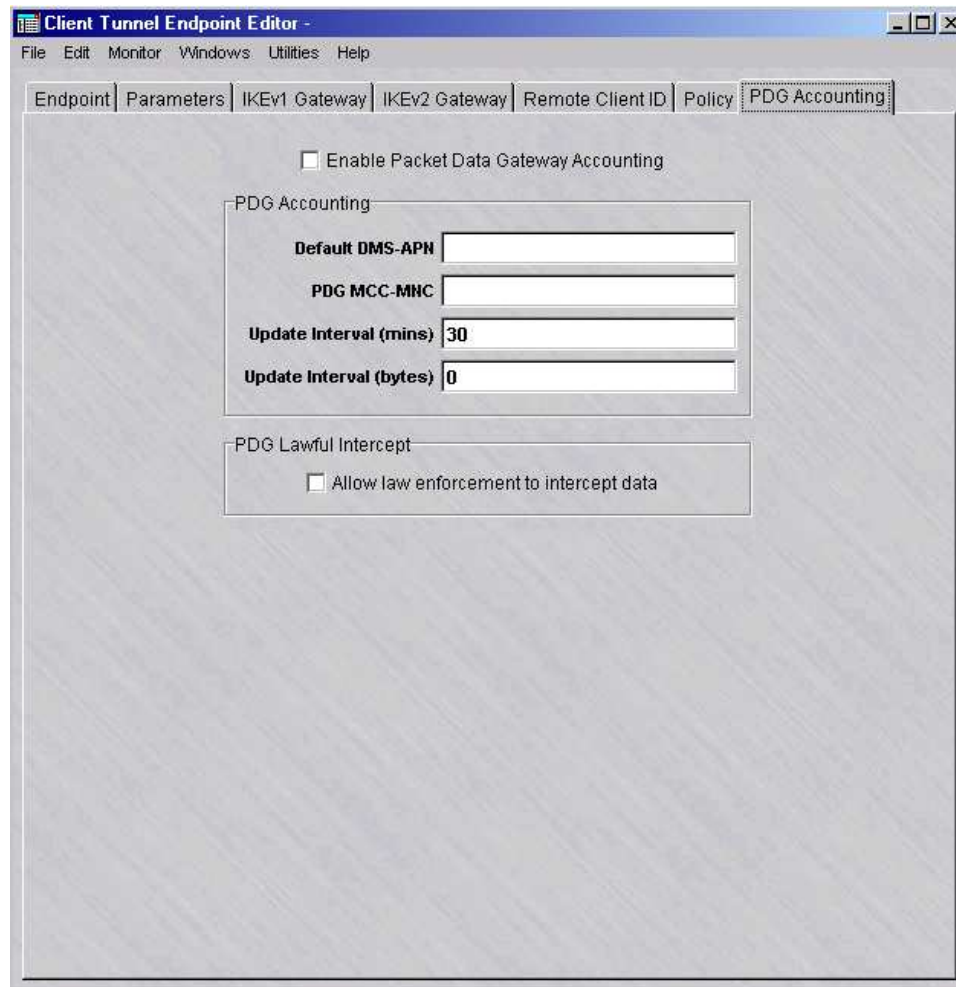
Note: before the Lawful Intercept feature is enabled, a rule must be created in the Brick zone ruleset that allows the intercepted data packets to be passed to the device port(s) (IP address(es)) being used by the law enforcement agency for monitoring purposes.

Complete the following steps to complete the PDG Accounting tab. The fields on this tab are used to enable or disable the Packet Data Gateway (PDG) Accounting feature, to configure the interim update intervals, in time period and traffic volume, and to enable or disable the PDG Lawful Intercept feature (if the PDG Lawful Intercept feature has been installed). The fields on this tab are optional. Refer to the [“Packet data gateway \(PDG\) accounting”](#) (p. 11-5) and [“PDG lawful intercept”](#) (p. 11-5) sections for additional information about these features.

-
- 1 Click **PDG Accounting** to display the PDG Accounting tab on the Client Tunnel Endpoint Editor.

Result The PDG Accounting tab is displayed (Figure 11-7, “Client Tunnel Endpoint Editor (PDG Accounting Tab)” (p. 11-30)).

Figure 11-7 Client Tunnel Endpoint Editor (PDG Accounting Tab)



-
- 2 Click the **Enable Packet Data Gateway Accounting** checkbox to enable the PDG Accounting feature. This checkbox is unchecked by default (the feature is disabled).

Result If the **Enable Packet Data Gateway Accounting** checkbox is checked, the remaining fields on the tab become editable. Go to Step 2.

-
- 3 In the **Default DMS-APN** field, enter the default Dual Mode Service (DMS) Access Point Name (APN) to be used to authenticate the originating access point of the traffic if the Dual Mode Handset (DMH) of the user does not provide this information as part of the IKE_AUTH_REQ message transmitted. This field accepts up to 100 characters.

- 4 In the **PDG MCC-MNC** field, enter the 5-6 digit Mobile Country Code (MCC) and the Mobile Network Code (MNC) of the network to which the Brick device belongs.
- 5 In the **Accounting Interval (mins)** field, enter the time period interval, in minutes, for accounting updates. The valid range for this field is 1-120 and the default value is 30.
- 6 In the **Accounting Interval (bytes)** field, enter the total traffic volume, in bytes, which, when reached, triggers an accounting update. The valid range for this field is 0-214783647. The default value for this field is 0, which means that no updates are performed based on the traffic volume, just the time period interval.
- 7 Click the **Allow law enforcement to intercept data** check box to enable the PDG Lawful Intercept feature on this client tunnel endpoint. The check box is unchecked by default (feature disabled).

END OF STEPS

Task: save and apply the tunnel

Once you have completed all tabs, you are ready to save and apply the tunnel. It is necessary to do this to make the tunnel operational. To save and apply the tunnel, follow the steps below:

- 1 From any tab in the Client Tunnel Endpoint Editor, display the File menu and select **Save and Apply**.

Result If an Alcatel-Lucent IPSec Client license limit has not been allocated for the TEP used for the tunnel, a warning dialog box is displayed advising you that this TEP cannot be used until the license limit is set. Refer to [“Allocating Alcatel-Lucent IPSec Client Licenses on Tunnel Endpoints”](#) (p. 11-35) for instructions on how to allocate an Alcatel-Lucent IPSec Client license limit on a tunnel endpoint. Click **OK** to close the warning dialog box.

A second warning dialog box is displayed, advising you that the **Save and Apply** action will recompile and apply the policy to the tunnel endpoint, and asks if you want to continue.
- 2 Click **OK** to proceed with recompiling and applying the policy to the tunnel endpoint.

Result The client TEP is created.

If any errors occur during processing, an error message dialog box is displayed with information about the error.

.....
E N D O F S T E P S



To Set Up a Client Tunnel with UDP Encapsulation

Changing the client tunnel transport mode

By default, traffic through Client tunnels is pure IPsec (type 50/51). However, you can change the transport mode to permit UDP encapsulated traffic instead, or to permit both pure IPsec and UDP encapsulation.

UDP encapsulation

When UDP encapsulation is employed, the original IP packet, including the entire IP header, is included as UDP payload data. The Brick can exchange these packets bi-directionally with other Bricks, and with hosts running the Alcatel-Lucent IPsec Client. The Brick will encapsulate the packets when the Security Association specifies, and decapsulate the packets upon receipt from the network.

UDP Encapsulation for IKEv1 key negotiation is an Alcatel-Lucent proprietary mechanism and only works with Bricks and Alcatel-Lucent IPsec Clients.

NAT traversal

NAT Traversal is the standard for performing UDP Encapsulation between IKEv2 tunnel endpoints. NAT Traversal for IKEv2 works with IPsec Client software deployed by other vendors.

If the **IKEv2** check box is checked in the **Allowed Protocols** box of the Client Tunnel Endpoints Editor, NAT Traversal is enabled automatically and is a read-only field on the **Parameters** tab of the Client Tunnel Endpoints Editor. If the **IKEv2** check box is unchecked, NAT Traversal is disabled and grayed out on the **Parameters** tab of the Client Tunnel Endpoints Editor.

If NAT Traversal is enabled, the SMS automatically adds the NAT Traversal ports to the **UDP_Encapsulation_Ports** service group. This service group is only visible on the Brick, not via the SMS GUI. Port 4500 is added to the UDP Encapsulation Ports list (which is also used for negotiating IKEv1 tunnels) and is used for negotiating tunnels using IKEv2.

Set Up UDP Encapsulation

The UDP encapsulation features only applies when the tunnel endpoint is a Brick. To set up UDP encapsulation, follow the steps below:

-
- 1 Configure the Client tunnel endpoint, just as you would any Client tunnel. Enter all the information required in the Main tab to set up the endpoint.

-
- 2 Click **Parameters** to display the Parameters tab, and then uncheck the check box labeled **Use Group Defaults**. This allows you to use the information you are about to enter in place of the default parameters.
-

- 3 In the **Allowed IPSec Transport Methods** box, select the transport methods to be allowed. You can select **Pure IPSec** only, **UDP Encapsulation** (for IKEv1 TEPs only), **NAT Traversal Port** (for IKEv2 TEPs only), or a combination of Pure IPSec and one of the other methods.

Important! If you want all your Client tunnels to use UDP encapsulation, you should change the default setting by clicking the **UDP Encapsulation Port(s)** check box on the Client Defaults Editor instead of the Client Tunnel Editor.

This way, every Client tunnel endpoint you set up will automatically handle UDP encapsulation, and you will not have to perform this procedure each time you set up a tunnel.

- 4 For UDP Encapsulation, the initiator destination port is 501.

For NAT Traversal, the initiator destination port is 4500.

However, when you clicked either **IKEv1** or **IKEv2** check box, a keyword called **UDP-Encapsulation-Ports** was added to the group in which the zone resides.

This keyword is automatically modified to include all the UDP destination ports specified in both the Client Tunnel Endpoint Editor and the LAN-LAN Tunnel Editor for all such entities in the zone.

END OF STEPS



What to Do Next

Post-tunnel setup

Once you have set up the tunnel, there are several other tasks that you will have to perform. The following explains:

Allocate Client Tunnel License Limit

The first thing you need to do after setting up the tunnel is allocate a client tunnel license limit. This procedure is explained in [“Allocating Alcatel-Lucent IPsec Client Licenses on Tunnel Endpoints”](#) (p. 11-35).

Allocating Alcatel-Lucent IPsec Client Licenses on Tunnel Endpoints

When the Brick is used to perform IKE negotiations, VPN Client users cannot establish tunnels to a TEP unless that endpoint has been allocated an Alcatel-Lucent IPsec Client License limit.

The SMS software installation key and feature keys determine how many client licenses are available on your SMS instance. You may distribute these licenses among the Groups on your SMS instance, then define user license limits for each Client TEP, up to the limit allocated for that Group.

Task

To allocate Alcatel-Lucent IPsec Client licenses:

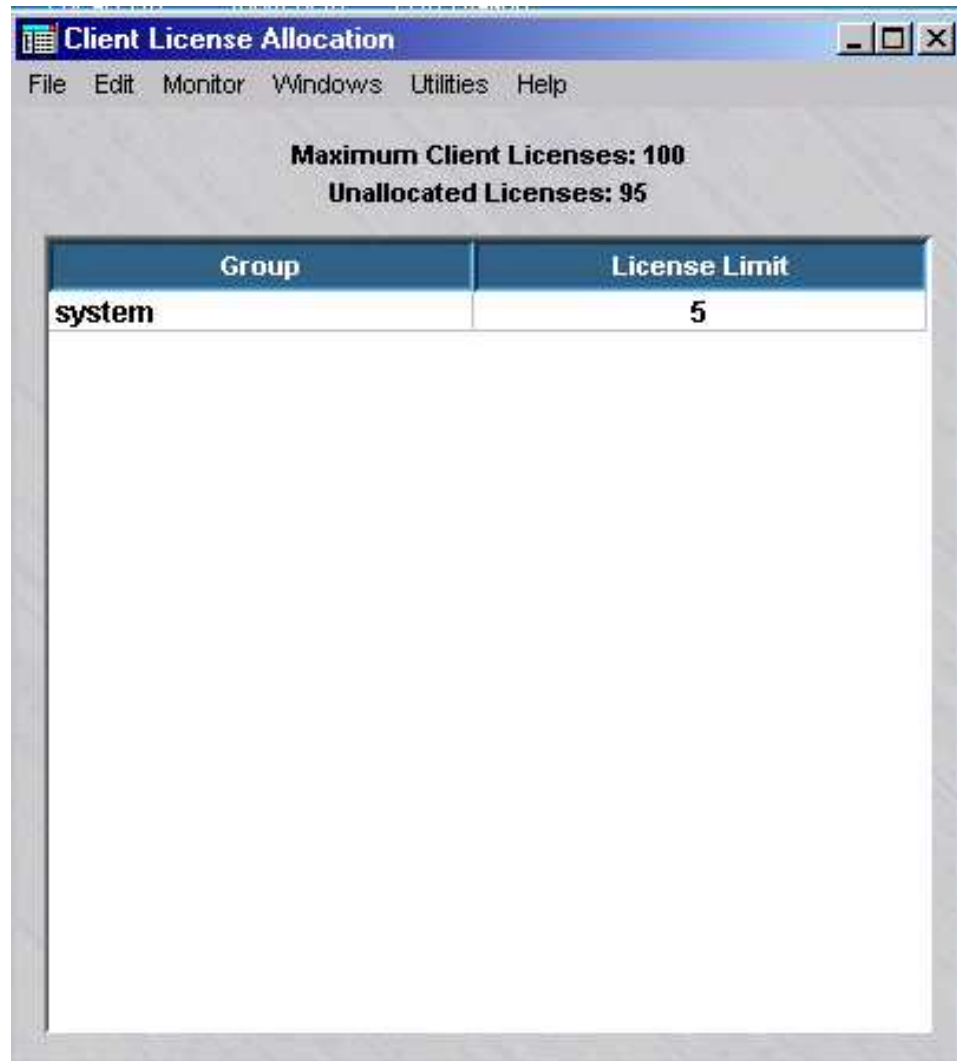
- 1 Log in using the SMS Navigator.

- 2 Click on the **Client Tunnel Endpoints** folder.
A list of Client tunnel endpoints is displayed in the Contents Panel.

- 3 Right-click on a Client tunnel endpoint and select **Allocate Licenses** from the pop-up menu.

The Client License Allocation window is displayed ([Figure 11-8, “Client License Allocation Window”](#) (p. 11-36) shows a sample window).

Figure 11-8 Client License Allocation Window



-
- 4 Right-click on a Group and select **Edit** or double-click on the Group.

The Group License Editor window is displayed ([Figure 11-9, “Group License Editor Window”](#) (p. 11-37) shows a sample window).

Figure 11-9 Group License Editor Window

Name	License Limit	TEP IP	Zone Ruleset	Device
mytonj	0	135.112.104.28	administrativezone	joesbrick

-
- 5 Right-click on each TEP in the list and select **Edit**.

A Tunnel Editor Endpoint window is displayed (Figure 11-10, “Tunnel Endpoint Editor Window” (p. 11-38) shows a sample window).

Figure 11-10 Tunnel Endpoint Editor Window

The screenshot shows a dialog box titled "TEP License Editor". It has a standard Windows-style title bar with a close button. The dialog contains five labeled text input fields:

- Name:** nytonj
- License Limit:** 0
- TEP IP:** 135.112.104.28
- Zone Ruleset:** administrativezone
- Device:** joesbrick

At the bottom of the dialog are two buttons: "OK" and "Cancel", each with a circular icon to its left.

.....

6 Change the value in the TEP License Limit field from **0** to the desired number of licenses for the TEP.

.....

7 Click the **OK** button.

The system returns to the Group License Editor window and shows the updated license limit for the TEP.

.....

8 Click the **Save and Apply** button.

.....

9 Repeat Steps 4-8 for each Group.

.....

END OF STEPS

.....

Set Up User Authentication

The next thing you should do after setting up the tunnel is to complete the set up of user authentication. Unless the client users are all coming from hosts with fixed IP addresses, you will have no way of preventing unauthorized access to the tunnel without some user authentication mechanism in place. SMS Release 9.4 supports four methods of authentication: Local Password, RADIUS, SecurID, and VPN Certificate.

At this point, you should have decided which method of authentication you are going to use, and you should have created the required authentication service, unless you are using the Local Password authentication service, which comes with the SMS application. The authentication service was required to set up the tunnel.

In addition, you have to do the following:

- **Local Password**

If you are using Local Password, you have to create an account for each user who will be authenticated. You may also want to create your own user groups to help manage this database of users. If not, you can use the **All_Users** group that is provided with the SMS application.

- **RADIUS and SecurID**

If you are using RADIUS or SecurID, you have to add a rule to the zone protecting the SMS to allow the SMS to communicate with the RADIUS or SecurID ACE/Server.

- **VPN Certificate**

If you are using VPN Certificate, you have to obtain and install an X.509 digital certificate using the Certificate Manager (see [Chapter 9, “Digital Certificates”](#) in the *SMS Policy Guide*).

For a detailed explanation of user authentication and instructions for creating the required rules, see [Chapter 8, “User Authentication”](#).

Create the User Access Rules

Every tunnel requires rules to permit the client’s traffic through the Brick, and to instruct the Brick when to encrypt the traffic and when to decrypt the traffic.

The following explains how to create a tunnel rule

Field	Explanation
Direction	The direction of the rule is usually into the zone, since the most Client sessions originates on a host outside the zone.
Source	The source has to be a user group. The user group can be one you have created, or it can be one of the system-created user groups (All_Users , Active_VPN_Users , Active_VPN_UserTEPs , Default_Auth_Users). See “SMS-defined user groups” (p. 8-19) in Chapter 8, “User Authentication” for definitions.
Destination	The destination has to be one or more hosts. It can be a specific IP address, a host group, or asterisk (all hosts in the zone).
Service	The service must match the application of the client user (HTTP, FTP, telnet).
Action	The action must be VPN or VPN Proxy.



Maintaining Client Tunnel Endpoints

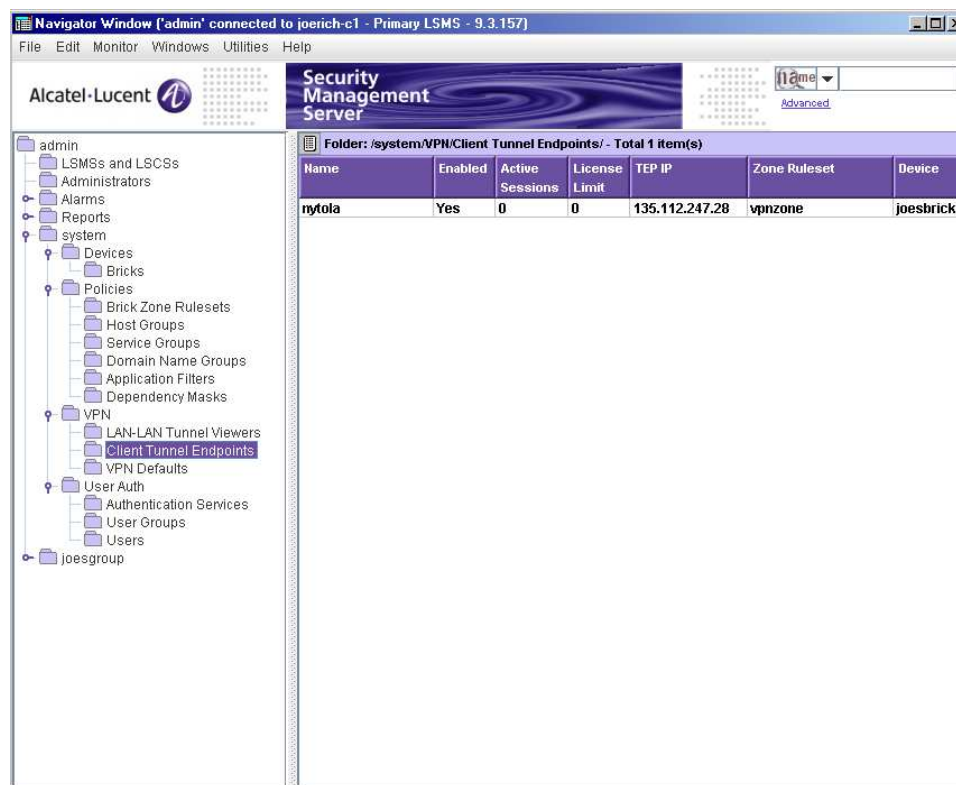
Overview

Once a client tunnel endpoint has been set up, it can be maintained by the administrator. An administrator can view the established client tunnel endpoints, display currently active tunnels by specifying search criteria, and enable and disable endpoints. You can also terminate individual existing sessions and all existing sessions.

To view existing tunnel endpoints

To view all the client tunnel endpoints that have been set up, open the appropriate group and VPN folders, and click the Client Tunnel Endpoints folder. The endpoints that have been configured will appear in the Navigator window. [Figure 11-11, “Client Tunnel Endpoints Window”](#) (p. 11-40) shows this display.

Figure 11-11 Client Tunnel Endpoints Window



The Client Tunnel Endpoints Window shows the following information for each client tunnel endpoint:

- Unique client tunnel endpoint name
- Client tunnel endpoint enabled indicator

- Number of currently active sessions
- License limit for the client tunnel endpoint
- Client tunnel endpoint IP address
- Zone ruleset for the client tunnel endpoint
- The associated device (Brick)

To enable and disable endpoints

To enable or disable a tunnel endpoint, right-click the endpoint in the Navigator window and select either **Enable** or **Disable** from the pop-up menu. No tunnels can be established to disabled endpoints. Existing sessions will be dropped if the endpoint is disabled.

You can also enable or disable an endpoint from the Main tab of the Client Tunnel Endpoint Editor (see [Figure 11-6, “Client Tunnel Endpoint Editor \(Endpoint Tab\)”](#) (p. 11-25)) by checking or unchecking the **Enable Tunnel Endpoint** check box.

To display active sessions

Use this procedure to search for and display active client tunnel sessions, based on the search criteria that you specify. You can search for and display all active client tunnel sessions by the originator’s user ID, IP address, or the local address to which the client is mapped.

Complete the following steps to search for and display active client tunnels.

- 1 Click on **Client Tunnel Endpoints** in the Folders panel of the Navigator.
Result A list of configured client tunnel endpoints is displayed in the Contents panel of the Navigator.
- 2 Right-click on an existing client tunnel endpoint and select **Search Client Tunnels** from the pop-up menu.

Result The Search Client Tunnels window is displayed (Figure 11-12, “Search Client Tunnels Window” (p. 11-42)

Figure 11-12 Search Client Tunnels Window



- 3 Select the search criteria to be used in displaying the active client tunnel sessions.

To	Do This
Search for active client tunnels by user ID	Click the User ID radio button (which is the default) and enter the full or partial user ID in the input field. An asterisk (*) can be used as a wildcard character to represent one or more characters of the ID (Examples: dav* , j*e . To search for active client tunnels for all user IDs, leave the asterisk (*) in the field.

To	Do This
Search for active client tunnels by IP address	Click the IP Address field and enter a single IP address, a range of IP addresses (Example: 10.10.10.1-10.10.10.99), or a subnet (Example: 10.10.10.1/24). To search for active client tunnels for all IP addresses, leave the asterisk (*) in the field.

- 4 After entering the search criteria, click the Search button.

Result The active client tunnel sessions that match the entered search criteria are displayed as a tabular listing in the window.

A maximum of 1000 matching client tunnel sessions are displayed.

If a client tunnel session was established using Alcatel-Lucent IPsec Client, and the endpoint has a Local Map IP Address (Local Presence is enabled), local map IP address and the actual client IP address are displayed. If there is no Local Map IP Address assigned to the client TEP for the active session, the client IP address is displayed in both columns of the table.

5

To	Do This
Terminate a specific client tunnel session in the listing	Right-click on the tunnel session and select Terminate Session from the pop-up menu.
Clear the listing of tunnel sessions	Right-click on any client tunnel session and select Clear Results from the pop-up menu.

END OF STEPS

To terminate all sessions

Complete the following steps to terminate all sessions on a client tunnel endpoint:

- 1 Click the **Client Tunnel Endpoints** folder in the Folders panel of the Navigator.

Result A list of defined client tunnel endpoints is displayed in the Contents panel of the Navigator.

- 2 Right-click the client tunnel endpoint in the Contents panel and select **Terminate All Sessions** from the pop-up menu.

Result A confirmation dialog box is displayed asking if you are sure that you want to terminate all sessions on this client tunnel endpoint.

- 3 Click **OK**.

Result An information dialog box is displayed, indicating the termination of all active sessions has been initiated and that it may take up to 30 seconds for the active sessions counter to be updated on the Navigator.

- 4 Click **OK** to close the information dialog box.

END OF STEPS

To modify an endpoint

Complete the following steps to modify a client tunnel endpoint:

- 1 With the Navigator window displayed, double-click the endpoint you want to modify. The Client Tunnel Endpoint Editor will appear (see [Figure 11-6, “Client Tunnel Endpoint Editor \(Endpoint Tab\)”](#) (p. 11-25)) with the fields populated with the endpoint you selected.
-

- 2 Make any necessary changes to the information in any of the tabs.
-

- 3 Display the File menu and select **Save and Apply**.

END OF STEPS


To delete an endpoint

To delete a client tunnel endpoint, follow the steps below:

- 1 With the Navigator window displayed, right-click the endpoint you want to delete and select **DELETE** from the pop-up menu.

A confirmation window similar to the one shown in [Figure 11-13, “Confirmation Window \(Client Tunnel Endpoint\)”](#) (p. 11-45) is displayed.

Figure 11-13 Confirmation Window (Client Tunnel Endpoint)



- 2 Click **OK** to recompile the policy assigned to the endpoint be recompiled and applied to the Brick and delete the tunnel endpoint.

END OF STEPS



To Create a Message for Alcatel-Lucent IPsec Client Users

When to use

Use this task to create a message that all Alcatel-Lucent IPsec Client users will see when they successfully enable a tunnel to an endpoint that has been configured on a Brick.

You can use this message facility to provide status messages, to indicate changes or additions to the hosts behind the tunnel that the client user can access, to warn users of potential problems, and other information.

Task

Complete the following steps to create a message:

- 1 Using a standard text editor such Notepad or vi, create a text file containing the message, and save it under the following name

`<group_name>.banner`

where `<group_name>` is the name of your group.

- 2 Move the file to the directory

`<$install_directory>\isms\lmf\vgc\IKEConfig`

where `<$install_directory>` is the directory in which you installed the SMS software (*users* on the Windows® or Vista® platform, or *opt* on the Solaris® platform if the defaults were used during installation).

This message will be displayed until you remove or update the file.

