



Lucent Security Management Server

Release 9.1

Policy Guide

260-100-016R9.1
Issue 1
August 2006

Lucent Technologies - Proprietary

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

Copyright © 2006 Lucent Technologies
Unpublished and Not for Publication
All Rights Reserved

This material is protected by the copyright and trade secret laws of the United States and other countries. It may not be reproduced, distributed, or altered in any fashion by any entity (either internal or external to Lucent Technologies), except in accordance with applicable agreements, contracts or licensing, without the express written consent of Lucent Technologies and the business management owner of the material.

Trademarks

All trademarks and service marks specified herein are owned by their respective companies.

Contents

About this information product

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

1 Lucent VPN Firewall *Brick*[®] Device Zone Rulesets

Overview	1-1
What is a Brick Zone Ruleset?	1-2
Create a Brick Zone Ruleset	1-10
To Create a Security Rule	1-17
To Add Time and Day Restrictions to a Rule	1-27
To Add Advanced Features to a Rule	1-30
To Add Bandwidth Management to a Rule	1-37
To Add TOS/Alarm Capability to a Rule	1-42
To Configure IP Address(es) for Rule-Based Routing	1-45
To Maintain Brick Zone Rulesets	1-47

	How to Maintain Security Rules	1-53
2	Host Groups	
	Overview	2-1
	What is a Host Group?	2-2
	How to Set Up a Host Group	2-4
	Global Host Groups	2-8
	Dynamic Host Groups	2-10
	Nested Host Groups	2-12
	How to Maintain a Host Group	2-15
	Host Groups Provided with the LSMS Application	2-21
3	Domain Name Groups	
	Overview	3-1
	What is a Domain Name Group?	3-2
	How to Set Up a Domain Name Group	3-3
	Global Domain Name Groups	3-8
	Nested Domain Name Groups	3-11
	How to Maintain a Domain Name Group	3-13
4	Service Groups	
	Overview	4-1
	What is a Service Group?	4-2
	How to Set Up a Service Group	4-3
	Global Service Groups	4-9
	Nested Service Groups	4-11
	How to Maintain a Service Group	4-14
	Service Groups Provided with the LSMS Application	4-21

5 Application Filters

Overview 5-1

Process Flow 5-3

DHCP Relay Application Filter 5-4

DNS Application Filters 5-7

ESP Application Filter 5-16

FTP Application Filter 5-19

GTP Application Filter 5-29

H.323 Application Filters 5-43

HTTP Application Filter 5-50

RPC Application Filters 5-58

SIP Application Filters 5-64

SMTP Application Filters 5-81

Global Application Filters 5-92

6 Network Address Translation

Overview 6-1

What is Network Address Translation? 6-2

How to Set Up Source Address Mapping 6-4

How to Set Up Destination Address Mapping 6-7

How to Set Up Destination Port Mapping 6-11

How to Perform Source Address Mapping with a Router 6-12

How to Perform Source Address Mapping without a Router 6-14

How to Perform Destination Address Mapping without a Router 6-16

Other Examples of a Brick Responding to ARPs 6-18

7 Dependency Masks

Overview 7-1

	What is a Dependency Mask?	7-2
	How to Set Up a Dependency Mask	7-3
	How to Maintain a Dependency Mask	7-10
	Example: RealAudio Session	7-16
8	Proxies	
	Overview	8-1
	How the Proxy Feature Works	8-2
	How to Make an Entry in the Proxy Table	8-4
	How to Maintain the Proxy Table	8-9
	How to Assign a Ruleset to the Brick Port	8-11
	How to Set Up Proxy Load Sharing	8-18
9	User Authentication	
	Overview	9-1
	What is User Authentication?	9-2
	How Authentication Works	9-4
	Before Setting Up User Authentication	9-8
	To Set Up Local Password Authentication	9-11
	How to Set Up RADIUS and SecurID Authentication	9-19
	How to Set Up VPN Certificate Authentication	9-29
10	Digital Certificates	
	Overview	10-1
	What is a Digital Certificate?	10-2
	What is the Certificate Manager?	10-3
	How to Start the Certificate Manager	10-5
	How to Obtain and Install a Server Certificate	10-7
	How to Obtain and Import a VPN Certificate	10-12

	Group Assignment for VPN Certificates	10-19
	Pending Certificate Signing Requests	10-23
	Server Versus VPN Certificate Signing Request	10-24
11	LAN-LAN Tunnels	
	Overview	11-1
	What is a LAN-LAN Tunnel?	11-2
	Set LAN-LAN Tunnel Defaults	11-5
	Set Up a LAN-LAN Tunnel	11-12
	Set Up a LAN-LAN Tunnel with UDP Encapsulation	11-23
	Redundant LAN-LAN Tunnels	11-25
	Maintain LAN-LAN Tunnels	11-26
	How to Set Up Service Level Agreements	11-30
12	Client Tunnel Endpoints	
	Overview	12-1
	What is a Client Tunnel?	12-2
	To Set the Client Tunnel Defaults	12-6
	To Set Up a Client Tunnel Endpoint	12-22
	To Set Up a Client Tunnel with UDP Encapsulation	12-29
	What to Do Next	12-31
	Maintain Client Tunnel Endpoints	12-37
	Create a Message for Lucent IPsec Client Users	12-43
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 Lucent VPN Firewall <i>Brick</i>[®] Device Zone Rulesets	
Overview	B-1
administrativezone	B-3
dhcpzone_on_inside_if	B-8
dhcpzone_on_internet_if	B-14
firewall	B-20
nocgwzone	B-27
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
IKEv1 Preshared Key	D-5
Timeout Information	D-6
Local Presence IP Information	D-8
Primary WINS/Secondary DNS	D-11
Sample Configuration	D-13

Index

List of tables

8 Proxies

8-1	Proxy Table without Load Sharing	8-19
8-2	Proxy Table with Load Sharing	8-19

9 User Authentication

9-1	LSMS-defined user groups	9-17
-----	--------------------------------	------

List of figures

1 Lucent VPN Firewall *Brick*[®] Device Zone Rulesets

1-1	Packet Filtering Process	1-3
1-2	Brick Zone Ruleset Editor	1-10
1-3	Brick Policy Assignment Editor	1-12
1-4	Apply Brick Window	1-15
1-5	Brick Zone Rule Editor (Basic Tab)	1-18
1-6	Browse: Select a Brick Window	1-20
1-7	Time and Day Fields	1-28
1-8	Rule Editor (Advanced Tab)	1-31
1-9	Brick Zone Rule Editor (Bandwidth Tab)	1-38
1-10	Asymmetric Bandwidth Configuration	1-40
1-11	Brick Zone Rule Editor (TOS/Alarms tab)	1-43
1-12	Brick Zone Rule Editor (Rules Based Routing tab)	1-45
1-13	View Brick Zone Rulesets	1-48
1-14	Copy Window	1-49
1-15	Confirmation Window (Brick Zone Rulesets)	1-50
1-16	Apply Policy Window	1-51
1-17	Brick Zone Ruleset Editor (Viewing Rules in a Zone Ruleset)	1-53
1-18	Confirmation Window (Rules)	1-56

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)	2-16
2-8	Entities Found Window	2-17
2-9	Copy To Window	2-18
2-10	Confirmation Window (Host Groups)	2-20

3 Domain Name Groups

3-1	Domain Name Group Editor	3-3
3-2	Domain Name Group Entry Window	3-5
3-3	Domain names in a Domain Name Group	3-6
3-4	Locate Original Window	3-9
3-5	Domain Name Group Entry Window (Nested Domain Name Groups)	3-12
3-6	Navigator Window (View Domain Name Groups)	3-13
3-7	Entities Found Window	3-14
3-8	Copy To Window	3-16
3-9	Confirmation Window (Domain Name Groups)	3-18

4 Service Groups

4-1	Service Group Editor	4-3
4-2	Service Editor	4-5
4-3	Service Group with Two Services	4-7
4-4	Locate Original Window (Service Groups)	4-10

4-5	Service Editor (Nested Service Groups)	4-12
4-6	Nested Service Groups	4-13
4-7	Navigator Window (View Service Groups)	4-15
4-8	Entities Found Window	4-16
4-9	Service Group Editor (Edit Mode)	4-17
4-10	Copy To Window	4-18
4-11	Confirmation Window (Service Groups)	4-20

5 Application Filters

5-1	DHCP Application Filter Editor	5-5
5-2	DNS Application Filter Editor (Protected Domain Names Tab)	5-8
5-3	DNS Application Filter Editor Screen (Names and Addresses Tab)	5-10
5-4	DNS Application Filter Editor Screen (RR Type Tab)	5-11
5-5	DNS Application Filter Editor (RR Class Tab)	5-12
5-6	DNS Application Filter Editor (Advanced Tab)	5-13
5-7	ESP Application Filter	5-16
5-8	FTP Application Filter Editor (Commands Filter Tab)	5-21
5-9	FTP Application Editor (Protocol Anomaly Check Tab)	5-23
5-10	FTP Application Filter Editor (Data Port & User Blocking Tab)	5-25
5-11	FTP Application Editor (Miscellaneous Options Tab)	5-27
5-12	GTP Application Filter -Stateful (GTP PDP Context Deletion Monitoring)	5-30
5-13	GTP Application Filter Editor	5-32
5-14	GTP Application Filter Editor (GP Interface Message List Tab)	5-34
5-15	GTP Application Filter Editor (Non-Gp Interface Message Tab)	5-35
5-16	GTP Application Filter Editor (IMSI) Prefix list Tab)	5-36
5-17	GTP Application Filter (MSISDN Prefix List Tab)	5-37
5-18	GTP Application Filter Editor(APN List Tab)	5-38
5-19	GTP Application Filter Editor (GTP Nesting Tab)	5-39

5-20	GTP Application Filter Editor (Stateful GTP Tab)	5-40
5-21	H.323_VOIP Application Filter (Address Translation Tab)	5-45
5-22	H.323_RAS Application Filter	5-46
5-23	Network Topology - Example 3	5-48
5-24	Application Filter Editor	5-51
5-25	URI Path Pattern Editor	5-54
5-26	Application Filter Editor (Keyword Tab)	5-55
5-27	Application Filter Editor (RPC Tab)	5-59
5-28	Application Filter Editor (Portmapper Tab)	5-61
5-29	SIP Proxy in the private network	5-67
5-30	SIP Proxy on public side of the Brick	5-67
5-31	SIP Proxy in the DMZ	5-68
5-32	SIP Application Filter Editor(Options Default Tab)	5-69
5-33	SIP Application Filter Editor (Options Tab)- both zones	5-71
5-34	SIP Application Filter Editor (Names and Other Addresses Tab)	5-73
5-35	SIP Editor (Names Entry Screen)	5-74
5-36	SIP Application Filter Editor (DNS Default Tab)	5-76
5-37	SIP Application Filter Editor (DNS Tab)-both zones	5-76
5-38	SIP Application Filter Editor (Methods Default Tab)	5-78
5-39	SIP Application Filter Editor (Methods Tab) - both zones	5-78
5-40	SMTP Application Filter Editor (Commands Tab)	5-82
5-41	SMTP Application Filter Editor (MIME Types Tab)	5-84
5-42	SMTP Application Editor (Attachment Tab)	5-86
5-43	SMTP Application Filter Editor (Anomalies Tab)	5-87
5-44	SMTP Application Filter Editor (Other Tab)	5-89
5-45	Locate Original Window (Application Filters)	5-93

6 Network Address Translation

6-1 Web Access Rule 6-4
6-2 Local Access Rule 6-8

7 Dependency Masks

7-1 Dependency Masks Editor 7-4
7-2 Brick Zone Ruleset Editor (Advanced Tab) 7-8
7-3 Navigator Window (View Dependency Masks) 7-11
7-4 Dependency Masks Editor (Edit Mode) 7-12
7-5 Copy To Window 7-13
7-6 Confirmation Window (Dependency Masks) 7-15
7-7 RealAudio Sessions 7-16

8 Proxies

8-1 Brick Editor (Proxies Tab) 8-5
8-2 Brick Proxy Editor 8-6
8-3 Confirmation Window (Proxies) 8-10
8-4 proxyzone Ruleset 8-12
8-5 Policy Assignment Tab 8-15
8-6 Brick Policy Assignment Editor 8-16

9 User Authentication

9-1 User Editor 9-12
9-2 Password Panel 9-13
9-3 User Group Editor 9-16
9-4 Authentication Service Editor 9-20
9-5 Authentication Service Editor (RADIUS Service Tab) 9-21
9-6 Authentication Service Editor (Attribute Codes Tab) 9-23
9-7 RADIUS Attributes Window 9-24

9-8	Authentication Service Editor	9-25
9-9	SecurID Options	9-26
9-10	Authentication Service Editor	9-29
9-11	VPN Certificate Options	9-30
9-12	Attributes	9-31
10	Digital Certificates	
10-1	Certificate Manager	10-3
11	LAN-LAN Tunnels	
11-1	Types of LAN-LAN Tunnel Endpoints	11-2
11-2	LAN-LAN Defaults Editor (Parameters Tab)	11-6
11-3	LAN-LAN Defaults Editor (Policy Tab)	11-8
11-4	All Group Tunnels Viewer	11-27
11-5	Default SLA Parameters	11-31
11-6	LAN-LAN Tunnel Editor (SLA Probes tab)	11-32
11-7	Tunnel Round Trip Delay Statistics	11-34
11-8	Group Round Trip Delay Statistics	11-34
12	Client Tunnel Endpoints	
12-1	Client Tunnels	12-3
12-2	Client Defaults Editor (Parameters Tab)	12-7
12-3	Client Defaults Editor (Remote Client ID Tab)	12-14
12-4	Client Defaults Editor (Policy Tab)	12-16
12-5	Apply Group Window	12-20
12-6	Client Tunnel Endpoint Editor (Endpoint Tab)	12-23
12-7	Client License Allocation Window	12-32
12-8	Group License Editor Window	12-33
12-9	Tunnel Endpoint Editor Window	12-34

List of figures

12-10	Client Tunnel Endpoints Window	12-37
12-11	Search Client Tunnels Window	12-39
12-12	Confirmation Window (Client Tunnel Endpoint)	12-42
12-13	Client Message	12-44

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

About this information product

Purpose

Welcome to the *LSMS Policy Guide*. This manual explains how to use Release 9.1 of the Lucent Security Management Server to create security policies and set up VPN tunnels on Lucent VPN Firewall *Brick*® devices.

Reason for reissue

Updated for Release 9.1 features.

Who Should Read This Book?

The *LSMS Policy Guide* is intended to be read by network administrators who will be using the LSMS 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 LSMS, these administrators are referred to as LSMS Administrators and Group Administrators, depending on the privileges they have been given when their profiles were created.

What Is in this Book

The *LSMS 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 *Policy Guide* covers the following topics:

Chapter	Purpose
1. Introduction	This chapter provides an introduction to this document.
2. Brick 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.
3. 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.
4. 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.
5. 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.
6. 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.
7. 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.
8. 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.
9. Proxies	This chapter explains how to set up a Brick device to proxy incoming and outgoing SMTP and HTTP sessions.
10. User Authentication	This chapter explains how to set up user authentication, utilizing either a database residing on the LSMS, a RADIUS or SecurID server, or X.509 digital certificates.
11. 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 LSMS.
12. 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
13. Client Tunnels	This chapter explains how to configure a Brick to serve as the endpoint of a client tunnel. A client tunnel is a tunnel between a host running the Lucent IPSec Client application and a Brick.
Appendix A. Local Presence	This appendix provides information on how to create a pool of addresses that can be used by IPSec Client users to gain access to a local LAN, known as <i>local presence</i> .
Appendix B. Pre-Configured Lucent VPN FirewallBrick® Device Zone Rulesets	This appendix contains a description of pre-configured Brick Zone rulesets.
Appendix C. Denial of Services Attacks	This appendix describes how to enable the Brick to fend off denial-of-service (DoS) attacks.
Appendix D. Radius Enhancements	This appendix describes the Remote Authentication Dial-in Service (RADIUS) protocol.

What is Not in this Book

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

- How to log on and off the LSMS
- 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 LSMS
- How to create groups and set up additional Administrator accounts

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

Supported Brick devices

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

- Model 20 Brick device
- Model 50 Brick device
- Model 150 Brick device
- Model 350 Brick device
- Model 1100/1100A Brick device

- Model 700 Brick device
- Model 1200 Standard and HS Brick devices

Some of the above Brick device models require a specific patch of the current LSMS release in order to be fully supported. For details about the LSMS patch release required for a specific Brick device model, refer to the *User's Guide* for the Brick device model or check out the **VPN Firewall Portfolio** link on the following website for specific LSMS patch release information: <http://www.lucent.com/security/>

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 LSMS.
- Model number and serial number of the Brick.
- Platform the LSMS runs on (*Windows*[™] or *Solaris*[®] operating systems).
- Description of problem.
- Layout of your network. For example, is the Brick connected to a device such as a hub or router? Is the Brick operating as a bridge or is it using static routes? What are the Brick ports connected to? What is the IP address range and VBA for each zone? What is the security policy for each port?

After gathering the information, contact Lucent Security Customer Care at 1-866-LUCENT-8.

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 email your comments to the Comments Hotline (comments@lucent.com).

1 Lucent VPN Firewall *Brick*[®] Device Zone Rulesets

Overview

Purpose

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, proxied or tunneled by the Brick device.

Contents

What is a Brick Zone Ruleset?	1-2
Create a Brick Zone Ruleset	1-10
To Create a Security Rule	1-17
To Add Time and Day Restrictions to a Rule	1-27
To Add Advanced Features to a Rule	1-30
To Add Bandwidth Management to a Rule	1-37
To Add TOS/Alarm Capability to a Rule	1-42
To Configure IP Address(es) for Rule-Based Routing	1-45
To Maintain Brick Zone Rulesets	1-47
How to Maintain Security Rules	1-53



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 Bricks. 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."

Packet Filtering

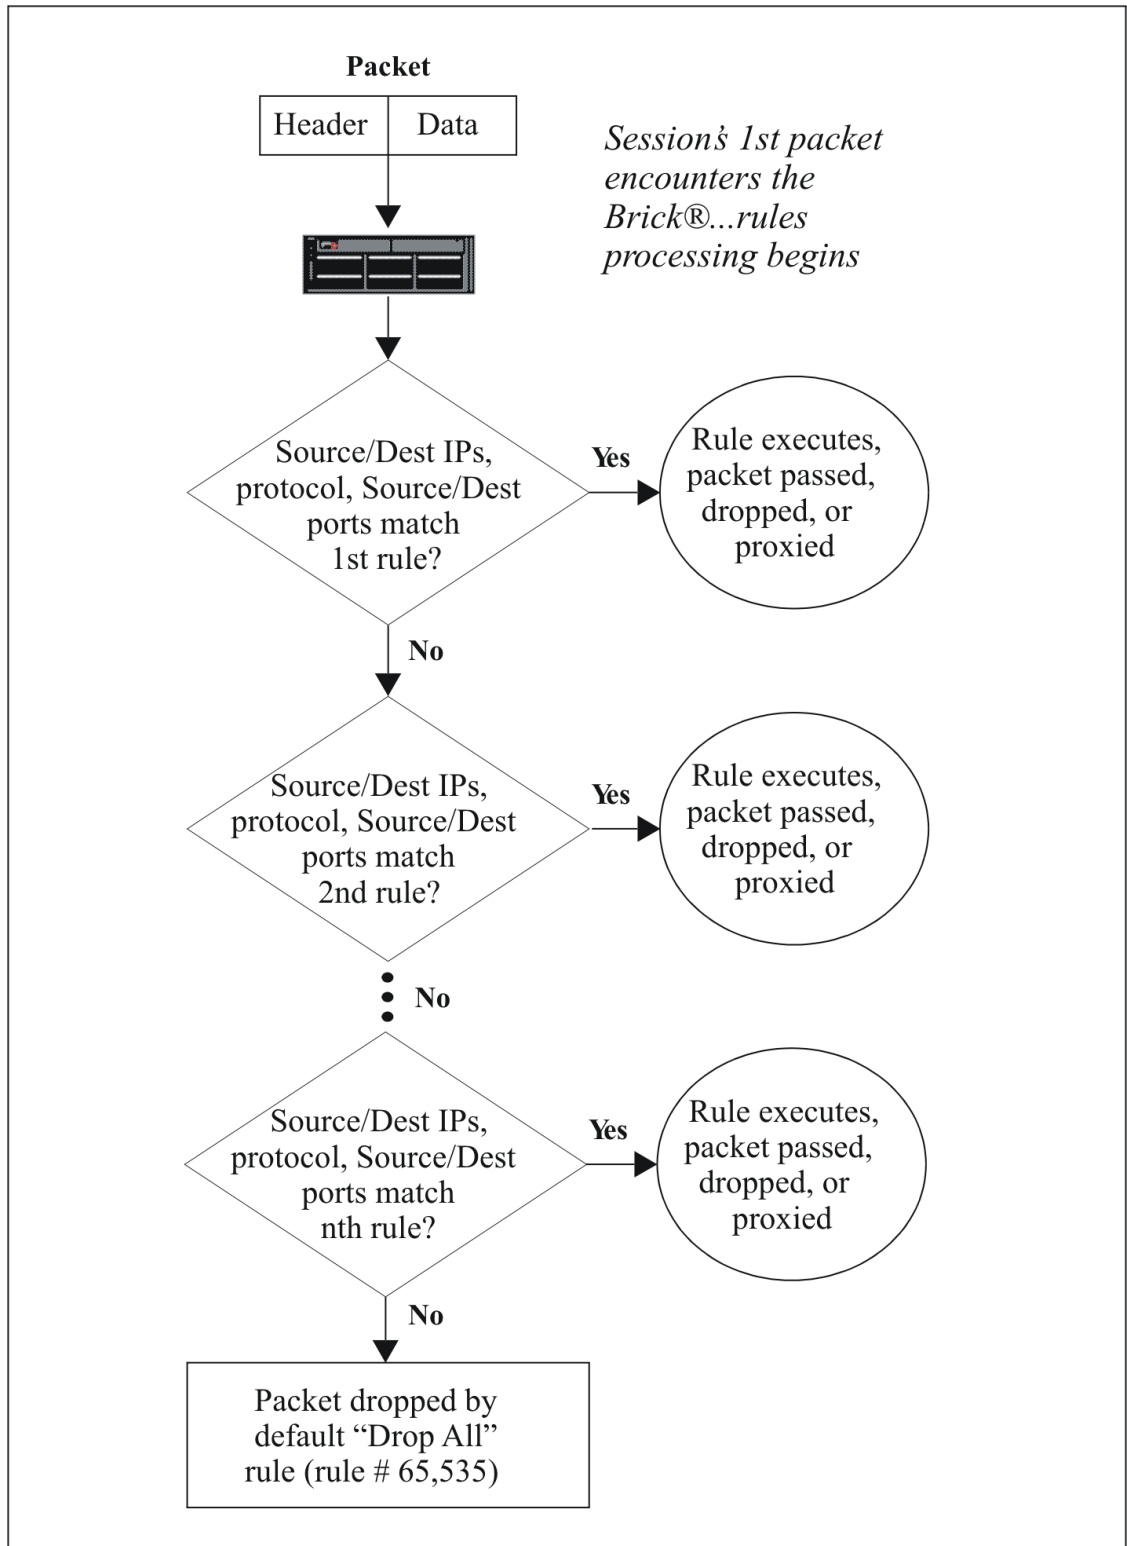
The Brick intercepts all inbound and outbound traffic traveling through each port and examine each packet header. The Brick then compares the information that it finds in the headers of the packets with the security rules in the zones that have been assigned to the ports.

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. LSMS 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

Starting with Release 9.1, the LSMS provides the option to configure a rule for HTTP, FTP, or SMTP 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 rule can be set up to route all packets of a certain traffic type (HTTP, FTP, SMTP) 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.

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.

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 LSMS.

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 LSMS 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.
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 LSMS or Group Administrators create. The numbers are assigned automatically by the LSMS 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 LSMS (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 (see [Figure 1-2, “Brick Zone Ruleset Editor”](#) (p. 1-10)).

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 LSMS 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 LSMS. 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.



Create a Brick Zone Ruleset

Overview

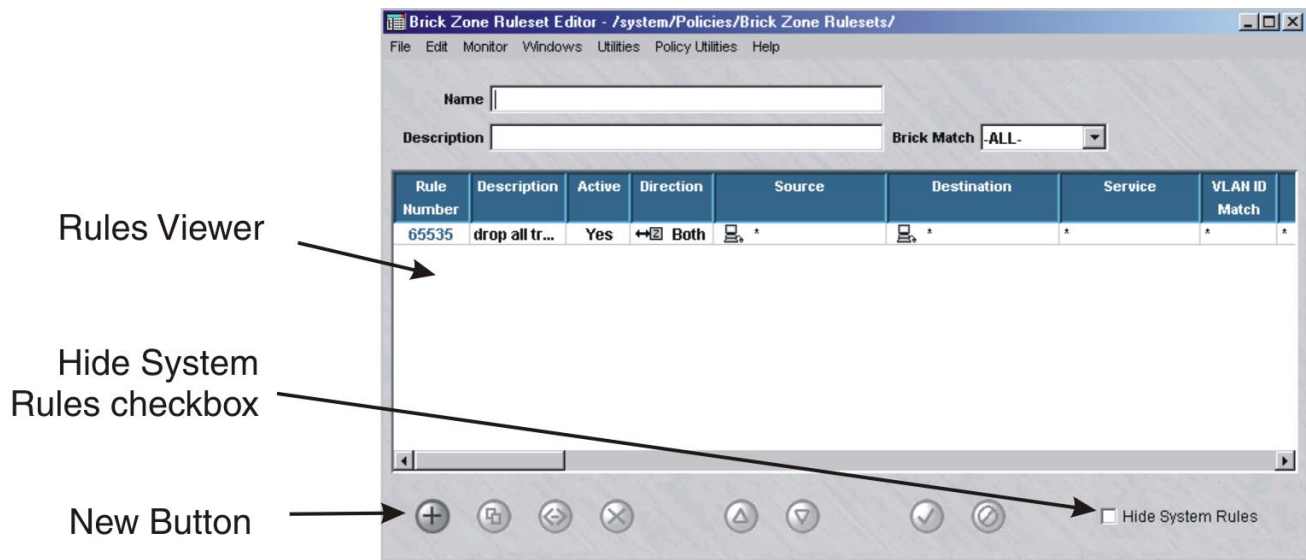
To set up a Brick zone ruleset, you have to display the Brick Zone Ruleset Editor and enter a zone name. To take effect, the ruleset then has to be assigned to a port on a Brick.

Display the Brick Zone Ruleset Editor

The easiest way to display the Brick Zone Ruleset Editor is to:

- 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. The Brick Zone Ruleset Editor is displayed (see [Figure 1-2, “Brick Zone Ruleset Editor”](#) (p. 1-10)).

Figure 1-2 Brick Zone Ruleset Editor



END OF STEPS

Create the Ruleset

To create a Brick zone ruleset, follow the steps below. Individual rules can be added using the procedure described in [“To Create a Security Rule” \(p. 1-17\)](#):

- 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 *Optional:* click the down arrow next to the **Brick Match** field to display a drop-down menu.

Select one of the following choices to filter the rules that are displayed in the Brick Editor:

- * - An asterisk (*). When this choice is selected, only generic rules in the zone ruleset that apply to all Bricks are displayed.
- **-ALL-** When this choice is selected, all rules in the zone ruleset for all Bricks are displayed. This is the default value.
- *Brickname* - If the zone has already been modified and contains rules that apply to a specific Brick, including generic rules, the name of the Brick appears in the drop-down list and can be selected.

A rule highlighted in yellow in the display is a Brick-specific rule.

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

END OF STEPS

Assign the Ruleset to a Physical Port

For the ruleset to take effect, you have to assign it to a port on at least one Brick. This is done using the Policy Assignment tab of the Brick Editor (see *Chapter 3*.

Configuring and Activating a Lucent VPN Firewall Brick® Device in the LSMS Administration Guide).

To assign a ruleset to a port, follow the steps below:

- 1 Open the Devices folder and click the Bricks folder to display all configured Bricks.
- 2 Double-click the Brick containing the port. The Brick Editor will appear.
- 3 Click **Policy Assignment** to display the Policy Assignment tab.
- 4 Double-click the port you want. The Brick Policy Assignment Editor is displayed, with the port you selected entered in the **Port** field (see [Figure 1-3, “Brick Policy Assignment Editor”](#) (p. 1-12)).

Figure 1-3 Brick Policy Assignment Editor

Brick Policy Assignment Editor - /system/Devices/Bricks/joesbrick

Basic Bandwidth TOS

Port ether2

Zone Ruleset dhcpzone_on_inside_if

TEP/VBA Addressing Method:
 Static Use DHCP
 Use PPPoE #1 Use PPPoE #2

Tunnel Endpoint /
Virtual Brick Address

Match Packets to or from this VBA

Hosts Behind Tunnel /
Zone IP Addresses *

Local Map Addresses

VPN Certificate

Authentication For External Users

Authentication Service

Authentication Timeout (mins) 480

Allowed Source IP Addresses *

OK Cancel

- 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 (see [“Overview” \(p. 6-1\)](#)).
- 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. See [“Overview” \(p. 2-1\)](#).
- 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 10, “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 LSMS host, you will have to enter an authentication service in the **Authentication Service for External Users** box.

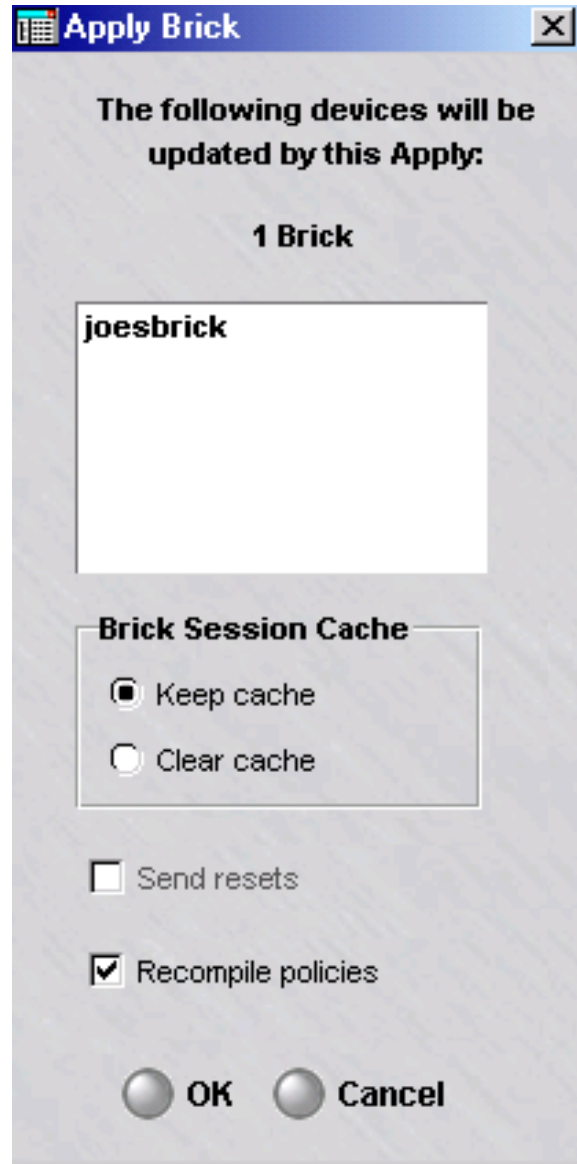
If the port will not be terminating a client tunnel, or if the user database resides on the LSMS host, you can skip this field. See [“Overview” \(p. 9-1\)](#) for a detailed explanation of user authentication.

- 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.

The Apply Brick window is displayed, with the associated Brick ([Figure 1-4, “Apply Brick Window”](#) (p. 1-15)).

Figure 1-4 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 5-23, “Network Topology - Example 3”](#) (p. 5-48)).

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

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



To Create a Security Rule

Overview

As [Figure 1-2, “Brick Zone Ruleset Editor” \(p. 1-10\)](#) shows, 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.

Task

Complete the following steps to create a new Brick zone ruleset.

- 1 Access the Brick Zone Ruleset Editor as described in [“Display the Brick Zone Ruleset Editor” \(p. 1-10\)](#).

The Brick Zone Ruleset Editor is displayed (see [Figure 1-2, “Brick Zone Ruleset Editor” \(p. 1-10\)](#)).

- 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 1-5, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-18)).

Figure 1-5 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 (i.e., 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 LSMS application.

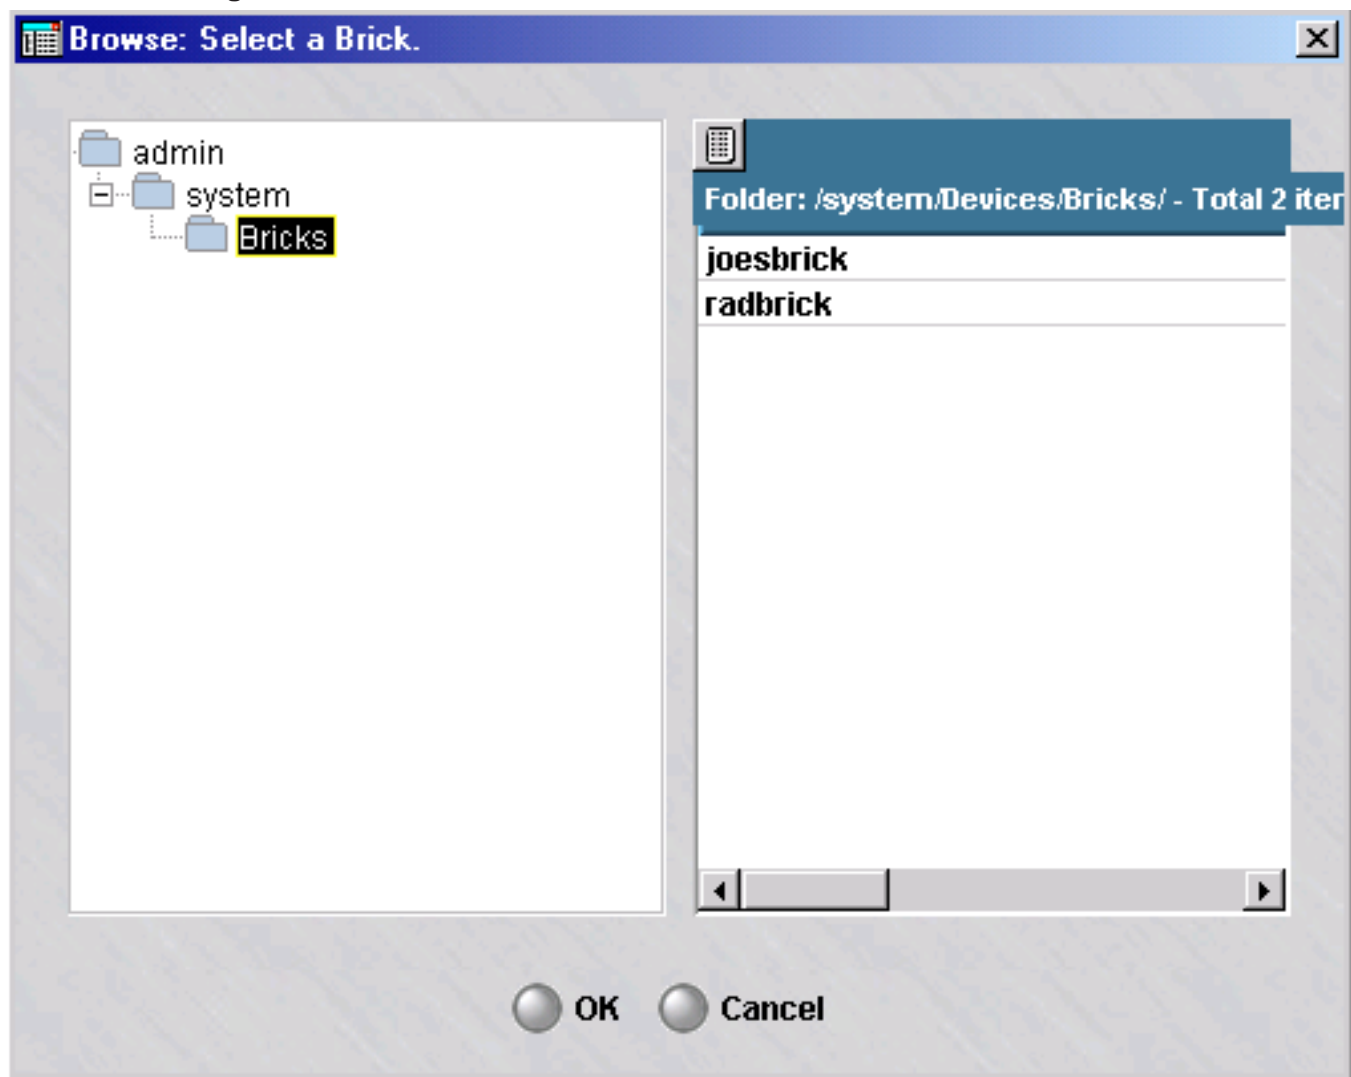
- 6 *Optional:* To make a rule specific to one or more Bricks, enter the name of a Brick device or click the down arrow next to the **Brick Match** field to display a drop-down list of choices.

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 1-6, “Browse: Select a Brick Window”](#) (p. 1-20)).

Figure 1-6 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. 1-42) on page 2-29 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	The Brick device will permit the session through the port.

Action	Explanation
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>To use the Proxy action, you must first put an entry that matches the service in the Proxy Table for this zone (see Chapter 8, “Proxies”).</p>
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. 11-1) and “Overview” (p. 12-1)).</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.

Application Layer Protocol	Basic Auditing	Detailed Auditing	Exception Auditing
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)
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)

Application Layer Protocol	Basic Auditing	Detailed Auditing	Exception Auditing
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
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 message 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 LSMS 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 LSMS time or a time offset relative to the LSMS 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 LSMS 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 LSMS-generated rules. If the color of the Rule Number is blue on the Brick Zone Ruleset Editor, it is an LSMS-generated rule; the Time and Day tab does not appear on the Brick Zone Rule Editor when you edit an LSMS-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 1-5, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-18)), click the **Time and Day** tab, and then click the **Time-of-Day**

Active checkbox. The time and day fields are displayed (see [Figure 1-7, “Time and Day Fields”](#) (p. 1-28)).

Figure 1-7 Time and Day Fields

The screenshot shows the 'Brick Zone Rule Editor' window with the 'Time and Day' tab selected. The 'Time-of-Day Active' checkbox is checked. The 'Start Time' field is set to '00:00:00' and the 'End Time' field is set to '23:59:59'. The 'Day of Week' section has checkboxes for Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday, all of which are checked. The 'Time Source' section has radio buttons for 'Brick' and 'LSMS', with 'LSMS' selected. At the bottom of the window are 'OK' and 'Cancel' buttons.

- 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

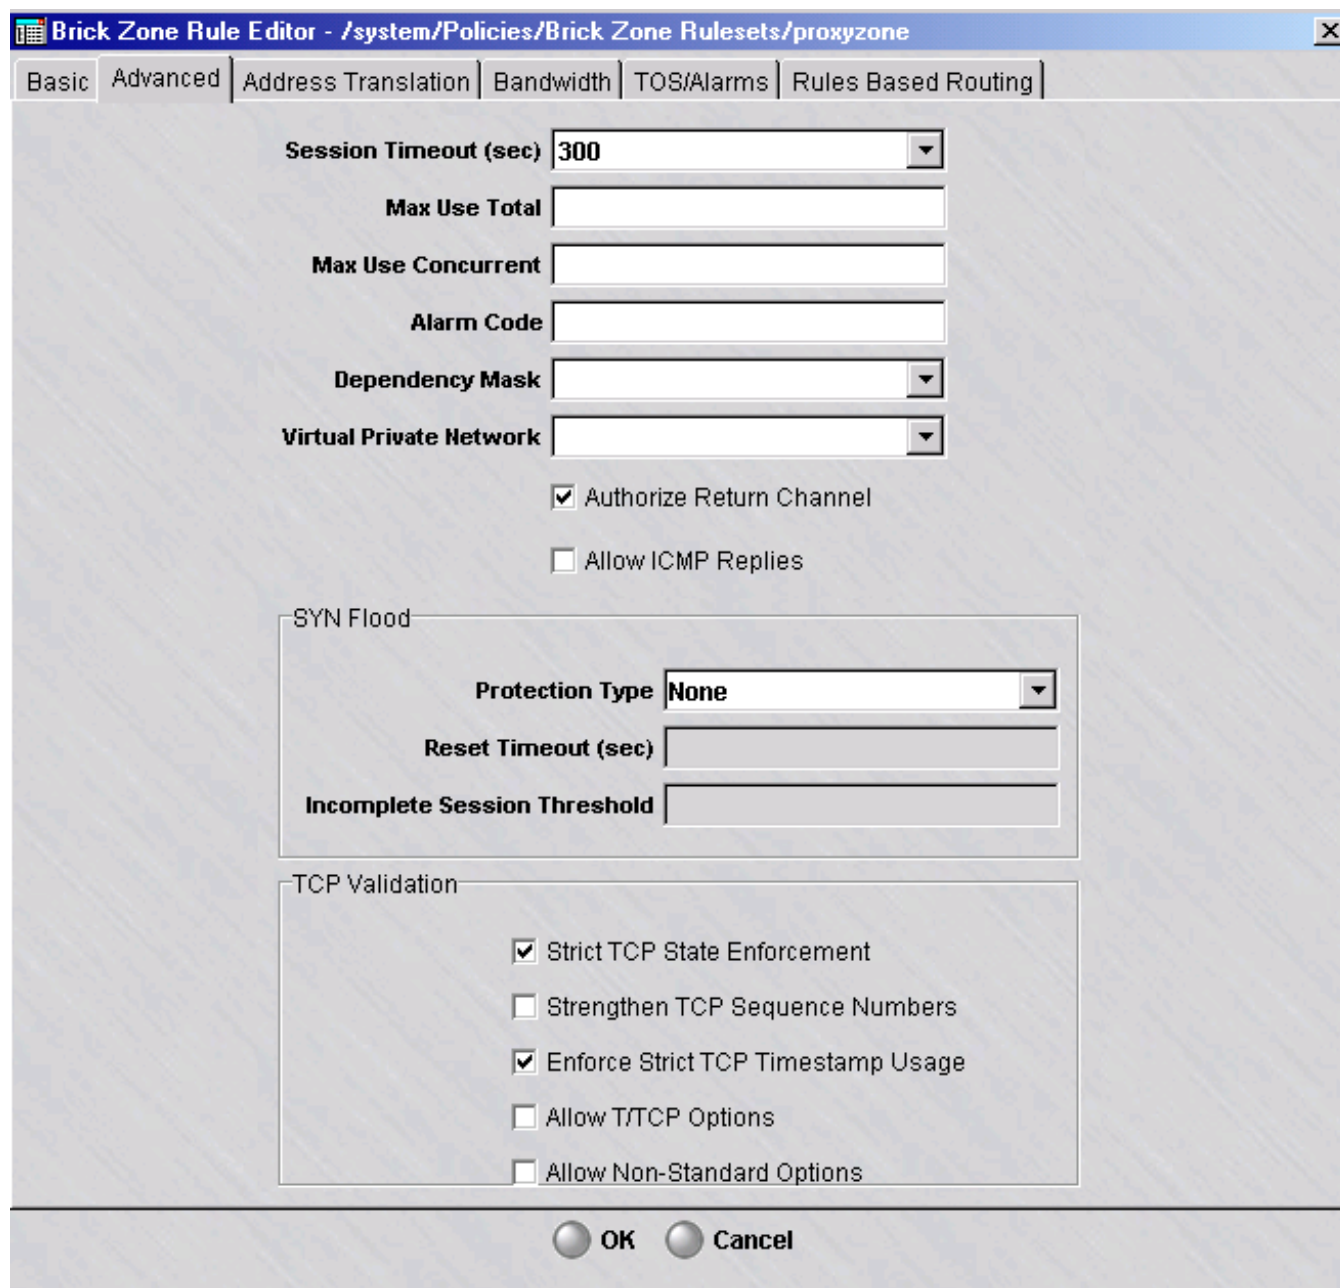
The LSMS provides a number of advanced features that you can add to a security rule. These features allow you to change the session timeout, set usage limitations on the rule, associate an alarm with the rule, and set up SYN flood protection and perform strict TCP validation.

These features can be accessed by clicking **Advanced** on the Rule Editor to display the Advanced tab, which is shown in [Figure 1-8, “Rule Editor \(Advanced Tab\)”](#) (p. 1-31) .

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 1-8 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.

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 (see *Chapter 10. Using the Configuration Assistant* in the *LSMS 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 *Chapter 11* in the *LSMS 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

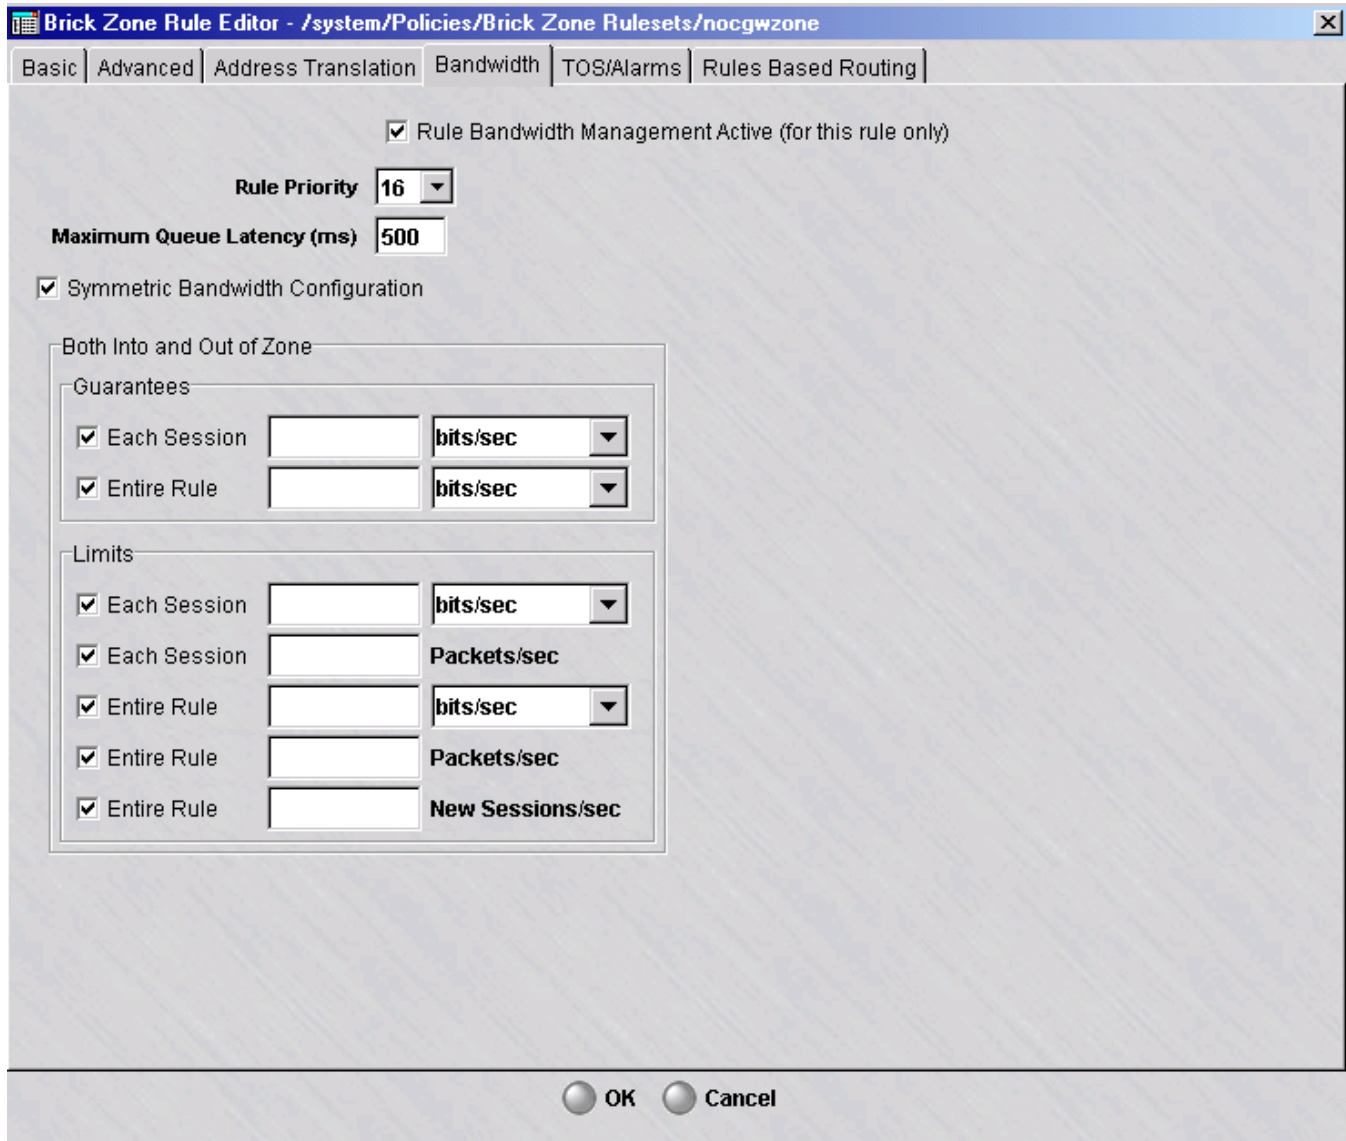
A guarantee sets the minimum acceptable bandwidth when traffic must be reduced in one place to accommodate higher traffic demands elsewhere. A limit caps the allowable bandwidth, packet rate, or new session rate under high traffic loads. These limits can defend against denial of service attacks, for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

To access the feature, follow the steps below:

-
- 1 In the Brick Zone Rule Editor ([Figure 1-2, “Brick Zone Ruleset Editor” \(p. 1-10\)](#)), 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 (see [Figure 1-9, “Brick Zone Rule Editor \(Bandwidth Tab\)”](#) (p. 1-38)).