END OF STEPS

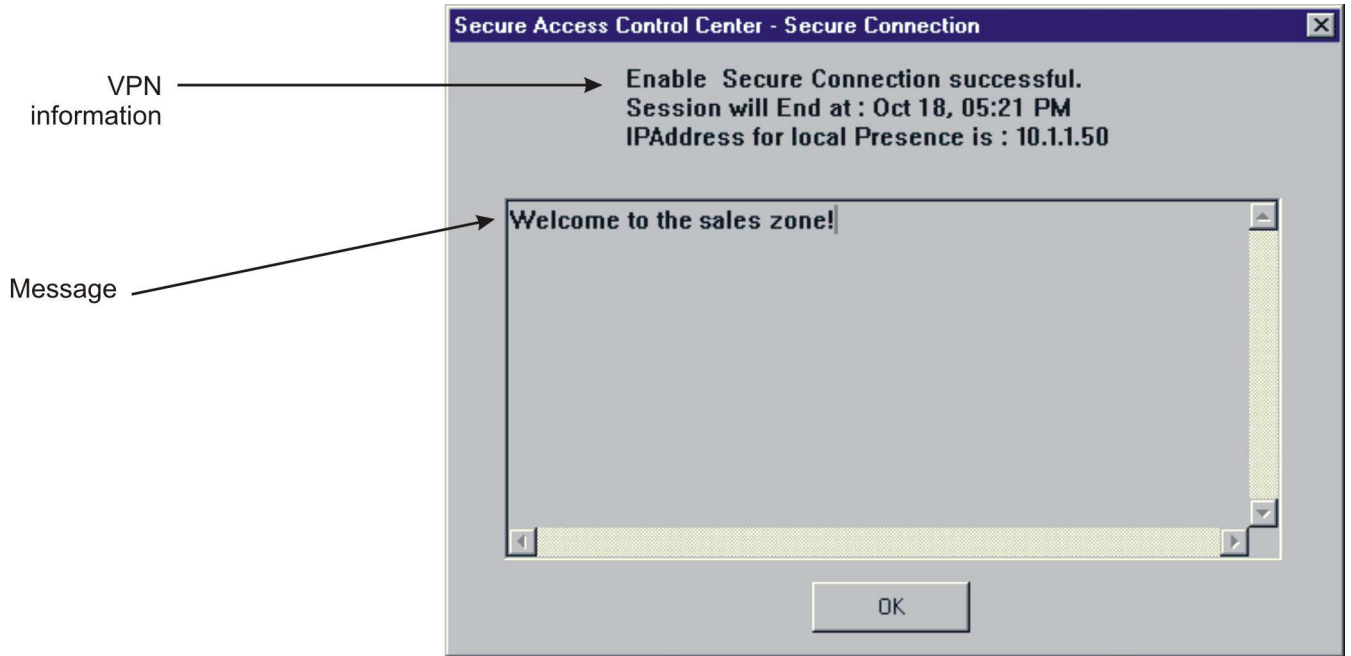
Sample Message

Figure 11-14, “Client Message” (p. 11-47) shows a typical message that an Alcatel-Lucent IPsec Client user might see after successfully enabling a tunnel.

Note that this window provides VPN information on top (the date and time the tunnel will terminate, the local address if one has been used) and the message in a text frame below. Scrollbars are provided to accommodate lengthy messages.

If you do not choose to post a message, the user will see a smaller pop-up window instead, with the VPN information only, and no text frame.

Figure 11-14 Client Message



Appendix A: Local Presence

Overview

Purpose

This appendix explains how to create a pool of addresses that can be used to provide users of the Alcatel-Lucent IPsec Client application with an IP address on a local LAN. This is referred to as a *local presence*.

The purpose of giving client users an IP address on a local LAN is to ensure that return packets to the client are routed to the specific Alcatel-Lucent *VPN Firewall Brick™* Security Appliance where the client's Security Association resides. The Security Association is required for encryption and decryption.

This feature is especially useful in solving traffic problems, and specifically Xwindows routing problems, in a network with asymmetric routes.

All administrators wishing to implement this feature must make changes to the Brick devices that pass the local presence packets. If some users are running older versions of Alcatel-Lucent IPsec Client, additional NAT rule changes must also be made.

Contents

How Local Presence Works	A-2
Why Set Up Local Presence?	A-3
How Local Presence Works	A-4
Local Presence - Scenario 1	A-5
Local Presence - Scenario 2	A-7



How Local Presence Works

Overview

The local presence feature uses a combination of:

- A field on the Brick Editor called **Local Map Addresses**
- A dynamic host group called **Active_VPN_Users**
- Rules that use that host group and (for earlier versions of the Alcatel-Lucent IPsec Client software) specify local presence mapping.

The host group **Active_VPN_Users** is created automatically by the SMS. When a VPN client is authenticated, the SMS looks at the list of available local presence addresses that you have created, and picks one for that user. It then adds that user's IP address and, if configured, performs a map-to address to **Active_VPN_Users**.

First, a rule can be created that uses the **Active_VPN_Users** host group. The following is an example of such a rule.

Direction	Source	Destination	Service	Action
IN	Active_VPN_Users	servers	FTP	Pass

Such a rule then specifies local address mapping under the Address Translation tab and an Action of "VPN" under the Basic tab.

The map-to address indicates the address to map the original IP address to. For or example, it might contain:

```
ip=135.92.38.112 map-to: 159.32.40.64
```

When a packet from 135.92.38.112 is allowed through the firewall because of that rule, its source IP address will be translated to 159.23.40.64. (Reverse packets will, of course, be unmapped just as they are for other types of network address translation.)

□

Why Set Up Local Presence?

Considerations

You should consider providing client users with a local presence if you have an environment consisting of multiple Bricks and asymmetric routes.

In a network with asymmetric routes, there is no guarantee that two successive packets, traveling between the same source and destination hosts, will take the same route to or from their destination.

However, for a tunnel to work, packets must enter and exit through the same Brick, because that device contains the IPsec Security Association and tunnel endpoint information that allows encrypted traffic to flow to and from the client.

By providing the client with a local IP address that is assigned to a specific port on a specific Brick, you are ensuring that the packet is directed back to the device that contains the correct Security Association and tunnel endpoint definition.

□

How Local Presence Works

Ways to implement local presence

There are two different implementations of the Local Presence feature. Your configuration will depend on what version of the Alcatel-Lucent IPsec Client package is in use by your remote users that access your network via a VPN tunnel.

1. Scenario 1 - All *Windows*[®] users are running Alcatel-Lucent IPsec Client 3.0.xxx or later.

For most users of the current SMS release, this is the expected scenario.

2. Scenario 2 - Some *Windows*[®] users are running an Alcatel-Lucent IPsec Client version earlier than 3.0.xxx.

Subscribers to SMS Software Support can easily download updated versions of the Alcatel-Lucent IPsec Client from the "private zone" on the VPN Firewall Product Registration and Support website (<https://www.lucent-ipsec.com>).



Local Presence - Scenario 1

Setting up local presence

To set up local presence, you have to display the Brick Policy Assignment Editor. This is the window used to assign a Brick zone ruleset to a port on a Brick. You can do this when you initially assign the ruleset to a port, or afterward by editing the port assignment.

(Refer to the *Configuring and Activating an Alcatel-Lucent VPN Firewall Brick® Security Appliance* chapter in the *SMS Administration Guide* for a more detailed discussion of port assignments.)

To set up local presence on the Brick

To set up local presence on the Brick, follow the steps below:

- 1 From the Navigator window, open the appropriate Group folder and Devices folder, and click the Brick folder to display all configured Bricks.

- 2 Right-click the Brick you want and select **Edit** from the pop-up menu. If you are configuring a new Brick, select **New** instead. The Brick Editor will appear with the Brick tab displayed. Refer to the *Configuring and Activating a Lucent VPN Firewall Brick® Security Appliance* chapter in the *SMS Administration Guide*.

- 3 Click **Policy Assignment** to display the Policy Assignment tab, and double click the port you want. The Brick Policy Assignment Editor will appear. Refer to the *Configuring Alcatel-Lucent VPN Firewall Brick® Security Appliance Ports* in the *SMS Administration Guide*.

- 4 In the **Local Map Addresses** field indicate the pool of addresses to be set aside for client users.

This can be done in any of the following ways:

- Enter a single IP address.
- Enter an address with subnet notation, for example, 10.10.10.10/24.
- Enter a range of IP addresses, separated by a hyphen.
- Display the drop-down list, select **Browse**, and select the host group you created for this purpose.

The **Local Map Addresses** field is only active when an IP address is entered in the **Tunnel Endpoint** field.

5 Click **OK** to dismiss the Brick Policy Assignment Editor.

6 Display the File menu and select **Save and Apply**.

Also, a change may be necessary to the Client Tunnel Endpoint configuration. The Hosts Behind Tunnel parameter should include a host group that you have defined consisting of:

- IP addresses of internal hosts and subnets to be accessed
- IP addresses entered in the "Local Map Addresses" field in the Brick Policy Assignment Editor.

For more information on editing Client Tunnels, please refer to *Client Tunnel Endpoints* chapter in this Guide.

END OF STEPS



Local Presence - Scenario 2

Setting up local presence

To give users of older versions of the Alcatel-Lucent IPsec Client an IP address on the local LAN, you must make use of the SMS's Network Address Translation (NAT) capability to automatically map the client's external IP address to an internal IP address. For more information, please refer to the *Network Address Translation* chapter in this Guide.

Set Up NAT Rule Changes for Older Alcatel-Lucent IPsec Client Versions

The objective is to force clear text packets traveling on the local LAN to the correct Brick where the IPsec processing will be performed. To accomplish this, you can make use of both source and destination mapping:

- **Source Address Mapping**

Source mapping is performed when the clear text packet originates at the Brick; in other words, when the client's address is the **source IP**.

In this case, the source address is mapped to an internal IP address. This internal address then becomes the destination IP address of the return packet. You must also set the **Type** to Local.

Since the internal IP address is associated with a particular Brick, this guarantees that the response will be routed back to the correct device.

- **Destination Address Mapping**

Destination mapping is performed for a clear text packet originating on a local host, that is, when the client's address is the **destination IP**. This is the case for X11 traffic. You must also set the **Type** to Local.

In this case, the destination address is mapped to the original client IP address from the internal address.

To permit the mapping to work, you need to set aside a pool of IP addresses to be used. Once you do this, it is a good idea to create a host group containing these addresses, so that you can easily enter them when required.

To set up local presence on the Brick

To set up local presence, you have to display the Brick Policy Assignment Editor. This is the window used to assign a Brick zone ruleset to a port on a Brick. You can do this when you initially assign the ruleset to a port, or afterward by editing the port assignment. (Refer to the *Configuring Alcatel-Lucent VPN Firewall Brick® Security Appliance Ports* chapter in the *SMS Administration Guide* for a more detailed discussion of port assignments.)

To set up local presence, follow the steps below:

- 1 From the Navigator window, open the appropriate Group folder and Devices folder, and click the Brick folder to display all configured Bricks.
- 2 Right-click the Brick you want and select **Edit** from the pop-up menu. If you are configuring a new Brick, select **New** instead. The Brick Editor will appear with the Brick tab displayed. Refer to the *Configuring and Activating an Alcatel-Lucent VPN Firewall Brick® Security Appliance* chapter in the *SMS Administration Guide*.
- 3 Click **Policy Assignment** to display the Policy Assignment tab, and double click the port you want. The Brick Policy Assignment Editor is displayed.
- 4 In the **Local Map Addresses** field indicate the pool of addresses to be set aside for client users.

This can be done in any of the following ways:

- Enter a single IP address.
- Enter an address with subnet notation, for example, 10.10.10.10/24.
- Enter a range of IP addresses, separated by a hyphen.
- Display the drop-down list, select **Browse**, and select the host group you created for this purpose.

The **Local Map Addresses** field is only active when an IP address is entered in the **Tunnel Endpoint** field.

- 5 Click **OK** to dismiss the Brick Policy Assignment Editor.
- 6 Display the File menu and select **Save and Apply**.

Also, a change may be necessary to the Client Tunnel Endpoint configuration. The Hosts Behind Tunnel parameter should include a host group that you have defined consisting of:

- IP addresses of internal hosts and subnets to be accessed
- IP addresses entered in the **Local Map Addresses** field in the Brick Policy Assignment Editor.

For more information on editing Client Tunnels, please refer to [Chapter 11, “Client Tunnel Endpoints”](#).

.....
E N D O F S T E P S



Appendix B: Pre-Configured Alcatel-Lucent *VPN Firewall Brick*TM Security Appliance Zone Rulesets

Overview

Purpose

This appendix explains each rule in the **administrativezone**, **dhcpzone_on_inside_if**, **dhcpzone_on_internet_if**, **firewall**, and **nocgwzone** rulesets. These are all pre-configured Brick zone rulesets that are provided with the SMS application. Rules that appear in the Brick Zone Ruleset Editor in black are general user rules for user traffic. Rules that appear in light blue print are generic to management traffic from Brick to SMS, and so forth.

The **administrativezone** ruleset is usually assigned to the port on the Brick that is connected to the SMS. It drops all traffic to the SMS except traffic from the Bricks it is managing. The **dhcpzone_on_inside_if** ruleset is the default zone for mobile Bricks for the protected side of the firewall and the **dhcpzone_on_internet_if** ruleset is the default zone for mobile Bricks for the internet side of the firewall.

If the SMS will be managing routers through a management tunnel, you usually assign the **nocgwzone** ruleset instead. It contains the rules necessary to establish the management tunnel to the routers.

The firewall ruleset is intended to protect the Brick, and is automatically assigned to a "local" interface.

There are two additional pre-configured rulesets. The **proxyzone** ruleset is used to protect the hosts running the Proxy Agent application (*Note: the Proxy Agent application has been replaced by the Rules-Based Routing feature and is no longer supported*). The **vpnzone** ruleset is used to set up LAN-LAN tunnels, and is described in [Chapter 10, "LAN-LAN Tunnels"](#).

Contents

administrativezone	B-3
dhcpzone_on_inside_if	B-11

dhcpzone_on_internet_if	B-17
firewall	B-23
nocgwzone	B-30



administrativezone

Task

To view the **administrativezone** ruleset:

- 1 With the Navigator window displayed, open the **System** group folder.
- 2 Open the Policies folder.
- 3 Click the Brick Zone Rulesets folder to display all the Brick zone rulesets in the Navigator window.
- 4 Double-click **administrativezone**.

END OF STEPS

Rulesets

The following sections describe each rule in the ruleset. The final rule in the ruleset is the standard drop-all packets rule.

Rules #200, #203, #206

The purpose of rules #200, #203, and #206 is to allow any Brick connected to the SMS to send audit data to the SMS and request the SMS to download the Brick. The rules differ only in their Source definition. The following explains how the rules work:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.

Field	Explanation
Source	<p>brickRemoteAddresses (#200)</p> <p>This host group is used when NAT is used for Bricks in a private address space relative to the SMS, when the addresses are fixed.</p> <p>bricks (#203)</p> <p>This is a host group that contains the IP addresses of all the Bricks that are being managed by this SMS. It was automatically created when the Administrative Zone was created.</p> <p>mobile_bricks (#206)</p> <p>This host group must be manually updated when mobile Bricks are created or fixed address Bricks are made mobile. This host group can define its hosts as * or a range, if known.</p>
Destination	<p>SMS</p> <p>This is the SMS host group that contains the IP addresses of all SMSs.</p>
Service	<p>brick_to_SMS_Services</p> <p>This is a service group that was created automatically to facilitate Brick to SMS communication. It consists of the following protocol/ports:</p> <p>TCP/900/*</p> <p>TCP/9000/*</p> <p>UDP/9014/*</p> <p>UDP/9020/911</p>
Action	<p>Pass</p> <p>The rule permits traffic matching the above fields, which can only be coming from a Brick, into the Administrative Zone.</p>
Audit	<p>Yes</p> <p>The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.</p>
Session Timeout	<p>300</p> <p>The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.</p>
Auth Return	<p>Yes</p> <p>The rule permits a return channel.</p>
Syn Flood Type	<p>None</p> <p>The rule does not protect against a Syn Flood attack.</p>

Rules #201, #204, #207

The purpose of rules #201, #204, and #207 is to allow the SMS to download a policy and send configuration information to any Brick connected to the SMS. The rules differ only in their Destination definition. The following explains how the rules work:

Field	Explanation
Direction	<p>Out Of Zone</p> <p>The purpose of the rule is to allow traffic from the SMS, which is inside the Administrative Zone, to reach the bricks, which are outside the zone.</p>
Source	<p>SMS</p> <p>This is the SMS host group that contains the IP addresses of all SMSs.</p>
Destination	<p>brickLocalAddresses (#201)</p> <p>This host group is used when NAT is used for Bricks in a private address space relative to the SMS, when the addresses are fixed.</p> <p>bricks (#204)</p> <p>This is a host group that contains the IP addresses of all the Bricks that are being managed by this SMS.</p> <p>mobile_bricks (#207)</p> <p>This host group must be manually updated when mobile Bricks are created or fixed address Bricks are made mobile. This host group can define its hosts as * or a range, if known.</p>
Service	<p>brick_from_SMS_Services</p> <p>This service group was created automatically to facilitate SMS to Brick communication. It consists of the following protocol/ports:</p> <p>TCP/910/*</p> <p>ICMP/8/0</p> <p>UDP/1024/*</p> <p>UDP/9014/*</p> <p>UDP/911/9020</p>
Action	<p>Pass</p> <p>The rule permits traffic matching the above fields, which can only be going to a Brick, out of the Administrative Zone.</p>

Field	Explanation
Audit	Yes The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rules #202, #205, #208

The purpose of rules #202, #205, and #208 is to allow the Brick to send policy download replies to the SMS when the clear cache option is used. These rules are automatically deleted a few seconds after they are loaded into the firewall. The rules differ only in their Source definition. The following explains how the rules work:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.
Source	brickRemoteAddresses (#202) This host group is used when NAT is used for Bricks in a private address space relative to the SMS, when the addresses are fixed. bricks (#205) This is a host group that contains the IP addresses of all the Bricks that are being managed by this SMS. mobile_bricks (#208) This host group must be manually updated when mobile Bricks are created or fixed address Bricks are made mobile. This host group can define its hosts as * or a range, if known.

Field	Explanation
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	TCP/*/910 The rule applies only to Brick traffic with a source port of 910.
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, into the Administrative Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Important! Rules #209, #210, and #211 are inactive rules intended as *sparcs*. Since you cannot create new System rules, these are available to be defined and activated for any purpose, such as for unusual cases involving NAT.

Rule 209

The purpose of Rule #209 is to allow any Brick connected to the SMS to send audit data to the SMS and request the SMS to download the Brick. The rules differ only in their Source definition. The following explains how the rules work:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.

Field	Explanation
Source (host group)	bricks This is a host group that contains the IP addresses of all the Bricks that are being managed by this SMS. It was automatically created when the Administrative Zone was created.
Destination (host group)	LSMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	brick_to_SMS_Services This is a service group that was created automatically to facilitate Brick to SMS communication. It consists of the following protocol/ports: TCP/900/* TCP/9000/* UDP/9014/* UDP/9020/911
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, into the Administrative Zone.
Session Audit	Yes The rule permits the Brick to audit all open and closed session packets. It will audit all passed packets.
Exception Audit	Yes The rule permits the Brick to audit all dropped packets.

Rule 210

The purpose of Rule #210 is to allow the SMS to download policy and configuration data to any Brick connected to that SMS. The following explains how the rules work:

Field	Explanation
Direction	Out of Zone The purpose of the rule is to allow traffic from the SMS, which is inside the Administrative Zone, to reach the Bricks, which are outside of the Administrative Zone.

Field	Explanation
Source (host group)	LSMS This is the SMS host group that contains the IP addresses of all SMSs.
Destination (host group)	bricks This is a host group that contains the IP addresses of all the Bricks that are being managed by this SMS. It was automatically created when the Administrative Zone was created.
Service	brick_from_SMS_Services This is a service group that was created automatically to facilitate SMS to Brick communication. It consists of the following protocol/ports: TCP/910/* ICMP/8/0 UDP/1024/* UDP/9014/* UDP/911/9020
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, into the Administrative Zone.
Session Audit	Yes The rule allows the Brick to audit all passed packets
Exception Audit	Yes The rule permits the Brick to audit all dropped packets.

Rule 211

The purpose of Rule #211 is send download replies from each Brick to the SMS when the Clear Cache option is used. The following explains how the rules work:

Field	Explanation
Direction	In to Zone The purpose of the rule is to allow traffic from each Brick, which is outside the Administrative Zone, to reach the SMS, which is inside the Administrative Zone.

Field	Explanation
Source (host group)	bricks This is a host group that contains the IP addresses of all the Bricks that are being managed by this SMS. It was automatically created when the Administrative Zone was created.
Destination (host group)	LSMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	tcp/*/910 This rule only applies to Brick traffic with a source port of 910.
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, into the Administrative Zone.
Session Audit	No The rule does not allow the Brick to audit all passed packets.
Exception Audit	Yes The rule permits the Brick to audit all dropped packets.



dhcpzone_on_inside_if

Task

To view the **dhcpzone_on_inside_if** ruleset:

- 1 With the Navigator window displayed, open the **System** group folder.

- 2 Open the Policies folder.

- 3 Click the Brick Zone Rulesets folder to display all the Brick zone rulesets in the Navigator window.

- 4 Double-click **dhcpzone_on_inside_if**.

END OF STEPS

Ruleset definitions

The following sections describe each rule in the ruleset. The final rule in the ruleset is the standard drop-all packets rule.

Rule #200

The purpose of rule #200 is to allow the Brick to send audit data to the SMS and request downloads. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.
Source	Virtual Brick Address The VBA is a dynamic address that the Brick has obtained from the DHCP server. This address is also used to perform NAT (Network Address Translation) for outbound traffic from the Brick to the internet.

Field	Explanation
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	brick_to_SMS_Services This is a service group that was created automatically to facilitate Brick to SMS communication. It consists of the following protocol/ports: TCP/900/* TCP/9000/* UDP/9014/* UDP/9020/911
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, out of the Firewall Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #201

The purpose of rule #201 is to allow SMS to download policy and configuration information to Bricks. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.

Field	Explanation
Source	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Destination	Virtual Brick Address The VBA is a dynamic address that the Brick has obtained from the DHCP server. This address is also used to perform NAT (Network Address Translation) for outbound traffic from the Brick to the internet.
Service	brick_from_SMS_Services This is a service group that was created automatically to facilitate Brick- to-SMS communication. It consists of the following protocol/ports: TCP/910/* ICMP/8/0 UDP/1024/* UDP/9014/* UDP/911/9020
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, out of the Firewall Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #202

The purpose of rule #202 is to allow policy download replies from Brick to SMS when Clear Cache option is used. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.
Source	Virtual Brick Address The VBA is a dynamic address that the Brick has obtained from the DHCP server. This address is also used to perform NAT (Network Address Translation) for outbound traffic from the Brick to the internet.
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	tcp/*/910
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, out of the Firewall Zone.
Audit	No
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #1000

The purpose of rule #1000 is to drop outbound traffic that uses IRC. The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow traffic from the SMS, which is inside the Administrative Zone, to reach the bricks, which are outside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all hosts.
Destination	Wildcard asterisk (*) The wildcard asterisk represents all hosts.
Service	irc
Action	Drop
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	10 The rule sets, to 10 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #1001

The purpose of rule #1001 is to allow user traffic to go out to the internet by using NAT. The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow traffic from the SMS, which is inside the Administrative Zone, to reach the Bricks, which are outside the zone.

Field	Explanation
Source	Wildcard asterisk (*) The wildcard asterisk represents all hosts.
Destination	Wildcard asterisk (*) The wildcard asterisk represents all hosts.
Service	Wildcard asterisk (*) The service can be any protocol or port.
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, out of the Firewall Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.



dhcpzone_on_internet_if

Task

To view the **dhcpzone_on_internet_if** ruleset:

- 1 With the Navigator window displayed, open the **System** group folder.

- 2 Open the Policies folder.

- 3 Click the Brick Zone Rulesets folder to display all the Brick zone rulesets in the Navigator window.

- 4 Double-click **dhcpzone_on_internet_if**.

END OF STEPS

Ruleset definitions

The following sections describe each rule in the ruleset. The final rule in the ruleset is the standard drop-all packets rule.

Rule #200

The purpose of rule #200 is to allow the Brick to send audit data to the SMS and request downloads. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.
Source	Virtual Brick Address The VBA is a dynamic address that the Brick has obtained from the DHCP server. This address is also used to perform NAT (Network Address Translation) for outbound traffic from the Brick to the internet.

Field	Explanation
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	brick_to_SMS_Services This is a service group that was created automatically to facilitate Brick-to-SMS communication. It consists of the following protocol/ports: TCP/900/* TCP/9000/* UDP/9014/* UDP/9020/911
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, out of the Firewall Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #201

The purpose of rule #201 is to allow SMS to download policy and configuration information to Bricks. The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow traffic from the SMS, which is inside the Administrative Zone, to reach the bricks, which are outside the zone.

Field	Explanation
Source	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Destination	Virtual Brick Address The VBA is a dynamic address that the Brick has obtained from the DHCP server. This address is also used to perform NAT (Network Address Translation) for outbound traffic from the Brick to the internet.
Service	brick_from_SMS_Services This is a service group that was created automatically to facilitate Brick-to-SMS communication. It consists of the following protocol/ports: TCP/910/* ICMP/8/0 UDP/1024/* UDP/9014/* UDP/911/9020
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, out of the Firewall Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #202

The purpose of rule #202 is to allow policy download replies from Brick to SMS when Clear Cache option is used. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.
Source	Virtual Brick Address The VBA is a dynamic address that the Brick has obtained from the DHCP server. This address is also used to perform NAT (Network Address Translation) for outbound traffic from the Brick to the internet.
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	tcp/*/910
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, out of the Firewall Zone.
Audit	No
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #1000

The purpose of rule #1000 is to drop outbound traffic that uses IRC. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all hosts.
Destination	Wildcard asterisk (*) The wildcard asterisk represents all hosts.
Service	irc
Action	Drop
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	10 The rule sets, to 10 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #1001

The purpose of rule #1001 is to allow user traffic to go out to the internet by using NAT. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the Administrative Zone, to reach the SMS, which is inside the zone.

Field	Explanation
Source	Wildcard asterisk (*) The wildcard asterisk represents all hosts.
Destination	Wildcard asterisk (*) The wildcard asterisk represents all hosts.
Service	Wildcard asterisk (*) The service can be any protocol or port.
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, out of the Firewall Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.



firewall

Task

To view the **firewall** ruleset:

- 1 With the Navigator window displayed, open the **System** group folder.
- 2 Open the Policies folder.
- 3 Click the Brick Zone Rulesets folder to display all the Brick zone rulesets in the Navigator window.
- 4 Double-click **firewall**.

END OF STEPS

Ruleset definitions

The following sections describe each rule in the ruleset. The final rule in the ruleset is the standard drop-all packets rule.

Rule #200

The purpose of rule #200 is to allow the Brick to send traffic to the SMS. The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow traffic from the Brick, which is inside the Firewall Zone, to reach the SMS, which is outside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.

Field	Explanation
Service	Wildcard asterisk (*) The service can be any protocol or port.
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, out of the Firewall Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #201

The purpose of rule #201 is to allow the Brick to receive traffic from the SMS. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from the SMS, which is outside the Firewall Zone, to reach the Brick, which is inside the zone.
Source	SMS This is the SMSHost group that contains the IP addresses of all SMSs.
Destination	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.

Field	Explanation
Service	brick_from_SMS_Services This is a service group that was created automatically to facilitate SMS-to-Brick communication. It consists of the following protocol/ports: TCP/910/* ICMP/8/0 UDP/1024/* UDP/9014/* UDP/911/9020
Action	Pass The rule permits traffic matching the above fields, which can only be coming from the SMS, into the Firewall Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #202

The purpose of rule #202 is to allow a drop notification through the Brick. If the **Drop Action** field in the Brick Zone Ruleset window is turned on, a drop notification is sent whenever the Brick drops a packet. Without this rule, the notification itself would be dropped by the Brick.

The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow drop notifications, which are originated by the Brick in the Firewall Zone, to leave the zone.

Field	Explanation
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	Wildcard asterisk (*) The destination can be any host on the Internet.
Service	1/3/13 These are the ports set aside for drop notifications.
Action	Pass This rule permits drop notifications to pass out of the Firewall Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #203

In a LAN-LAN VPN, either another Brick or another IPSec-compliant device must initiate IKE negotiations with the Brick that is serving as the tunnel endpoint.

In a Client to LAN VPN, an Alcatel-Lucent IPSec Client has to initiate IKE negotiations with the Brick that is serving as the tunnel endpoint.

In both cases, the IKE negotiation packets must be permitted into the Firewall Zone so they can reach the Brick that is serving as the tunnel endpoint. The purpose of rule #203 is to allow these packets into the Firewall Zone. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The rule allows IKE negotiation packets initiated by a client to reach the appropriate Brick in the Firewall Zone.

Field	Explanation
Source	Wildcard asterisk (*) The source (the VPN client) can be any host on the Internet.
Destination	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Service	UDP/500/* These are the protocol and ports to be used for IKE negotiation.
Action	Pass The traffic in the above fields can only be IKE negotiations because of the protocol and ports. The rule allows this traffic to enter the zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	59 The rule sets, to 59 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #204

In a LAN-LAN VPN, IKE negotiations can be initiated by a Brick. The purpose of rule #204 is to allow these IKE negotiation packets to leave the Firewall Zone.

The rule also has a second purpose — to allow rekey notification in both Client to LAN and LAN-LAN VPNs.

The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The rule allows IKE negotiation packets initiated by a Brick out of the Firewall Zone.

Field	Explanation
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	Wildcard asterisk (*) The destination (VPN client) can be any host on the Internet.
Service	UDP/*/500 These are the protocol and ports to be used for IKE negotiations.
Action	Pass The traffic in the above fields can only be IKE negotiations because of the protocol and ports. The rule allows this traffic to leave the zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	59 The rule sets, to 59 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #205

The purpose of rule #205 is to allow queries from proxy hosts about reflection to reach the Brick. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The proxy hosts querying the Brick are located outside the Firewall Zone.
Source	Wildcard asterisk (*) The source can be any host on the Internet.
Destination	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.

Field	Explanation
Service	UDP/1024/* These are the protocol and ports to be used for the query.
Action	Pass The traffic in the above fields can only be a reflection query because of the protocol and ports. The rule allows this traffic to enter the zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.



nocgwzone

Task

To view the **nocgwzone** ruleset:

- 1 With the Navigator window displayed, open the **System** group folder.
- 2 Open the Policies folder.
- 3 Click the Brick Zone Rulesets folder to display all the Brick zone rulesets in the Navigator window.
- 4 Double-click **nocgwzone**.

END OF STEPS

Ruleset definition

The following sections describe each rule in the ruleset. The final rule in the ruleset is the standard drop-all packets rule.

Rule #200

The purpose of rule #200 is to allow any Brick connected to the SMS to send traffic to the SMS. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the NOC Gateway Zone, to reach the SMS, which is inside the zone.
Source	bricks This is a host group that contains the IP addresses of all the Bricks that are being managed by this SMS. It was automatically created when the NOC Gateway Zone was created.
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.

Field	Explanation
Service	brick_to_SMS_Services This is a service group that was created automatically to facilitate Brick to SMS communication. It consists of the following protocol/ports: TCP/900/* TCP/9000/* UDP/9014/* UDP/9020/911
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, into the NOC Gateway Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #201

The purpose of rule #201 is to allow the SMS to send traffic to any Brick connected to the SMS. The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow traffic from the SMS, which is inside the NOC Gateway Zone, to reach the Bricks, which are outside the zone.
Source	SMS This is the SMS host group that contains the IP addresses of all SMSs.

Field	Explanation
Destination	bricks This is a host group that contains the IP addresses of all the Bricks that are being managed by this SMS.
Service	brick_from_SMS_Services This is a service group that was created automatically to facilitate SMS-to-Brick communication. It consists of the following protocol/ports: TCP/910/* ICMP/8/0 UDP/1024/* UDP/9014/* UDP/911/9020
Action	Pass The rule permits traffic matching the above fields, which can only be going to a Brick, out of the NOC Gateway Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #202

The purpose of rule #202 is to suppress an error message from appearing when the NOC Gateway Zone is loaded on a Brick. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the NOC Gateway Zone, to reach the SMS, which is inside the zone.
Source	bricks This is a host group that contains the IP addresses of all the Bricks that are being managed by this SMS. It was automatically created when the NOC Gateway Zone was created.
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	TCP/*/910 The rule applies only to Brick traffic with a source port of 910.
Action	Pass The rule permits traffic matching the above fields, which can only be coming from a Brick, into the NOC Gateway Zone.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #203

The purpose of rule #203 is to permit tunneled, ICMP traffic from any Brick into the NOC Gateway Zone and then onto the SMS. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the NOC Gateway Zone, to reach the SMS, which is inside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	ICMP Traffic sent from any Brick is using the ICMP protocol and can be sent from any source port.
Action	VPN The rule permits tunneled traffic matching the above fields, which can only be coming from a Brick, into the NOC Gateway Zone, and arriving at the SMS.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	10 The rule sets, to 10 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
VPN	External The rule causes packets leaving the zone to be <i>encrypted</i> and packets coming into the zone to be <i>decrypted</i> .
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #204

The purpose of rule #204 is to permit tunneled, TCP traffic from port 9017 of any Brick into the NOC Gateway Zone and then onto the SMS. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the NOC Gateway Zone, to reach the SMS, which is inside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	TCP/9017 Traffic sent from any Brick is using the TCP protocol and must use port 9017.
Action	VPN The rule permits tunneled traffic matching the above fields, which can only be coming from a Brick, into the NOC Gateway Zone, and arriving at the SMS.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	3600 The rule sets, to 3600 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
VPN	External The rule causes packets leaving the zone to be <i>encrypted</i> and packets coming into the zone to be <i>decrypted</i> .
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #205

The purpose of rule #205 is to permit tunneled, TCP traffic from port 9018 of any Brick into the NOC Gateway Zone and then onto the SMS. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the NOC Gateway Zone, to reach the SMS, which is inside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	TCP/9017 Traffic sent from any Brick is using the TCP protocol and must use port 9017.
Action	VPN The rule permits tunneled traffic matching the above fields, which can only be coming from a Brick, into the NOC Gateway Zone, and arriving at the SMS.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	3600 The rule sets, to 3600 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
VPN	External The rule causes packets leaving the zone to be <i>encrypted</i> and packets coming into the zone to be <i>decrypted</i> .
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #206

The purpose of rule #206 is to permit tunneled, UDP traffic from port 69 of any Brick into the NOC Gateway Zone and then onto the SMS. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks, which are outside the NOC Gateway Zone, to reach the SMS, which is inside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Service	UDP/69 Traffic sent from any Brick is using the UDP protocol and must use port 69.
Action	VPN The rule permits tunneled traffic matching the above fields, which can only be coming from a Brick, into the NOC Gateway Zone, and arriving at the SMS.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
VPN	External The rule causes packets leaving the zone to be <i>encrypted</i> and packets coming into the zone to be <i>decrypted</i> .
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #207

The purpose of rule #207 is to allow the SMS to send tunneled, TCP traffic to port 9017 of any Brick. The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow traffic from the SMS, which is inside the NOC Gateway Zone, to reach the Bricks, which are outside the zone.
Source	SMS This is the SMS host group that contains the IP addresses of all SMSs.
Destination	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Service	TCP/9017 Traffic sent to any Brick is using the TCP protocol and must be sent to port 9017.
Action	VPN The rule permits tunneled traffic matching the above fields, which can only be coming from the SMS and going to any Brick.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	3600 The rule sets, to 3600 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
VPN	External The rule causes packets leaving the zone to be <i>encrypted</i> and packets coming into the zone to be <i>decrypted</i> .
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #208

The purpose of rule #208 is to allow the SMS to send tunneled, TCP traffic to port 9018 of any Brick. The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow traffic from the SMS, which is inside the NOC Gateway Zone, to reach the Bricks, which are outside the zone.
Source	LSMS This is the SMS host group that contains the IP addresses of all SMSs.
Destination	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Service	TCP/9018 Traffic sent to any Brick is using the TCP protocol and must be sent to port 9018.
Action	VPN The rule permits tunneled traffic matching the above fields, which can only be coming from the SMS and going to any Brick.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	3600 The rule sets, to 3600 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
VPN	External The rule causes packets leaving the zone to be <i>encrypted</i> and packets coming into the zone to be <i>decrypted</i> .
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #420

The purpose of rule #420 is to permit UDP traffic from port 500 of any Brick to the NOC Gateway. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks to reach the NOC Gateway.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	Virtual Brick Address (VBA) of the NOC Gateway. The destination has to be the Virtual Brick Address (VBA) of the NOC Gateway. This is the IP address that was entered on the Brick Policy Assignment Editor when assigning the NOC Gateway Zone to the port on the Brick configured as the NOC Gateway.
Service	UDP/500/* Traffic sent to any Brick is using the TCP protocol and must be sent to port 9018.
Action	Pass The rule permits traffic matching the above fields, which can be coming from any Brick and arriving at the SMS.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #430

The purpose of rule #430 is to allow any Brick to send encrypted IPsec traffic to the NOC Gateway. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from Bricks to reach the NOC Gateway.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	Virtual Brick Address (VBA) of the NOC Gateway. The destination has to be the Virtual Brick Address (VBA) of the NOC Gateway. This is the IP address that was entered on the Brick Policy Assignment Editor when assigning the NOC Gateway Zone to the port on the Brick configured as the NOC Gateway.
Service	ipsec Traffic sent to any Brick is using the ipsec protocol from any port.
Action	VPN The rule permits tunneled traffic matching the above fields, which can be coming from any Brick and arriving at the Brick configured as the NOC Gateway.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
VPN	External The rule causes packets leaving the zone to be <i>encrypted</i> and packets coming into the zone to be <i>decrypted</i> .
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.

Rule #440

The purpose of rule #440 is to allow any Brick to send decrypted IPsec traffic to the NOC Gateway. The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow traffic from the Brick to reach the NOC Gateway.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	Virtual Brick Address (VBA) of the NOC Gateway. The destination has to be the Virtual Brick Address (VBA) of the NOC Gateway. This is the IP address that was entered on the Brick Policy Assignment Editor when assigning the NOC Gateway Zone to the port on the Brick configured as the NOC Gateway.
Service	ipsec Traffic sent to any Brick is using the ipsec protocol from any port.
Action	VPN The rule permits tunneled traffic matching the above fields, which can be coming from any Brick and arriving at the Brick configured as the NOC Gateway.
Audit	Yes. The rule permits the Brick to audit all open and closed session packets. It will audit both dropped and passed packets.
Session Timeout	300 The rule sets, to 300 seconds, the length of time for a particular packet of any protocol to remain in the session cache.
VPN	Internal The rule causes packets leaving the zone to be <i>encrypted</i> and packets coming into the zone to be <i>decrypted</i> .
Auth Return	Yes The rule permits a return channel.
Syn Flood Type	None The rule does not protect against a Syn Flood attack.



Appendix C: Denial of Service Attacks

Overview

Purpose

This appendix describes three features that enable the Alcatel-Lucent *VPN Firewall Brick™* Security Appliance to fend off denial-of-service (DoS) attacks.

Contents

Syn Flood Protection	C-2
Intelligent Cache Management	C-4
Robust Fragment Reassembly	C-6



Syn Flood Protection

Overview

Syn flood attacks are one of the most common types of DoS attacks.

Syn Flood Attack

A syn flood attack takes advantage of the three-way handshake that characterizes TCP/IP connections. The following explains how the handshake works:

- When a client attempts to establish a connection with a server, the client sends a TCP packet to the server with the synchronization (SYN) flag set.
- The server then sends a packet back to the client with both the SYN and acknowledgment (ACK) flags set.
- The client then responds to the SYN-ACK packet by sending an ACK packet back to the server...completing the three-way handshake and establishing the connection.

In a syn flood attack, an attacker, using a fake source IP address, sends a large number of connection requests (SYN packets) to a specific server. The server sets aside memory for each request, and then sends a SYN-ACK packet back to the fake source IP address. When the SYN-ACK message times out, the server re-sends it, forcing the server to keep the memory available.

If enough connections are kept half open, the server will run out of memory and will not be able to respond to legitimate connection requests. Ultimately, the server could crash for lack of sufficient memory.

Flood Protection

The Brick has a rule-based syn flood protection feature that can be activated from the SMS. This feature allows you to create a rule allowing certain sessions through the Brick, and then set:

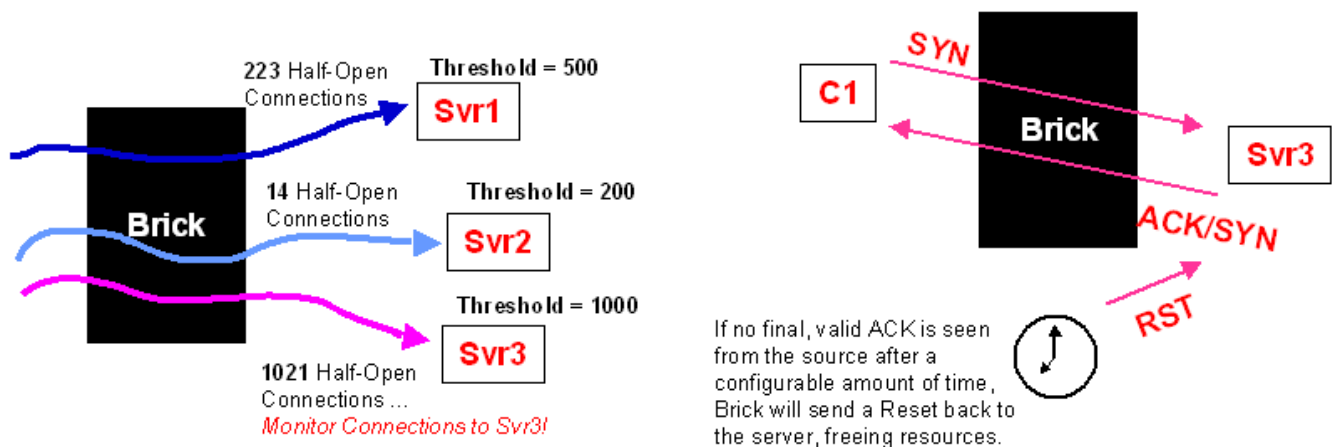
- A threshold that determines the number of half open sessions that will be allowed by this rule
and
- A countdown timer that determines how long these sessions can remain half open after the threshold is reached before the Brick sends a reset (RST) packet to the destination IP address. The RST packet causes the destination server to remove the session from its memory.

Important! The following are some important things to consider when setting the countdown timer:

1. Many servers have their own syn flood protection mechanism. Applying the Brick's syn flood defense to such servers could interfere with the server's mechanisms. Since the host knows how much memory is in use for half open sessions, if its defenses are adequate, it is best to not enable syn flood protection on the Brick for those hosts.
2. The following trade-off applies to setting the countdown timer:
 - Setting the countdown timer to a high value gives time for very remote hosts and/or hosts on very congested networks to complete the three-way handshake.
 - Setting the timer to a low value helps reduce the memory consumed by a syn flood attack.

The diagram in [Figure C-1, "Syn Flood Protection"](#) (p. C-3) illustrates the Brick's syn flood protection feature. For instructions on activating this feature, see [Step 10](#).

Figure C-1 Syn Flood Protection



□

Intelligent Cache Management

Overview

The SMS provides an intelligent cache management feature that has been specifically designed to deal with flood attacks that can saturate the Brick's physical memory.

When this feature is activated, the Brick will periodically scan its session cache and remove established sessions, as needed, to free session cache memory. You can configure the order in which the sessions are removed.

Session Cache

Each model of the Brick contains a session cache in which entries are made for every session currently active through the Brick. Since the Brick's memory consists of physical RAM, there are limits on the amount set aside for the session cache. For a model 300 Brick, this limit is about 400,000 simultaneous session entries.

In the event that all session cache memory is used, the Brick will not be able to create cache entries for any further new sessions. Packets in existing sessions (for which cache entries exist) will be allowed through the Brick as usual.

Packets for new sessions will be processed according to the rules; but since the Brick cannot create cache entries for those sessions (that is, keep state information), it is very likely that reverse packets in the new sessions will not be allowed to pass. Therefore, some traffic in those new sessions may not pass through the Brick. (TCP sessions, in particular, may not be allowed to become established.)

A "session" is not just a TCP connection. The Brick also treats UDP and ICMP packets as sessions. Some services (such as FTP) require multiple sessions.

Flood Protection


If too many cache entries are being established, the Brick may be under an attack of some sort. In that case, it is necessary to make some determination about which cache entries are being created because of that attack.

The individual cache class thresholds provide a means of identifying the "likely suspects." An Administrator makes an educated guess about the maximum cache memory that is likely to be used by certain types of sessions under normal circumstances. Anything that exceeds those guesses is suspect.

Too many half-open TCP sessions is indicative of a TCP SYN-flood attack; too many dropped sessions could indicate some sort of a probe for open ports; etc. The intelligent cache management process begins by removing sessions that are in the suspect class; if that doesn't free enough memory, it goes after other cache classes, starting with the ones whose entries can be removed with the least amount of disturbance to any valid established sessions.

The diagram below illustrates the intelligent cache management feature. For instructions on activating this feature, see *Chapter 4. Configuring Lucent VPN Firewall Brick® Device Ports* in the *SMS Administration Guide*.

Figure C-2 Intelligent Cache Management



□

Robust Fragment Reassembly

Overview


The Brick has been designed to protect against a sustained flood of fragments. Fragments are checked for duplication and overlap, and invalid fragments are immediately dropped.

In addition, a robust fragment reassembly feature has been built into the Brick. If the total number of outstanding fragments exceeds a threshold, or if a new packet requires an unavailable reassembly queue, this feature is activated.

When the feature is activated, it frees a single reassembly queue, chosen pseudo-randomly. This ensures that only a single IP packet is affected each time.

The diagram below illustrates the robust fragment reassembly feature.

Figure C-3 Robust Fragment Reassembly



□

Appendix D: RADIUS Attributes

Overview

Purpose

RADIUS (Remote Authentication Dial-in Service) is a protocol that enables the SMS to communicate with a database server that provides user authentication, access control, and accounting functionality.

It is possible to configure the SMS to use some of the information contained in the RADIUS protocol messages that the RADIUS server sends to the SMS when authenticating users. To do this, you have to assign RADIUS attribute codes to certain SMS parameters, and then indicate the RADIUS attribute data type.

The following lists the SMS parameters you can use and indicates the RADIUS attribute codes and data types that can be assigned to them.

Contents

User Group Information	D-2
Admin Key	D-4
Client Program Information	D-5
IKEv1 Preshared Key	D-7
Timeout Information	D-8
Local Presence IP Information	D-10
Primary WINS/Secondary DNS	D-13
Sample Configuration	D-15



User Group Information

Overview

An administrator can configure the SMS to process user group information contained in the RADIUS protocol message.

How to Configure the SMS and RADIUS Server

To configure the SMS and RADIUS server, follow the steps below:

-
- 1 Using the interface on the RADIUS server, configure a RADIUS attribute type to contain user group information.
-
- 2 Using the SMS, create a user group and assign that user group to one or more rules in one or more zones assigned to one or more Alcatel-Lucent *VPN Firewall Brick™* Security Appliances. The name of the user group must be exactly the same as the name of the user group you configured on the RADIUS server. The SMS user group can be empty.

The procedure for creating an SMS user group is described in [“Overview” \(p. 8-1\)](#) (see [“To create a user group” \(p. 8-17\)](#)).

-
- 3 Create a RADIUS authentication service. On the **Attribute Codes** tab of the Authentication Service Editor, set the SMS Parameter to **User Group** and select one of the following RADIUS Attribute Codes. The RADIUS Attribute Type (string or text) will be entered automatically when you enter the code.

- | | |
|-------------------|--------------------------|
| • User-Name | • Called-Station-id |
| • User- Password | • Calling-Station-id |
| • CHAP-password | • NAS-Identifier |
| • Filter-id | • Proxy-State |
| • Reply-Message | • Login-LAT-Service |
| • Callback-Number | • Login-LAT-Node |
| • Callback-id | • Login-LAT-Group |
| • Framed-Route | • Framed-Apple-Talk-Zone |
| • State | • CHAP-Challenge |
| • Class | • Login-LAT-Port |

The procedure for creating a RADIUS authentication service is described in “Overview” (p. 8-1) (see “To Set Up RADIUS and SecurID Authentication” (p. 8-21)).

.....
E N D O F S T E P S
.....

How the Feature Works

Once the RADIUS server and SMS have been properly configured, the user group feature will work with both application users and Alcatel-Lucent IPSec Client users. The following explains how it works:

1. When a user whose authentication service is configured to accept RADIUS user groups is authenticated, the SMS will scan the RADIUS “Access-Accept” message for attributes of the type specified in the authentication service. (This is also true when the authentication service for external users is configured to accept RADIUS User Groups.)
2. If any of these attribute types are found in the RADIUS message, the SMS will parse the string or text accompanying each attribute as a UTF-8 string. The SMS will perform a character-by-character match on all user groups configured in the SMS group to which this user is a member.
3. If a user group is found in the SMS group that matches the string described above, the SMS will apply the IP address associated with the end user performing the authentication to the SMS user group matching that in the RADIUS response. This user group will be applied to every Brick and zone in which that user group is used.
4. The process described in [Step 2](#) and [Step 3](#) will be repeated for every attribute of the appropriate type found in the RADIUS message. This allows a user to be a member of multiple user groups. (Note that in any case, the user group must be pre-configured on the SMS.)
5. Regardless of whether or not a user group is not found in the SMS group that matches the string described above, the SMS will always apply the IP address to the *All_Users* user group. This user group will be applied to every Brick and zone in which that user group is used.

□

Admin Key

Definition

When a RADIUS or SecurID authentication service is assigned to an SMS or Group administrator, they must be assigned an Admin Key, which is a fixed "password" that is required to log in, in addition to the RADIUS password or SecurID token code+pin. If the Admin Key is configured as an SMS parameter in the RADIUS attributes, RADIUS can be configured to return the Admin Key in a RADIUS response attribute so that the administrator does not have to enter 2 passwords when they log in. The RADIUS Attribute Type (string or text) will be entered automatically when you enter the code.

The following table shows the Admin Key parameters.

- User-Name
- User- Password
- CHAP-password
- Filter-id
- Reply-Message
- Callback-Number
- Callback-id
- Framed-Route
- State
- Class
- Called-Station-id
- Calling-Station-id
- NAS-Identifier
- Proxy-State
- Login-LAT-Service
- Login-LAT-Node
- Login-LAT-Group
- Framed-Apple-Talk-Zone
- CHAP-Challenge
- Login-LAT-Port



Client Program Information

Definition

The SMS can be set up to receive information from RADIUS and passed to the Alcatel-Lucent IPsec Client about the program to be run after the tunnel is established and the action to be taken if the program returns an error or cannot be found.

Client Program

The value of the RADIUS attribute assigned to this parameter is sent to the Alcatel-Lucent IPsec Client as the program to be run after the tunnel is established. Other RADIUS attributes can be assigned to this parameter. This parameter must be assigned to a string attribute.

The default value of this attribute is **26/831/13**, which is in the format `26/<vendorID>/<vendorType>`.

Client Program Error Action

The value of the RADIUS attribute assigned to this parameter indicates whether the Alcatel-Lucent IPsec Client should drop the tunnel if an error is encountered when running the Client Program, or just ignore the error. Other RADIUS attributes can be assigned to this parameter. This parameter must be assigned to an integer attribute.

The default value of this attribute is **26/831/14**, which is in the format `26/<vendorID>/<vendorType>`. The integer values for the **Lucent-AAA-Client-Error-Action** attribute returned by RADIUS are 1 = ignore, 2 = drop.

How the feature works

Once the RADIUS server and SMS have been properly configured, the Client Program feature will work with Alcatel-Lucent IPsec Client users. The following explains how it works:

1. If the RADIUS attribute assigned to the Client Program parameter is received, the Alcatel-Lucent IPsec Client runs the specified program when the tunnel is established. The attribute contains the full path to the program to be run. If the path contains `%DC%`, the Alcatel-Lucent IPsec Client will substitute the IP address of the Domain Controller in its place.
2. If the value of the RADIUS attribute assigned to the Client Program Error Action parameter is 2(drop), the Alcatel-Lucent IPsec Client drops the tunnel if the run program returns an error or is not found. Otherwise, if the attribute is not received or if the value is 1(ignore), the Alcatel-Lucent IPsec Client will not drop the tunnel if the run program returns an error or is not found.

3. If the Reply-Message attribute is received from RADIUS, it is used in place of a banner message displayed to the Alcatel-Lucent IPsec Client user in a pop-up window, if the authentication and tunnel setup succeeds; otherwise, the banner message configured on the SMS for the Group of the Zone assigned to the client tunnel endpoint is displayed in a pop-up window. If the authentication or tunnel setup fails, an error message is displayed.
4. If the Client Program does not complete within a configured timeout value (default is 180 seconds), the Alcatel-Lucent IPsec Client considers the program hung, and attempts to run it a second time. If the program still does not run successfully, the Alcatel-Lucent IPsec Client considers this a failure and takes the specified error action (ignore or drop).



IKEv1 Preshared Key

Definition

IKEv1 Preshared Key is only used when the Unlicensed Mobile Access (UMA) feature is enabled. It defines the unique preshared key that is used for negotiating Phase 1 IKE when that user connects to the Client Tunnel Endpoint on the Brick. The RADIUS Attribute Type for this field is `string` or `text`.

The following table shows the IKEv1 Preshared Key parameters.

- User-Name
- User- Password
- CHAP-password
- Filter-id
- Reply-Message
- Callback-Number
- Callback-id
- Framed-Route
- State
- Class
- Called-Station-id
- Calling-Station-id
- NAS-Identifier
- Proxy-State
- Login-LAT-Service
- Login-LAT-Node
- Login-LAT-Group
- Framed-Apple-Talk-Zone
- CHAP-Challenge
- Login-LAT-Port



Timeout Information

Overview

An administrator can configure the SMS to process user timeout and idle timeout information contained in the RADIUS protocol message.

How to Configure the SMS and RADIUS Server

To configure the SMS and RADIUS server, follow the steps below:

- 1 Using the interface on the RADIUS server, configure a RADIUS attribute type to contain timeout information.
- 2 Create a RADIUS authentication service. On the **Attribute Codes** tab of the Authentication Service Editor, set the SMS Parameter to **User Timeout** or **Client-Tunnel Idle Timeout** and select one of the following RADIUS Attribute Codes. The RADIUS Attribute Type (always integer) will be entered automatically when you enter the code.

- NAS-Port
- Service-Type
- Framed-Protocol
- Framed-Routing
- Framed-MTU
- Framed-Compression
- Login-Service
- Login-TCP-Port
- Session-Timeout
- Idle-Timeout
- Termination-Action
- Framed-Apple-Talk-Link
- Framed-Apple-Talk-Network
- NAS-Port-Type
- Port-Limit

The procedure for creating a RADIUS authentication service is described in [“Overview” \(p. 8-1\)](#) (see [“To Set Up RADIUS and SecurID Authentication” \(p. 8-21\)](#)).