Figure 1-9 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 (i.e., 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 1-10, “Asymmetric Bandwidth Configuration”](#) (p. 1-40).

Figure 1-10 Asymmetric Bandwidth Configuration

The screenshot shows the 'Brick Zone Rule Editor' window with the 'Bandwidth' tab selected. The window title is '/system/Policies/Brick Zone Rulesets/nocgwzone'. The 'Basic' tab is active, showing the following settings:

- Rule Bandwidth Management Active (for this rule only)
- Rule Priority: 16
- Maximum Queue Latency (ms): 500
- Symmetric Bandwidth Configuration

The configuration is split into two sections: 'Into Zone' and 'Out of Zone'. Each section has 'Guarantees' and 'Limits' sub-sections.

Into Zone Guarantees:

- Each Session [] bits/sec
- Entire Rule [] bits/sec

Into Zone Limits:

- Each Session [] bits/sec
- Each Session [] Packets/sec
- Entire Rule [] bits/sec
- Entire Rule [] Packets/sec
- Entire Rule [] New Sessions/sec

Out of Zone Guarantees:

- Each Session [] bits/sec
- Entire Rule [] bits/sec

Out of Zone Limits:

- Each Session [] bits/sec
- Each Session [] Packets/sec
- Entire Rule [] bits/sec
- Entire Rule [] Packets/sec
- Entire Rule [] New Sessions/sec

At the bottom of the window are 'OK' and 'Cancel' buttons.

- 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

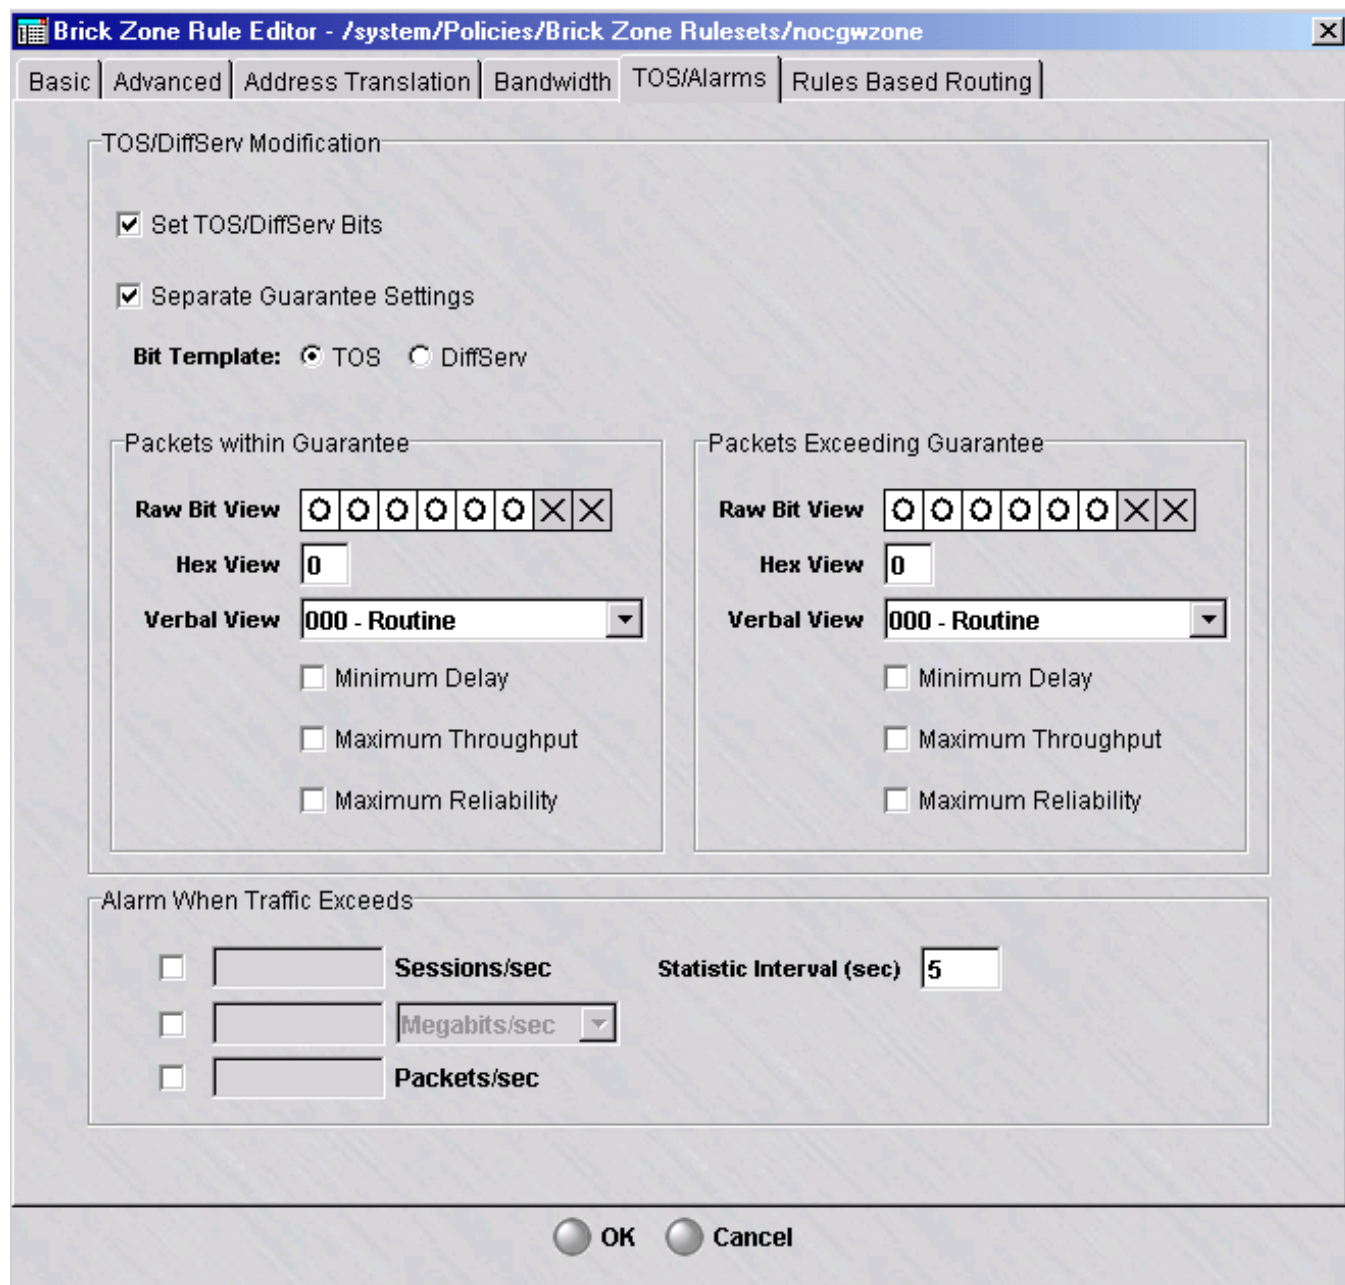
Overview

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” \(p. 1-10\)](#)) to display the

TOS/Alarms tab (see Figure 1-11, “Brick Zone Rule Editor (TOS/Alarms tab)” (p. 1-43)).

Figure 1-11 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 1-11, “Brick Zone Rule Editor \(TOS/Alarms tab\)”](#) (p. 1-43)) 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 Rule-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.

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. 1-10)), 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 1-12, “Brick Zone Rule Editor (Rules Based Routing tab)” (p. 1-45)).

Figure 1-12 Brick Zone Rule Editor (Rules Based Routing tab)



2 To	Do This
Configure an IP address for forward routing of packets that match the rule	Enter an IP address in the Forward Route Address field.
Configure an IP address for reverse routing of packets that match the rule	Enter an IP address in the Reverse Route Address field.

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

END OF STEPS



To Maintain 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 below entitled “[Apply a Ruleset](#)” (p. 1-50) explains how to modify and delete rules in the Brick zone ruleset.

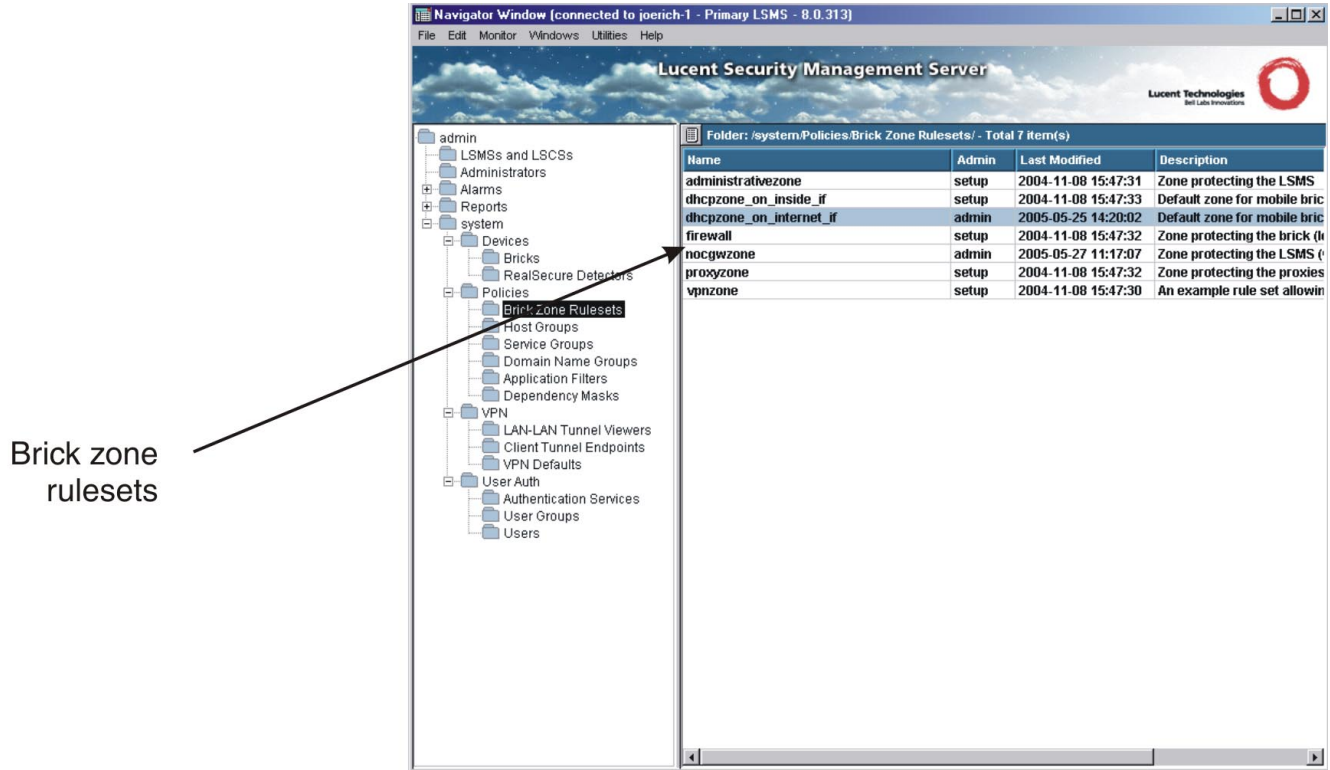
View Existing Zone Rulesets

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 1-13, “View Brick Zone Rulesets”](#) (p. 1-48)).

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 1-13 View Brick Zone Rulesets



END OF STEPS

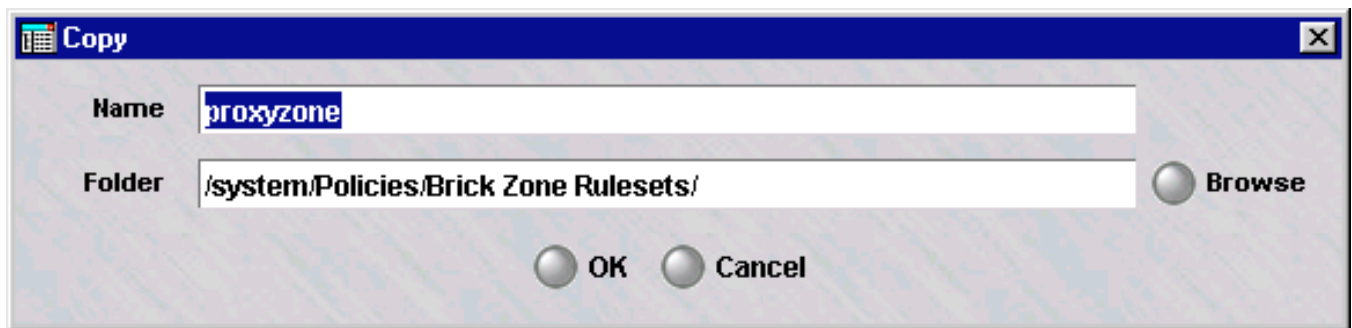
Copy a Brick Zone Ruleset

Complete the following steps to copy a Brick zone ruleset.

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

Result A Copy window is displayed (Figure 1-14, “Copy Window” (p. 1-49)).

Figure 1-14 Copy Window



- 2 In the **Name** field, enter the name of the copy.
- 3 In the **Folder** field, click **Browse** and select the folder you want to copy this Brick zone ruleset to.
- 4 Click **OK** to dismiss the Copy window. The ruleset will be copied to the folder you indicated.

END OF STEPS

Move a Ruleset

To move a ruleset, follow the steps below:

- 1 With the existing rulesets displayed in the Navigator window, right-click the ruleset 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 the ruleset to, and click **OK** to dismiss the Browse window. The ruleset will be moved to the folder you indicated.

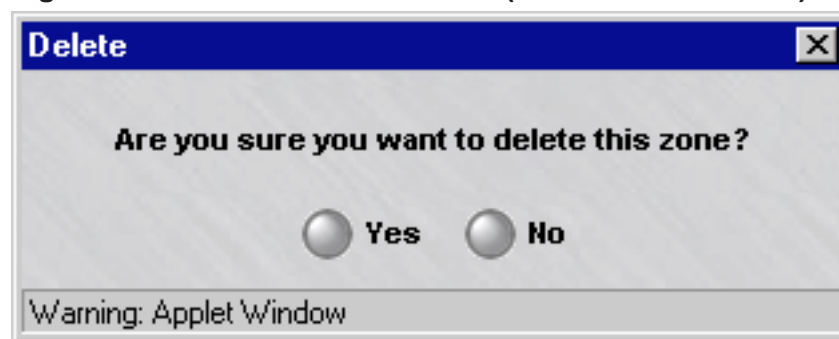
END OF STEPS

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 1-15, “Confirmation Window \(Brick Zone Rulesets\)”](#) (p. 1-50) will appear.

Figure 1-15 Confirmation Window (Brick Zone Rulesets)



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

END OF STEPS

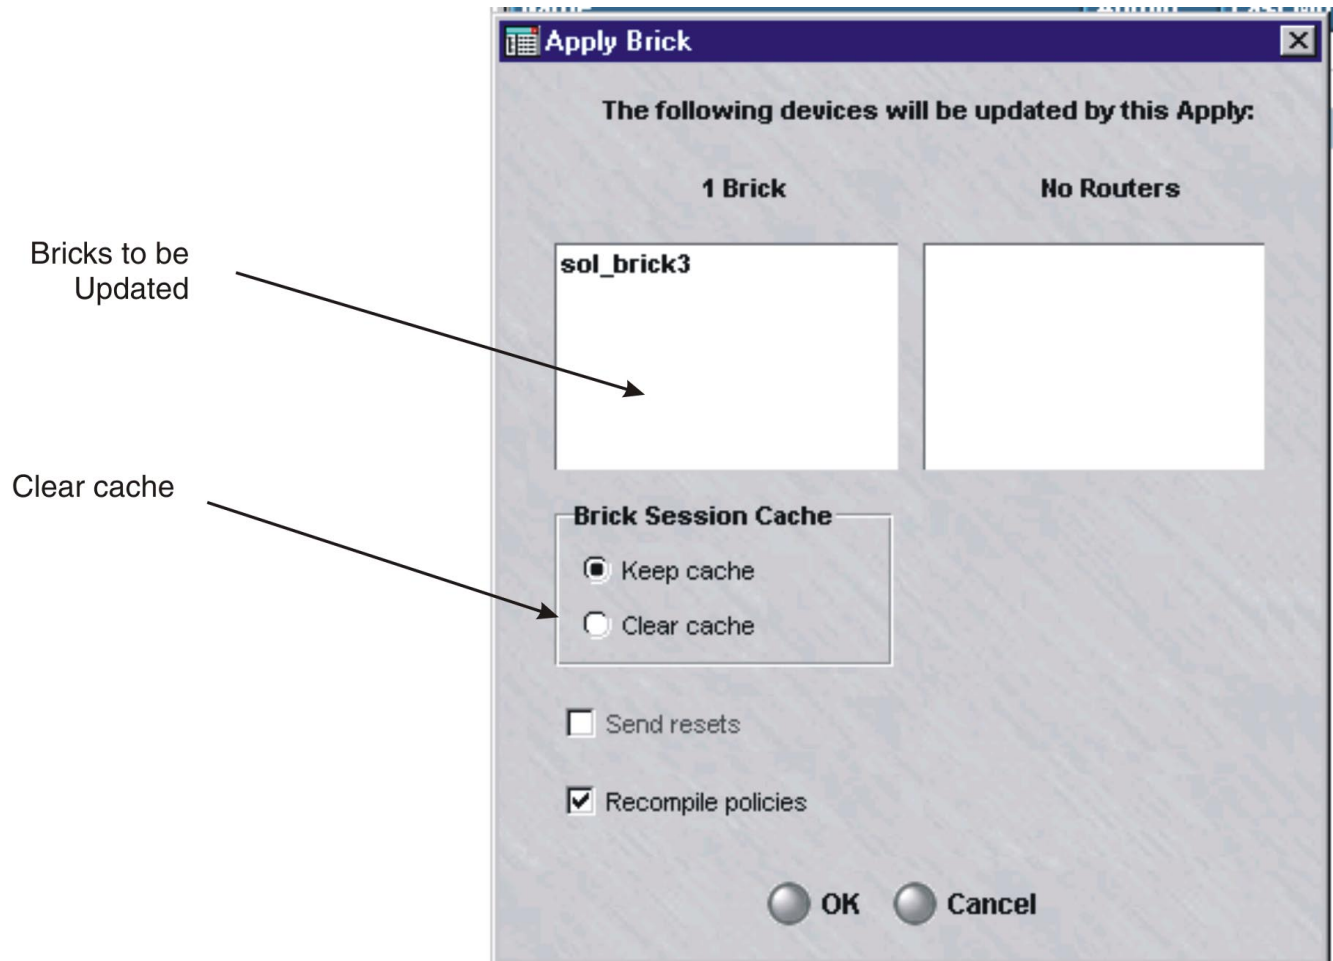
Apply a Ruleset

Whenever changes are made to a ruleset that has been assigned to a port on one or more Bricks, the ruleset has to be applied to the Bricks. 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 1-16, “Apply Policy Window”](#) (p. 1-51), indicates the Bricks that will be updated by this apply.

Figure 1-16 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.

□

How to Maintain 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.

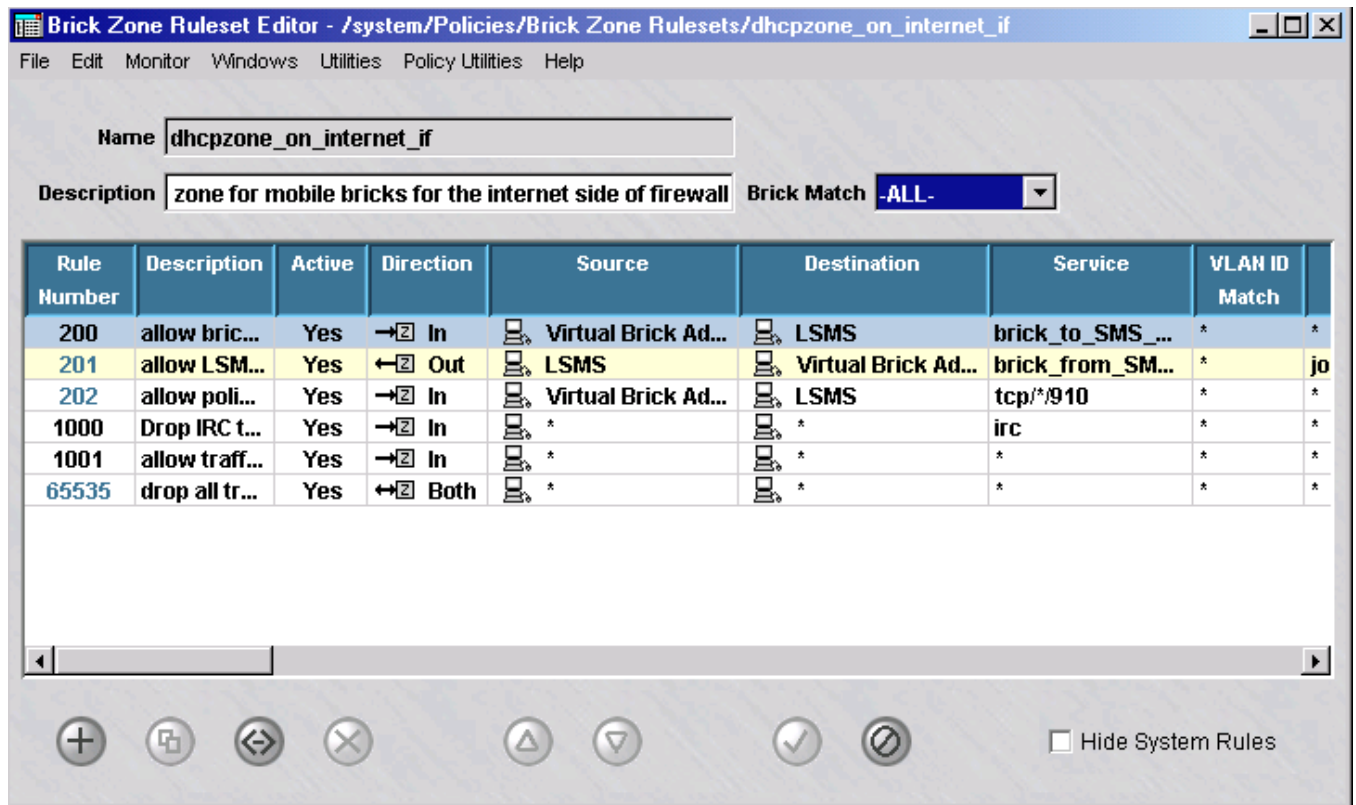
View Security Rules

To view the security rules in a ruleset, do the following:

- 1 Access the Brick Zone Ruleset Editor as described in “Display the Brick Zone Ruleset Editor” (p. 1-10).

The Brick Zone Ruleset Editor (see Figure 1-17, “Brick Zone Ruleset Editor (Viewing Rules in a Zone Ruleset)” (p. 1-53)) shows all the security rules that have been created to date for that ruleset, in ascending order.

Figure 1-17 Brick Zone Ruleset Editor (Viewing Rules in a Zone Ruleset)



- 2 To filter the display of rules in the zone ruleset, 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 only the generic rules in the zone ruleset that apply to all Bricks are displayed.
- **-ALL-**, which means that all rules in the zone ruleset are displayed.
- A *Brickname*, which means that all the rules (including the generic rules) in the zone ruleset that apply to the selected Brick are displayed.











A rule highlighted in yellow in the display is a Brick-specific rule.