END OF STEPS

How the Feature Works

Once the RADIUS server and SMS have been properly configured, the User Timeout feature applies to both application users and Alcatel-Lucent IPSec Client users and the Client-Tunnel Idle Timeout feature only applies to IPSec client users. The following explains how it works:

1. When a user whose authentication service is configured to accept RADIUS timeouts is authenticated, the SMS will scan the RADIUS "Access-Accept" message for attributes of the type specified in the authentication service. (This is also true when the authentication service for external users is configured to accept RADIUS timeouts.)
2. If any of these attribute types are found in the RADIUS message, the SMS will parse the integer accompanying the attribute as a timeout value, in seconds. If multiple attributes of the same type are found, only the first attribute will be parsed.
The value returned from RADIUS is in seconds. The SMS converts the timeout value to minutes because the User Authentication Timeout and Client-Tunnel Idle Timeout fields are in minutes.
3. If a timeout attribute is returned from RADIUS, that value is used and overrides the default timeout specified for that VBA/Client Tunnel Endpoint. It also overrides any timeout value that might be assigned to that User even if that User is defined locally on the SMS.

□

Local Presence IP Information

Overview

An administrator can configure the SMS to process any local presence IP address contained in the RADIUS protocol message. A local presence IP address is an address on the local LAN that is temporarily given to Alcatel-Lucent IPsec Client users who are enabling a client tunnel.

How to Configure the SMS and RADIUS Server

To configure the SMS and RADIUS server, follow the steps below:

- 1 Using the interface on the RADIUS server, configure a RADIUS attribute type to contain the local IP address.
- 2 Create a RADIUS authentication service. On the **Attribute Codes** tab of the Authentication Service Editor, set the SMS Parameter to **Local Presence IP** and select one of the following RADIUS Attribute Codes. The RADIUS Attribute Type (address, string or text) will be entered automatically when you enter the code.

- User-Name
- User-Password
- CHAP-Password
- NAS-IP-Address
- Framed-IP-Address
- Framed-IP-Netmask
- Filter-id
- Login-IP-Host
- Reply-Message
- Callback-Number
- Callback-id
- Framed-Route
- Framed-IPX-Network
- State
- Class
- Called-Station-id
- Calling-Station-id
- NAS-Identifier
- Proxy-State
- Login-LAT-Service
- Login-LAT-Node
- Login-LAT-Group
- Framed-Apple-Talk-Zone
- CHAP-Challenge
- Login-LAT-Port

The procedure for creating a RADIUS authentication service is described in “Overview” (p. 8-1) (see “To Set Up RADIUS and SecurID Authentication” (p. 8-21)).

.....
E N D O F S T E P S
.....

How the Feature Works

Once the RADIUS server and SMS have been properly configured, the local presence IP feature will work with Alcatel-Lucent IPsec Client users. The following explains how it works:

1. When a user whose authentication service is configured to accept the RADIUS local IP address is authenticated, the SMS will scan the RADIUS “Access-Accept” message for attributes of the type specified in the authentication service. (This is also true when the authentication service for external users is configured to accept the RADIUS local IP address.)
To use this feature correctly, you must configure the Local Map Pool to NOT intersect with any IP addresses that may be coming from RADIUS. This will prevent the SMS from handing out an IP address that may be preassigned to an individual who is simply inactive at that time.
2. To allow the RADIUS server the flexibility of choosing a different local IP address (depending on the tunnel endpoint to which the client connected), the SMS will send as part of its RADIUS “Access-Request” message the “Tunnel-Server-Endpoint” attribute (type #67) and include a dotted-decimal representation of the actual endpoint to which the client connected (since the attribute is a string type).
3. If any of these attribute types are found in the RADIUS message, the SMS will parse the value accompanying the attribute as an IP address (directly if the attribute is an address type, or via string conversion if the attribute is string or text type). If multiple attributes of the same type are found, only the first attribute will be parsed.
4. If the SMS finds a local IP address in the RADIUS message, the SMS will present that address to the Brick to use as the local presence address, regardless of whether or not any other local presence addresses are configured in the SMS for that tunnel.
5. If the SMS determines that the local IP address in the RADIUS message has already been allocated to another tunnel, the SMS will NOT allow that client to authenticate. It will also log an error message indicating that a local IP address assigned by RADIUS had already been given out by the SMS.

6. If the SMS determines that the local IP address in the RADIUS message intersects with the configured Local Map Pool, but that the address has NOT been assigned to a particular user, the SMS will allow that client to authenticate. However, the SMS will log a warning message indicating that a local IP address assigned by RADIUS conflicts with the configured Local Map Pool. The SMS will NOT re-assign that particular IP address while the existing tunnel (using that same address) remains up.
7. The Brick, upon receiving a Local IP address for a tunnel that is NOT in its pre-configured Local Map Pool, will respond to ARP messages on behalf of that address automatically. When such a tunnel is disabled, the Brick will cease to respond to such ARP messages.



Primary WINS/Secondary DNS

Overview

An administrator can configure the SMS to process any WINS/DNS information contained in the RADIUS protocol message.

How to Configure the SMS and RADIUS Server

To configure the SMS and RADIUS server, follow the steps below:

- 1 Using the interface on the RADIUS server, configure a RADIUS attribute type to contain the primary WINS or secondary DNS address.
- 2 Create a RADIUS authentication service. On the **Attribute Codes** tab of the Authentication Service Editor, set the SMS Parameter to **Primary WINS, Secondary WINS, Primary DNS** or **Secondary DNS** and select one of the following RADIUS Attribute Codes. The RADIUS Attribute Type (address, string or text) will be entered automatically when you enter the code.

- User-Name
- User-Password
- CHAP-Password
- NAS-IP-Address
- Framed-IP-Address
- Framed-IP-Netmask
- Filter-id
- Login-IP-Host
- Reply-Message
- Callback-Number
- Callback-id
- Framed-Route
- Framed-IPX-Network
- Primary DNS Address¹
- Secondary DNS Address²
- State
- Class
- Called-Station-id
- Calling-Station-id
- NAS-Identifier
- Proxy-State
- Login-LAT-Service
- Login-LAT-Node
- Login-LAT-Group
- Framed-Apple-Talk-Zone
- CHAP-Challenge
- Login-LAT-Port
- Primary WINS Address³
- Secondary WINS Address⁴

Notes:

1. The default RADIUS attribute is **26/311/28 - MS-Primary-DNS-Server**
2. The default RADIUS attribute is **26/311/29 - MS-Secondary-DNS-Server**
3. The default RADIUS attribute is **26/311/30 - MS-Primary-NBNS-Server**
4. The default RADIUS attribute is **26/311/31 - MS-Secondary-NBNS-Server**

The procedure for creating a RADIUS authentication service is described in the procedural section [“To Set Up RADIUS and SecurID Authentication”](#) (p. 8-21).

END OF STEPS



Sample Configuration

RADIUS attribute codes

The following is an example of a typical SMS-RADIUS configuration:

Parameter	RADIUS Attribute Code	RADIUS Attribute Type
User Group	77	Text
User Timeout	27	Integer
Local Presence IP	8	Address
Idle Timeout	34	Integer
Primary WINS Address	34	String
Secondary WINS Address	36	String

Important! Configurations on the RADIUS server may make use of the fact that the Called-Station-Id attribute is set to the client TEP in the protocol message that the SMS sends to the RADIUS server for authentication.

□

Appendix E: Configuring the Brick for VoIP/NOE Traffic Using Pre-Defined SMS Templates

Overview

Purpose

This appendix provides step-by-step instructions and illustrations on how to set up a Brick device to serve as an Application Layer Gateway (ALG) in various VoIP deployment scenarios, using the pre-defined Brick zone ruleset, host group, service group, and application filter templates that are provided with the SMS application.

The following sections in the SMS documentation have useful background information about the deployment of a Brick device as an ALG in a VoIP network and the pre-defined templates that are provided with the SMS application to configure the Brick for this purpose:

- *Deployment of a Brick device as an application layer gateway (ALG) for VoIP/NOE phone communications* in the *SMS Administration Guide*
- [“Brick Zone Ruleset Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 1-66) in the *SMS Policy Guide*
- [“Host Group Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 2-25) in the *SMS Policy Guide*
- [“Service Group Templates Provided with the SMS Application for VoIP/NOE Traffic”](#) (p. 4-27) in the *SMS Policy Guide*
- [“TFTP Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-123) in the *SMS Policy Guide*
- [“NOE Application Filter Templates Provided With the SMS Application for VoIP/NOE Traffic”](#) (p. 5-68) in the *SMS Policy Guide*
- Chapter 5, [“Application Filters”](#) in the *SMS Policy Guide*

Contents

Configuring the Bricks and SMS Locations	E-3
VPN Hub and Spoke Configurations	E-7

Brick VPN Configurations	E-19
Multiple Zones Per Site	E-29
NATed Configurations	E-33

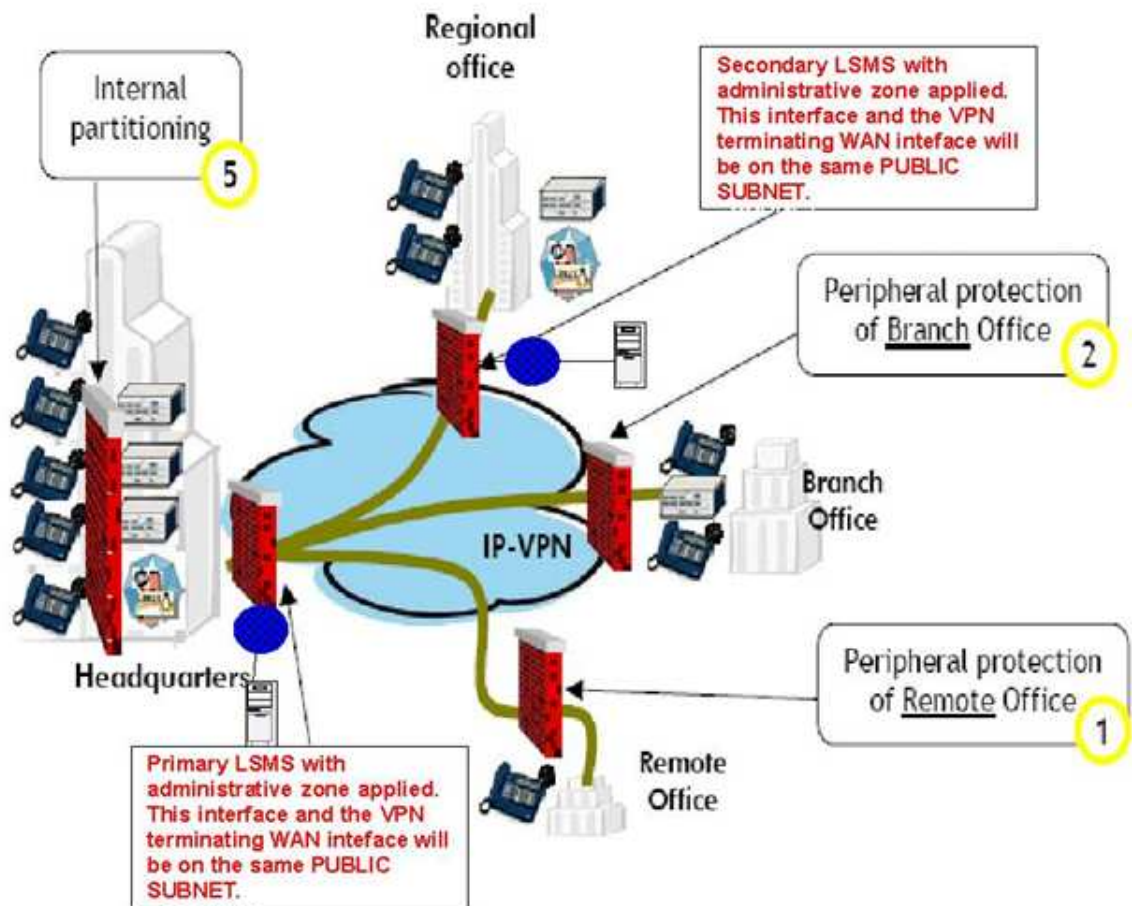


Configuring the Bricks and SMS Locations

Background

A Brick device, serving as an Application Layer Gateway (ALG) in a service provider's IP-based telecommunications network to secure VoIP communications, and its managing SMSs, can be deployed in any of the logical vantage point in a customer's enterprise VoIP network, as shown in [Figure E-1, "Deployment of Brick devices and SMSs in a Customer's Enterprise VoIP Network"](#) (p. E-3).

Figure E-1 Deployment of Brick devices and SMSs in a Customer's Enterprise VoIP Network



The Primary SMS (and Secondary SMS, if a redundant pair has been installed), must have public IP addresses and reside on the same subnet as the VPN terminating IP address. All of the Brick devices should also have public IP addresses. Assume, in this example, that the public Internet is a 192.168.168.3.0/24 network and the private network IP address is 10.1.1.0/24.

After installing a Primary SMS with a public IP address (in this example, 192.168.3.155), a Brick device is configured by assigning the proper public IP address (in this example, 192.168.3.1). Note that Ether0 is the WAN interface on the 192.168.3.0 subnet; Ether2 is the Brick interface to which the SMS will be connected. Ether1 is the 10.1.1.0/24 private network where the LAN zone will be applied. The Brick can now be flopped without assigning any zones (refer to [Figure E-2, “Brick Editor \(Assigning Public IP Address for Brick <-> SMS Communications\)”](#) (p. E-4) and [Figure E-3, “Brick Editor \(Physical Ports Tab, Showing Private IP Address Assigned to Ether1\)”](#) (p. E-5)).

Figure E-2 Brick Editor (Assigning Public IP Address for Brick <-> SMS Communications)

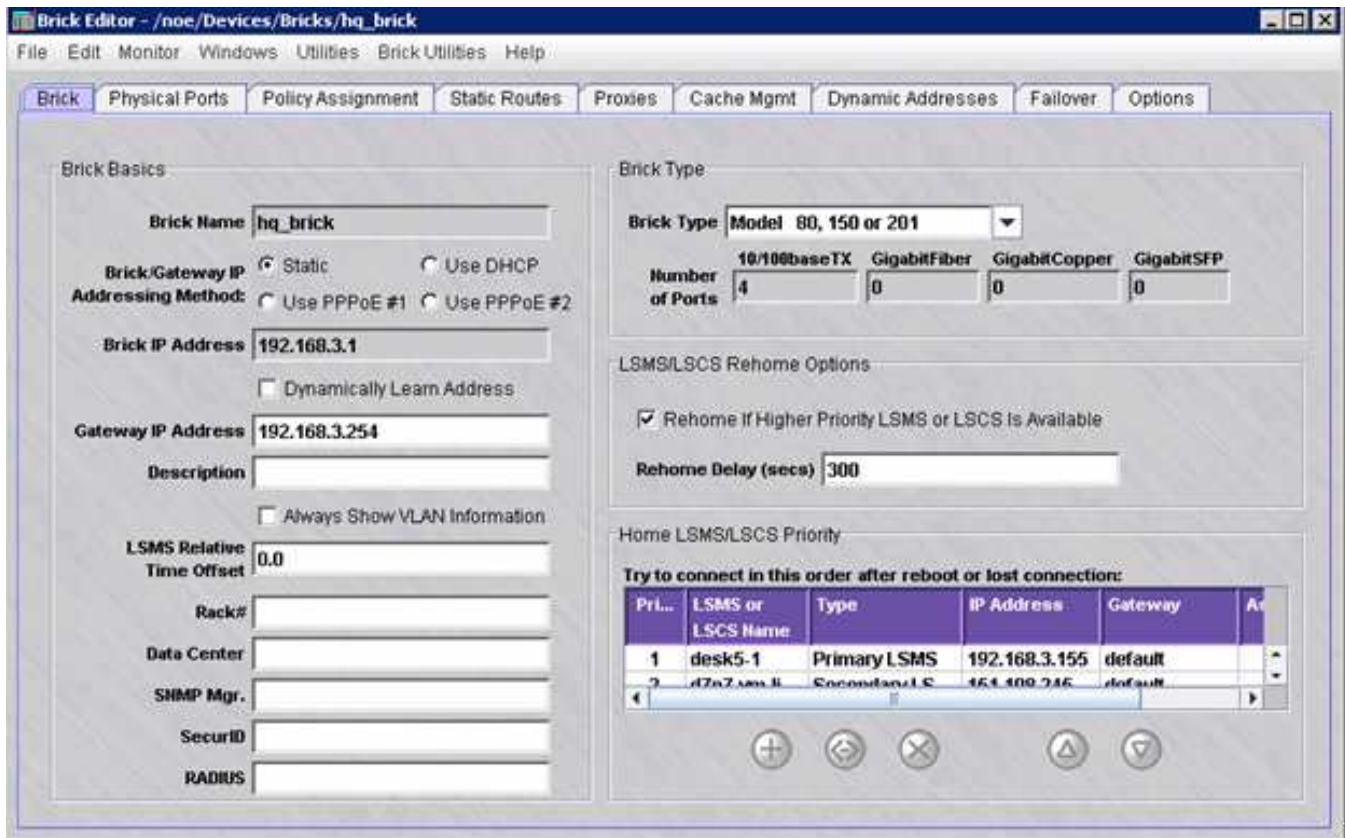
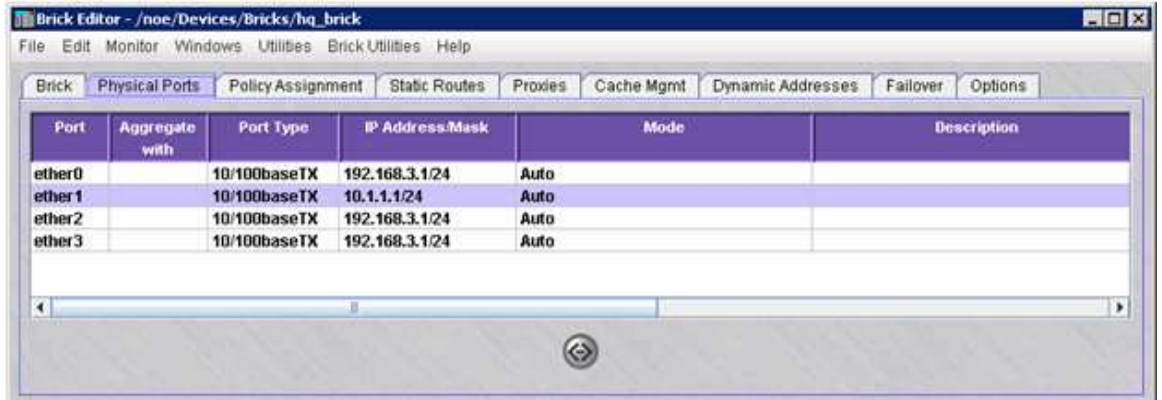


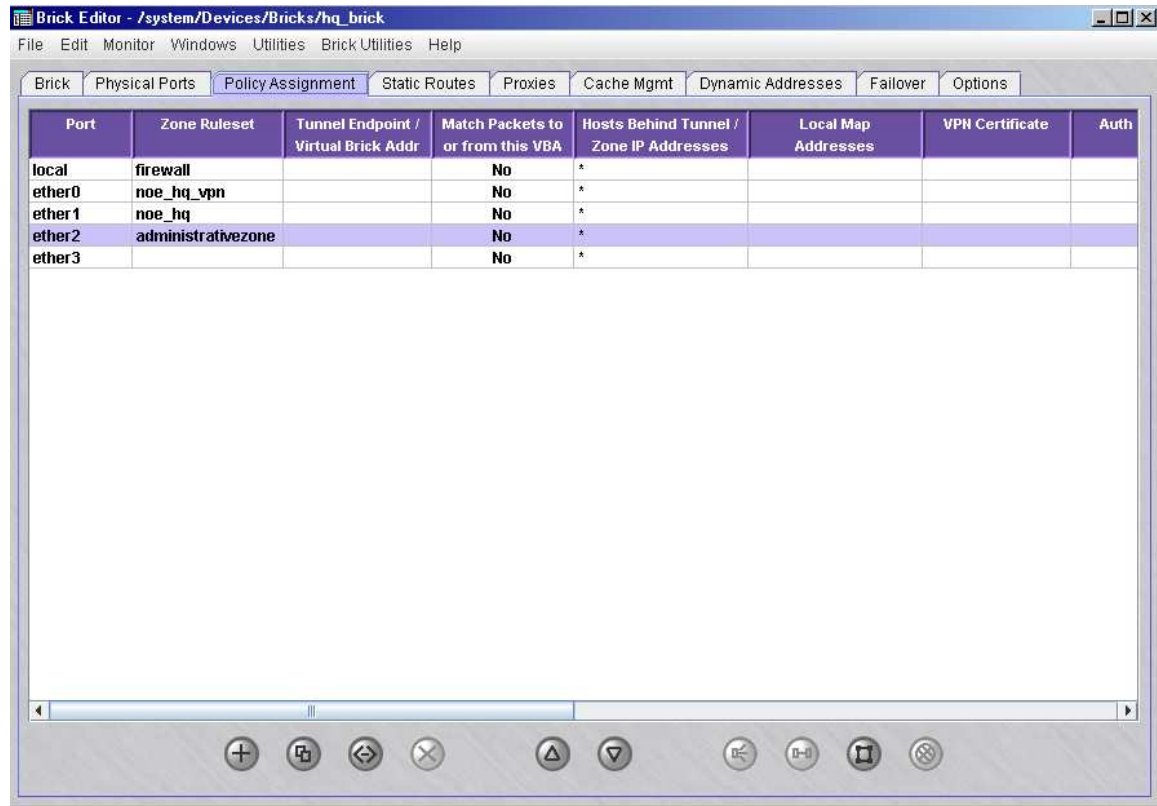
Figure E-3 Brick Editor (Physical Ports Tab, Showing Private IP Address Assigned to Ether1)



The same procedure is used to create all of the Bricks in the OXE VoIP deployment. Once all of the Bricks are up and the SMS is communicating with the Bricks, the pre-configured **administrativezone** zone ruleset can be assigned as shown in [Figure E-4, “Brick Editor \(Policy Assignment Tab, Showing administrativezone Ruleset and LSMS Host Group Assigned\)”](#) (p. E-6) and then the rest of the zone rulesets applied as indicated in the section [“VPN Hub and Spoke Configurations”](#) (p. E-7).

Note that for the **administrativezone**, in the **Hosts Behind Tunnel / Zone IP Addresses** field on the Policy Assignment Editor, the pre-defined host group **LSMS** should be selected, which contains the private IP address (and public IP addresses, if assigned) of the SMS(s) (refer to [Figure E-4, “Brick Editor \(Policy Assignment Tab, Showing administrativezone Ruleset and LSMS Host Group Assigned\)”](#) (p. E-6)).

Figure E-4 Brick Editor (Policy Assignment Tab, Showing administrativezone Ruleset and LSMS Host Group Assigned)



The same procedure is applied to the Regional Office Brick when assigning the Secondary SMS. For details about configuring a Primary SMS and Secondary SMS, refer to the *SMS Administration Guide*.

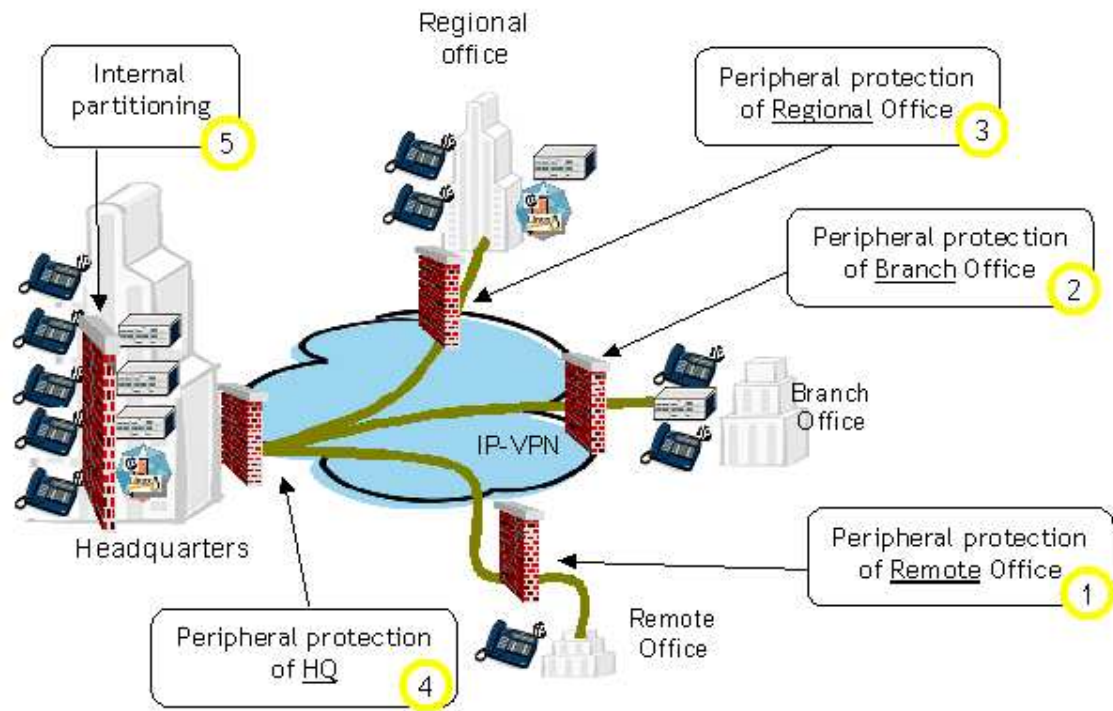
□

VPN Hub and Spoke Configurations

Overview


This section explains how to configure Brick devices that are positioned to protect VoIP equipment at various sites and are connected through an IP-VPN or IPsec VPN in a “hub and spoke” network architecture.

Refer to the following diagram when following the steps outlined for the various Brick deployment scenarios.



For each deployment scenario, the main provisioning task for setting up the Brick is to enter the correct IP addresses of the equipment that is located behind the Brick at the related site into the associated pre-defined host groups provided with the SMS. Pre-defined host groups have been created to identify the equipment at each possible deployment site (for example, **NOE_PHONES_Branch_Office** is a pre-defined host group in which you would enter the IP addresses of the IP phone devices located behind the Brick at a Branch Office site).

Brick VPN zones will always terminate on the WAN side. All other zones will terminate on the LAN side of the Brick (as depicted below).



When following the steps for configuring the Brick in the following sections, it is recommended that the Primary SMS have a public IP address and located at the Headquarters site; a secondary SMS can be located at the Regional Office site. Both SMSs should be using a dedicated interface on the Brick with the **administrativezone** ruleset applied.

It is assumed that the Bricks have been correctly configured and downloaded with the corrected SMS release software and that the SMS is able to communicate with the Bricks.

Headquarters site deployment

The following pre-defined host groups, available in the **NOE** subfolder under the **Host Groups** folder in the SMS Navigator, must be edited as follows before the pre-defined **noe_hq** ruleset can be assigned and applied to the Brick. One these host groups are edited, the other zones are also ready to be applied.

- **NOE_TFTP_Server_HQ**— this host group must be filled in with the IP address of the TFTP server located behind the Brick on the LAN side at the Headquarters site. If there are two servers (no spatial redundancy), then enter the IP addresses of both servers.
- **NOE_TFTP_Server_Regional**— this host group must be filled in with the IP address of the TFTP server located at the Regional Office, if one is deployed (WAN side).
- **NOE_Call_Server_HQ**— this host group must be filled in with the IP address of the call server located behind the Brick on the LAN side at the Headquarters site. If there are two servers (no spatial redundancy), then enter the IP addresses of both servers.
- **NOE_Call_Server_Regional**— this host group must be filled in with the IP address of the call server located at the Regional Office, if one is deployed (WAN side).

- **NOE_PHONES_Headquarters**— this host group must be filled in with the IP addresses of all IP phone devices located behind the Brick on the LAN side at the Headquarters site.
- **NOE_PHONES_Regional_Office**— this host group must be filled in with the IP addresses of all IP phone devices located at the Regional Office site.
- **NOE_PHONES_Branch_Office**— this host group must be filled in with the IP addresses of all IP phone devices located at the Branch Office site.
- **NOE_PHONES_Remote_Office**— this host group must be filled in with the IP addresses of all IP phone devices located at the Remote Office site.
- **NOE_GD_IPs_Headquarters**— this host group must be filled in with the IP addresses of all GD boards of the MGW unit(s) located behind the Brick at the Headquarters site.
- **NOE_GA_IPs_Headquarters**— this host group must be filled in with the IP addresses of all GA boards of the MGW unit(s) located behind the Brick at the Headquarters site.
- **NOE_GD_IPs_Regional_Office**— this host group must be filled in with the IP addresses of all GD boards of the MGW unit(s) located behind the Brick at the Regional Office site.
- **NOE_GA_IPs_Regional_Office**— this host group must be filled in with the IP addresses of all GA boards of the MGW unit(s) located behind the Brick at the Regional Office site.
- **NOE_GD_IPs_Branch_Office**— this host group must be filled in with the IP addresses of all GD boards of the MGW unit(s) located behind the Brick at the Branch Office site.
- **NOE_GA_IPs_Branch_Office**— this host group must be filled in with the IP addresses of all GA boards of the MGW unit(s) located behind the Brick at the Branch Office site.
- **NOE_Presentation_Server**— this host group must be filled in with the IP address of the Presentation Server.
- **NOE_PCS**— this host group must be filled in with the IP address of the Passive Communication Server (PCS).

There are pre-defined service groups and host groups for 4670 server communications to clients and call servers. However, there are no rules in the pre-defined zone rulesets to allow these communications. If this traffic must be supported, then the appropriate pre-defined host groups must be filled in with the IP addresses of this equipment and assigned to rules within the appropriate zone ruleset with the proper service groups and direction. Use the existing rules in the ruleset as an example.

If the 4760 server is in the Headquarters zone, then the pre-defined host group **NOE_Omni_Vista_4760_Server** must be filled in with the IP address of this server and added (nested) within the pre-defined host group **NOE_VPN_HQ_HUB**. Any

new hosts that are added to the Headquarters zone must also be added to the proper host group nested within the **NOE_VPN_HQ_HUB** host group. Use the existing rules in the ruleset as an example.

The proper communication for the **NOE_PCS_to_CS** and **NOE_CS_to_PCS** service groups must be selected from the rules within the zone ruleset. Deactivate the rules that are not being used. If secured communication is used, choose the secured rules and deactivate the non-secured rules. The opposite is true if non-secured communication is used.

Regional office site deployment

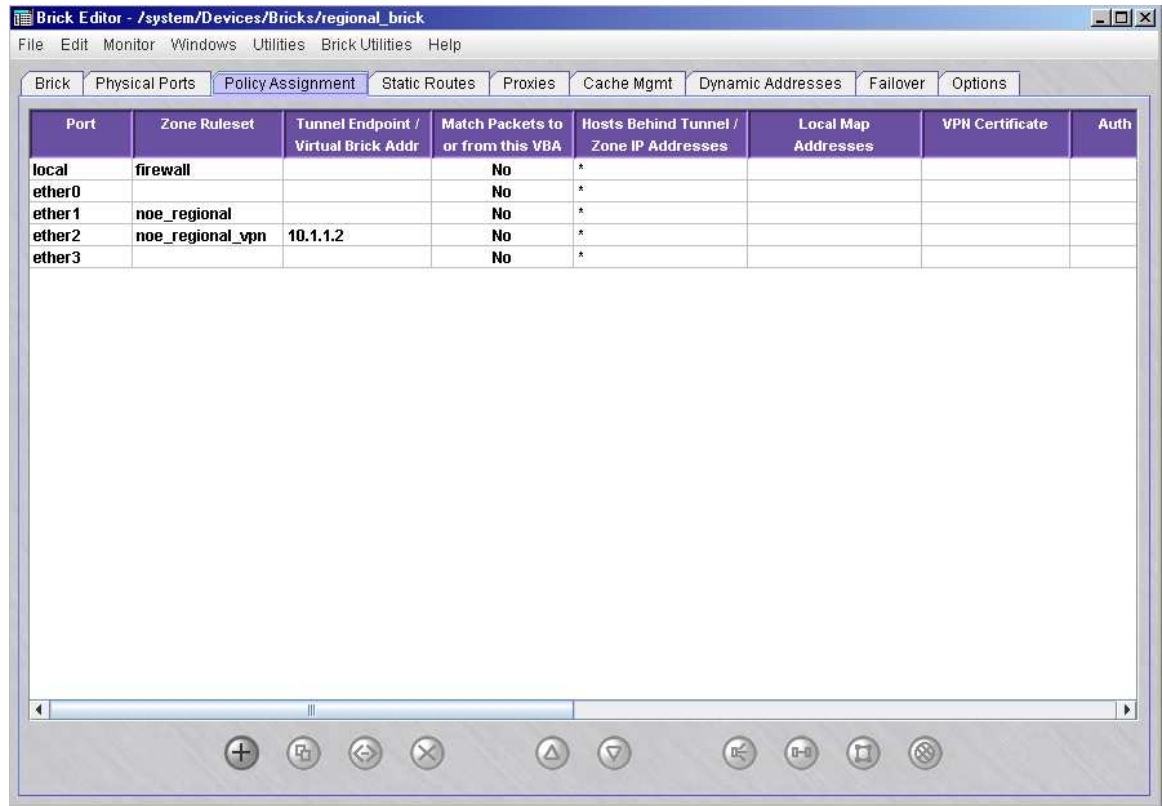
All pre-defined host groups indicated in the section [“Headquarters site deployment” \(p. E-8\)](#) must be completed before applying the Regional Office zone rulesets. Once these host groups have been edited, the zones can be applied to the Regional Office Brick in a manner similar to what is shown in [Figure E-5, “Brick Editor \(Policy Assignment Tab, Showing Pre-Defined Regional Office Zone Rulesets Applied\)” \(p. E-11\)](#). The interfaces, Brick names, and Brick models may be different than the example shown. Note that a Virtual Brick Address (VBA) is required for the VPN zone on the WAN side.

There are pre-defined service groups and host groups for 4670 server communications to clients and call servers. However, there are no rules in the pre-defined zone rulesets to allow these communications. If this traffic must be supported, then the appropriate pre-defined host groups must be filled in with the IP addresses of this equipment and assigned to rules within the appropriate zone ruleset with the proper service groups and direction. Use the existing rules in the ruleset as an example.

If the 4760 server is in the Regional zone, then the pre-defined host group **NOE_Omni_Vista_4760_Server** must be filled in with the IP address of this server and added (nested) within the pre-defined host groups **NOE_VPN_Regional**, **NOE_VPN_Outside_HQ**, and **NOE_VPN_HQ_HUB**. Any new hosts that are added to the Headquarters zone must also be added to these host groups. Use the existing rules in the ruleset as an example.

The proper communication for the **NOE_PCS_to_CS** and **NOE_CS_to_PCS** service groups must be selected from the rules within the zone ruleset. Deactivate the rules that are not being used. If secured communication is used, choose the secured rules and deactivate the non-secured rules. The opposite is true if non-secured communication is used.

Figure E-5 Brick Editor (Policy Assignment Tab, Showing Pre-Defined Regional Office Zone Rulesets Applied)



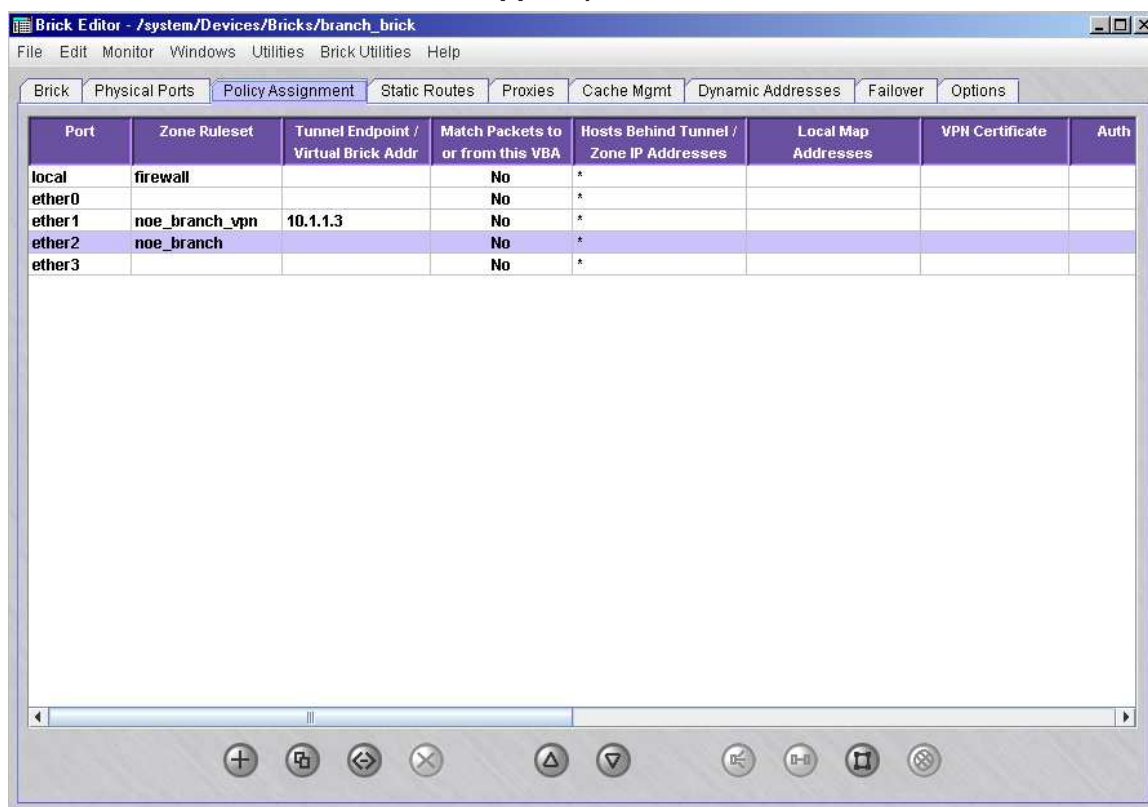
Branch office site deployment

All pre-defined host groups indicated in the section [“Headquarters site deployment”](#) (p. E-8) must be completed before applying the Branch Office zone rulesets. Once these host groups have been edited, the zones can be applied to the Branch Office Brick in a manner similar to what is shown in [Figure E-6, “Brick Editor \(Policy Assignment Tab, Showing Pre-Defined Branch Office Zone Rulesets Applied\)”](#) (p. E-12). The interfaces, Brick names, and Brick models may be different than the example shown. Note that a Virtual Brick Address (VBA) is required for the VPN zone on the WAN side.

Any new hosts that are added to the Branch zone must also be added to the **NOE_VPN_Branch**, **NOE_VPN_Outside_HQ**, and **NOE_VPN_HQ_HUB** host groups and proper rules added to the zone to allow the desired communication. Use the existing rules in the ruleset as an example.

The proper communication for the **NOE_PCS_to_CS** and **NOE_CS_to_PCS** service groups must be selected from the rules within the zone ruleset. Deactivate the rules that are not being used. If secured communication is used, choose the secured rules and deactivate the non-secured rules. The opposite is true if non-secured communication is used.

Figure E-6 Brick Editor (Policy Assignment Tab, Showing Pre-Defined Branch Office Zone Rulesets Applied)



Duplicating the noe_branch zone templates for multiple branch office sites

The **noe_branch** zone ruleset template can be copied and configured for Brick deployments at multiple Branch Office sites. The related host group copies of the duplicated **noe_branch** zone ruleset must be edited by adding the IP addresses of the equipment at the Branch Office site (call servers, IP phone devices, GA/GD boards of MGW units) as explained in the section [“Headquarters site deployment”](#) (p. E-8).

The **noe_branch** zone ruleset and the templates for its related “components” (host groups, application filters, service groups) can be copied multiple times and renamed as needed for Branch Office site, using the Component Copy tool, as explained in the [“To copy a Brick zone ruleset”](#) (p. 1-55) section in Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”.

When using the Component Copy tool to copy the **noe_branch** ruleset, make sure that:

- You give the **noe_branch** zone ruleset a new name if it is in the same group (Example: **noe_branch_1**).
- The **Copy components when copying within a Group** checkbox on the Copy window is checked, to make copies of all of its related components.
- The **Component Suffix** field on the Copy window is filled in with a text string that will be appended to the name of each component, to associate it with the zone ruleset itself (Example: **_X**, where **_X** is entered as **_1** to label all components associated with the **noe_branch_1** ruleset, **_2** to rename and label all components associated with the **noe_branch_2** zone ruleset, and so forth).

Note: the Component Suffix used must be unique and cannot be reused across zone rulesets. For example, if you enter “_2” as the Component Suffix for all component copies associated with the noe_branch_2 zone ruleset, you would have to enter a different Component Suffix for all component copies associated with the noe_regional_2 ruleset; e.g., “_2a” or “_2x”, or similarly distinguished numbering/labeling scheme for the suffix applied to component copies of each zone ruleset.

- The **Component Folder** field on the Copy window is entered as **NOE**.

After the Copy is performed, some additional manual updates have to be made to the newly created components as follows:

1. Four new host groups will be created: **NOE_Phones_Branch_Office_X**, **NOE_GD_IPs_Branch_Office_X**, **NOE_GA_IPs_Branch_Office_X**, **NOE_GD_Public_IPs_Branch_Office_X**, and **NOE_GA_Public_IPs_Branch_Office_X**, where **X** is the suffix used in the Component Copy command.
2. A new host group **NOE_MGWs_Branch_Office_X** is created. This host contains the host groups **NOE_GD_IPs_Branch_Office_X** and **NOE_GA_IPs_Branch_Office_X**.
3. The host groups **NOE_Phones_Branch_Office_X** and **NOE_MGWs_Branch_Office_X** must be manually added to the host groups **NOE_Phones_MGW_outside_HQ** and **NOE_Phones_MGW_outside_Regional**.
*Note: if using NAT at the branch office site, then **NOE_GD_Public_IPs_Branch_Office_X** and **NOE_GA_Public_IPs_Branch_Office_X** must be added to the outside Regional and outside HQ host groups, instead of **NOE_MGWs_Branch_Office_X**.*
4. The host group **NOE_GD_IPs_Branch_Office_X** must be manually added to the host group **NOE_GD_outside_Regional**.
*Note: if using NAT at the branch office site, then **NOE_GD_Public_IPs_Branch_Office_X** must be added to the outside Regional and outside HQ host groups instead.*

5. The host group **NOE_GA_IPs_Branch_Office_X** must be manually added to the host groups **NOE_GA_IPs_outside_HQ** and **NOE_GA_IPs_outside_Regional**
Note: if using NAT at the branch office site, then **NOE_GA_Public_IPs_Branch_Office_X** must be added to the outside Regional and outside HQ host groups instead.
6. The host groups **NOE_Phones_Branch_Office_X** and **NOE_MGWs_Branch_Office_X** must be manually added to the host group **NOE_VPN_HQ_HUB**.
7. The host groups **NOE_Phones_Branch_Office_X** and **NOE_MGWs_Branch_Office_X** must be manually added to the host group **NOE_VPN_Outside_HQ**.
8. A new host group **NOE_PHONES_MGW_inside_Branch_X** is created. This host group contains the host groups **NOE_Phones_Branch_Office_X** and **NOE_MGWs_Branch_Office_X**.
9. Create a regular non-component copy of the host group **NOE_VPN_Branch** and name it **NOE_VPN_Branch_X**. Manually add the host groups **NOE_Phones_Branch_Office_X** and **NOE_MGWs_Branch_Office_X** and remove the original entries.
10. A new NOE application filter **NOE_NAT_Branch_X** is created. The application filter copy entries are the same as the original one, except that the Hosts Behind Brick value is **NOE_PHONES_MGW_inside_Branch_X** and the MGW IPLink entries are **NOE_GD_IPs_Branch_Office_X**, **NOE_GA_IPs_Branch_Office_X**, **NOE_GD_Public_IPs_Branch_Office_X**, and **NOE_GA_Public_Branch_Office_X**.
11. A new NOE application filter **NOE_Standard_inside_Branch_X** is created. The application filter copy entries are the same as the original one, except that the Hosts Behind Brick value is **NOE_PHONES_MGW_inside_Branch_X** and the IPLink entries are **NOE_GD_IPs_Branch_Office_X** and **NOE_GA_IPs_Branch_Office_X**.
12. A new TFTP application filter **TFTP_with_NOE_NAT_Branch_X** is created. Entries in this copy are the same as the original, except that it calls the NOE application filter copy **NOE_NAT_Branch_X**.
13. A new TFTP application filter **TFTP_with_NOE_inside_Branch_X** is created. Entries in this copy are the same as the original one, except that it calls the NOE application filter copy **NOE_Standard_inside_Branch_X**.
14. A new service group **NOE_Service_NAT_Branch_X** is created. Entries in this copy are the same as the original one, except that it calls the application filter copy **TFTP_with_NOE_NAT_Branch_X**.
15. A new service group **NOE_Service_Standard_inside_Branch_X** is created. Entries in this copy are the same as the original one, except that it calls the application filter copy **TFTP_with_NOE_inside_Branch_X**.

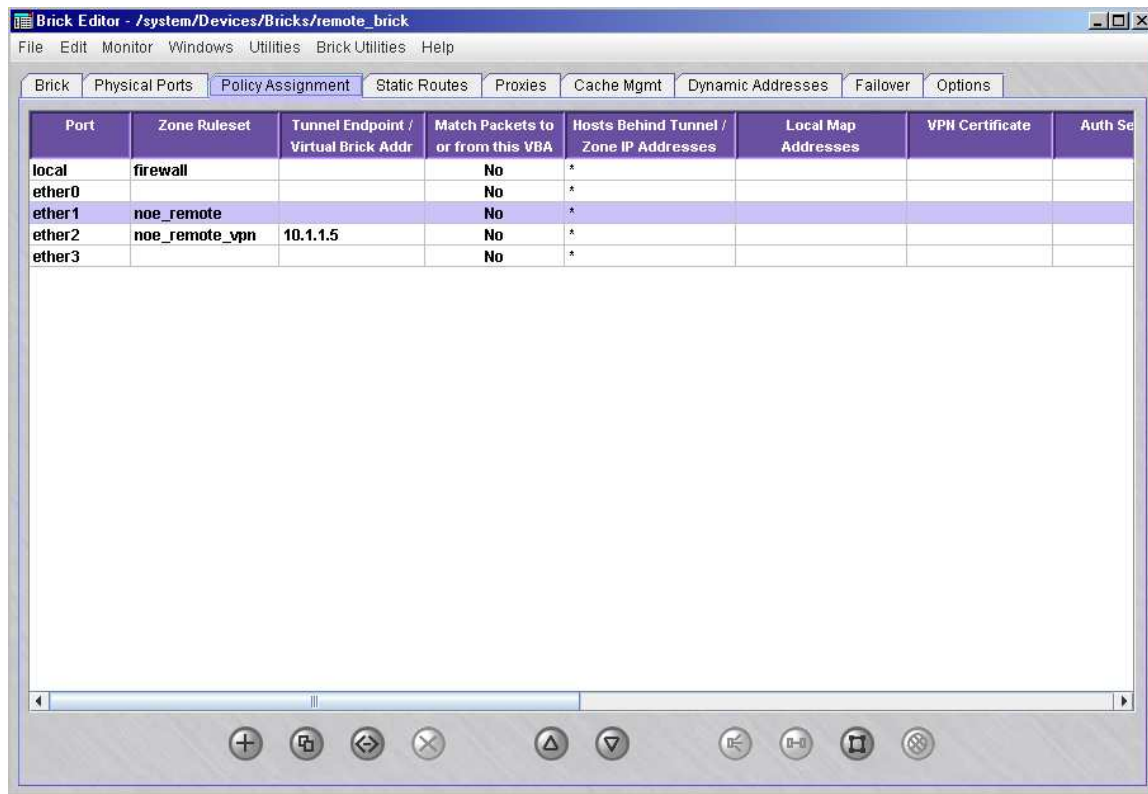
16. A new zone ruleset **noe_branch_x** is created (applying a different name than the original one if it is in the same group). This zone ruleset copy has all of the same rules as the original one, except that all references to the **noe_branch** host groups and service groups have been replaced with the **noe_branch_x** groups and services just created.
17. Manually perform a non-component copy of the existing **noe_branch_vpn** zone ruleset and call it **noe_branch_vpn_X**. This will have all the same rules as the previous zone ruleset. There are no changes to this zone ruleset and it can be copied exactly and renamed.

Remote office site deployment

All pre-defined host groups indicated in the section [“Headquarters site deployment” \(p. E-8\)](#) must be completed before applying the Remote Office zone rulesets. Once these host groups have been edited, the zones can be applied to the Remote Office Brick in a manner similar to what is shown in [Figure E-7, “Brick Editor \(Policy Assignment Tab, Showing Pre-Defined Remote Office Zone Rulesets Applied\)” \(p. E-16\)](#). The interfaces, Brick names, and Brick models may be different than the example shown. Note that a Virtual Brick Address (VBA) is required for the VPN zone on the WAN side.

Any new hosts that are added to the Remote zone must also be added to the **NOE_VPN_Remote**, **NOE_VPN_Outside_HQ**, and **NOE_VPN_HQ_HUB** host groups and proper rules added to the zone to allow the desired communication. Use existing rules in the ruleset as examples.

Figure E-7 Brick Editor (Policy Assignment Tab, Showing Pre-Defined Remote Office Zone Rulesets Applied)



Duplicating the noe_remote zone templates for multiple remote office sites

The **noe_remote** zone ruleset template can be copied and configured for Brick deployments at multiple Branch Office sites. The related host group copies of the duplicated **noe_remote** zone ruleset must be edited by adding the IP addresses of the equipment at the Branch Office site (call servers, IP phone devices, GA/GD boards of MGW units) as explained in the section [“Headquarters site deployment”](#) (p. E-8).

The **noe_remote** zone ruleset and the templates for its related “components” (host groups, application filters, service groups) can be copied multiple times and renamed as needed for Branch Office site, using the Component Copy tool, as explained in the [“To copy a Brick zone ruleset”](#) (p. 1-55) section in Chapter 1, *“Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”*.

When using the Component Copy tool to copy the **noe_branch** ruleset, make sure that:

- You give the **noe_remote** zone ruleset a new name if it is in the same group (Example: **noe_remote_1**).
- The **Copy components when copying within a Group** checkbox on the Copy window is checked, to make copies of all of its related components

- The **Component Suffix** field on the Copy window is filled in with a text string that will be appended to the name of each component, to associate it with the zone ruleset itself (Example: **_X**, where **_X** is entered as **_1** to label all components associated with the **noe_remote_1** ruleset, **_2** to rename and label all components associated with the **noe_remote_2** zone ruleset, and so forth.
Note: the Component Suffix used must be unique and cannot be reused across zone rulesets. For example, if you enter “_2” as the Component Suffix for all component copies associated with the noe_branch_2 zone ruleset, you would have to enter a different Component Suffix for all component copies associated with the noe_remote_2 ruleset; e.g., “_2a” or “_2x”, or similarly distinguished numbering/labeling scheme for the suffix applied to component copies of each zone ruleset.
- The **Component Folder** field on the Copy window is entered as **NOE**.

After the Copy is performed, some additional manual updates have to be made to the newly created components as follows:

1. A new host group **NOE_Phones_Remote_Office_X** will be created, where **X** is the Component Suffix used in the copy command.
2. You must manually add the host group **NOE_Phones_Remote_Office_X** to the host groups **NOE_Phones_MGW_outside_HQ** and **NOE_Phones_MGW_outside_Regional**.
3. The host group **NOE_Phones_Remote_Office_X** must be manually added to the host group **NOE_VPN_HQ_HUB**.
4. The host group **NOE_Phones_Remote_Office_X** must be manually added to the host group **NOE_VPN_Outside_HQ**.
5. Create a regular non-component copy of the host group **NOE_VPN_Remote** and name it **NOE_VPN_Remote_X**. Manually add the host group **NOE_Phones_Remote_Office_X** to the host group **NOE_VPN_Remote_X**.
6. A new application filter **NOE_NAT_Remote_X** is created. The entries are the same as the original one, except that the Hosts Behind Brick value is **NOE_Phones_Office_X**.
7. A new application filter **NOE_Standard_inside_Remote_X** is created. The entries are the same as the original one, except that the Hosts Behind Brick value is **NOE_Phones_Remote_Office_X**.
8. A new TFTP application filter **TFTP_with_NOE_NAT_Remote_X** is created. The entries are the same as the original one, except that it calls the application filter copy **NOE_NAT_Remote_X**.
9. A new TFTP application filter **TFTP_with_NOE_inside_Remote_X** is created. The entries are the same as the original one, except that it calls the application filter copy **NOE_Standard_inside_Remote_X**.
10. A new service group **NOE_Service_NAT_Remote_X** is created. The entries are the same as the original one, except that it calls the application filter copy **TFTP_with_NOE_NAT_Remote_X**.

11. A new service group **NOE_inside_Remote_X** is created. The entries are the same as the original one, except that it call the application filter copy **TFTP_with_NOE_inside_Remote_X**.
12. A new zone ruleset **noe_remote_X** is created. This will have all the same rules as the previous zone ruleset, except all references to **noe_remote** are now replaced with the **noe_remote_X** groups and services just created.
13. Manually perform a non-component copy of the zone ruleset **noe_remote_vpn** and call it **noe_remote_vpn_X**. This will have all the same rules as the previous zone ruleset. There are no changes to this zone ruleset and it can be copied exactly and renamed.