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  (green)	Drop  (red)	Proxy  (yellow)	VPN  (purple)	VPN Proxy  (purple/yellow)

END OF STEPS

Modify a Rule

To modify a security rule, follow the steps below:

- 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. 1-10)).
- 3 Double-click the rule you want to modify. The Brick Zone Rule Editor will appear (see [Figure 1-5, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-18)) 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 [“Apply a Ruleset”](#) (p. 1-50) above).

END OF STEPS

Duplicate a Rule

To duplicate a security rule, follow the steps below:

- 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 1-5, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-18)) 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

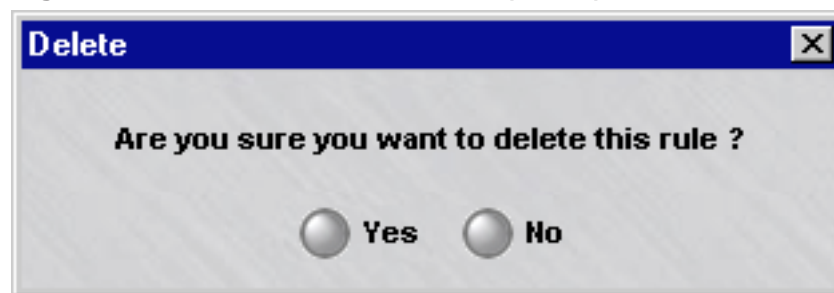
Delete a Rule

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 1-18, “Confirmation Window \(Rules\)”](#) (p. 1-56) will appear

Figure 1-18 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 LSMS 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. 1-10)).

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.



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 LSMS application.

Contents

What is a Host Group?	2-2
How to Set Up a Host Group	2-4
Global Host Groups	2-8
Dynamic Host Groups	2-10
Nested Host Groups	2-12
How to Maintain a Host Group	2-15
Host Groups Provided with the LSMS Application	2-21



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 instances where host groups can be used:

- To identify the source and destination hosts in a security rule in a Lucent VPN Firewall *Brick*[®] device 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 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 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.



How to Set Up a Host Group

Overview

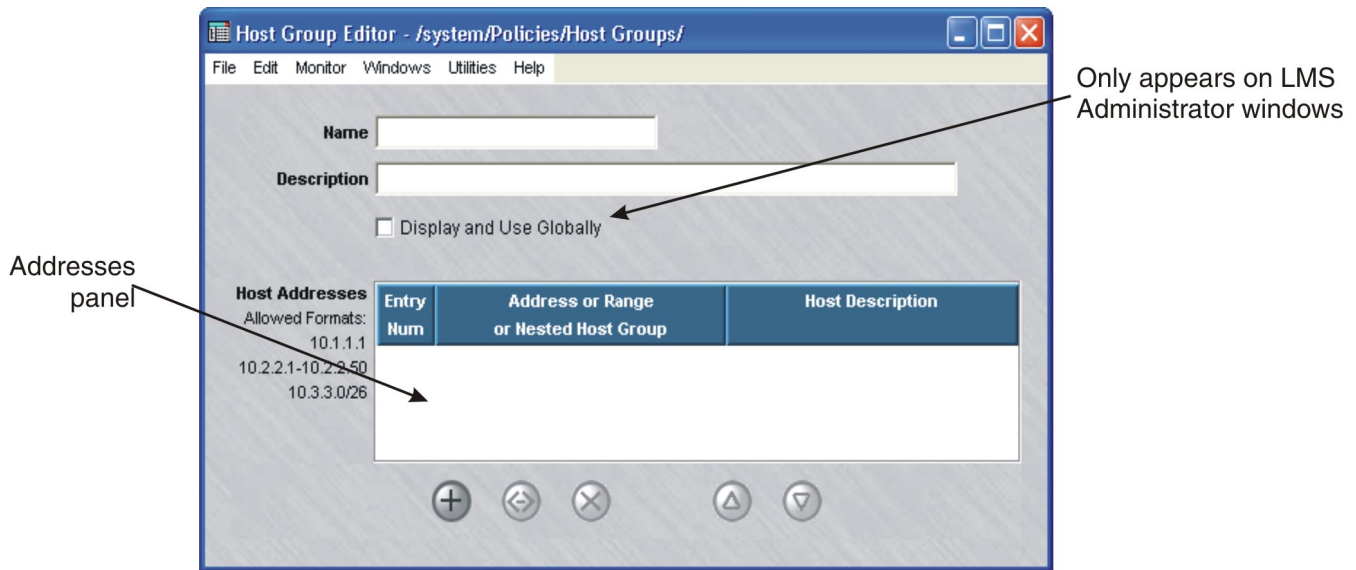
To set up a host group, you have to display the Host Groups Editor and enter the name of the host group and the IP address(es) to be included in the host group.

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 will appear. It is shown in [Figure 2-1, “Host Group Editor”](#) (p. 2-4).

Figure 2-1 Host Group Editor



END OF STEPS

Create a Host Group

To create a host group, follow the steps below.

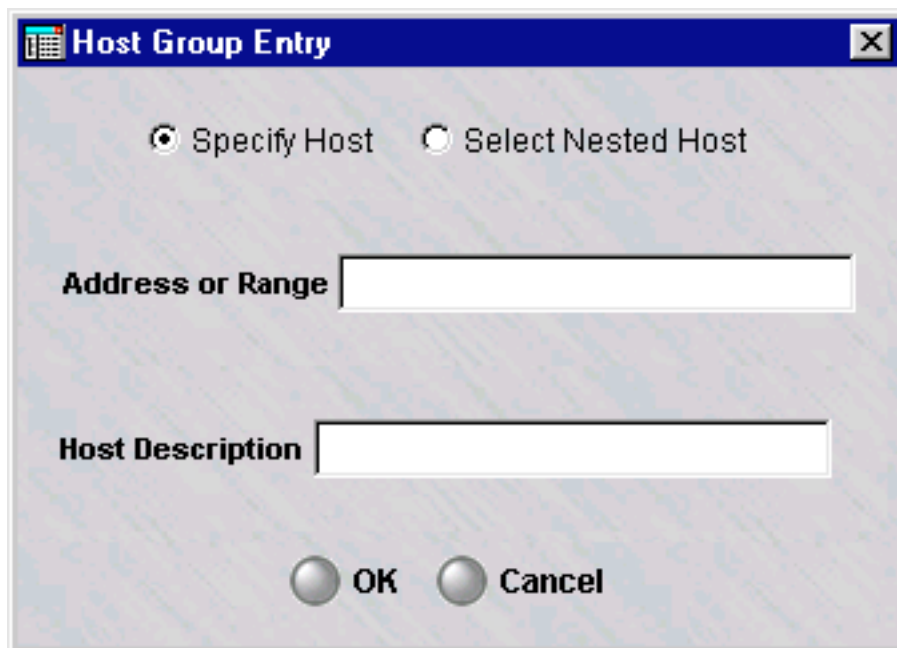
- 1 In the **Name** field, enter a unique name to identify this host group. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers and certain special characters.
- 2 In the **Description** field, you can enter an optional description of the host group. The 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 “[View Host Groups](#)” (p. 2-15).

Important! If you are an LSMS 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 LSMS group, but can be seen and used in every other group. Only LSMS Administrators can create global host groups. Global host groups are dynamically linked, and will update everywhere when modified.

Create a Global Host Group

If you are an LSMS 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 [“Modify a Host Group”](#) (p. 2-17) 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.

View Global Host Groups

Global host groups appear in the Navigator window, just like standard host groups (see [“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.

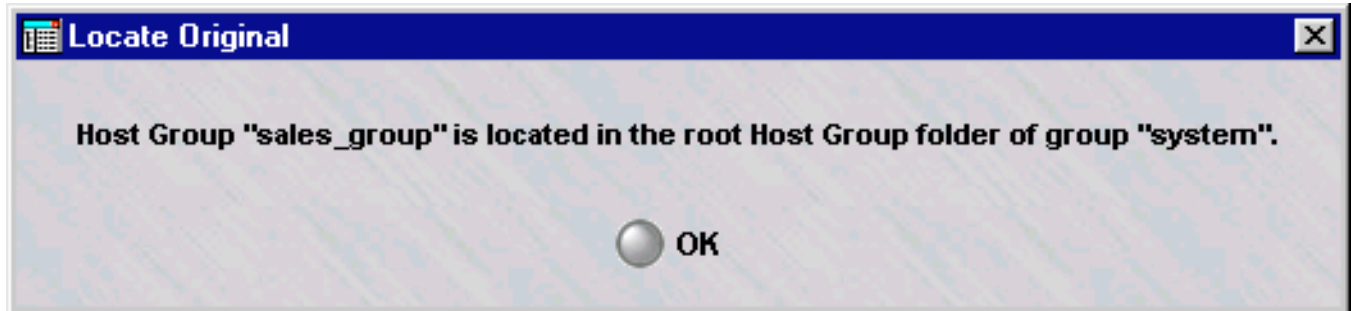
However, any global host groups created in *a different group* can be identified by a globe icon that appears to the left of the entry, as shown below.

 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



Removing 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 [“Delete a Host Group”](#) (p. 2-19) below).

Similarly, only the original source of a global host group can be moved (see [“Move a Host Group”](#) (p. 2-19) 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 LSMS 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 LSMS to the Release 9.1 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.

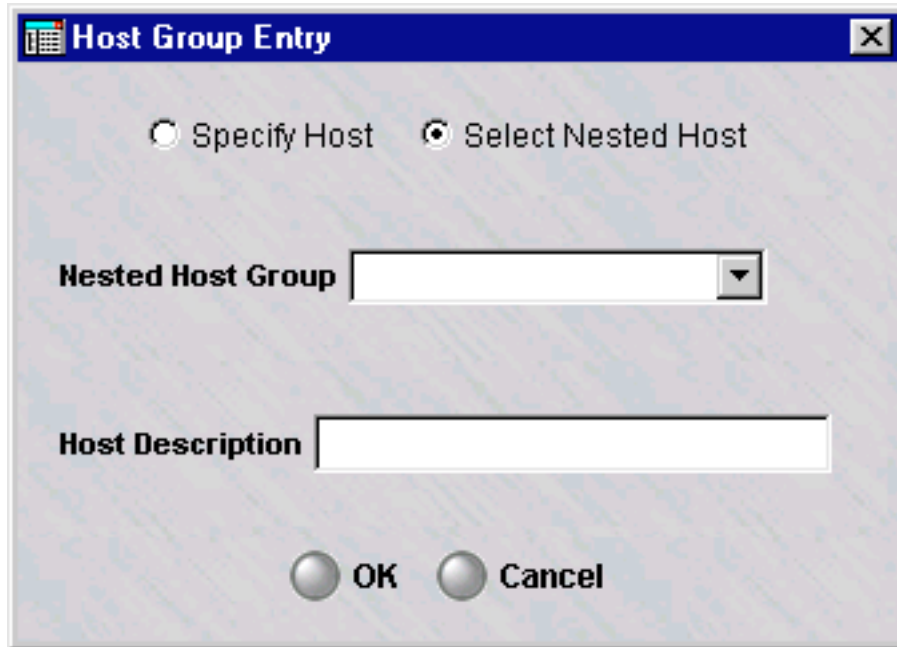
Task

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

-
- 1 Create a host group as previously explained (see [“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)



-
- 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. 4-13) shows an example of two host groups nested within a another host group.

Figure 2-6 Nested Host Groups

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

END OF STEPS



How to Maintain a Host Group

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.

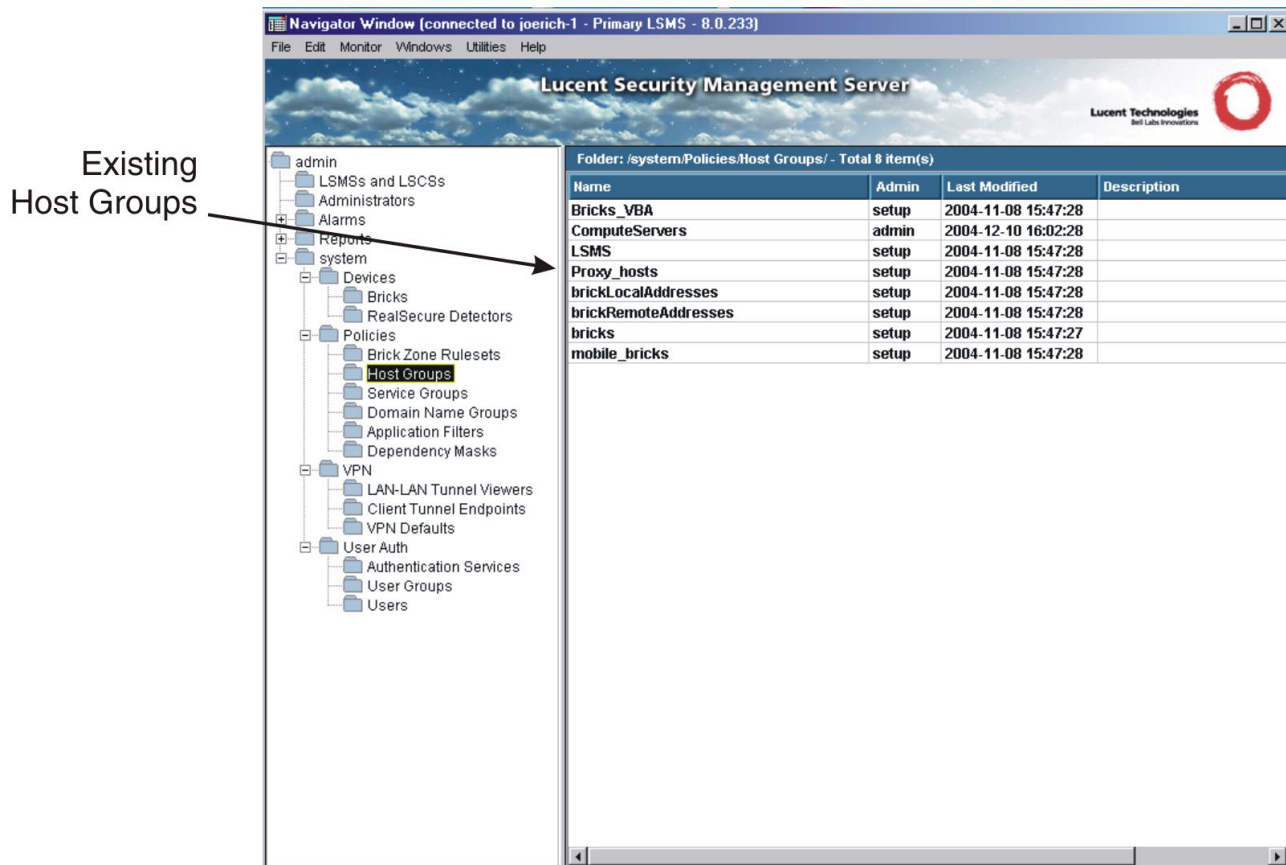
View Host Groups

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.

- 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. 2-16).

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. 2-16) shows, there are eight host groups that are provided with the LSMS application (**bricks**, **mobile_bricks**, **brickRemoteAddresses**, **brickLocalAddresses**, **Proxy_hosts**, **Bricks_VBA**, **LSMS** and **ComputeServers**). 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).

END OF STEPS

Find Entities Using a Host Group

- 1 It is possible to identify all the entities (such as Brick zone rulesets) using a host group. With the host groups displayed in the Navigator window, right-click the host group you want and select **Find Entities Using this Host Group** from the pop-up menu. A window similar to the one shown in [Figure 2-8, “Entities Found Window”](#) (p. 2-17) will appear and list all the entities using this host group.

Figure 2-8 Entities Found Window



As [Figure 2-8, “Entities Found Window”](#) (p. 2-17) shows, the window gives the entity type (in the example, a Brick zone ruleset), the group and subfolder in which the entity is found, and details about the entity (if the entity is a Brick zone ruleset, it gives the specific rules in which the global host group is used). You may bring up an editor for the entity by right-clicking on the entry and choosing **Edit**.

END OF STEPS

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.

END OF STEPS

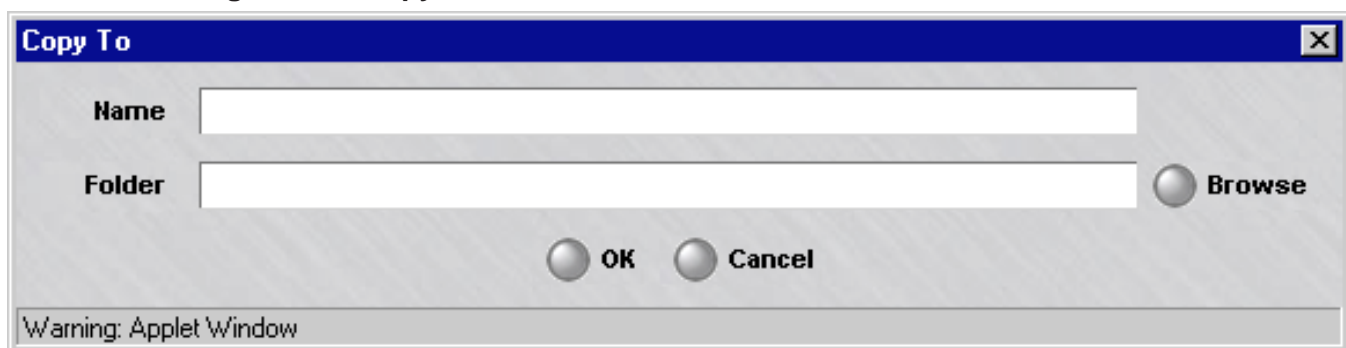
Copy a Host Group

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. However, you can only copy it to the folder labeled "Host Groups" or to a subfolder under that folder. If you try to copy it to another folder, you will get an error message. If a host group is copied to another group, and the host group has nested host groups within it, the nested host groups are copied (if one with the same name and content does not exist) and possibly renamed (if one with the same name exists, but has different content).

To copy 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 copy and select **Copy** from the pop-up menu. A Copy To window will appear (see [Figure 2-9, "Copy To Window"](#) (p. 2-18)).

Figure 2-9 Copy To Window



-
- 2 In the **Name** field, enter the name you want to give the copy. If you are copying the host group to the same group, you must assign the copy a new name.

-
- 3 In the **Folder** field, click **Browse** and select the folder you want to copy this host group to.
-
- 4 Click **OK** to copy the host group and dismiss the Copy To window. You will be returned to the Navigator window.

END OF STEPS

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 "["Find Entities Using a Host Group" \(p. 2-17\)](#)" 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

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 “[Find Entities Using a Host Group](#)” (p. 2-17) above).

To delete a host group, follow the steps below:

- 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. 2-20) 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 LSMS Application

Overview

Three host groups are automatically created when the LSMS 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 LSMS application. The following explains:

- *administrativezone*
The purpose of the *administrativezone* ruleset is to protect an LSMS that is managing Bricks, but not routers (or is managing routers in the clear). It is assigned to the port connecting the LSMS to the Brick, and it only allows the LSMS 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 LSMS (Rule #203), and the destination of the rule allowing traffic to the Bricks from the LSMS (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 LSMS.
- *nocgwzone*
The purpose of the *nocgwzone* ruleset is to protect an LSMS 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 Lucent 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.
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 LSMS 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 LSMS, 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 (as of R8.0) and, therefore, appears in every group. This group contains the IP address of a single LSMS (in a standalone configuration) or two LSMSs (in a redundant configuration). The public IP addresses of these LSMSs, 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 *LSMS* host group. This group contains the IP address of every Compute Server (LSCS) 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 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 *LSMS Administration Guide*.)

The *brickLocalAddresses* group should be populated with the "public" address that the LSMS 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.



3 Domain Name Groups

Overview

Purpose

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

Contents

What is a Domain Name Group?	3-2
How to Set Up a Domain Name Group	3-3
Global Domain Name Groups	3-8
Nested Domain Name Groups	3-11
How to Maintain a Domain Name Group	3-13



What is a Domain Name Group?

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.



How to Set Up a Domain Name Group

Overview

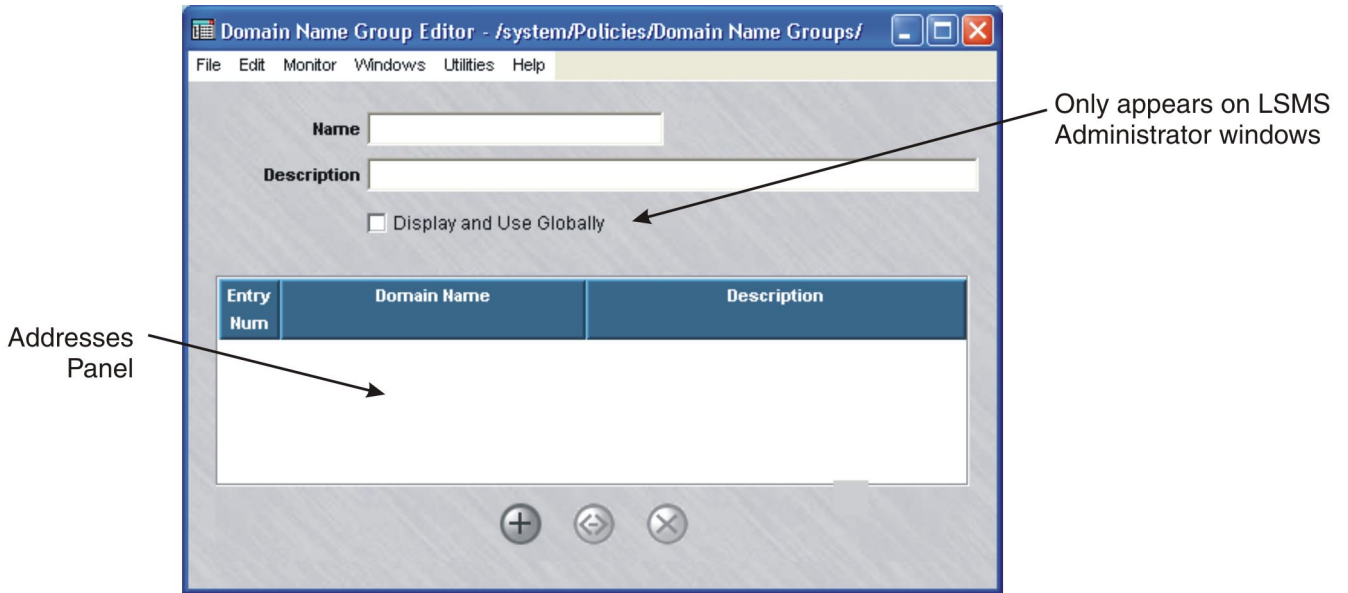
To set up a domain name group, you have to display the Domain Name Groups Editor and enter the name of the domain name group and an optional description.

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 will appear. It is shown in [Figure 3-1, “Domain Name Group Editor”](#) (p. 3-3).

Figure 3-1 Domain Name Group Editor



END OF STEPS

Create a Domain Name Group

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

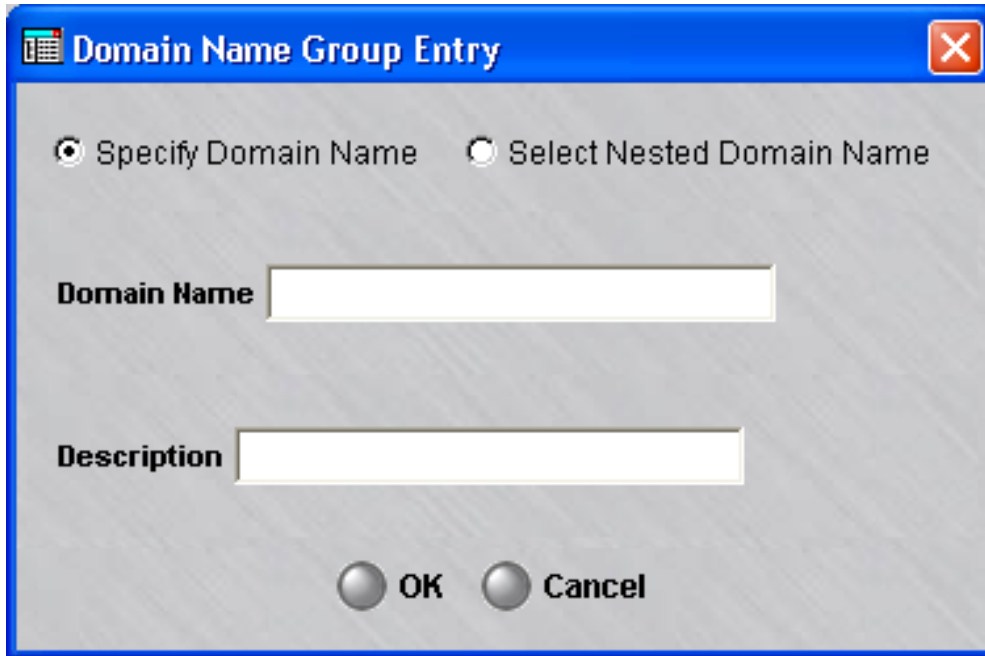
- 1 In the **Name** field, enter a unique name to identify this domain name group. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers and certain special characters.
- 2 In the **Description** field, you can enter an optional description of the domain name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

Important! If you are an LSMS 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. See "Global Domain Name Groups" below 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. 3-5).