□

Brick VPN Configurations

Overview

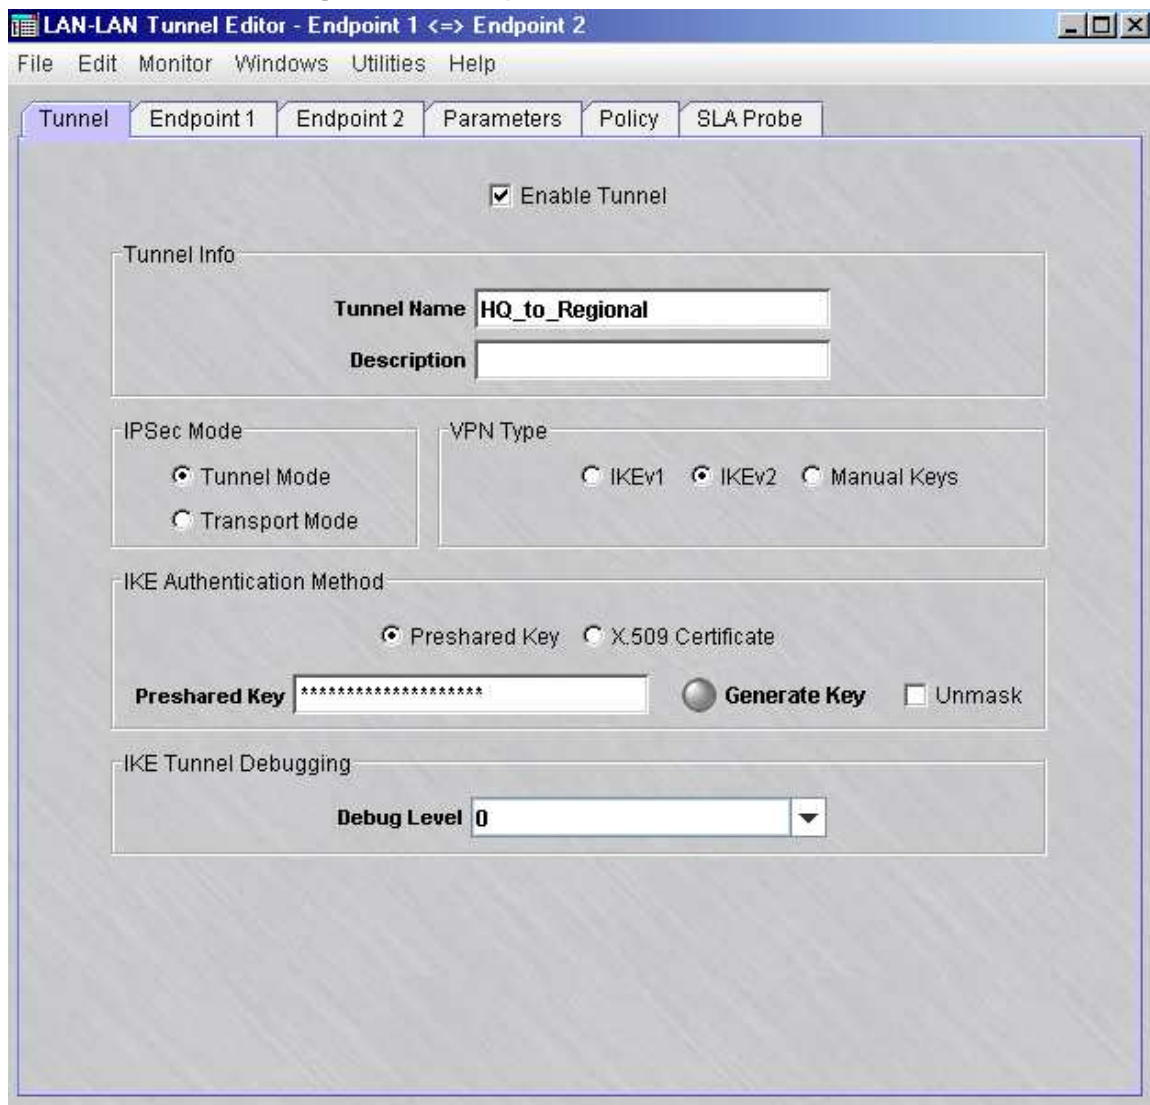
Once all of the Bricks have been applied and configured as shown in the previous sections, VPN tunnels can be established between the various sites. Open the LAN-LAN Tunnel Editor and configure tunnels as shown below. *Note: In each case, make sure that IKEv2 is selected for the VPN Type and allow the Brick to generate the Preshared Key.*

To create a headquarters-to-regional office tunnel

Complete the following steps:

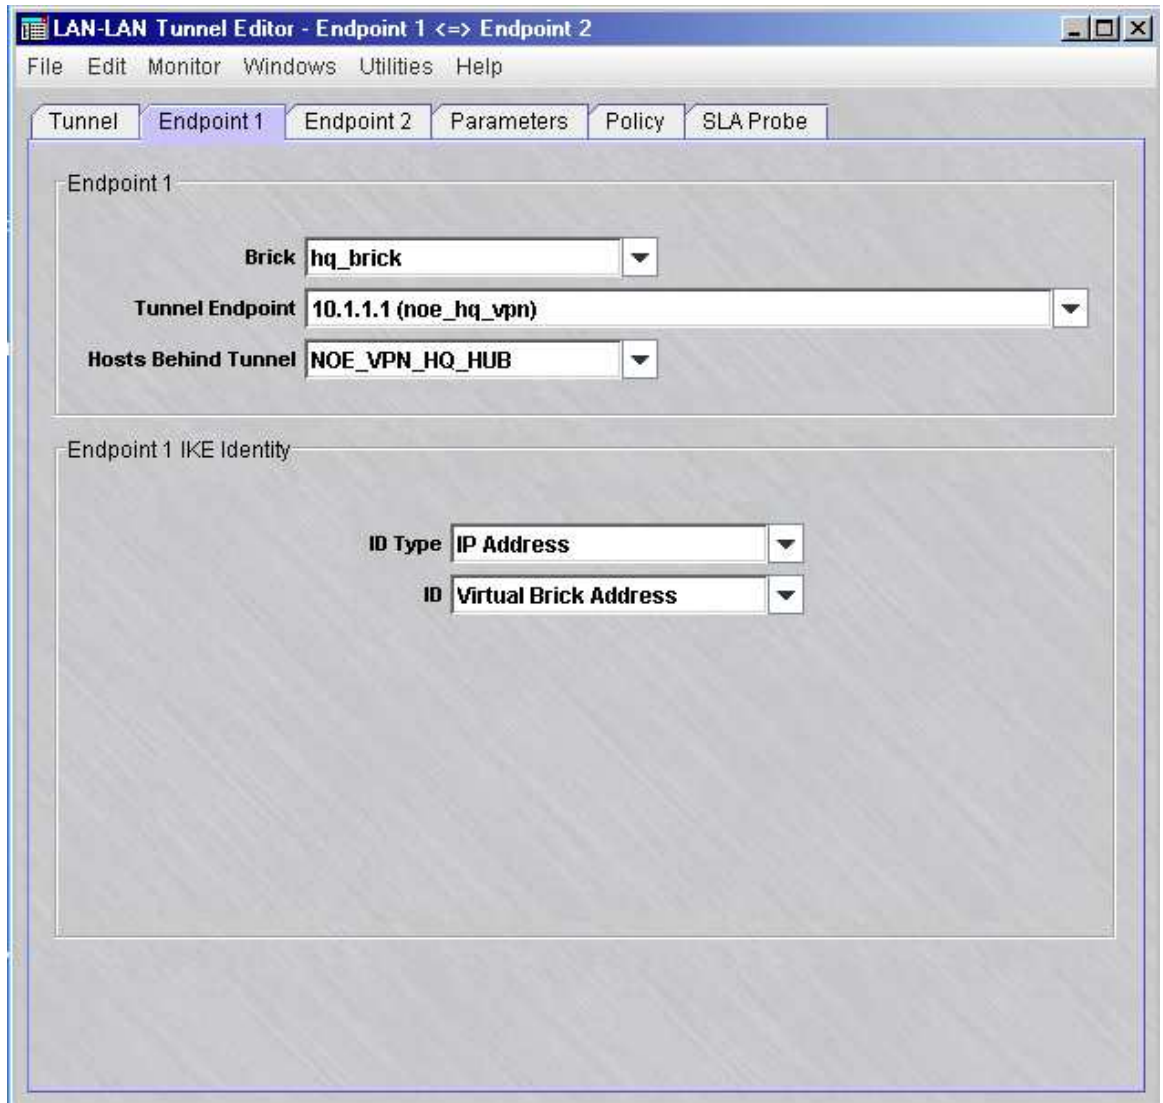
- 1 Open the LAN-LAN Tunnel Editor and enter a Tunnel Name for the tunnel ([Figure E-8, “LAN-LAN Tunnel Editor \(Showing Parameter Settings for HQ_to_Regional Tunnel\)”](#) (p. E-20)).

**Figure E-8 LAN-LAN Tunnel Editor (Showing Parameter Settings for
HQ_to_Regional Tunnel)**



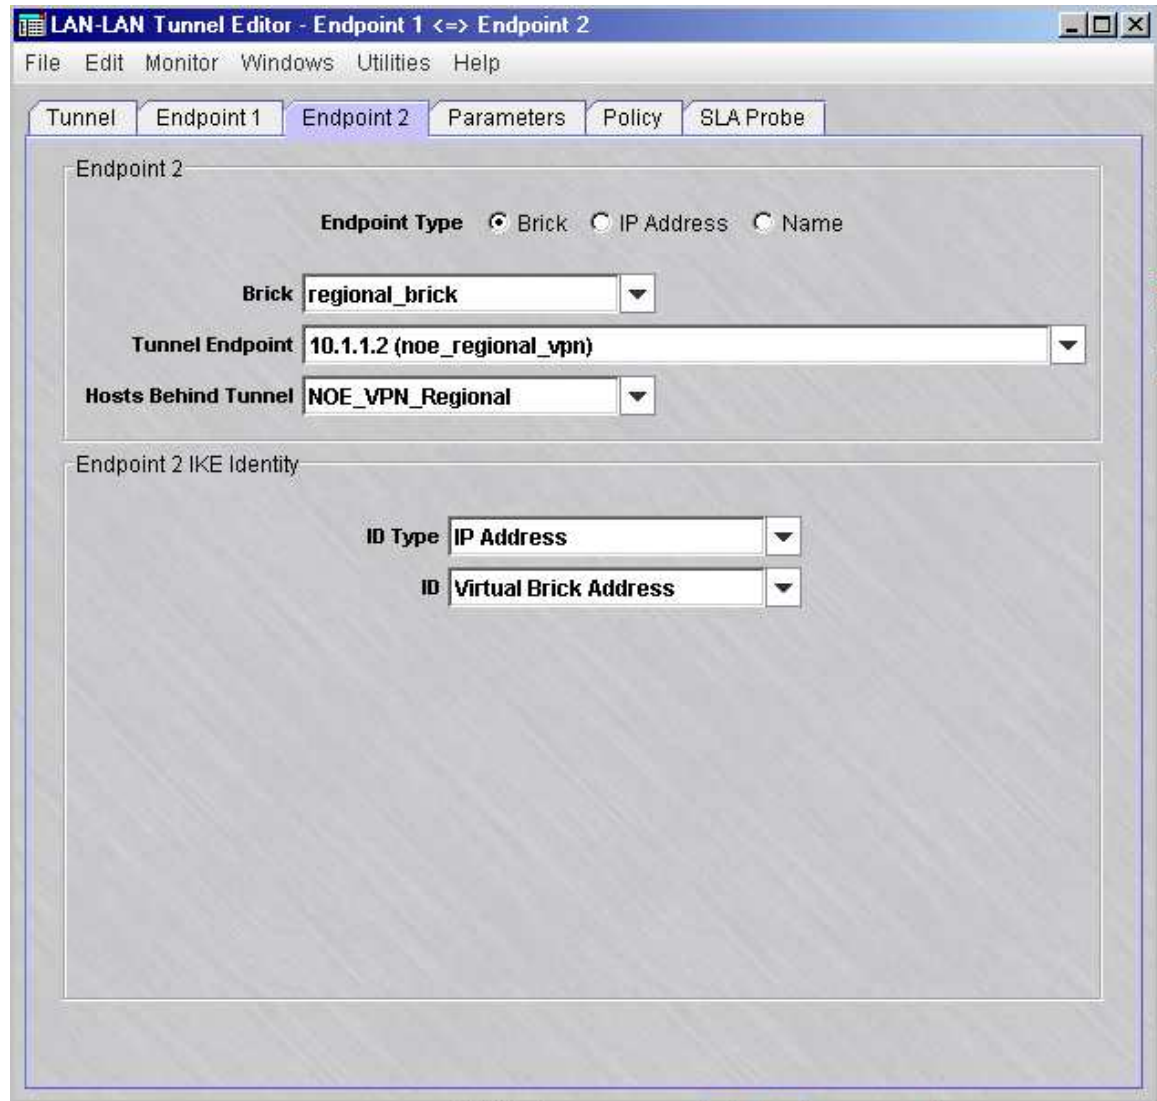
- 2 Choose the Headquarters site Brick as Endpoint1, **noe_hq_vpn** as the tunnel endpoint, and the pre-defined host group **NOE_VPN_HQ_HUB** in the **Hosts Behind Tunnel** field (Figure E-9, “LAN-LAN Tunnel Editor (Endpoint1 Tab, Showing Headquarters Brick and Endpoint 1 Assignment)” (p. E-21)).

Figure E-9 LAN-LAN Tunnel Editor (Endpoint1 Tab, Showing Headquarters Brick and Endpoint 1 Assignment)



- 3 Choose the Regional Brick as Endpoint2, **noe_regional_vpn** as the tunnel endpoint, and **NOE_VPN_Regional** in the **Hosts Behind Tunnel** field (Figure E-10, “LAN-LAN Tunnel Editor (Endpoint2 Tab, Showing Regional Brick and Endpoint 2 Assignment)” (p. E-22)).

Figure E-10 LAN-LAN Tunnel Editor (Endpoint2 Tab, Showing Regional Brick and Endpoint 2 Assignment)



-
- 4 Save and apply the tunnel configuration to the Brick.

This same process is repeated for each of the tunnels in the Hub and Spoke office site configuration.

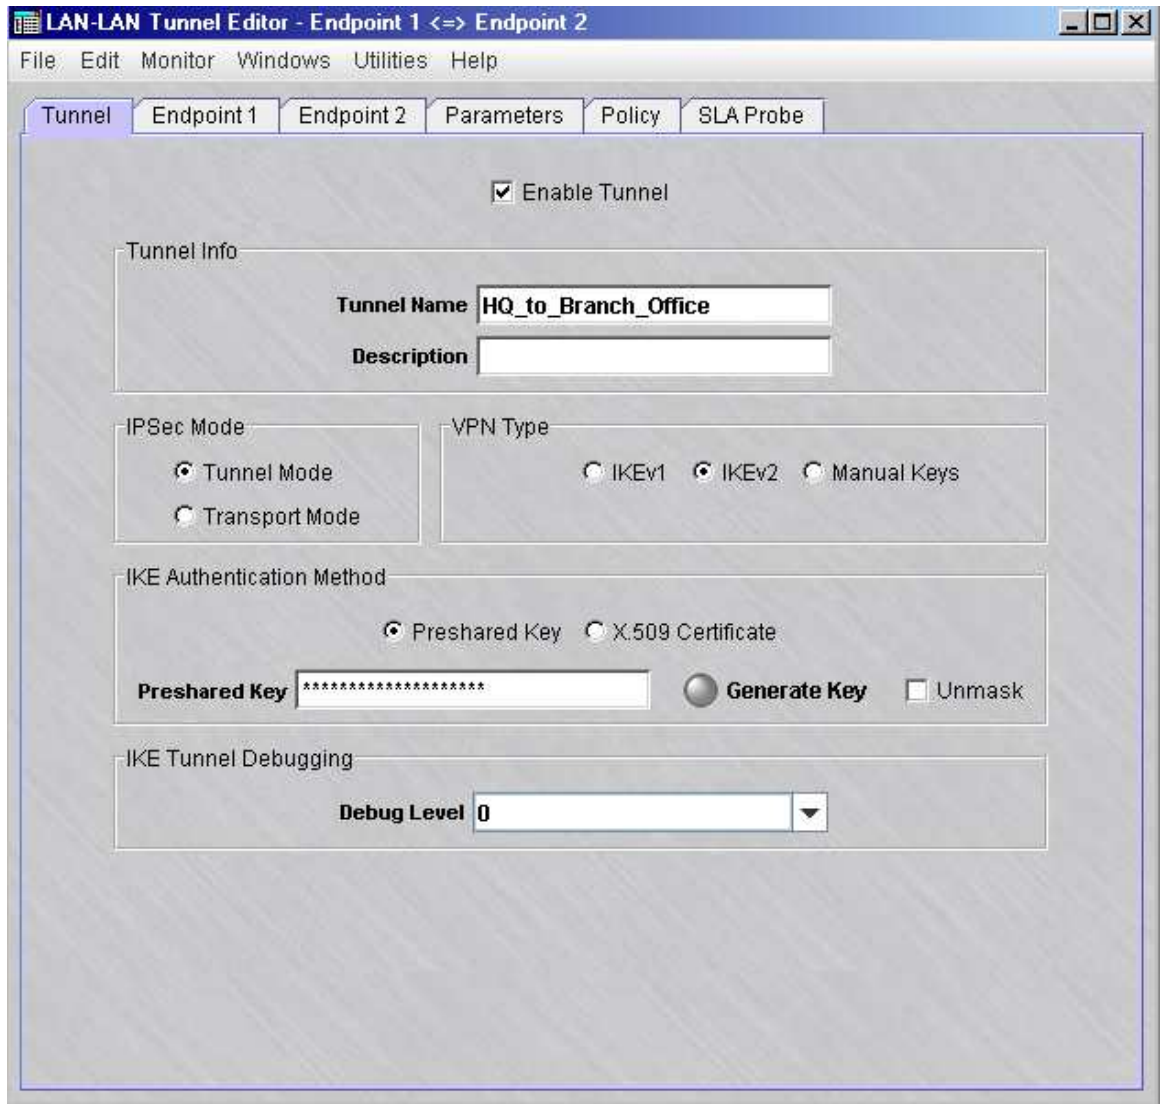
END OF STEPS

To create a headquarters-to-branch office tunnel

Complete the following steps:

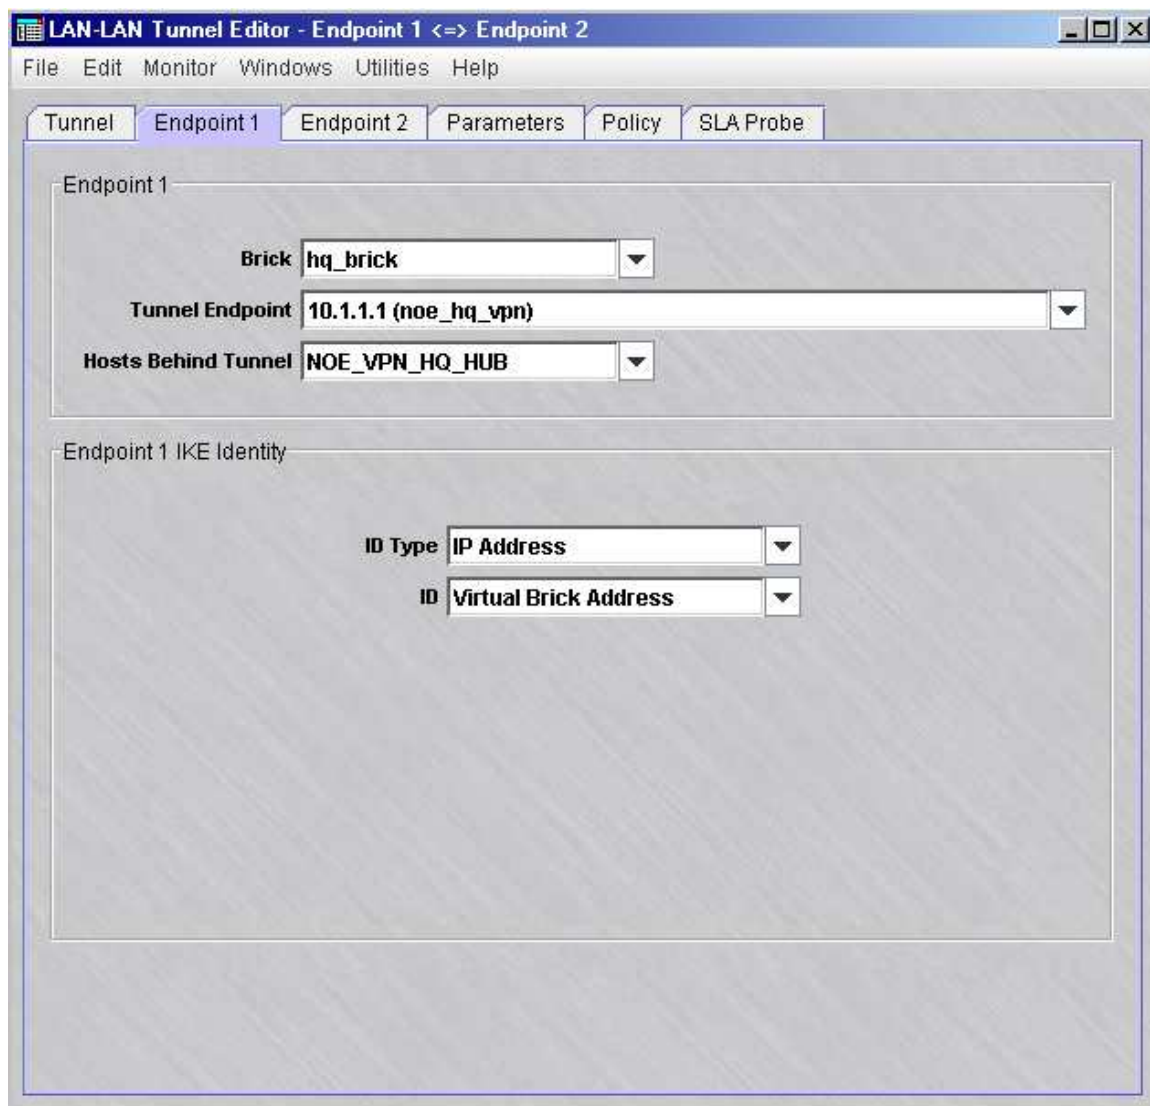
- 1 Open the LAN-LAN Tunnel Editor and enter a Tunnel Name for the tunnel (Figure E-11, “LAN-LAN Tunnel Editor (Showing Parameter Settings for HQ_to_Branch_Office Tunnel)” (p. E-23)).

Figure E-11 LAN-LAN Tunnel Editor (Showing Parameter Settings for HQ_to_Branch_Office Tunnel)



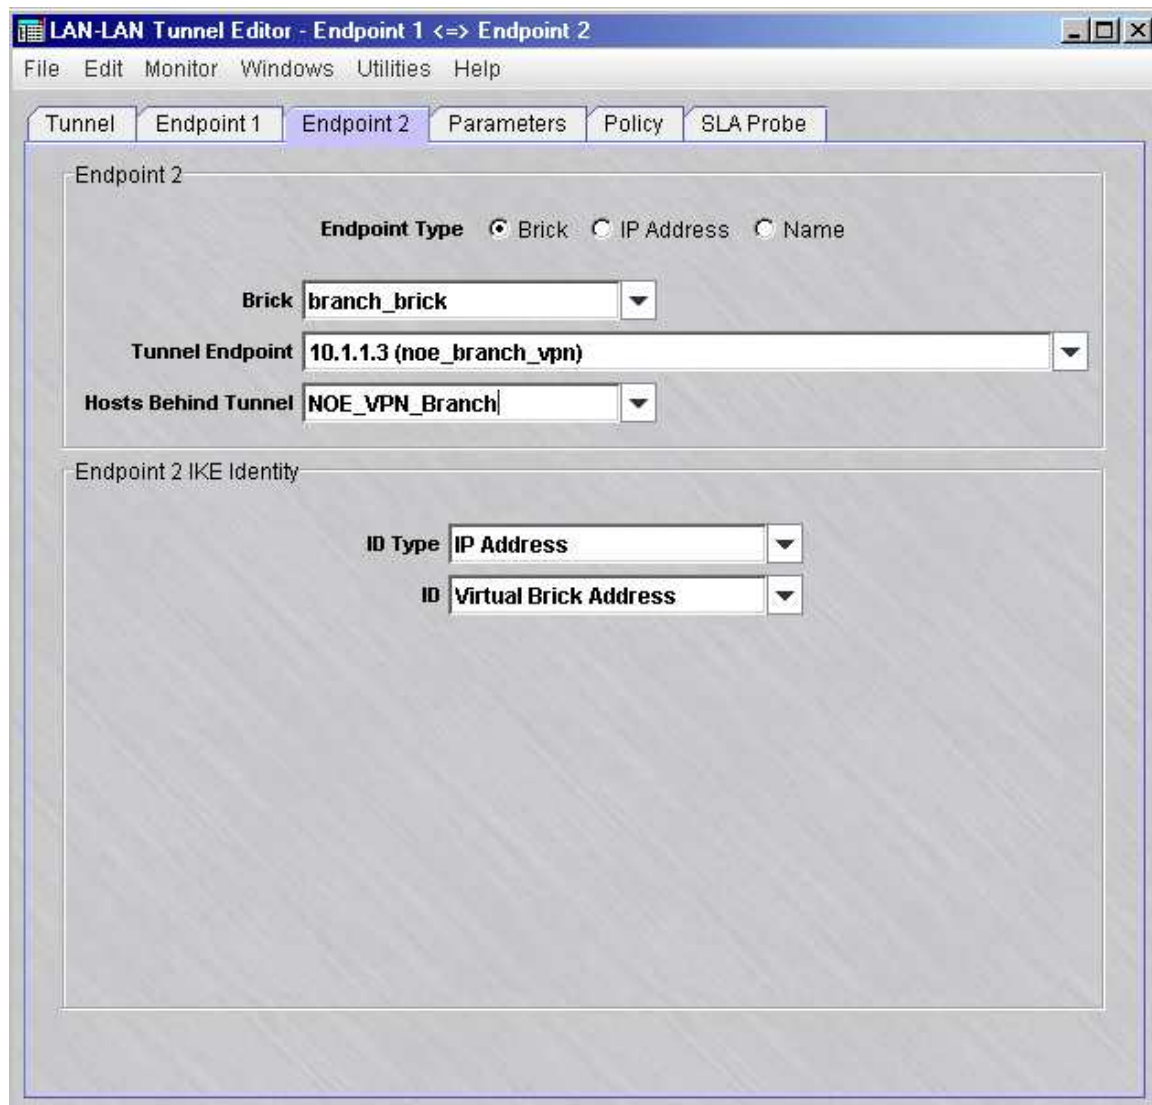
- 2 Choose the Headquarters Brick as Endpoint1, **noe_hq_vpn** as the tunnel endpoint, and the pre-defined host group **NOE_VPN_HQ_HUB** in the **Hosts Behind Tunnel** field (Figure E-12, “LAN-LAN Tunnel Editor (Endpoint 1 Tab, Showing Headquarters Brick and Endpoint 1 Assignment)” (p. E-24)).

Figure E-12 LAN-LAN Tunnel Editor (Endpoint 1 Tab, Showing Headquarters Brick and Endpoint 1 Assignment)



- 3 Choose the Branch Office Brick as Endpoint2, **noe_branch_vpn** as the tunnel endpoint, and **NOE_VPN_Branch** as the hosts behind the tunnel (Figure E-13, “LAN-LAN Tunnel Editor (Endpoint 2 Tab, Showing Branch Office Brick and Endpoint 2 Assignment)” (p. E-25)).

**Figure E-13 LAN-LAN Tunnel Editor (Endpoint 2 Tab, Showing Branch Office
Brick and Endpoint 2 Assignment)**



-
- 4 Save and apply the tunnel configuration to the Brick.

This same process is repeated for each of the tunnels in the Hub and Spoke office site configuration.

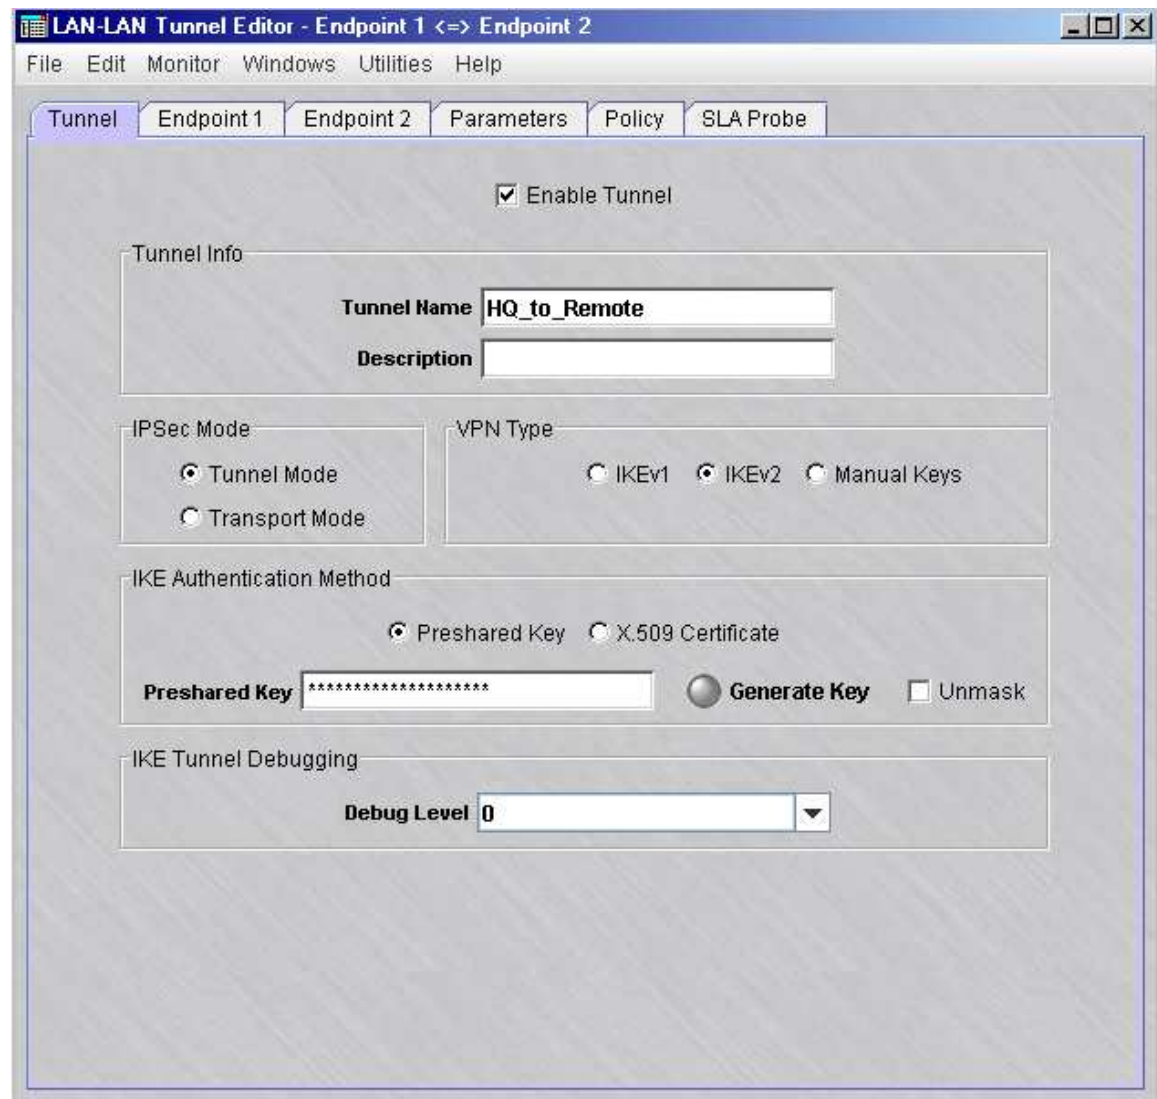
END OF STEPS

To create a headquarters-to-remote office tunnel

Complete the following steps:

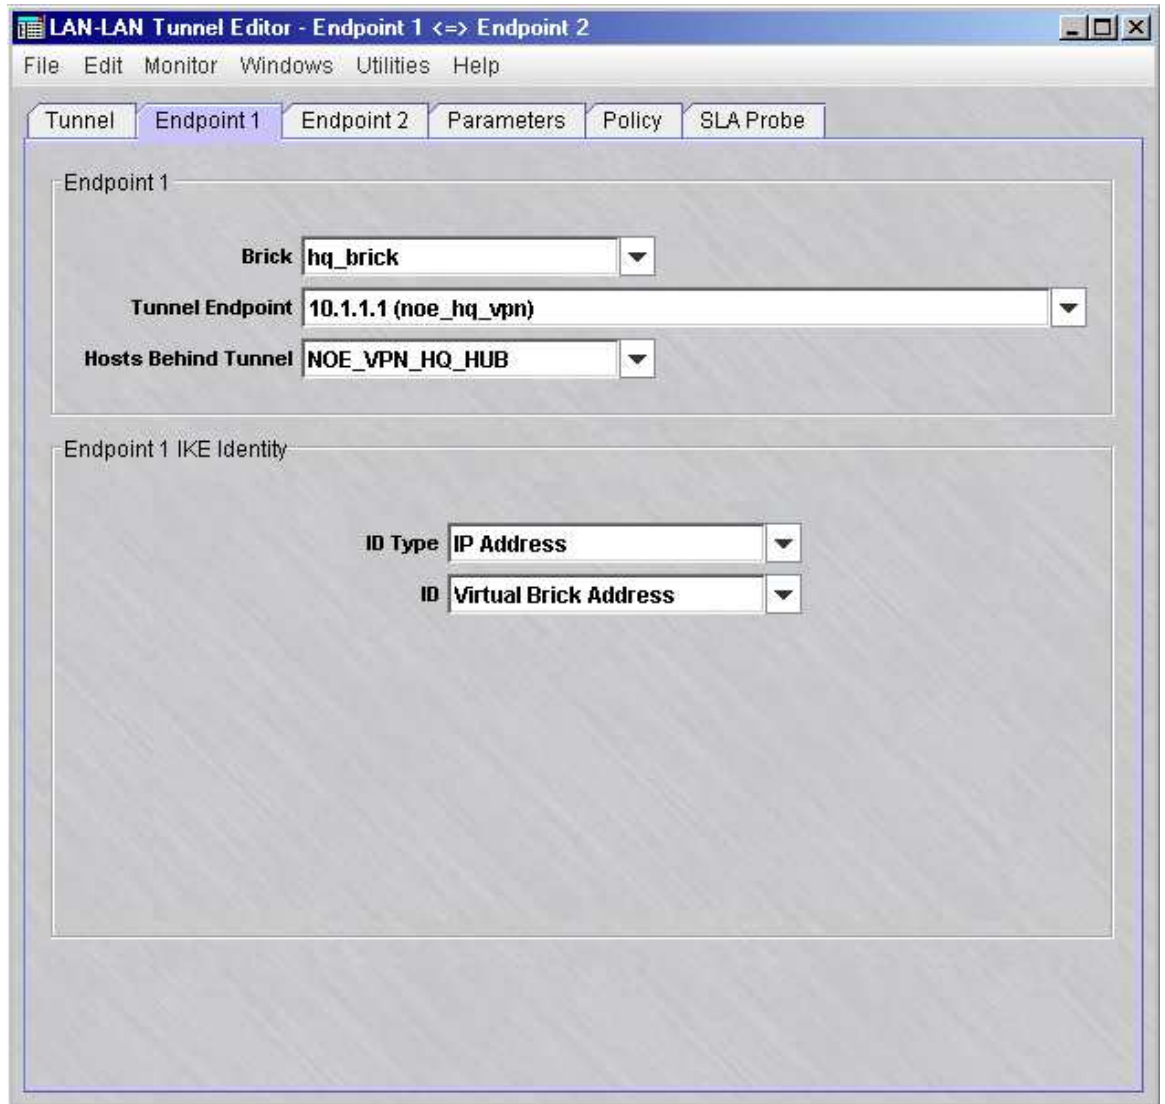
- 1 Open the LAN-LAN Tunnel Editor and enter a Tunnel Name for the tunnel (Figure E-14, “LAN-LAN Tunnel Editor (Showing Parameter Settings for HQ_to_Remote Tunnel)” (p. E-26)).

Figure E-14 LAN-LAN Tunnel Editor (Showing Parameter Settings for HQ_to_Remote Tunnel)



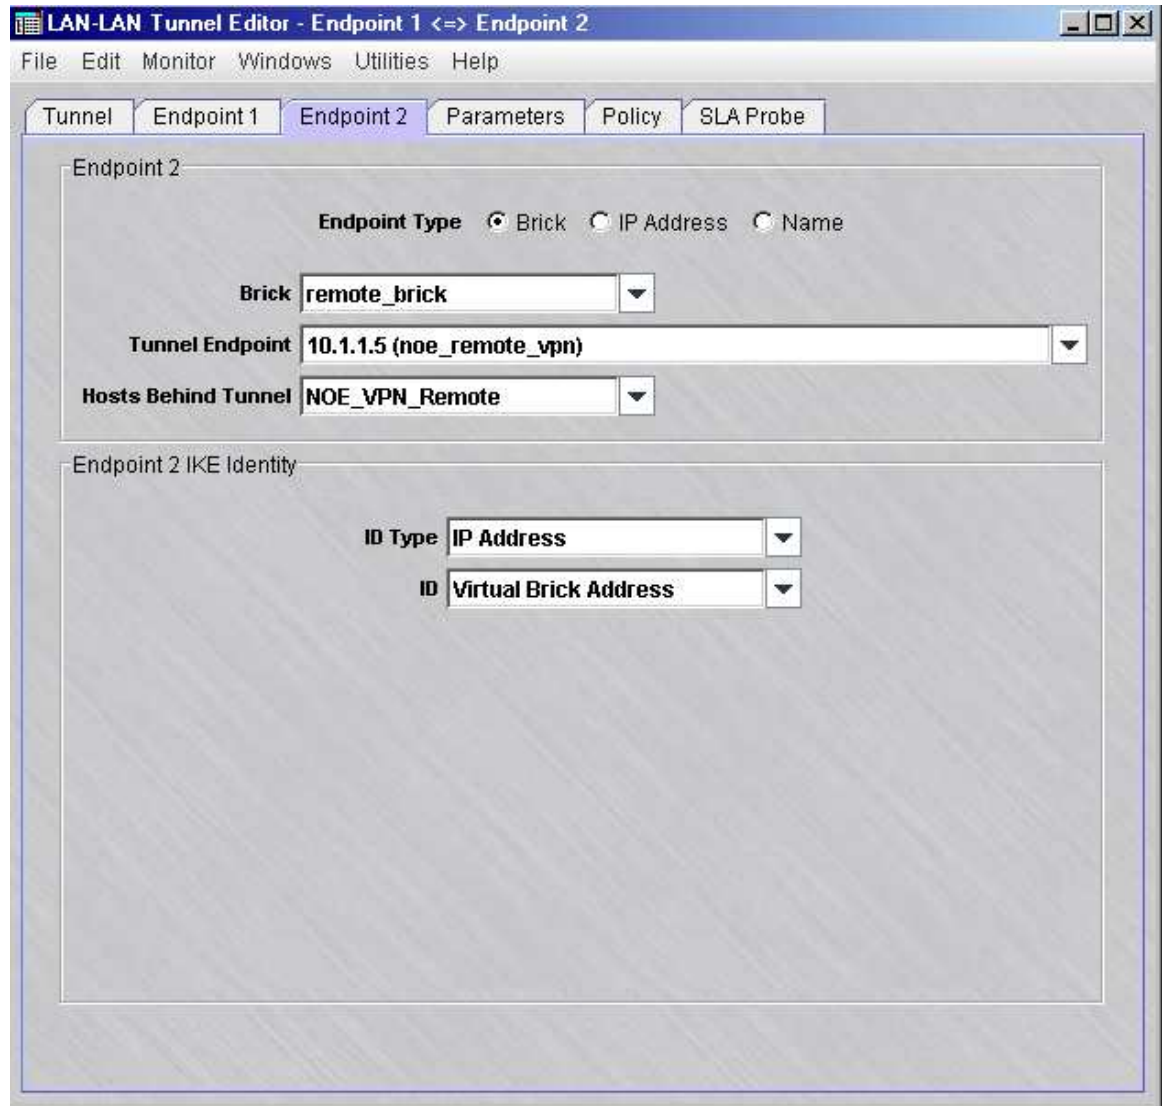
- 2 Choose the Headquarters Brick as Endpoint1, **noe_hq_vpn** as the tunnel endpoint, and **NOE_VPN_HQ_HUB** in the **Hosts Behind Tunnel** field (Figure E-15, “LAN-LAN Tunnel Editor (Endpoint 1 Tab, Showing Headquarters Brick and Endpoint 1 Assignment)” (p. E-27)).

Figure E-15 LAN-LAN Tunnel Editor (Endpoint 1 Tab, Showing Headquarters Brick and Endpoint 1 Assignment)



- 3 Choose the Branch Office Brick as Endpoint2, **noe_remote_vpn** as the tunnel endpoint, and **NOE_VPN_Remote** as the hosts behind the tunnel (Figure E-16, “LAN-LAN Tunnel Editor (Endpoint 2 Tab, Showing Branch Office Brick and Endpoint 2 Assignment)” (p. E-28)).

**Figure E-16 LAN-LAN Tunnel Editor (Endpoint 2 Tab, Showing Branch Office
Brick and Endpoint 2 Assignment)**



-
- 4 Save and apply the tunnel configuration to the Brick.

All tunnels should now be up and communications can be tested for proper operation.

END OF STEPS



Multiple Zones Per Site

Overview

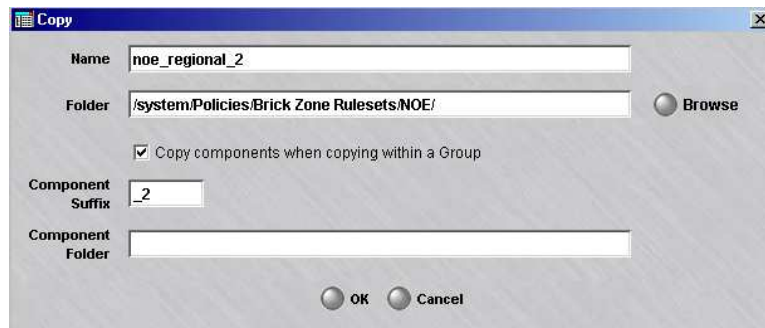
A Brick, acting as a VoIP ALG, supports up to 2,000 phone calls per zone. Multiple zones will have to be used at each site where there can be more than 2,000 IP phone devices. In this case, all that is required is to make a copy of the zone and place this new zone as another zone instance on the LAN interface. The following procedure shows you how to approach this situation, using the Component Copy tool as described in the “To copy a Brick zone ruleset” (p. 1-55) section in [Chapter 1, “Alcatel-Lucent VPN Firewall Brick™ Security Appliance Zone Rulesets”](#) to make multiple copies of a zone ruleset that can be applied at the Brick deployment site.

- 1 Right-click the zone ruleset that needs to be copied and select **Copy** from the pop-up menu.

Result The Copy window is displayed.

- 2 Enter a new name for the duplicate zone ruleset, for example, **noe_regional_2** ([Figure E-17, “Copy Window \(Renaming Duplicate Zone Ruleset\)”](#) (p. E-29)).

Figure E-17 Copy Window (Renaming Duplicate Zone Ruleset)



It is recommended that you copy and rename the associated components of the zone ruleset template by checking the **Copy components when copying within a Group** checkbox. It is also recommended that you add a text string in the **Component Suffix** field to associate the copied components with the copied zone ruleset.

- 3 Click **OK** to apply the change.

4 Open the Brick Editor for the appropriate Brick (in this case, the Regional Office Brick) and access the Policy Assignment tab.

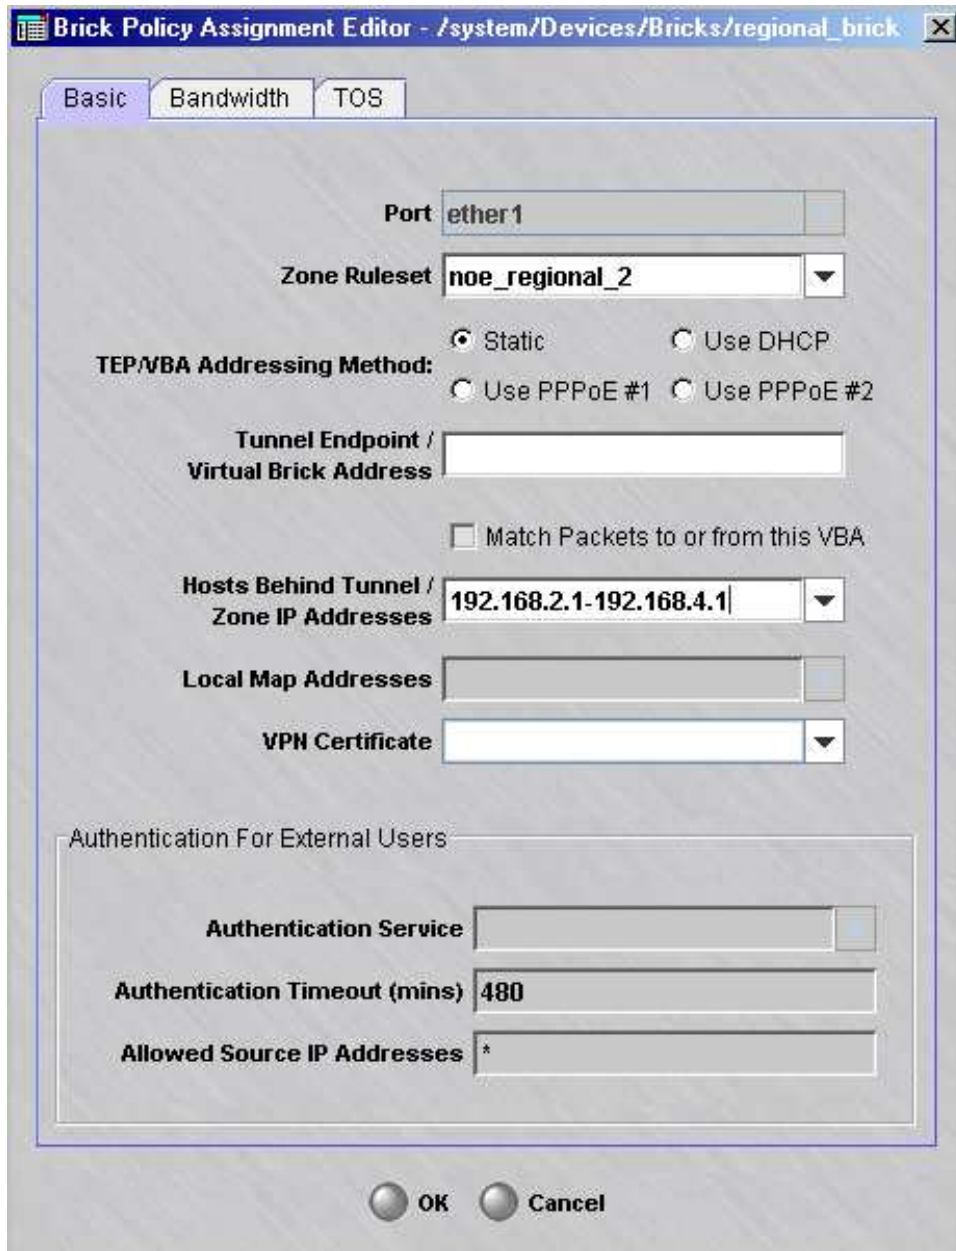
5 Right-click on the interface that has the **noe_regional** ruleset assigned and select **Edit** from the pop-up menu.

Result The Brick Policy Assignment Editor is displayed.

6 In the **Zone Ruleset** field, select the new/duplicate zone ruleset, in this case, **noe_regional_2**.

7 In the **Hosts Behind Tunnel** field, enter the range of IP addresses for the phones that will be directed to the new/duplicate zone ruleset that has just been created ([Figure E-18, “Brick Policy Assignment Editor \(Showing Range of Phone IP Addresses for Duplicate Zone Ruleset in Hosts Behind Tunnel Field\)”](#) (p. E-31)).

Figure E-18 Brick Policy Assignment Editor (Showing Range of Phone IP Addresses for Duplicate Zone Ruleset in Hosts Behind Tunnel Field)



The **Hosts Behind Tunnel** field of the original zone ruleset will contain an asterisk (*) for the hosts behind the tunnel. Any IP address that does not fall within the range specified for the new zone ruleset will use the original zone ruleset.

Repeat these steps as required for additional zone rulesets.

The following shows an example of five zone rulesets assigned to the same Brick interface.

Port	Zone Ruleset	Tunnel Endpoint / Virtual Brick Addr	Match Packets to or from this VBA	Hosts Behind Tunnel / Zone IP Addresses	Local Map Addresses	VPN Certificate	Auth
local	firewall		No	*			
ether0			No	*			
ether1	noe_regional		No	*			
ether1	noe_regional_5		No	*			
ether1	noe_regional_4		No	*			
ether1	noe_regional_3		No	*			
ether1	noe_regional_2		No	*			
ether2	noe_regional_vpn	10.1.1.2	No	*			
ether3			No	*			

END OF STEPS



NATed Configurations

Overview

The only two VoIP configurations that support the Brick as a NATing device are the Branch Office and Remote Office cases, where the call servers and MGW equipment are “hosted” at the Telephony Service Provider’s premises in a multi-company arrangement, as shown in [Figure E-19, “Hosted PCX with Local PSTN Access”](#) (p. E-33) and [Figure E-20, “Hosted PCX With Full VoIP Deployment”](#) (p. E-33).

Figure E-19 Hosted PCX with Local PSTN Access

Telephony Service Provider

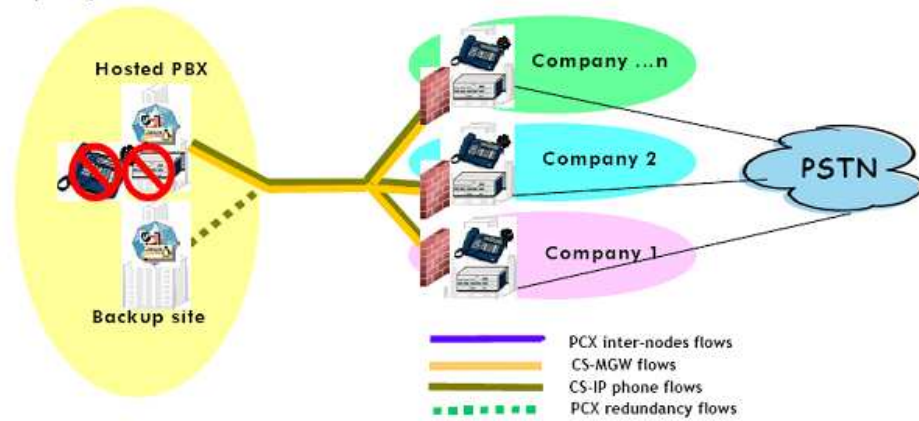
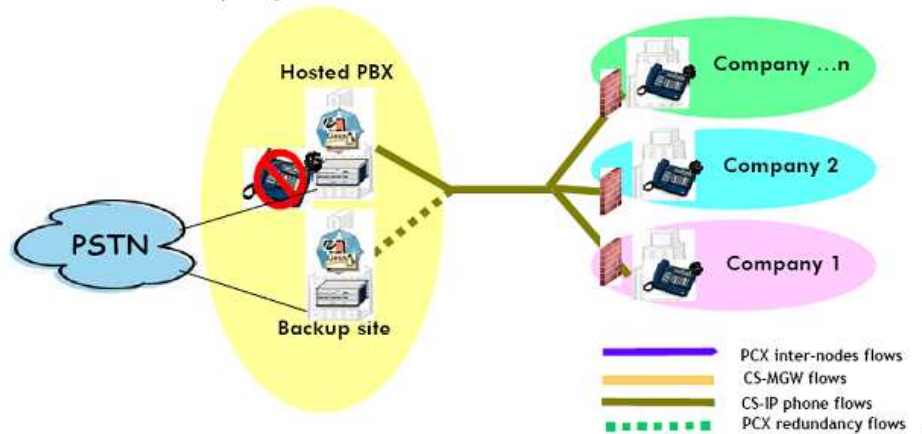


Figure E-20 Hosted PCX With Full VoIP Deployment

Telephony Service Provider



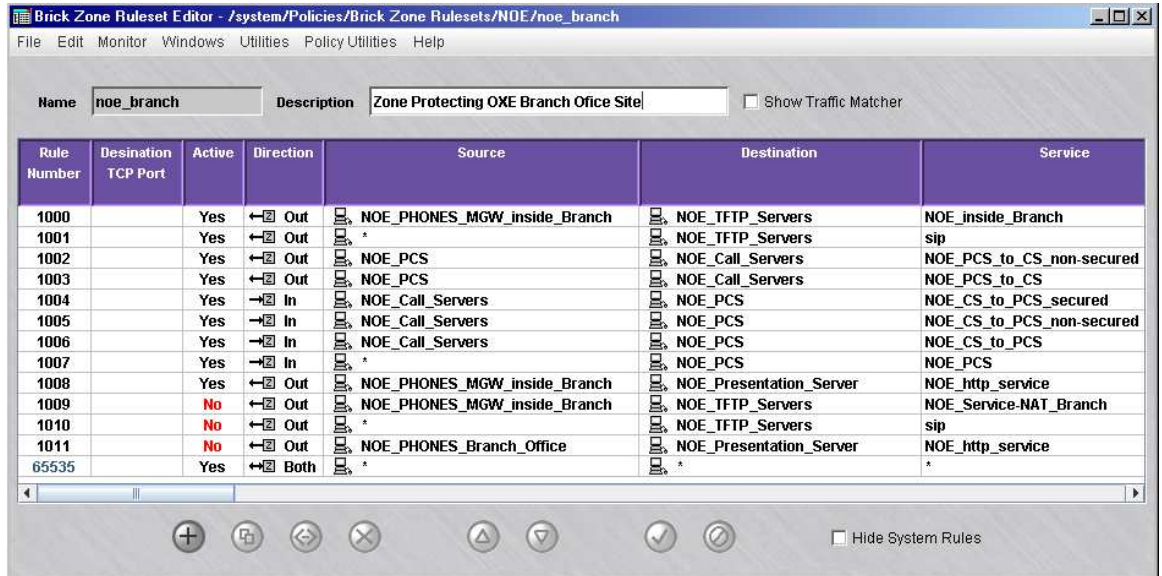
Hosted PCX with local PSTN access

The following pre-defined host groups, available in the **NOE** subfolder under the **Host Groups** folder in the SMS Navigator, must be edited as follows before the pre-defined **noe_branch** ruleset can be assigned and applied to the Brick:

- **NOE_TFTP_Server_HQ**— this host group must be filled in with the public IP address of the TFTP server at the TSP site
- **NOE_TFTP_Server_Regional**— this host group must be filled in with the public IP address of the TFTP server that will be acting as the backup TFTP server
- **NOE_Call_Server_HQ**— this host group must be filled in with the public IP address of the call server at the TSP site
- **NOE_PHONES_Branch_Office**— this host group must be filled in with the IP addresses of all IP phone devices located at the Branch office site
- **NOE_GD_IPs_Branch_Office**— this host group must be filled in with the IP addresses of the GD boards of the MGW unit(s) located at the Branch Office site
- **NOE_GA_IPs_Branch_Office**— this host group must be filled in with the IP addresses of the GA boards of the MGW unit(s) located at the Branch Office site
- **NOE_Presentation_Server**— this host group must be filled in with the IP address of the Presentation Server
- **NOE_GD_Public_IPs_Branch_Office**— this host group must be filled in with the public IP addresses that have been set aside for the GD boards in the MGW unit(s) located behind the Brick at the Branch Office site
- **NOE_GA_Public_IPs_Branch_Office**— this host group must be filled in with the public IP addresses that have been set aside for the GA boards in the MGW unit(s) located behind the Brick located at the Branch Office site

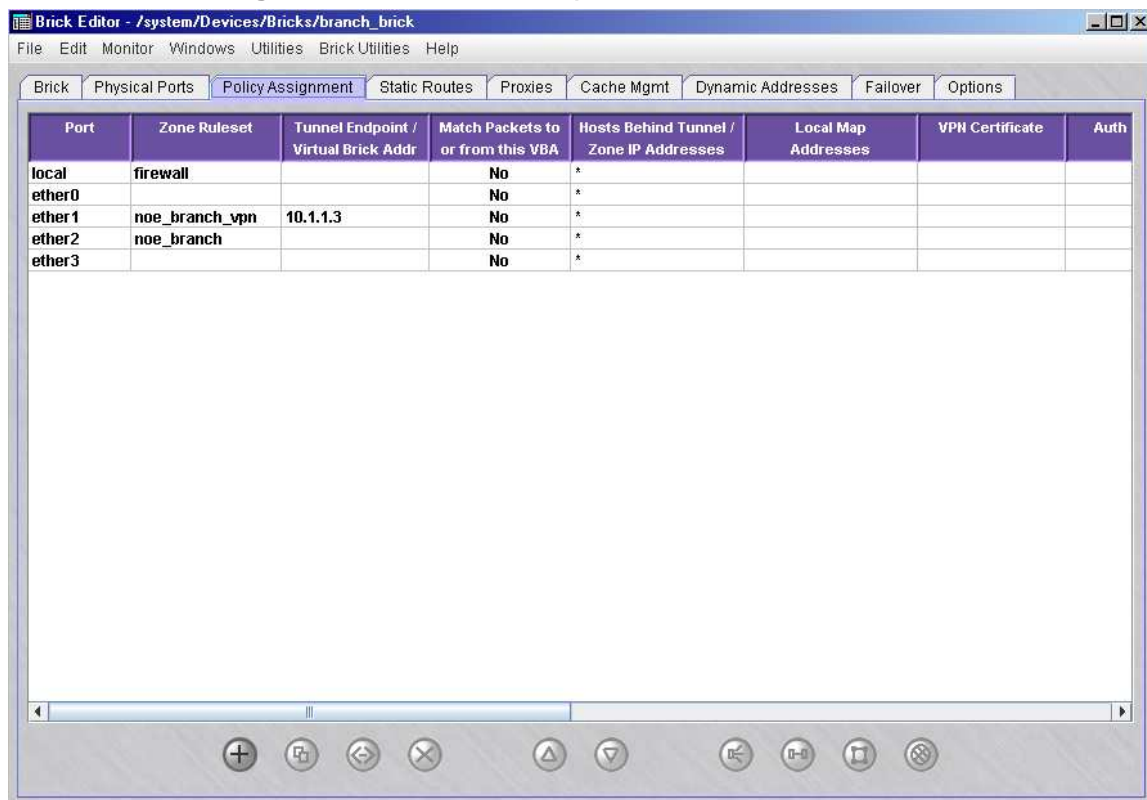
Once the above host groups have been completed, the zone ruleset can be applied to the Branch Office Brick ([Figure E-21, “Brick Zone Ruleset Editor \(Showing Pre-Defined Host Groups Assigned within noe_branch Ruleset\)”](#) (p. E-35) shows an example).