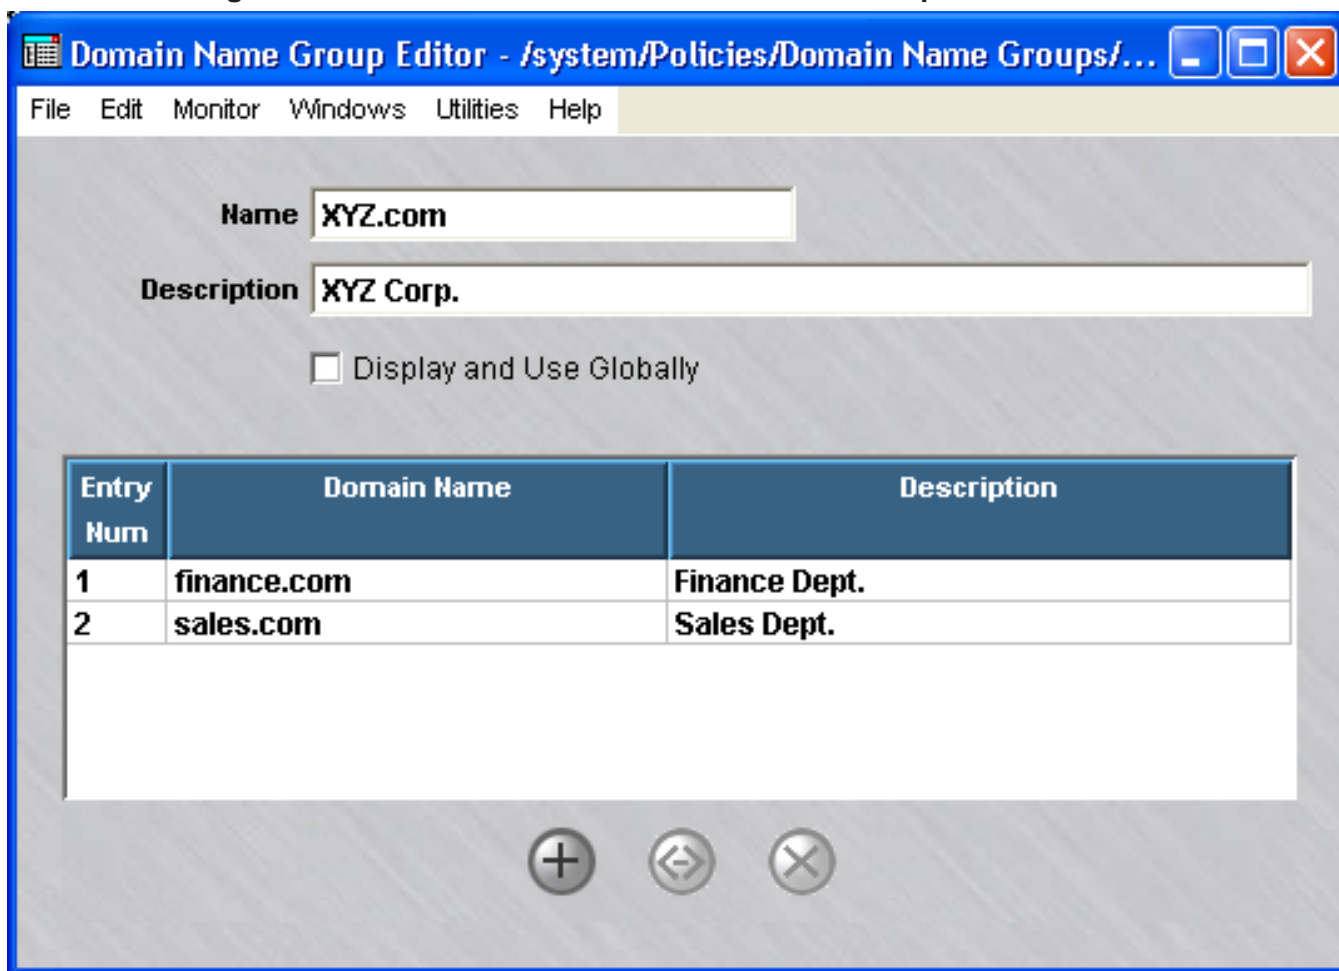
Figure 3-2 Domain Name Group Entry Window



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, see “Nested Domain Name Groups” below.

-
- 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” \(p. 3-6\)](#) 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]**.

END OF STEPS



Global Domain Name Groups

Definition

A global domain name group is a domain name group that is created in one LSMS group, but can be seen and used in every other group. Only LSMS 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 LSMS 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 (see “View Domain Name Groups” below). 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. 3-9) will appear.

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 (see “Delete a Domain Name Group” below).

Similarly, only the original source of a global domain name group can be moved (see “Move a Domain Name Group” below).

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.

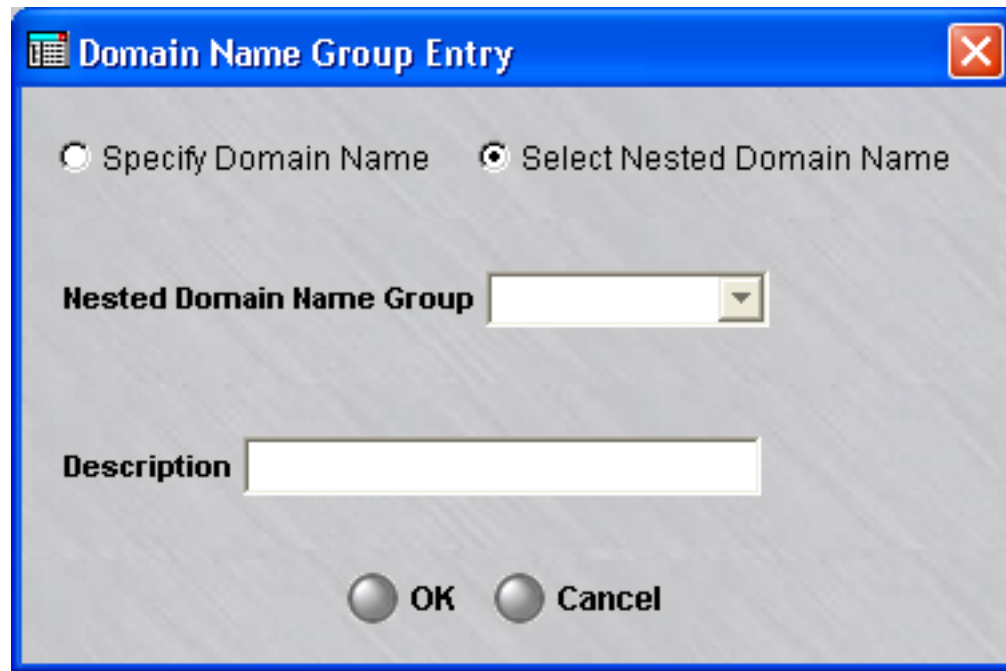
Task

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. 3-5)).
 - 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. 3-5)). 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



How to Maintain a Domain Name Group

Overview

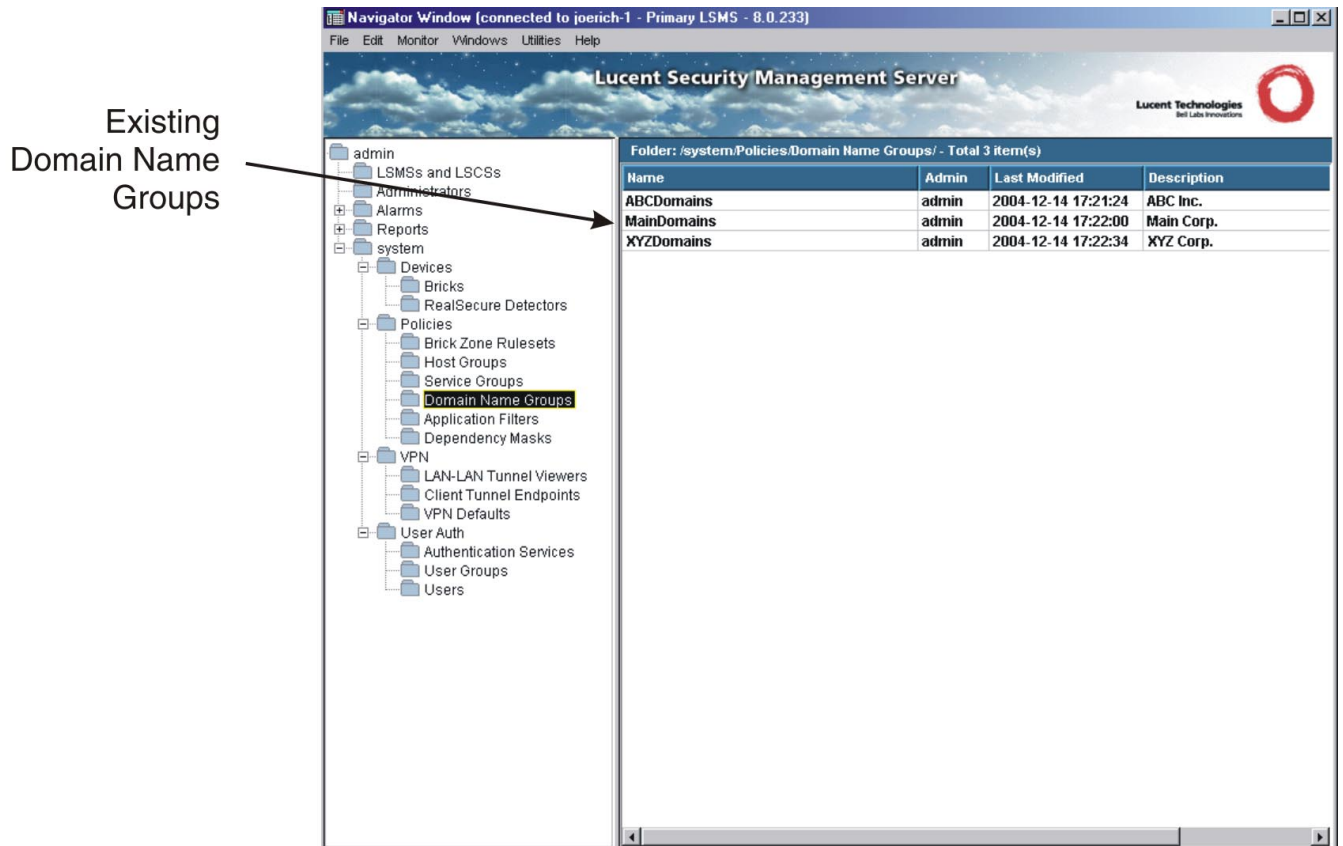
Once a domain name group has been created and saved, it can be viewed, modified, copied, or moved. If the domain name group is no longer needed, it can be deleted.

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. 3-13).

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



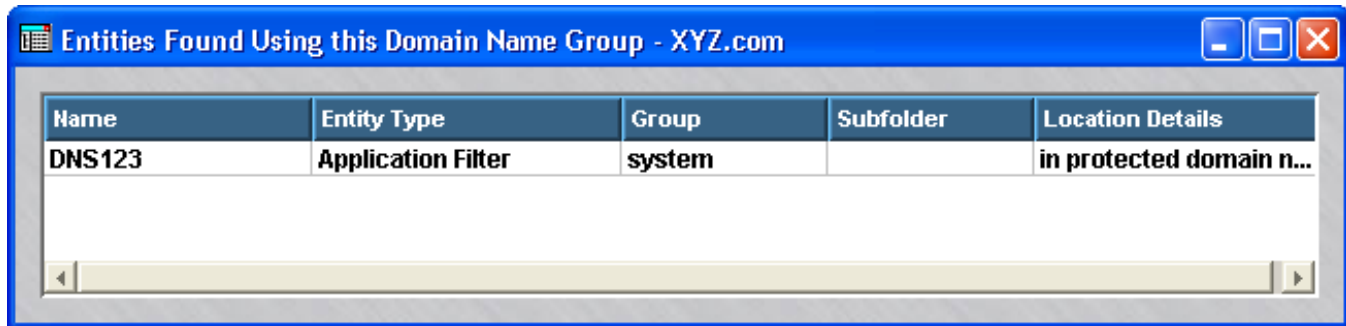
As you can see, for each domain name group, the Navigator window shows the Administrator who created the domain name group, the date and time the group was created, and a brief description, if one was entered when the group was created.

END OF STEPS

Find Entities Using a Domain Name Group

- 1 It is possible to identify all the entities using a domain name group. With the domain name groups displayed in the Navigator window, right-click the domain name group you want and select **Find Entities Using this Domain Group** from the pop-up menu. A window similar to the one shown in [Figure 3-7, “Entities Found Window”](#) (p. 3-14) will appear and list all the entities using this domain name group.

Figure 3-7 Entities Found Window



As [Figure 3-7, “Entities Found Window”](#) (p. 3-14) shows, the window gives the entity type (in the example, an application filter), the group and subfolder in which the entity is found, and details about the entity.

END OF STEPS

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. 3-6).

- 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

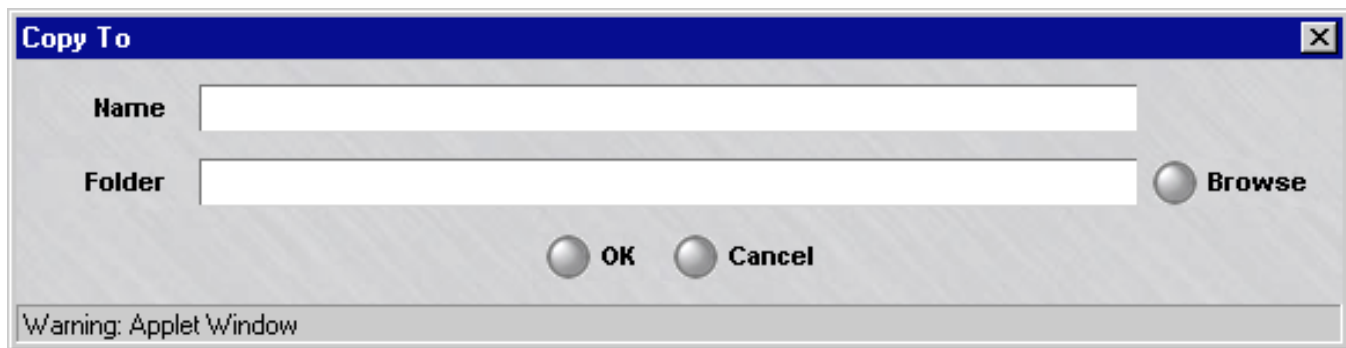
Copy a Domain Name Group

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. However, you can only copy it to the folder labeled “Domain Name Groups” or to a subfolder under that folder. If you try to copy it to another folder, you will get an error message. If a domain name group is copied to another group, and the domain name group has nested domain name groups within it, the nested groups are copied (if one with the same name and content does not exist) and possibly renamed (if one with the same name exists, but has different content).

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

- 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. A Copy To window will appear (see [Figure 3-8, "Copy To Window"](#) (p. 3-16)).

Figure 3-8 Copy To Window



- 2 In the **Name** field, enter the name you want to give the copy. If you are copying the domain name group to the same group, you must assign the copy a new name.
- 3 In the **Folder** field, click **Browse** and select the folder to which you want to copy this domain name group.
- 4 Click **OK** to copy the domain name group and dismiss the Copy To window. You will be returned to the Navigator window.

END OF STEPS

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.

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

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.

A confirmation window similar to the one shown in [Figure 3-9, “Confirmation Window \(Domain Name Groups\)”](#) (p. 3-18) will appear.

Figure 3-9 Confirmation Window (Domain Name Groups)



-
- 2 Click **Yes** to delete the domain name group and dismiss the Confirmation window. The domain name group will be 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 LSMS application.

Contents

What is a Service Group?	4-2
How to Set Up a Service Group	4-3
Global Service Groups	4-9
Nested Service Groups	4-11
How to Maintain a Service Group	4-14
Service Groups Provided with the LSMS Application	4-21



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, and you can make use of the 60 service groups that have been provided with the LSMS application. These 60 service groups cover most of the important services used by Internet applications today.

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 LSMS
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 Lucent VPN Firewall *Brick*® device 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.



How to Set Up a Service Group

Overview

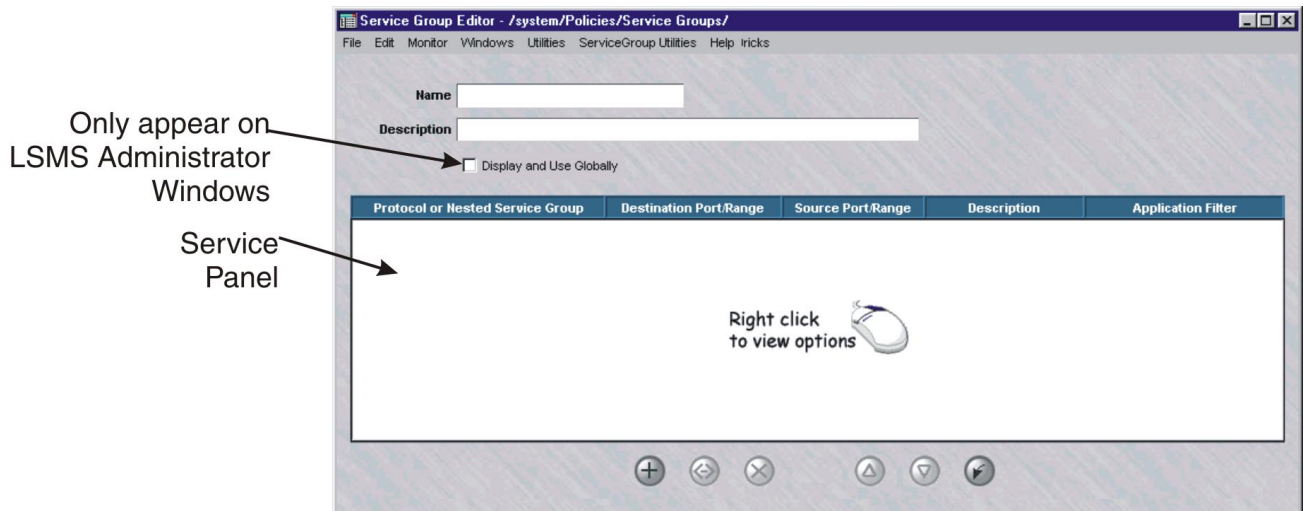
To set up a service group, you have to display the Service Groups Editor and enter the name of the service group and the services (protocol/destination port/source port) to be included in the group.

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. The Service Groups Editor will appear. It is shown in [Figure 4-1, “Service Group Editor”](#) (p. 4-3).

Figure 4-1 Service Group Editor



END OF STEPS

Create a Service Group

To create a service group, follow the steps below.

- 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 [“View Service Groups”](#) (p. 4-14).

Important! If you are an LSMS 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). The Service Editor, which is shown in [Figure 4-2, “Service Editor”](#) (p. 4-5), will appear.

Figure 4-2 Service Editor

Service Editor - /system/Policies/Service Groups/

Description

Specify Service Select Nested Service Group

Protocol * ▼

Destination Port or Range (TCP/UDP) * -
Type (ICMP) * -

Source Port or Range (TCP/UDP) * -
Code (ICMP) * -

Session Timeout (seconds)

Application Filter ▼

OK Cancel

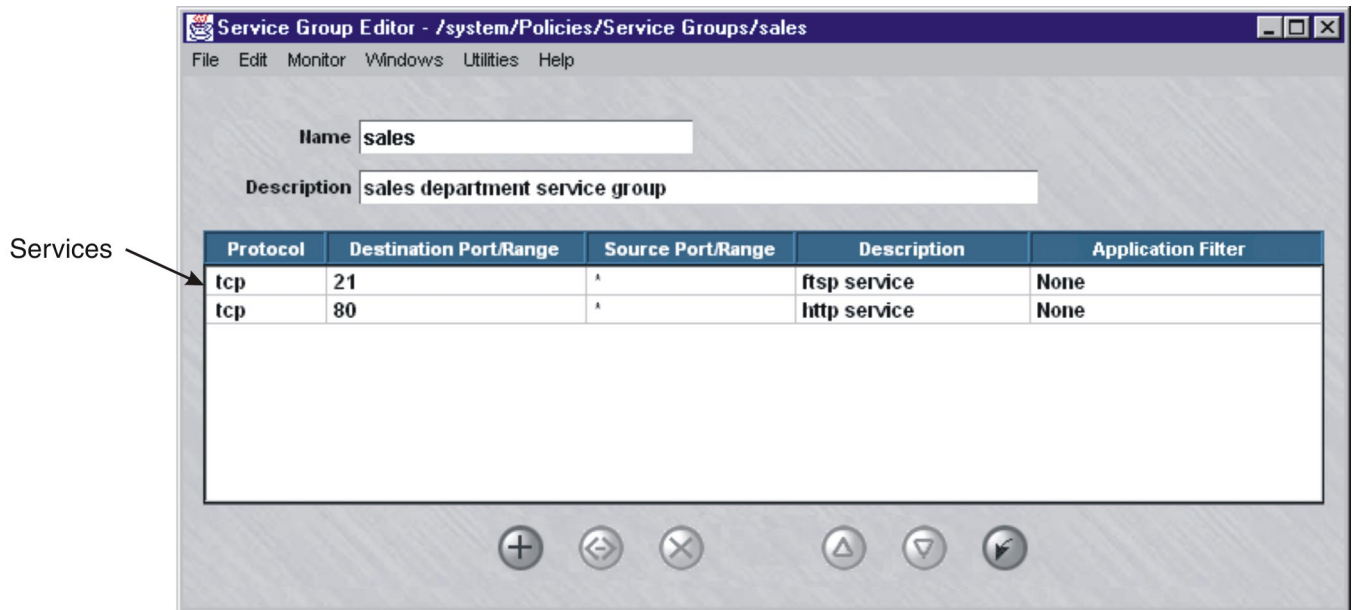
- 4 The purpose of the Service Editor is to allow you to add services to the service group.

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. 4-3)).
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 Host Groups”](#) (p. 2-12).
2. In the **Protocol** field, select a protocol from the drop-down list. This field is required. The choices are:
 - Asterisk (any protocol)
 - TCP
 - UDP
 - ICMP
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 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”](#).

7. 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.
8. If you need to add other services to this service group, repeat 1 - 7 for each additional service. Figure 4-3, “Service Group with Two Services” (p. 4-7) shows a service group called *sales* that contains two services (HTTP and FTP).

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 “Create a Service Group” (p. 4-4) above), select **Import Services** instead. A browse window will appear instead of the Service Editor (Figure 4-2, “Service Editor” (p. 4-5)) 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. 4-3)). 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.



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 LSMS Administrators can create global service groups.

Create a Global Service Group


If you are an LSMS 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. 4-3)). 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 [“Modify a Service Group”](#) (p. 4-16) 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.

View Global Service Groups

Global service groups appear in the Navigator window, just like standard service groups (see [“View Service Groups”](#) (p. 4-14) 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.

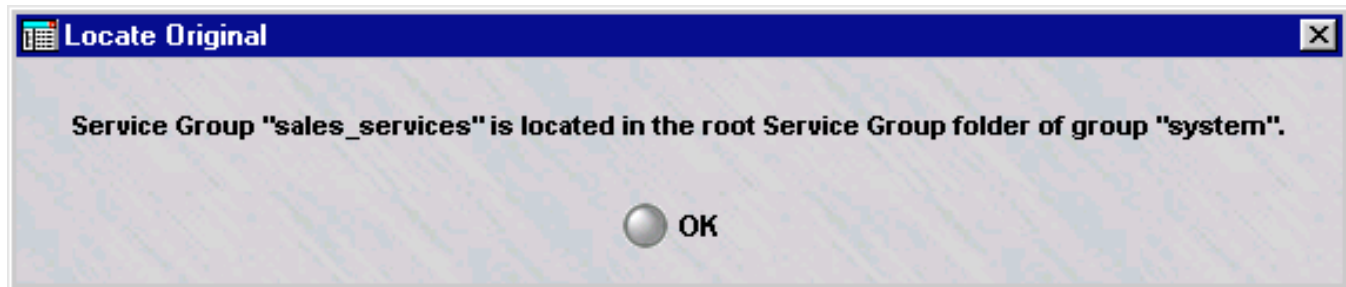
However, any global service groups created in *a different group* can be identified by a globe icon that appears to the left of the entry, as shown below.

 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)



Removing 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. 4-3)), 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 [“Delete a Service Group”](#) (p. 4-19) below).

Similarly, only the original source of a global service group can be moved (see [“Move a Service Group”](#) (p. 4-19) 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.

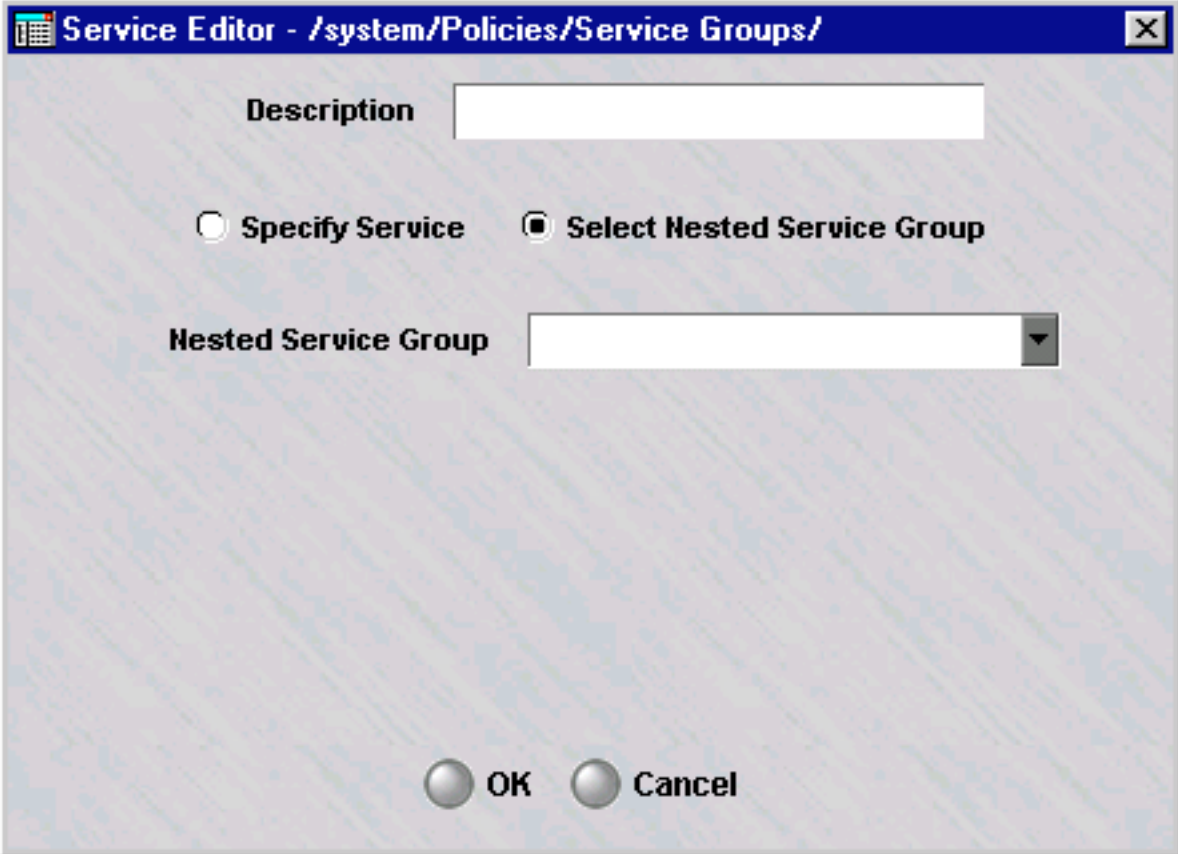
Task

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

-
- 1 Create a service group as previously explained (see [“Create a Service Group”](#) (p. 4-4) above). Follow the steps until the Service Editor appears ([Figure 4-2, “Service Editor”](#) (p. 4-5)).
 - 2 Click the **Select Nested Service Group** checkbox in the Service Editor (see [Figure 4-2, “Service Editor”](#) (p. 4-5)). 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)



The screenshot shows a dialog box titled "Service Editor - /system/Policies/Service Groups/". It contains a "Description" text field, two radio buttons labeled "Specify Service" and "Select Nested Service Group" (the latter is selected), and a "Nested Service Group" drop-down menu. At the bottom are "OK" and "Cancel" buttons.

-
- 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” \(p. 4-13\)](#) 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	Description	Application Filter
ftp			file transfer services	
http			web services	
smtp			email services	

END OF STEPS



How to Maintain a Service Group

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.

View 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:

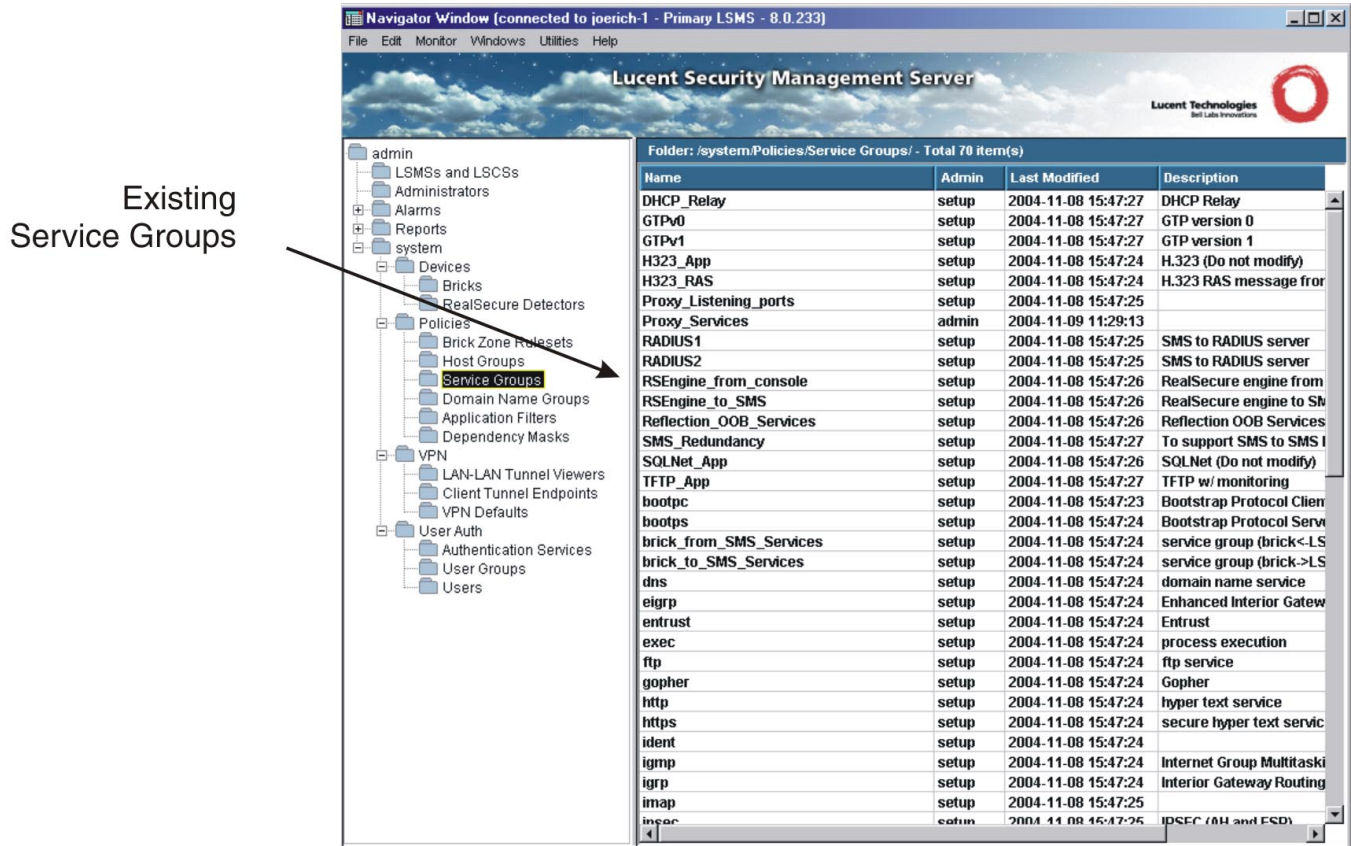
- 1 Open the appropriate group folder, and then open the Policy Components folder.
- 2 Click the Service Groups folder. All existing service groups, including those provided with the LSMS application, will be displayed in the Navigator window (see [Figure 4-7, “Navigator Window \(View Service Groups\)”](#) (p. 4-15)).

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 LSMS 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 “Modify a Service Group” (p. 4-16) for instructions.

Figure 4-7 Navigator Window (View Service Groups)



END OF STEPS

Find Entities Using a Global Service Group

- 1 It is possible to identify all the entities (such as Brick zone rulesets) using a global service group. With the service groups displayed in the Navigator window, right-click the service group and select **Find Entities Using this Service Group** from the pop-up menu. A window similar to the one shown in [Figure 4-8, “Entities Found Window”](#) (p. 4-16) will appear and list all the entities using this service group.

Figure 4-8 Entities Found Window



As [Figure 4-8, “Entities Found Window”](#) (p. 4-16) shows, the window gives the entity type (in the example, a Brick zone ruleset), the group and subfolder in which the entity is found, and details about the entity (if the entity is a Brick zone ruleset, it gives the specific rules in which the global service group is used).

END OF STEPS

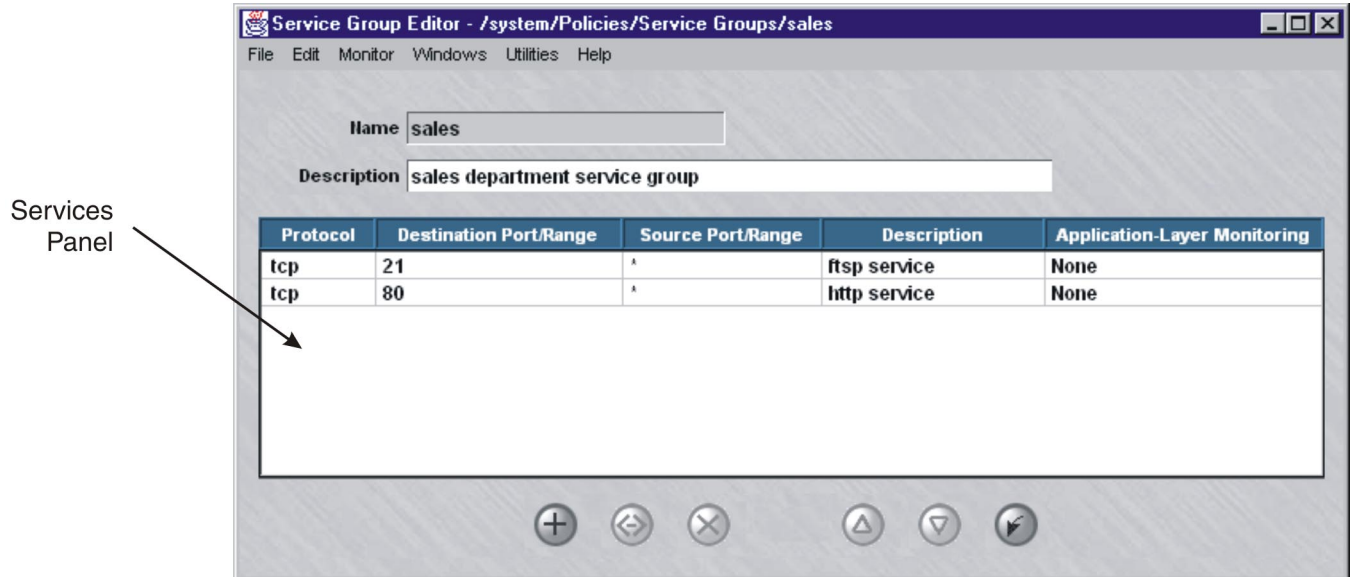
Modify a Service Group

You can modify any field in a service group except the **Name** field. To modify a service group, follow the steps below:

- 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 4-9, “Service Group Editor \(Edit Mode\)”](#) (p. 4-17).

Figure 4-9 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. 4-5)) 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. 4-5)) 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

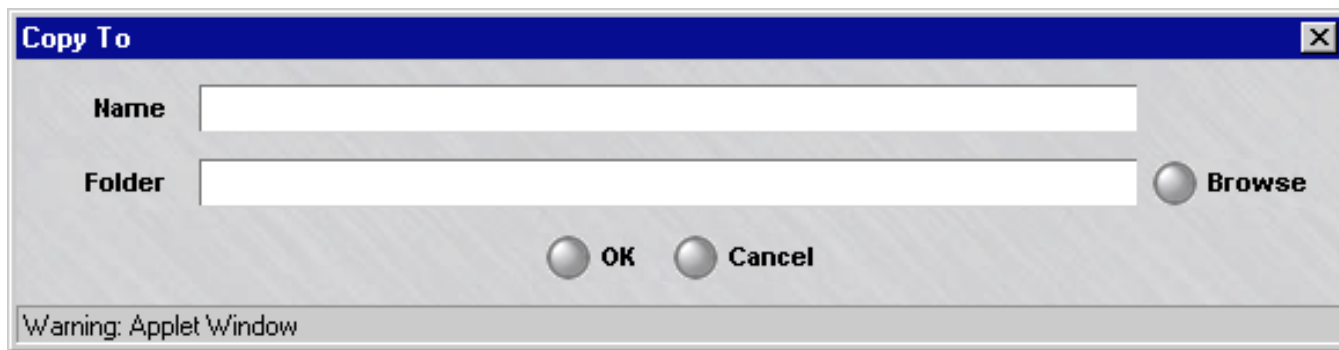
Copy a Service Group

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 folder. However, you can only copy it to the folder labeled "Service Groups" or to a subfolder under that folder. If you try to copy it to another folder, you will get an error message. If a service group is copied to another group, and the service group has nested service groups within it, the nested service groups are copied (if one with the same name and content does not exist) and possibly renamed (if one with the same name exists, but has different content).

To copy 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 copy and select **Copy** from the pop-up menu. A Copy To window will appear (see [Figure 4-10, "Copy To Window"](#) (p. 4-18)).

Figure 4-10 Copy To Window



-
- 2 In the **Name** field, enter the name you want to give the copy. If you are copying the service group to the same group, you must assign the copy a new name.
 - 3 In the **Folder** field, click **Browse** and select the Service Groups folder you want to copy this service group to.

-
- 4 Click **OK** to copy the service group and dismiss the Copy To window. You will be returned to the Navigator window.

END OF STEPS

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 ["Find Entities Using a Global Service Group"](#) (p. 4-16) 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

Delete a Service Group

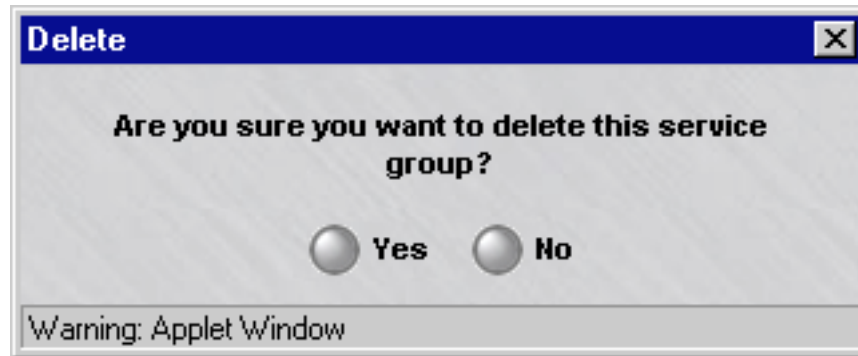
Before deleting a service group, make sure that it is not currently in use. If you attempt to delete a service group that is in use, you will get an error message. If you receive an error message when attempting to delete a service group, you can run the "find entity" utility to discover where it is in use (see ["Find Entities Using a Global Service Group"](#) (p. 4-16) 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 4-11, “Confirmation Window \(Service Groups\)”](#) (p. 4-20) will appear.

Figure 4-11 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 LSMS Application

Overview

More than 60 service groups are automatically created when the LSMS 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 LSMS	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<-LSMS) for managing Brick devices	TCP	910	*
		ICMP	8	0
		UDP	911	9020
		UDP	1024	*
		UDP	9014	*
Brick_to_SMS_Services	service group (brick->LSMS) 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 <-LSMS) for managing router	UDP	9014	*
		UDP	69	
router_from_SMS_tcp	service group (router <-LSMS) for managing router	TCP	9017-9018	*

Service Group Name	Description	Protocol	Dest Port	Source Port
router_to_SMS_other	service group (udp & icmp) (router ->LSMS) for managing router	ICMP	*	*
		UDP	69	*
		UDP	9014	*
router_to_SMS_tcp	Service group (tcp) (router -> LSMS) 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 LSMS	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	*



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 Lucent VPN Firewall *Brick*[®] device.

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.
- 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.
- SIP - Session Initiation Protocol 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.

In addition, the application layer monitoring reasons previously available are now present as application filters. The affected protocols are H.323_RAS, Oracle SQL*Net, FTP and TFTP. 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.

Future releases will both expand the capability of the HTTP application filter and add new application filters for other protocols.

Contents

Process Flow	5-3
DHCP Relay Application Filter	5-4
DNS Application Filters	5-7
ESP Application Filter	5-16
FTP Application Filter	5-19
GTP Application Filter	5-29
H.323 Application Filters	5-43
HTTP Application Filter	5-50
RPC Application Filters	5-58
SIP Application Filters	5-64
SMTP Application Filters	5-81
Global Application Filters	5-92



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.

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.

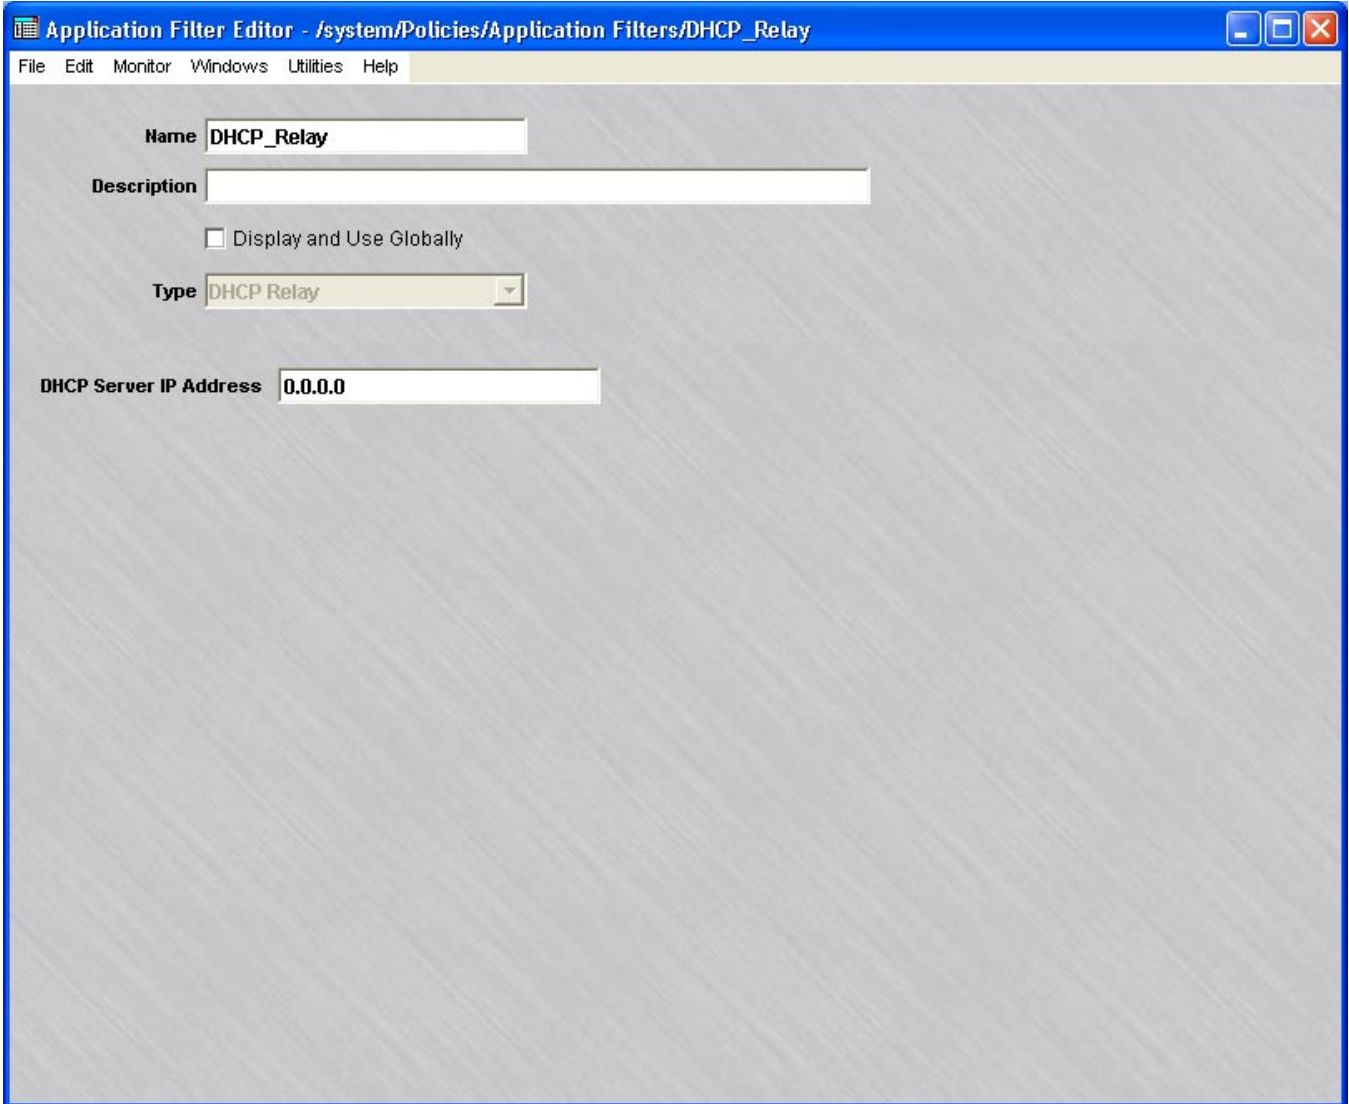
Task

To configure this application filter, follow the steps below:

- 1 From the LSMS 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. 5-5) is displayed.