Figure E-21 Brick Zone Ruleset Editor (Showing Pre-Defined Host Groups Assigned within noe_branch Ruleset)



Once the host groups have been configured for the zone ruleset, it can be assigned to a Brick interface (Figure E-22, “Brick Editor (Policy Assignment Tab, Showing noe_branch Ruleset Assigned to Brick Interface)” (p. E-36) shows an example).

Figure E-22 Brick Editor (Policy Assignment Tab, Showing noe_branch Ruleset Assigned to Brick Interface)



The zone ruleset can now be saved and applied to the Brick.

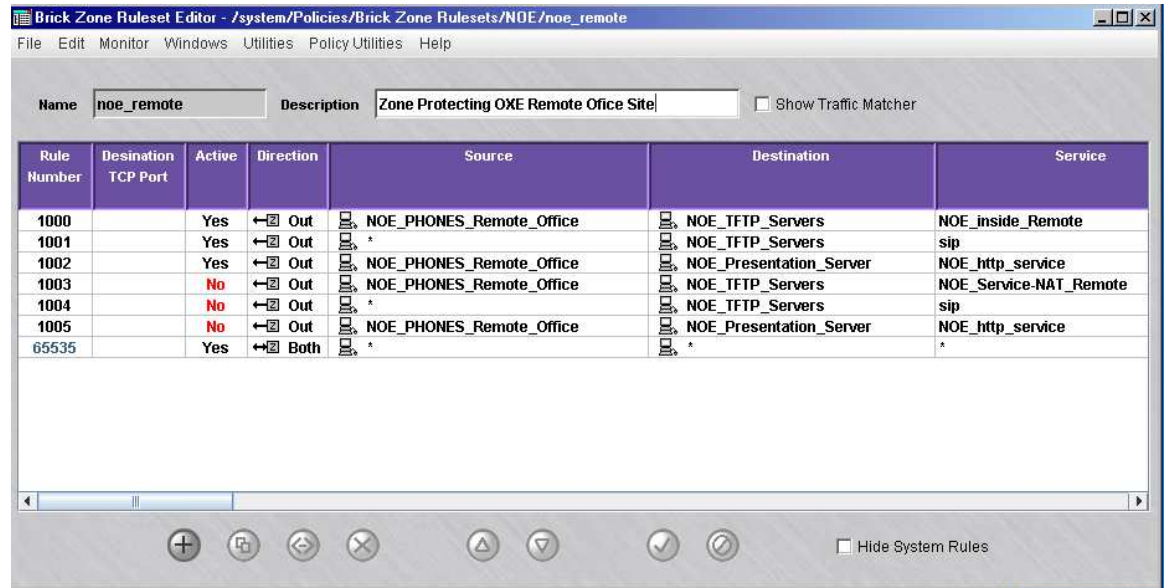
Hosted PCX with full VoIP deployment

The following pre-defined host groups, available in the **NOE** subfolder under the **Host Groups** folder in the SMS Navigator, must be edited as follows before the pre-defined **noe_remote** ruleset can be assigned and applied to the Brick:

- **NOE_TFTP_Server_HQ**— this host group must be filled in with the public IP address of the TFTP server at the TSP site
- **NOE_TFTP_Server_Regional**— this host group must be filled in with the public IP address of the TFTP server that will be acting as the backup TFTP server
- **NOE_Call_Server_HQ**— this host group must be filled in with the public IP address of the call server at the TSP site
- **NOE_Call_Server_Regional**— this host group must be filled in with the public IP address of the call server that will be acting as the backup call server
- **NOE_PHONES_Remote_Office**— this host group must be filled in with the IP addresses of the IP phone devices located at the Remote Office site
- **NOE_Presentation_Server**— this host group must be filled in with the IP address of the Presentation Server

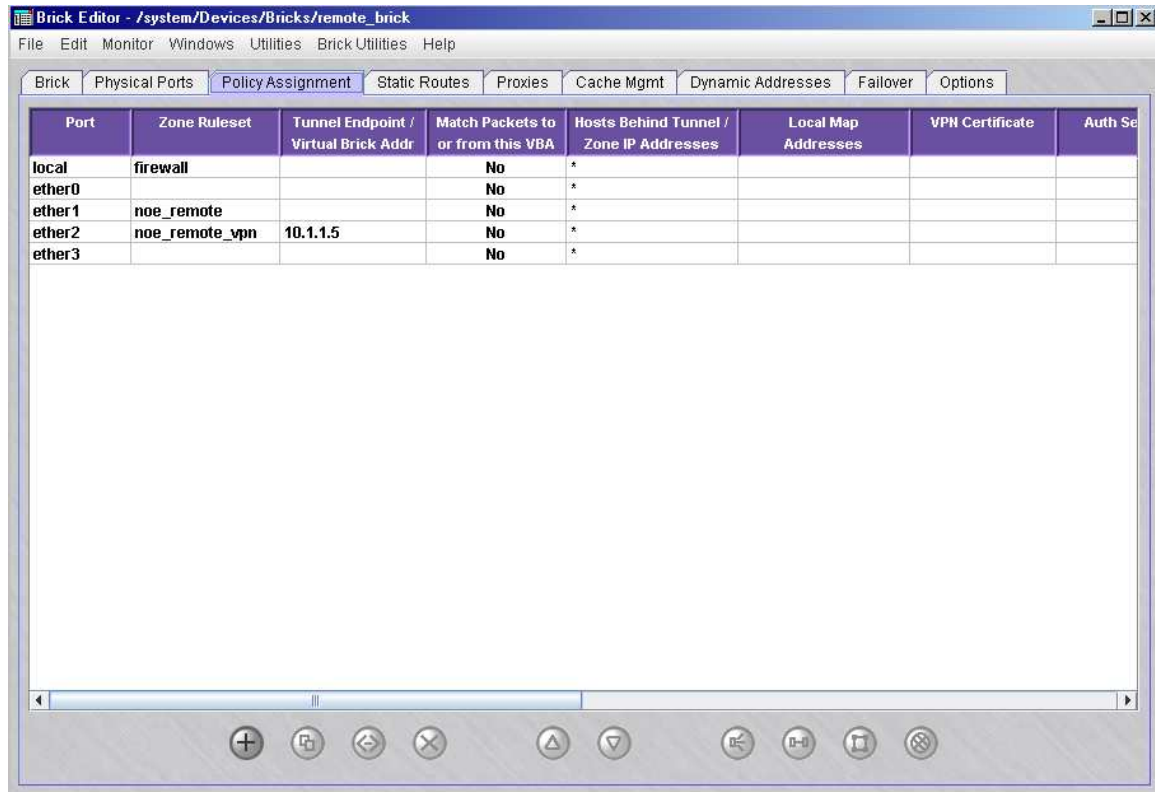
Once the above host groups have been completed, the zone ruleset can be applied to the Remote Office Brick (shows an example).

Figure E-23 Brick Zone Ruleset Editor (Showing Pre-Defined Host Groups Assigned within noe_remote Ruleset)



Once the host groups have been configured for the zone ruleset, it can be assigned to a Brick interface (shows an example).

Figure E-24 Brick Editor (Policy Assignment Tab, Showing noe_remote Ruleset Assigned to Brick Interface)



The zone ruleset can now be saved and applied to the Brick.



Index

A Accounting shared secret

- Accounting shared secret, [8-23](#)
- ACE/Server see [SecurID, 8-21](#)
- Activate a rule, [1-64](#)
- administrativezone ruleset
 - administrativezone, [B-3](#)
- Application filters
 - copy, [5-129](#)
- Application users, [8-2, 8-4](#)
 - Login procedure, [8-4](#)
- Apply
 - Brick zone ruleset, [1-57](#)
 - LAN-LAN tunnel, [10-23](#)
- Authenticatio service, [11-23](#)
- Authentication
 - See: [User authentication](#)
- Authentication port, [8-23](#)
- Authentication service, [8-21, 8-28, 8-30](#)
 - RADIUS, [8-21](#)
 - SecurID, [8-28](#)
 - VPN certificate, [8-30](#)
- Authentication Services
 - copy, [8-35](#)
- Authentication shared secret, [8-23, 8-23](#)
- Authentication timeout, [8-6, 8-30](#)
- Authorize return channel, [1-30](#)

Automated key exchange, [10-3](#)

B Brick, [10-2, 11-2](#)

- LAN-LAN tunnel endpoint, [10-2](#)
- Brick devices
 - supported, [xxv](#)
- Brick session cache see [Session cache, 1-8](#)
- Brick Zone Ruleset Editor, [1-14](#)
- Brick zone rulesets, [1-1, 1-64](#)
 - add rules-based routing IP addresses, [1-40](#)
 - Advanced features, [1-28](#)
 - Apply, [1-57](#)
 - assign to physical port, [1-46](#)
 - Brick Zone Ruleset Editor, [1-14](#)
 - Copy, [1-55](#)
 - create, [1-14](#)
 - Create, [1-15](#)
 - customized zones, [1-4](#)
 - Definition, [1-2](#)
 - Delete, [1-57](#)
 - finding rules that match traffic patterns in, [1-7](#)
 - Move, [1-56](#)
 - Packet filtering, [1-2](#)
 - pre-defined templates for VoIP/NOE traffic, [1-66](#)
 - rules-based routing and, [1-4](#)

- SCTP port filtering and, [1-6](#), [1-42](#)
- View, [1-51](#)
- Brick zone traffic
 - finding rules that match, [1-7](#)
- Bricks
 - assign zone ruleset to physical port, [1-46](#)
 - configure as Application Layer Gateway (ALG) for VoIP/NOE traffic, [E-1](#)
-
- C** CA
 - See: Certificate Authorities(CAs)
- Certificate Authorities (CAs), [9-1](#)
 - selecting multiple CAs for client tunnel endpoint, [11-3](#)
- Certificate manager, [9-7](#)
 - Create certificate signing request, [9-7](#)
 - Start on Windows NT, [9-5](#)
- Certificate signing request
 - Create, [9-7](#), [9-7](#)
 - Submit to certificate authority, [9-9](#)
- Change session timeout, [1-28](#)
- Client Defaults Editor, [11-8](#)
 - Parameters tab, [11-8](#)
- Client firewall, [11-11](#)
- Client tunnel endpoint
 - set up with UDP encapsulation, [11-33](#)
- Client tunnel endpoints
 - maintain, [11-40](#)
 - set up, [11-24](#)
- Client tunnels, [11-1](#)
 - authentication service, [11-6](#)
 - Client Defaults Editor, [11-8](#)
 - Client firewall, [11-11](#)
 - defaults, [11-6](#)
 - display active sessions, [11-41](#)
 - Host behind tunnel, [11-23](#)
 - host groups and, [11-6](#)
 - Keep alive interval, [11-10](#)
 - missed keepalive limit, [11-10](#)
 - Save password, [11-11](#)
 - Scenarios, [11-2](#)
 - Set the defaults, [11-8](#)
- Client users, [8-2](#), [8-4](#)
 - Login procedure, [8-4](#)
- Configuration Assistant
 - user authentication parameters, [8-11](#)
- Copy
 - authentication service, [8-35](#)
 - Brick zone ruleset, [1-55](#)
 - Dependency mask, [7-11](#)
 - Domain name group, [3-14](#)
 - Host group, [2-18](#)
 - Service group, [4-19](#)
 - user account, [8-15](#)
 - user group, [8-18](#)
- Copy/paste
 - rule, [1-62](#)
- Create, [8-11](#), [8-11](#)
 - Brick zone ruleset, [1-14](#), [1-15](#)
 - Certificate signing request, [9-7](#)
 - Dependency mask, [7-4](#)
 - Host group, [2-5](#), [3-4](#)
 - RADIUS authentication service, [8-21](#)
 - Rule, [1-16](#)
 - SecurID authentication service, [8-28](#)
 - Service group, [4-5](#)
 - User account, [8-11](#)
 - User group, [8-17](#)

CSR see Certificate Signing Request, [9-7](#)

D Delete

Brick zone ruleset, [1-57](#)

Dependency mask, [7-14](#)

Domain name group, [3-17](#)

Host group, [2-21](#)

LAN-LAN tunnel, [10-30](#)

Rule, [1-64](#)

Service group, [4-21](#)

Dependency masks, [7-1](#), [7-18](#)

Associate the with a rule, [7-7](#)

Copy, [7-11](#)

Create, [7-4](#)

Definition, [7-2](#)

Delete, [7-14](#)

Dependency Masks Editor, [7-3](#)

find entities of, [7-10](#)

Modify, [7-10](#)

Move, [7-13](#)

RealAudio, [7-15](#)

View, [7-8](#)

Dependency Masks Editor, [7-3](#)

Destination address mapping, [6-5](#), [6-20](#), [A-7](#)

Direct, [6-10](#)

Pool, [6-10](#)

Destination port mapping, [6-12](#)

Deactivate a rule, [1-64](#)

Diffie-Hellman group, [10-9](#)

Digital Certificates, [9-1](#)

Digital certificates, [9-26](#)

Domain name groups

Copy, [3-14](#)

Delete, [3-17](#)

Move, [3-16](#)

Duplicate

rule, [1-61](#)

E Enable compression, [11-20](#)

Enable LAN-LAN tunnel;Disable LAN-LAN tunnel;LAN-LAN tunnels

Enable, [10-30](#)

Enable perfect forward secrecy, [10-10](#), [11-20](#)

Endpoints see Tunnel Endpoints, [10-2](#)

F firewall ruleset

firewall, [B-6](#)

FTP traffic

content filtering for, [1-4](#)

H Host group

client tunnels and, [11-6](#)

Host group templates

pre-defined for VoIP/NOE traffic, [2-25](#)

Host groups, [2-1](#), [2-23](#), [3-1](#)

Benefits, [2-3](#)

copy, [2-18](#)

Create, [2-5](#), [3-4](#)

Definition, [2-2](#), [3-2](#)

delete, [2-21](#)

find entities of, [2-16](#)

find overlapping IP address of hosts, [2-17](#)

Host Groups Editor, [2-4](#)

modify, [2-18](#)

Modify, [3-14](#)

move, [2-20](#)

Provided with the SMS application, [2-9](#)

provided with the SMS application, [2-22](#)

- Provided with the SMS application, [3-8](#), [5-128](#)
 - rules-based routing and, [1-4](#), [1-40](#)
 - Uses, [2-2](#), [3-2](#)
 - View, [2-15](#), [3-12](#)
 - Host Groups Editor, [2-4](#)
 - Hosts behind tunnel, [11-23](#)
 - Client tunnel, [11-23](#)
 - HTTP application filter template
 - pre-defined for VoIP/NOE traffic, [5-60](#)
 - HTTP traffic
 - content filtering for, [1-4](#)
-
- I**
 - IKE see Internet Key Exchange, [10-3](#)
 - IKEv1
 - See: Internet Key Exchange Version 1 (IKEv1)
 - client tunnel endpoint defaults, [11-20](#)
 - IKEv2
 - See: Internet Key Exchange Version 2 (IKEv2)
 - client tunnel endpoint defaults, [11-20](#)
 - Internet Key Exchange, [10-3](#)
 - Internet Key Exchange Version 1 (IKEv1), [10-4](#), [D-7](#)
 - Internet Key Exchange Version 2 (IKEv2), [10-4](#), [11-3](#)
 - IPSec proposal, [10-10](#)
 - Authentication type, [10-10](#), [11-19](#)
 - Enable compression, [11-20](#)
 - Enable compression;Enable compression, [10-11](#)
 - Enable perfect forward secrecy, [11-20](#)
 - Enable perforward secrecy, [10-10](#)
 - Encryption type, [10-10](#), [11-19](#)
 - Protocol, [10-10](#), [11-19](#)
 - SA lifetime (kbytes), [10-10](#)
 - SA lifetime (Kbytes), [11-20](#)
 - SA lifetime (secs), [11-20](#)
 - SA lifetime(secs), [10-10](#)
 - IPSec SA proposal, [11-19](#)
-
- K**
 - Keepalive interval, [11-10](#)
-
- L**
 - LAN-LAN Defaults Editor, [10-6](#)
 - Parameters tab, [10-6](#)
 - LAN-LAN Tunnel Editor
 - Main tab, [10-14](#)
 - LAN-LAN Tunnel EditorLAN-LAN tunnels
 - LAN-LAN Tunnel Editor, [10-14](#)
 - LAN-LAN tunnels, [10-1](#), [10-14](#), [10-35](#)
 - Automated key exchange, [10-3](#)
 - Brick, [10-2](#)
 - defaults, [10-4](#)
 - Delete, [10-30](#)
 - Disable, [10-30](#)
 - Endpoints, [10-2](#)
 - IPSec transport method, [10-8](#)
 - keepalive interval, [10-8](#)
 - LAN-LAN Defaults Editor, [10-6](#)
 - maintain, [10-28](#)
 - Modify, [10-29](#)
 - Receive any proposals, [10-8](#)
 - Save and apply, [10-23](#)
 - Set defaults, [10-6](#)
 - set up, [10-13](#)
 - set up with UDP encapsulation, [10-25](#)
 - View, [10-28](#)
 - Lawful intercept feature
 - Packet Data Gateway (PDG), [11-5](#)
 - Local password, [8-3](#), [8-11](#), [8-11](#), [8-17](#)
 - Create user account, [8-11](#), [8-11](#)

- Create user group, [8-17](#)
 - Set up, [8-11](#)
 - Local Presence, [A-1](#)
 - Local presence
 - Destination address mapping, [A-7](#)
 - Source address mapping, [A-7](#)
 - Lucent IPsec
 - create message for users, [11-46](#)
 - Lucent Netcare Professional Services, [xxvi](#)
-

M Max use concurrent, [1-29](#)

- Max use total, [1-29](#)
- Modify
 - Dependency mask, [7-10](#)
 - Host group, [2-18](#), [3-14](#)
 - LAN-LAN tunnel, [10-29](#)
 - rule, [1-61](#)

Move

- Brick zone ruleset, [1-56](#)
 - Dependency mask, [7-13](#)
 - domain name group, [3-16](#)
 - Host group, [2-20](#)
 - Service group, [4-20](#)
-

N NAT

- See: Network address translation (NAT)
 - Network address translation, [6-20](#)
 - Destination address mapping, [6-5](#), [6-20](#)
 - Destination port mapping, [6-12](#)
 - Source address mapping, [6-18](#)
 - With a router, [6-16](#)
 - Without a router, [6-18](#), [6-20](#)
 - Network address translation (NAT), [6-1](#)
-

- nocgwzone ruleset
 - nocgwzone, [B-30](#)
 - NOE application filter
 - pre-defined templates for VoIP/NOE traffic, [5-68](#)
-

P Packet Data Gateway (PDG) accounting, [11-5](#)

- Packet Data Gateway (PDG) lawful intercept feature, [11-5](#)
 - Packet filtering, [1-2](#)
 - Pool, [6-6](#), [6-10](#)
 - Ports
 - assign zone ruleset to, [1-46](#)
 - Pre-configured brick zone rulesets, [B-1](#)
 - Pre-configured Brick zone rulesets, [B-3](#), [B-6](#), [B-30](#)
 - Pre-configured brick zone rulesets, [B-30](#)
 - Pre-configured Brick zone rulesets, [B-41](#)
 - Proxy servers
 - rules-based routing and, [1-4](#)
 - Proxy Table, [8-6](#)
 - Destination port, [8-6](#)
 - Proxy port, [8-6](#)
-

R RADIUS, [8-3](#), [8-21](#), [8-21](#), [8-23](#), [8-23](#), [8-23](#), [8-23](#), [8-23](#), [8-23](#)

- attribute codes, [D-1](#)
 - Authentication port, [8-23](#)
 - Authentication service, [8-21](#)
 - Required rule, [8-8](#)
 - RealAudio, [7-15](#)
 - Receive any proposals, [10-8](#)
 - Renumber
 - Rules, [1-64](#)
-

- Required attributes
 - Required attributes, [8-32](#)
- Resource encrypted file, [8-9](#), [8-28](#)
- Routers, [11-2](#)
- RTSP traffic
 - content filtering for, [1-4](#)
- Rules
 - Activate, [1-64](#)
 - Associate alarm, [1-29](#)
 - Brick-specific, [1-4](#)
 - copy/paste, [1-62](#)
 - Create, [1-16](#)
 - Deactivate, [1-64](#)
 - Delete, [1-64](#)
 - duplicate, [1-61](#)
 - modify, [1-61](#)
 - Numbering, [1-8](#)
 - Renumber, [1-64](#)
 - Types, [1-8](#)
 - View, [1-59](#)
- Rules-based routing, [1-4](#)
 - add IP addresses for, [1-40](#)
 - host groups and, [1-4](#), [1-40](#)
-
- S** SA lifetime (Kbytes), [10-11](#), [11-20](#)
- SA lifetime (secs), [10-10](#), [10-11](#), [11-20](#), [11-20](#)
- SCTP
 - See: Stream Control Transport Protocol (SCTP)
- SecurID, [8-3](#), [8-9](#), [8-28](#), [8-28](#), [8-28](#)
 - Authentication service, [8-28](#)
 - Required, [8-8](#)
 - Resource encrypted file, [8-9](#), [8-28](#)
- Security Associations (SAs)
 - See: configuring multiple SA proposals for client tunnel endpoint
- Service Group Editor, [4-4](#)
- Service group templates
 - pre-defined for VoIP/NOE traffic, [4-27](#)
- Service groups, [4-1](#), [4-23](#)
 - Benefits, [4-2](#)
 - Copy, [4-19](#)
 - Create, [4-5](#)
 - Definition, [4-2](#)
 - Delete, [4-21](#)
 - Move, [4-20](#)
 - Provided with SMS application, [4-23](#)
- service groups
 - Provided with the SMS application, [4-10](#)
- Service groups
 - Service Group Editor, [4-4](#)
 - Uses, [4-2](#)
 - View, [4-13](#)
- Services file
 - Services file, [8-10](#)
- Session timeout, [1-28](#)
- SMS parameters
 - Client Program, [D-5](#)
- SMTP traffic
 - content filtering for, [1-4](#)
- Source address mapping, [6-18](#), [A-7](#)
 - Direct, [6-6](#)
 - Pool, [6-6](#)
- Spoofing attacks
 - prevention of, [1-12](#)
- Stream Control Transport Protocol (SCTP), [1-6](#), [1-42](#)
- Strong password (SOX compliance), [8-11](#)

SYN flood protection, [1-31](#)

T Technical support, [xxvi](#)

TFTP application filter

pre-defined templates for VoIP/NOE traffic,
[5-123](#)

Traffic matcher tool, [1-7](#), [1-53](#)

Tunnel endpoints, [10-2](#), [10-2](#), [11-2](#), [11-2](#)

Brick, [10-2](#), [11-2](#)

Router, [11-2](#)

Types, [10-2](#)

Tunnels

see Client Tunnels, [10-1](#)

see LAN-LAN Tunnels, [10-1](#)

U UDP encapsulation

set up client tunnel endpoint with, [11-33](#)

set up LAN-LAN tunnel with, [10-25](#)

UMA

See: Unlicensed Mobile Access (UMA)

Unlicensed Mobile Access (UMA), [11-13](#), [11-16](#),
[D-7](#)

User account, [8-11](#), [8-11](#)

User accounts

copy, [8-15](#)

Create, [8-11](#)

User authentication, [8-1](#), [8-10](#), [8-32](#)

Application users, [8-2](#), [8-4](#)

Authentication process, [8-4](#)

Authentication server, [8-5](#), [8-6](#)

Authentication timeout, [8-6](#), [8-30](#)

Client users, [8-2](#), [8-4](#)

How it works, [8-4](#)

Local password, [8-3](#), [8-11](#)

Proxy Table, [8-6](#)

RADIUS, [8-3](#)

SecurID, [8-3](#)

Types of authentication, [8-3](#)

Types of users, [8-2](#)

VPN certificate, [8-3](#)

User groups, [8-17](#)

copy, [8-18](#)

Create, [8-17](#)

SMS-defined, [8-19](#)

Users, [8-2](#), [8-2](#)

Application users, [8-2](#)

Client users, [8-2](#)

V View

Brick zone rulesets, [1-51](#)

Dependency masks, [7-8](#)

Host groups, [2-15](#), [3-12](#)

LAN-LAN tunnels, [10-28](#)

rules, [1-59](#)

Service groups, [4-13](#)

Virtual brick address, [6-5](#), [6-9](#), [6-18](#)

VPN certificate, [8-3](#), [8-30](#), [8-32](#), [8-32](#)

Authentication service, [8-30](#)

How to set up, [8-32](#)

VPN certificate, [8-32](#)