Figure 5-1 DHCP Application Filter Editor



The screenshot shows a window titled "Application Filter Editor - /system/Policies/Application Filters/DHCP_Relay". The window has a menu bar with "File", "Edit", "Monitor", "Windows", "Utilities", and "Help". The main area contains the following fields:

- Name:** DHCP_Relay
- Description:** (empty text box)
- Display and Use Globally
- Type:** DHCP Relay (dropdown menu)
- DHCP Server IP Address:** 0.0.0.0

Important! If you are an LSMS 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. 5-58) 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 firewall will ensure that an outside DNS server will not be able to tell 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

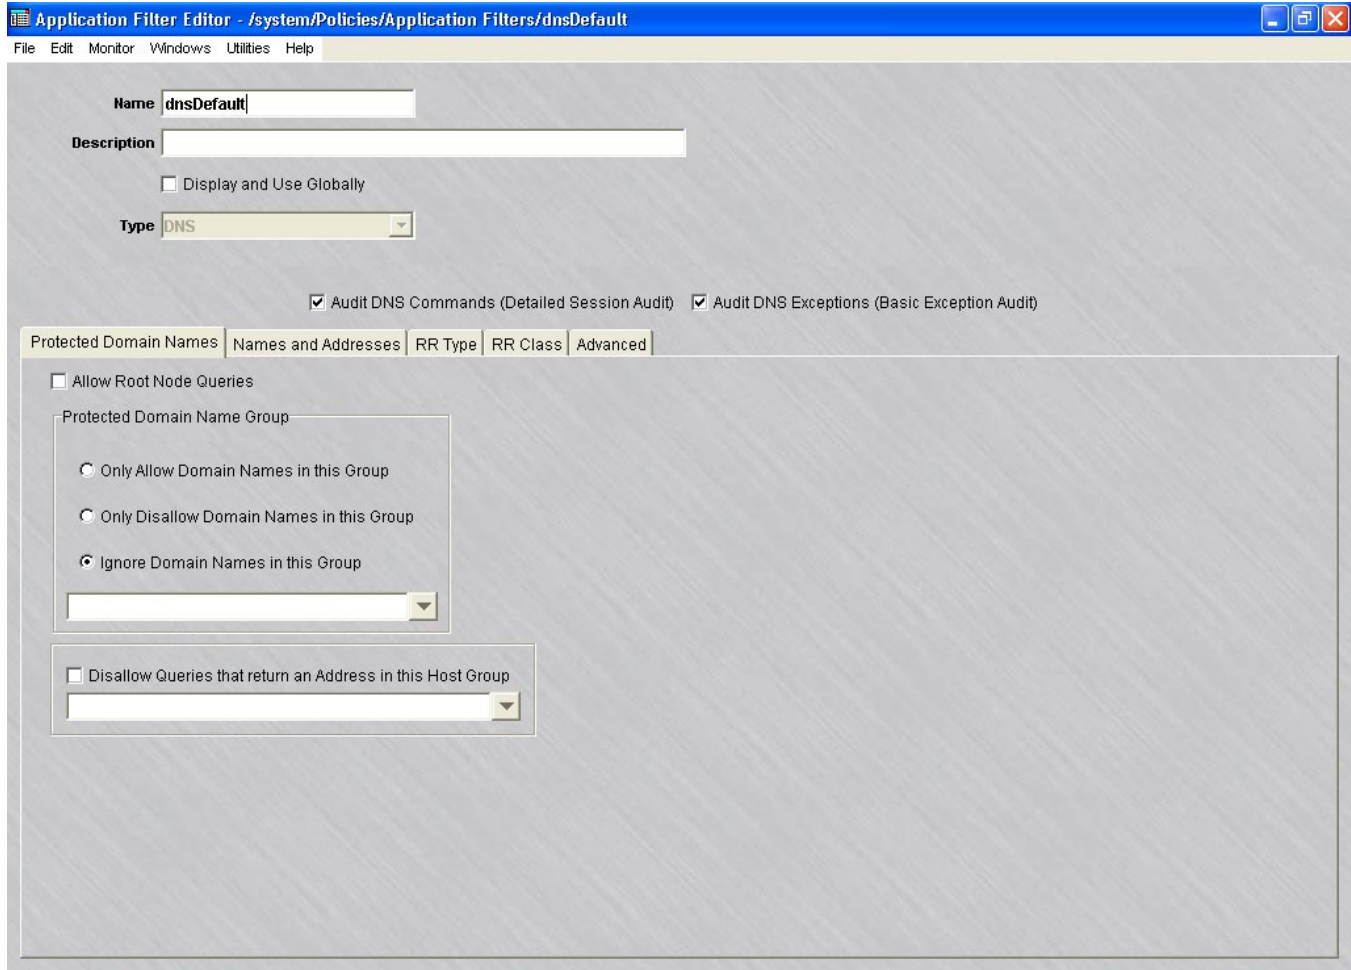
A DNS application filter called `dnsDefault` is created automatically in the system group for a newly installed LSMS, or for each pre-existing group for an upgraded LSMS. For a newly installed LSMS and for newly created groups, `dnsDefault` is additionally assigned to the UDP entry of the `dns` service 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 LSMS Navigator, open the Policies folder in the desired group, and click **Application Filters**.

-
- 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. 5-8) will appear.

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



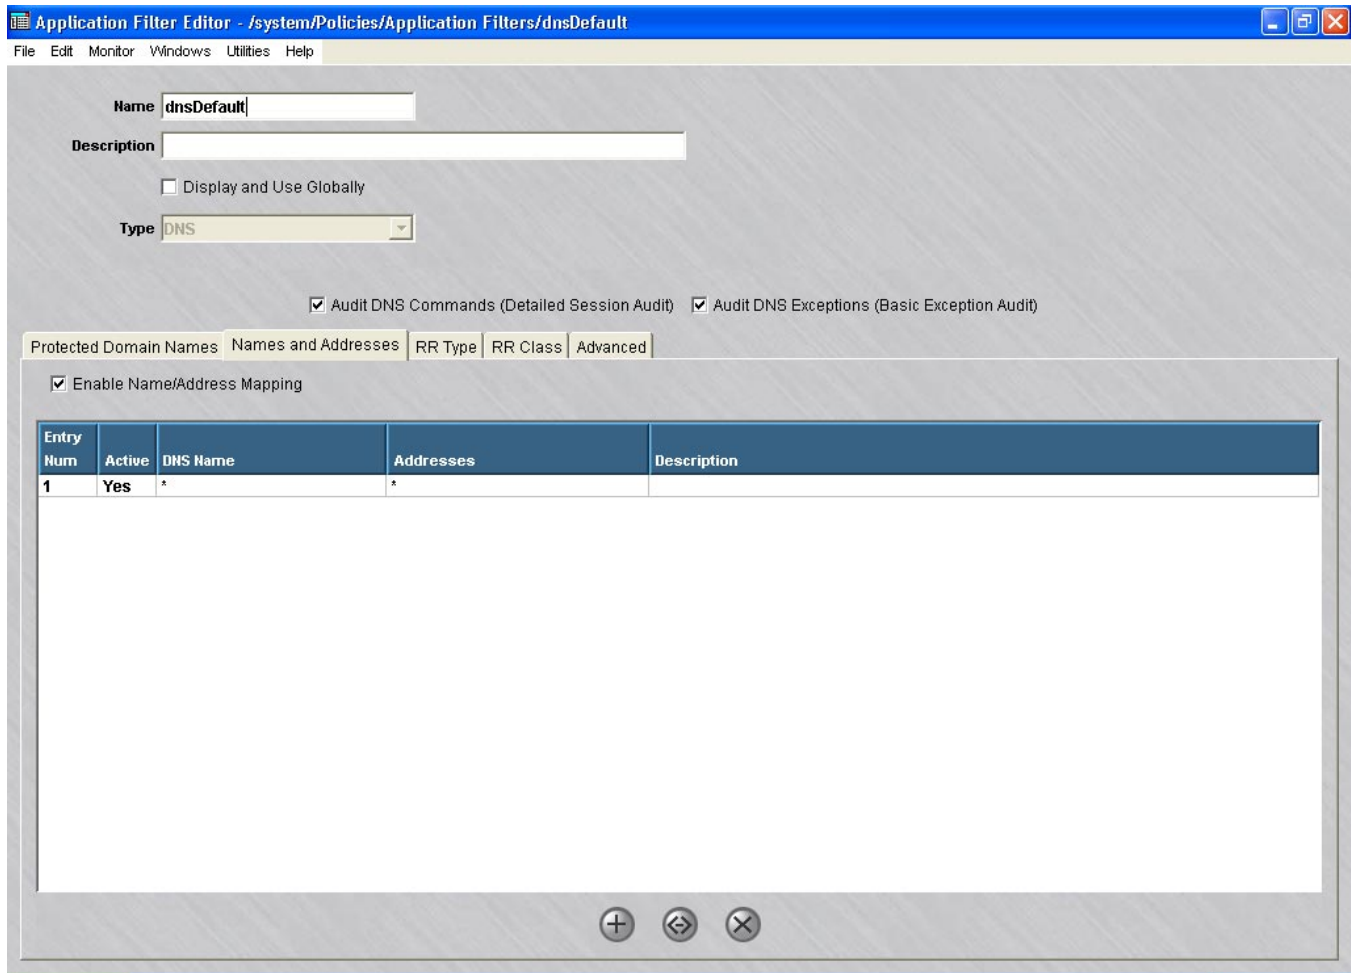
-
- In the **Name** and **Description** fields, enter a name (default is dnsDefault) for the application filter and a brief description. The name is required, but the description is optional.

Important! If you are an LSMS 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 “Global Application Filters” below 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. 5-10)) 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

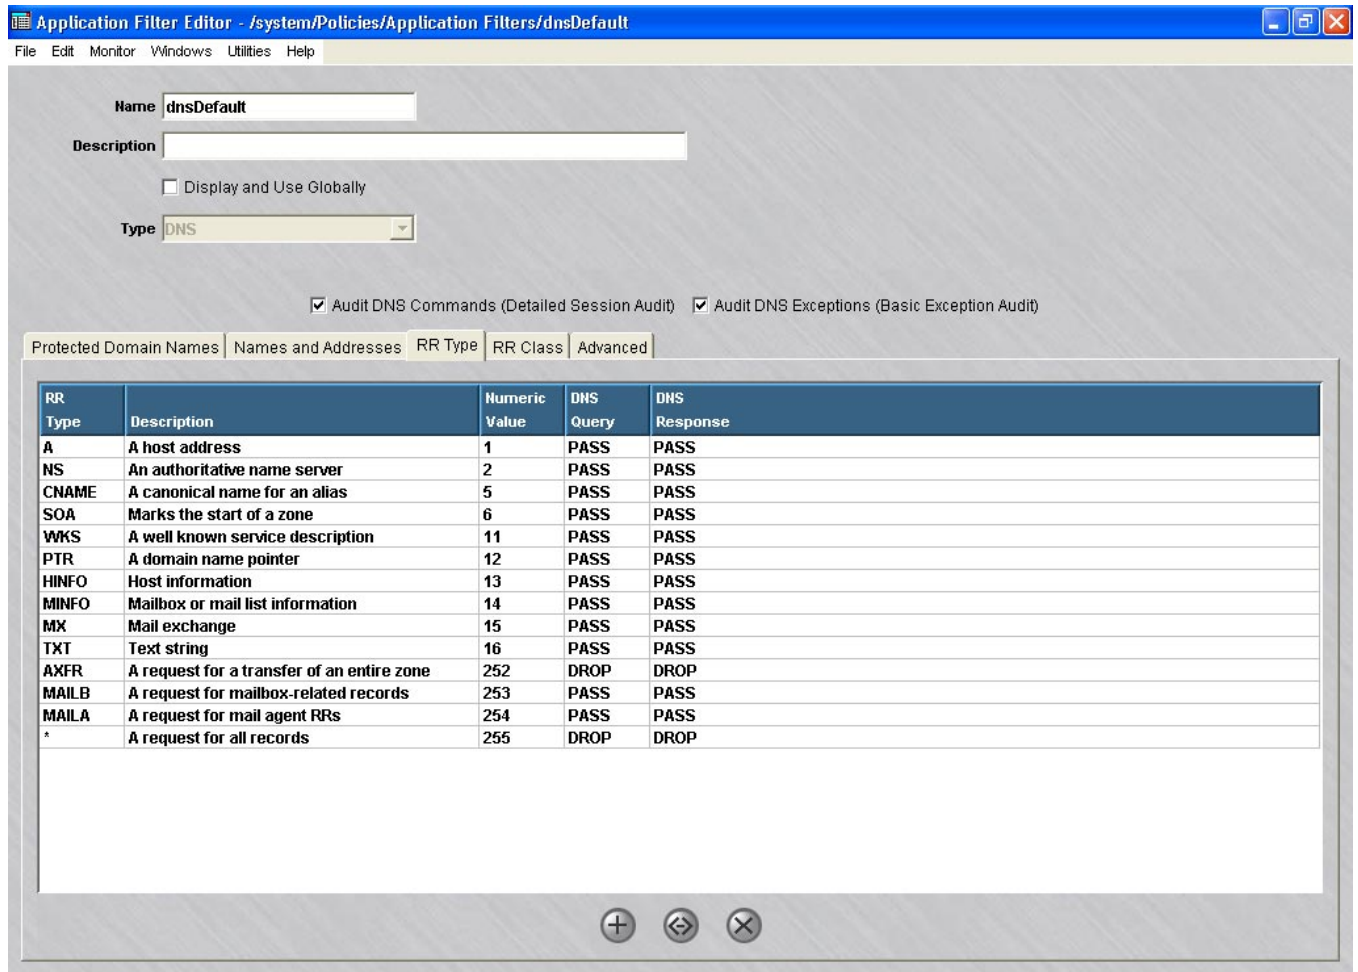
names such as *.lucent.com. The address can be ranges of IP's or a specific IP. Host groups can be used in the usual manner.

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. 5-11):

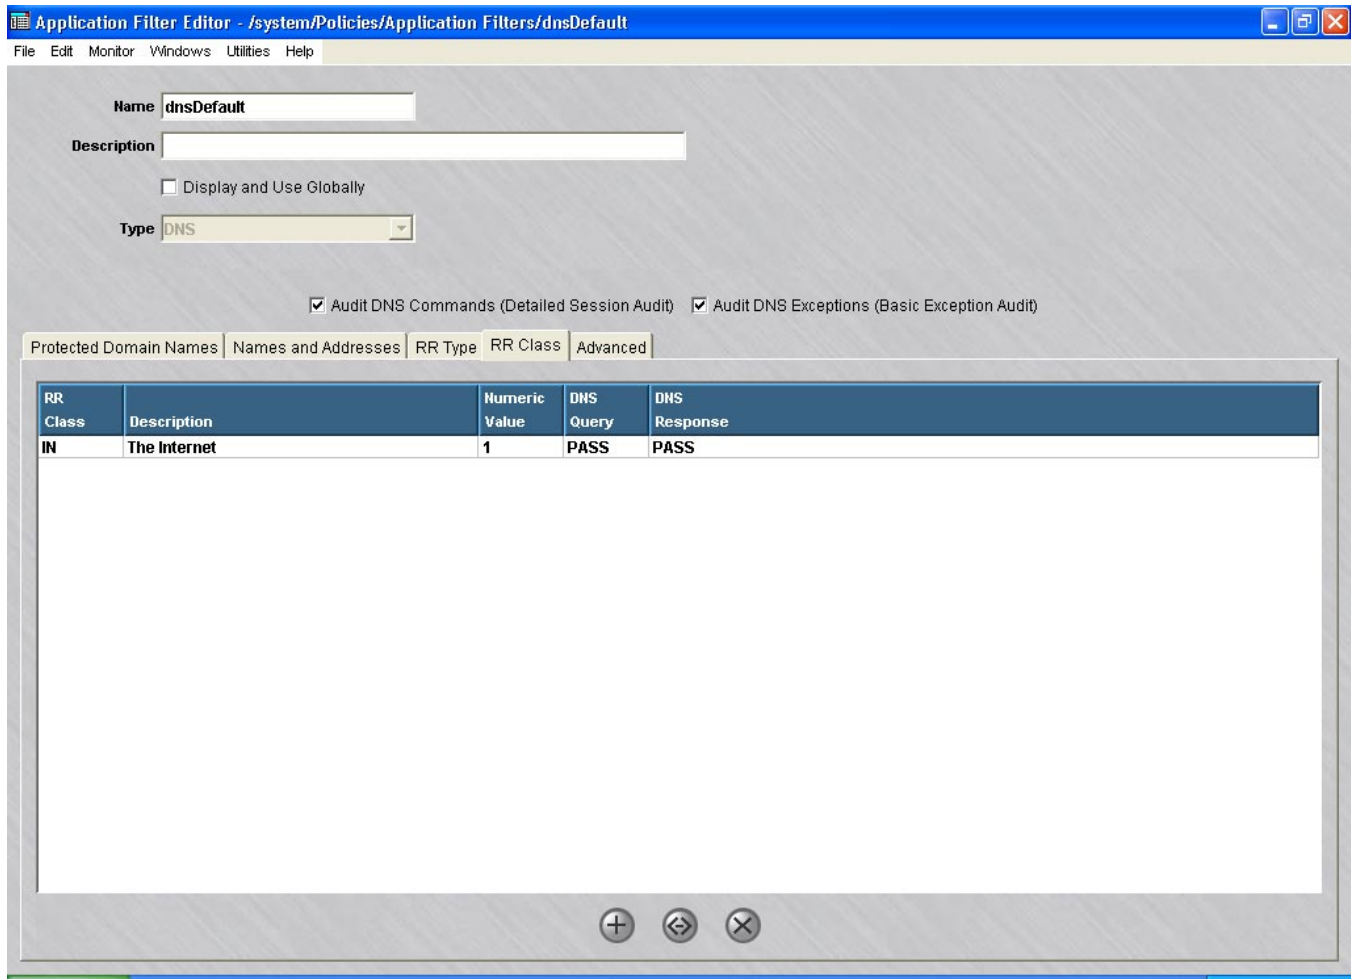
Figure 5-4 DNS Application Filter Editor Screen (RR Type Tab)



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:

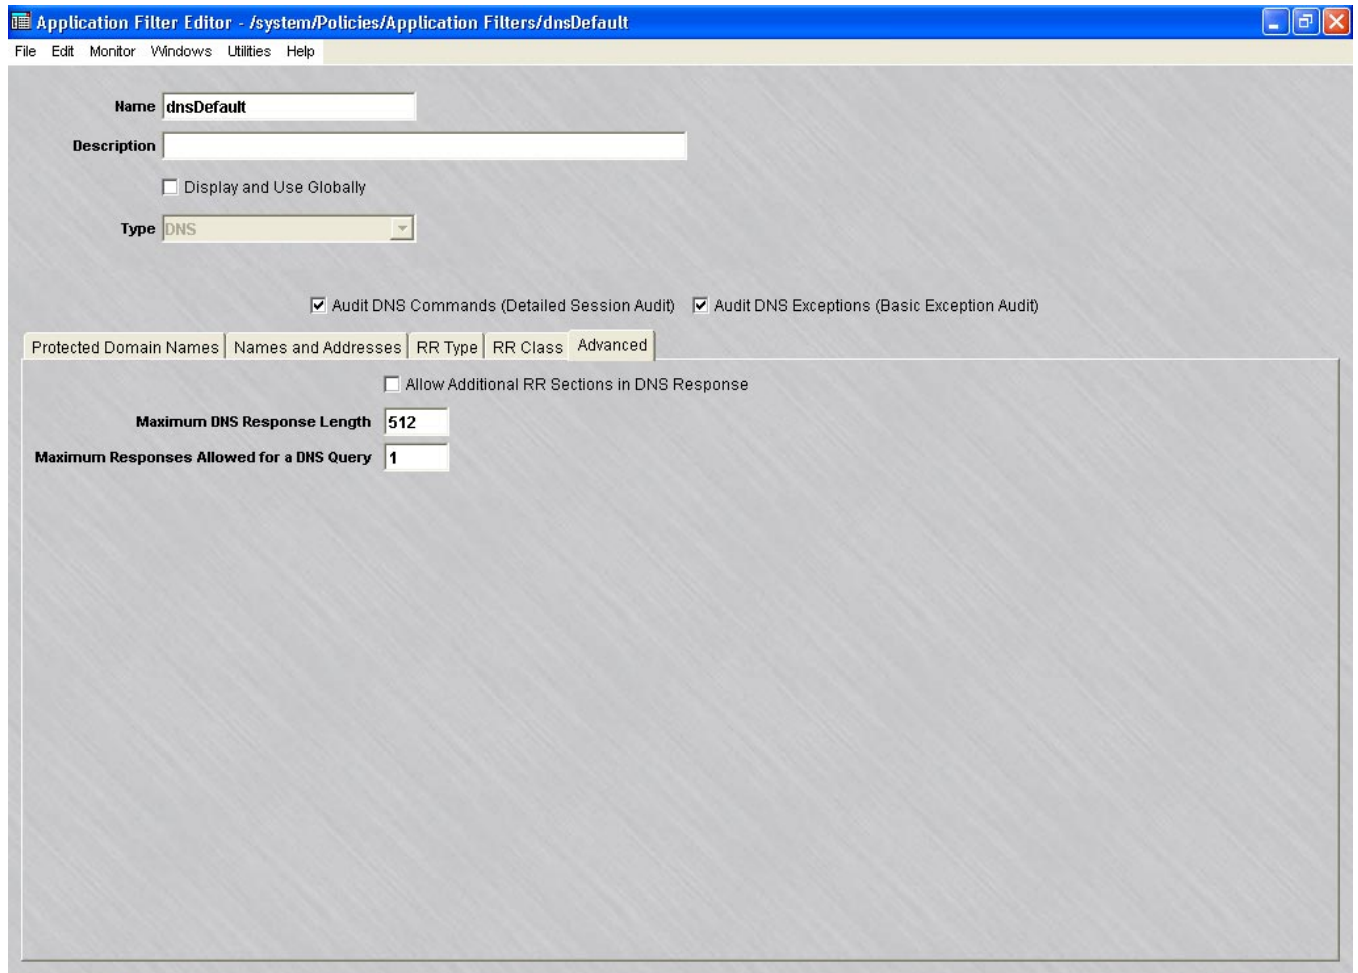
Figure 5-5 DNS Application Filter Editor (RR Class Tab)



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 **Advance** 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.
- 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

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, “Lucent VPN Firewall Brick® Device Zone Rulesets”](#). For additional information on service groups, see [Chapter 4, “Service Groups”](#).

END OF STEPS



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 LSMS Navigator, open the Policies folder in the desired group and click **Application Filters**.
-

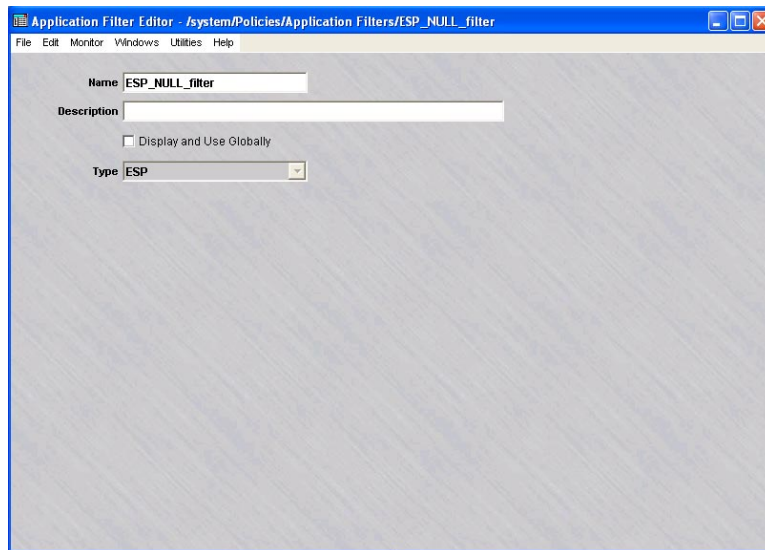
- 2 In the right hand column, right-click the mouse and select **New Application Filter** from the pop-up menu.

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” \(p. 5-16\)](#)).

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 LSMS 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” \(p. 5-92\)](#) 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

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, “Lucent VPN Firewall Brick® Device Zone Rulesets”](#). For additional information on service groups, see [Chapter 1, “Lucent VPN Firewall Brick® Device 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 LSMS 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 LSMS and, once installed, is created automatically for newly created groups. Similarly, an *ftpDefault* filter is created when an LSMS 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.

How to Configure an FTP Application Filter

To configure an FTP application filter, follow the steps below:

- 1 From the LSMS 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 \(Commands Filter Tab \)” \(p. 5-21\)](#) 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 'Name' and 'Description' fields are empty. The 'Type' dropdown is set to 'FTP'. The 'Display and Use Globally' checkbox is unchecked. Both 'Audit FTP Commands (Detailed Session Audit)' and 'Audit FTP Exceptions (Basic Exception Audit)' checkboxes are checked. The 'Commands Filter' tab is active, showing a 'Pass List' with items: mkd, port, user, quit, stat, stor, type, pasv. The 'Block List' has one item: site. Below each list is an 'Add:' field and a close button.

- 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 LSMS 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. 5-92) 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**

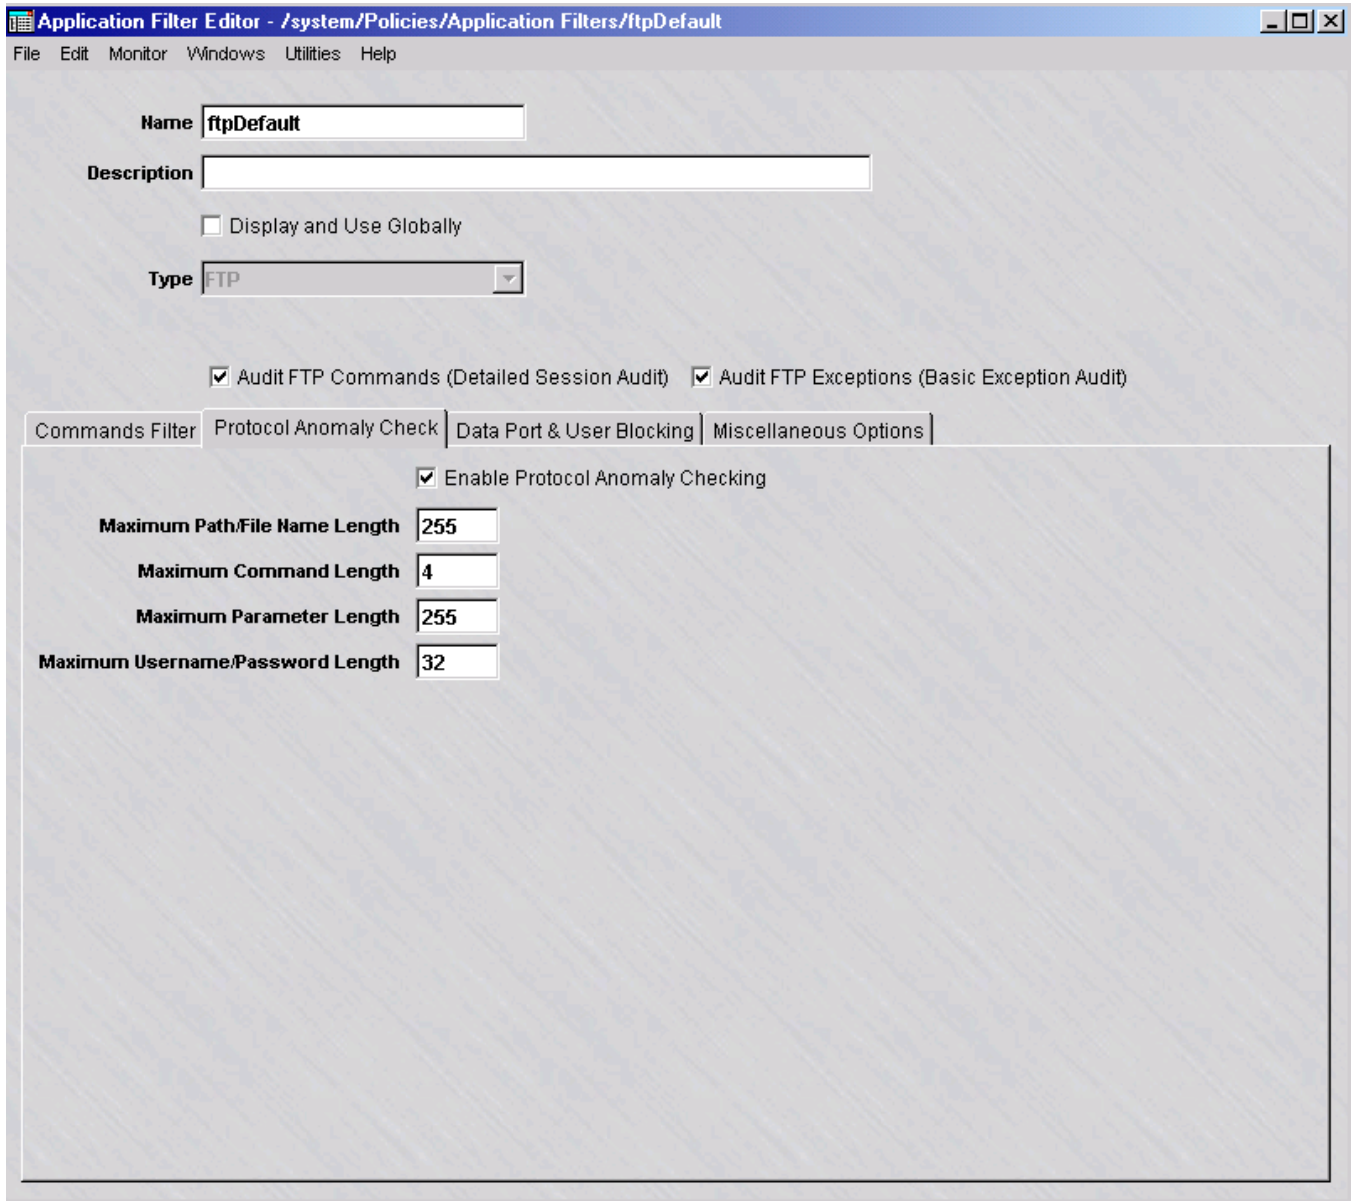
Command filtering is accomplished in one of two separate modes: pass or block. If you select the **Pass List** radio button, only those commands in the pass list are allowed. If you select the **Block List** radio button, commands in the block list are prohibited.

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. 5-23) shows a sample screen):

Figure 5-9 FTP Application Editor (Protocol Anomaly Check Tab)

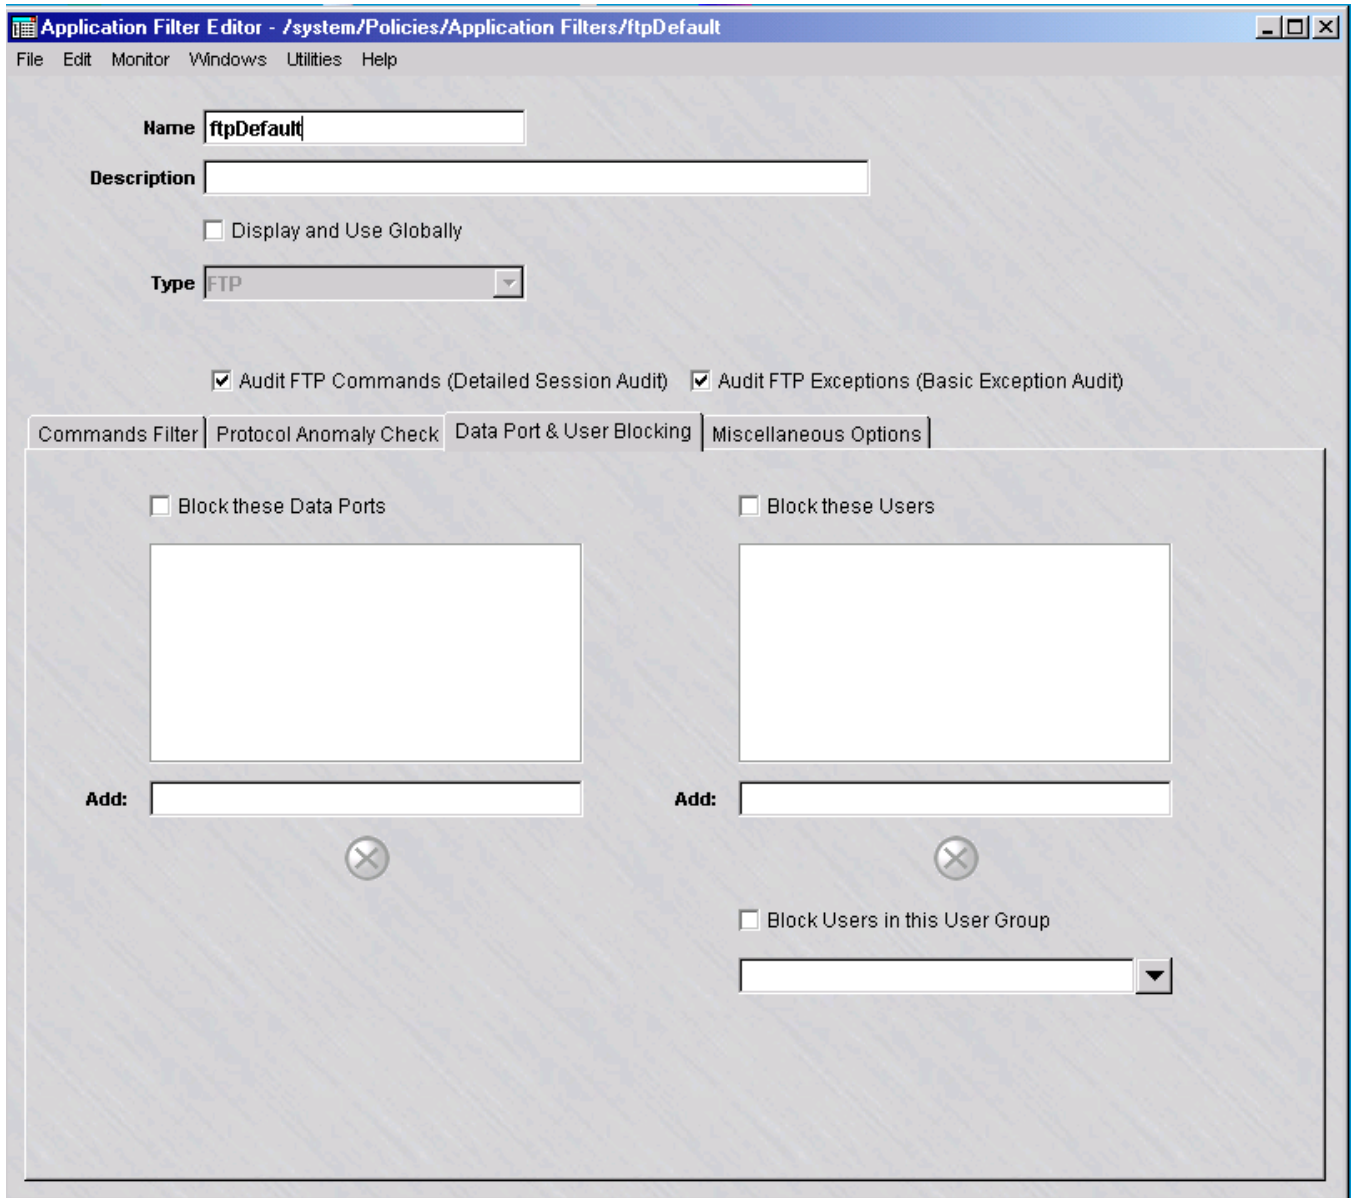


The following explains these fields:



- **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. 5-25) shows a sample screen):

Figure 5-10 FTP Application Filter Editor (Data Port & User Blocking Tab)

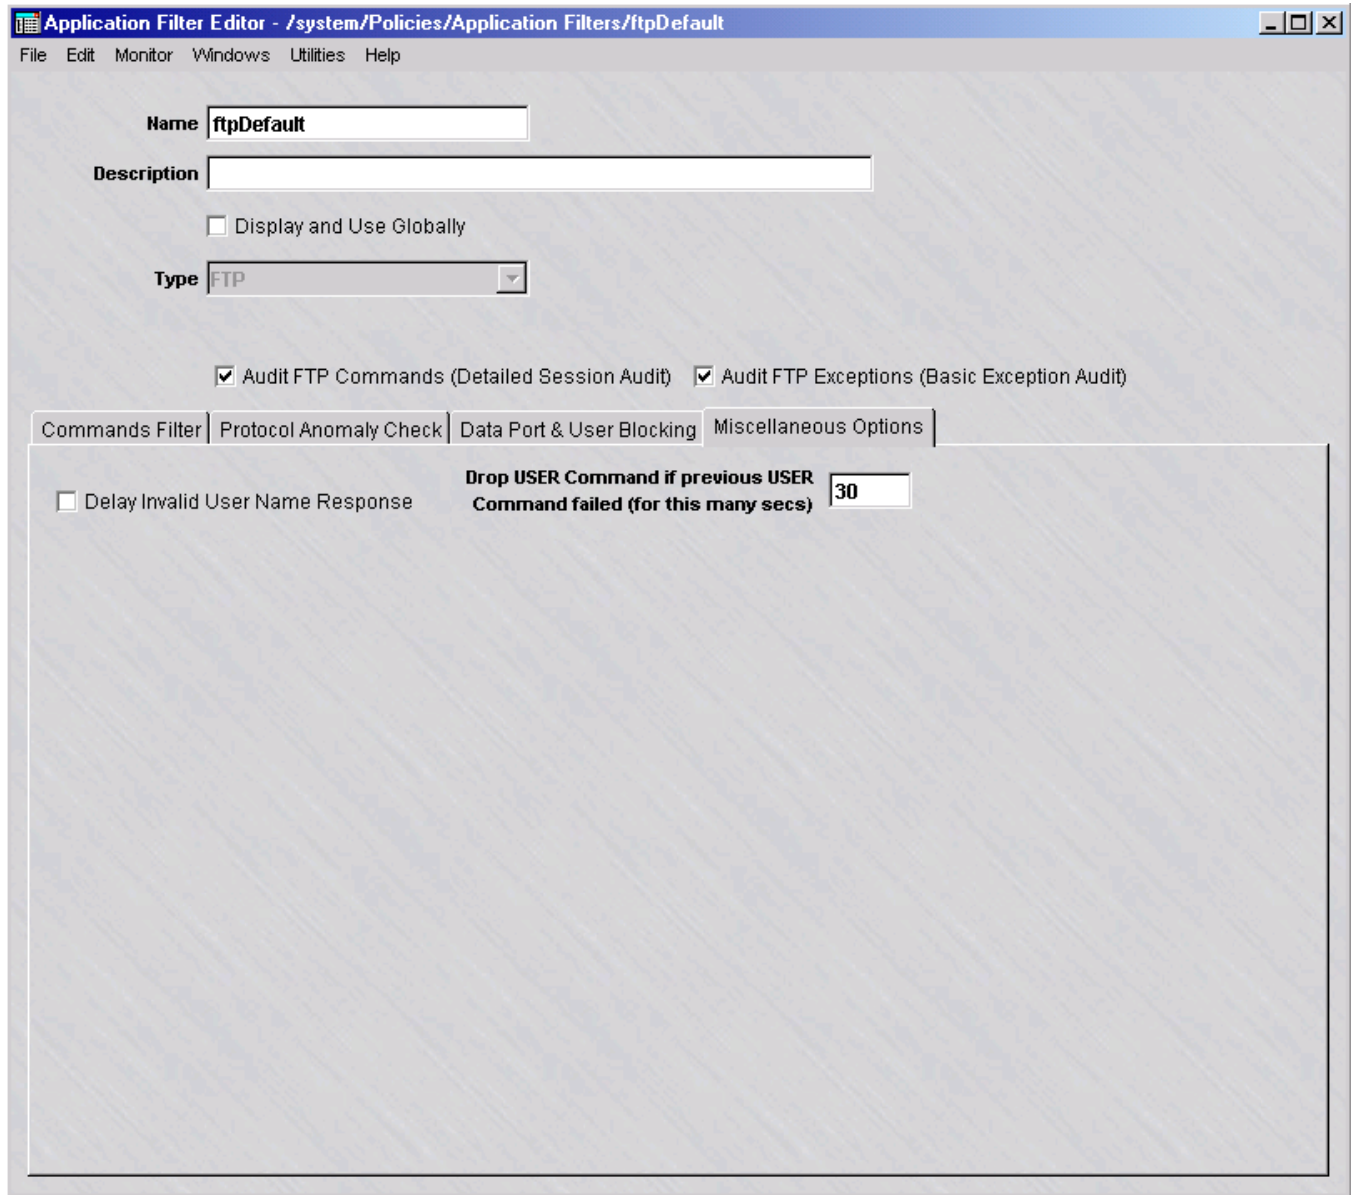


The following explains these fields:

- **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. 5-27) 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.

-
- 9 From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.

.....
 END OF STEPS

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, “Lucent VPN Firewall *Brick*® Device Zone Rulesets”](#). For additional information on service groups, see [Chapter 1, “Lucent VPN Firewall *Brick*® Device 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 LSMS 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

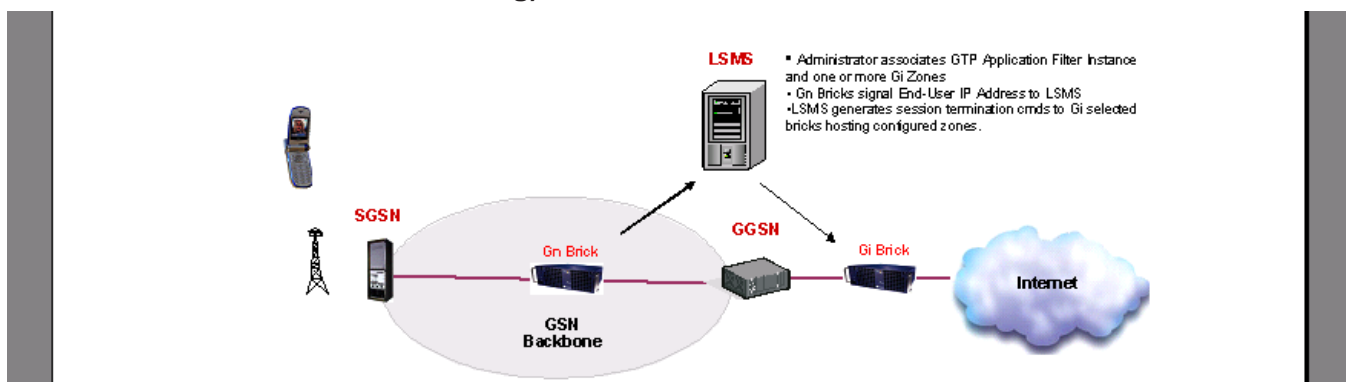
The increased security mechanisms provided by stateful GTP application filtering is further enhanced by a built-in GTP PDP Context Deletion Monitor that coordinates the termination of firewall sessions on the Gi interface of the Brick with the termination of the associated GTP PDP Contexts on the GPRS/UMTS 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 LSMS 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 LSMS 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 LSMS indicating the deletion of the last PDP context and release of the associated IP address. The LSMS 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. 5-30).

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 “[Configure a GTP Application Filter](#)” (p. 5-31).

Configure a GTP Application Filter

To configure a GTP application filter, follow the steps below:

- 1 From the LSMS 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**. The GTP application filter window is displayed (see [Figure 5-13](#), “GTP Application Filter Editor” (p. 5-32)).

Figure 5-13 GTP Application Filter Editor

Message Name	Numerical Value	Direction
Create AA PDP Context Request	22	Both
Create AA PDP Context Response	23	Both
Create PDP Context Request	16	Both
Create PDP Context Response	17	Both
Delete AA PDP Context Request	24	Both
Delete AA PDP Context Response	25	Both
Delete PDP Context Request	20	Both
Delete PDP Context Response	21	Both
Echo Request	1	Both
Echo Response	2	Both
Error Indication	26	Both
Identification Request	48	Both
Identification Response	49	Both

- 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 LSMS 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. 5-92) 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 LSMS. Additional details about configuring Stateful GTP Filtering are provided in [Step 11](#) 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. 5-34) shows a sample screen):

Figure 5-14 GTP Application Filter Editor (GP Interface Message List Tab)

Message Name	Numerical Value	Direction
Data Record Transfer Request	240	None
Data Record Transfer Request	241	None
Failure Report Request	34	None
Failure Report Response	35	None
Node Alive Request	4	None
Node Alive Response	5	None
Note MS GPRS Present Request	36	None
Note MS GPRS Present Response	37	None
Redirection Request	6	None
Redirection Response	7	None
Send Routing Information for GP...	32	None
Send Routing Information for GP...	33	None

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.

- 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. 5-35) shows a sample screen):

Figure 5-15 GTP Application Filter Editor (Non-Gp Interface Message Tab)

Message Name	Numerical Value	Direction
Data Record Transfer Request	240	None
Data Record Transfer Request	241	None
Failure Report Request	34	None
Failure Report Response	35	None
Node Alive Request	4	None
Node Alive Response	5	None
Note MS GPRS Present Request	36	None
Note MS GPRS Present Response	37	None
Redirection Request	6	None
Redirection Response	7	None
Send Routing Information for GP...	32	None
Send Routing Information for GP...	33	None

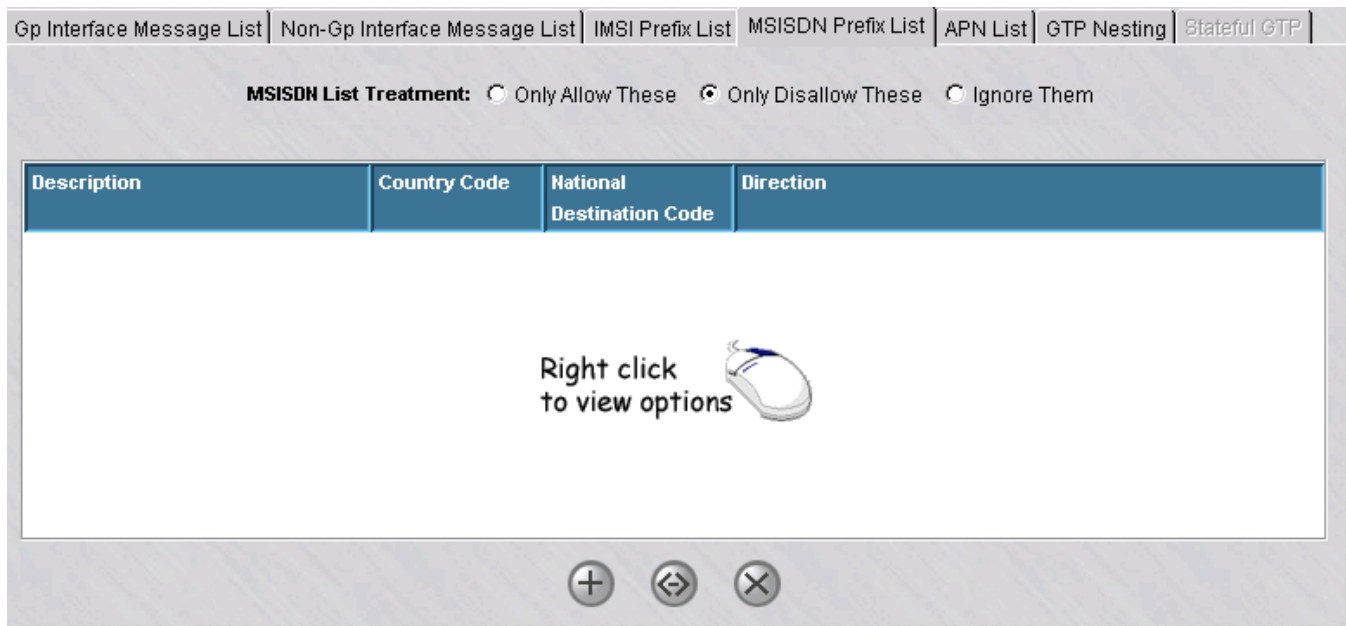
The following explains these fields:

- **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. 5-36) shows a sample screen):

Figure 5-16 GTP Application Filter Editor (IMSI) Prefix list Tab



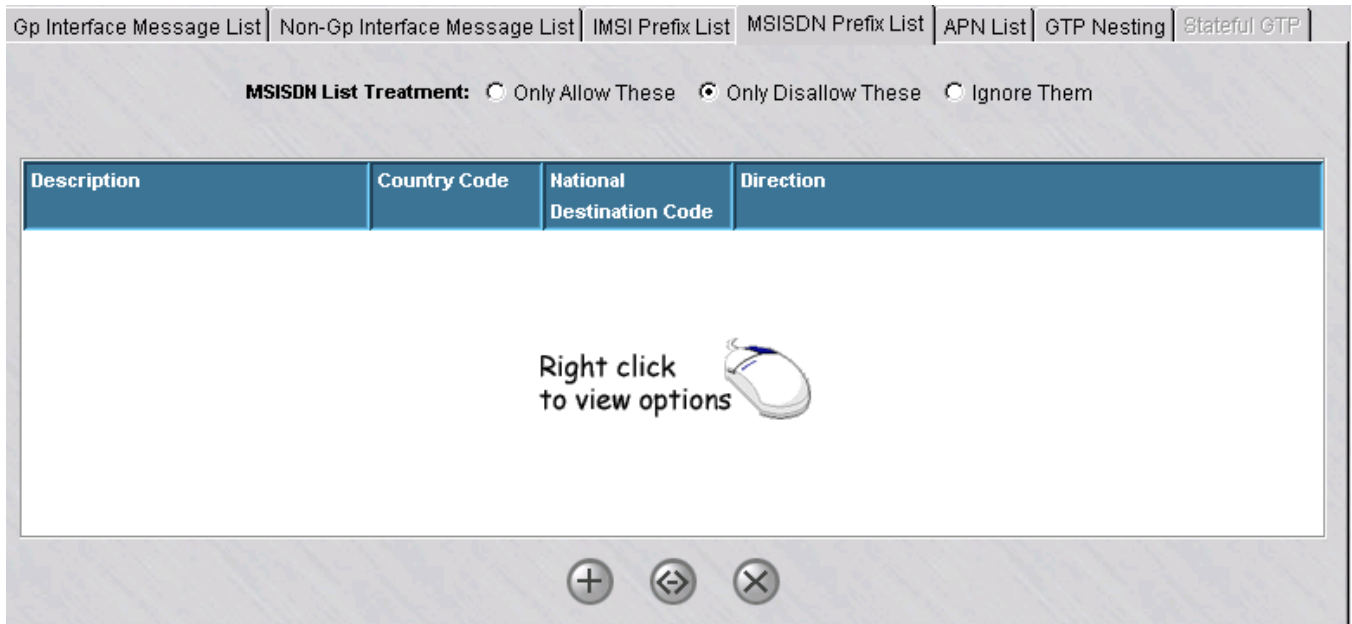
The following explains these fields:

- **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. 5-37) shows a sample screen):

Figure 5-17 GTP Application Filter (MSISDN Prefix List Tab)



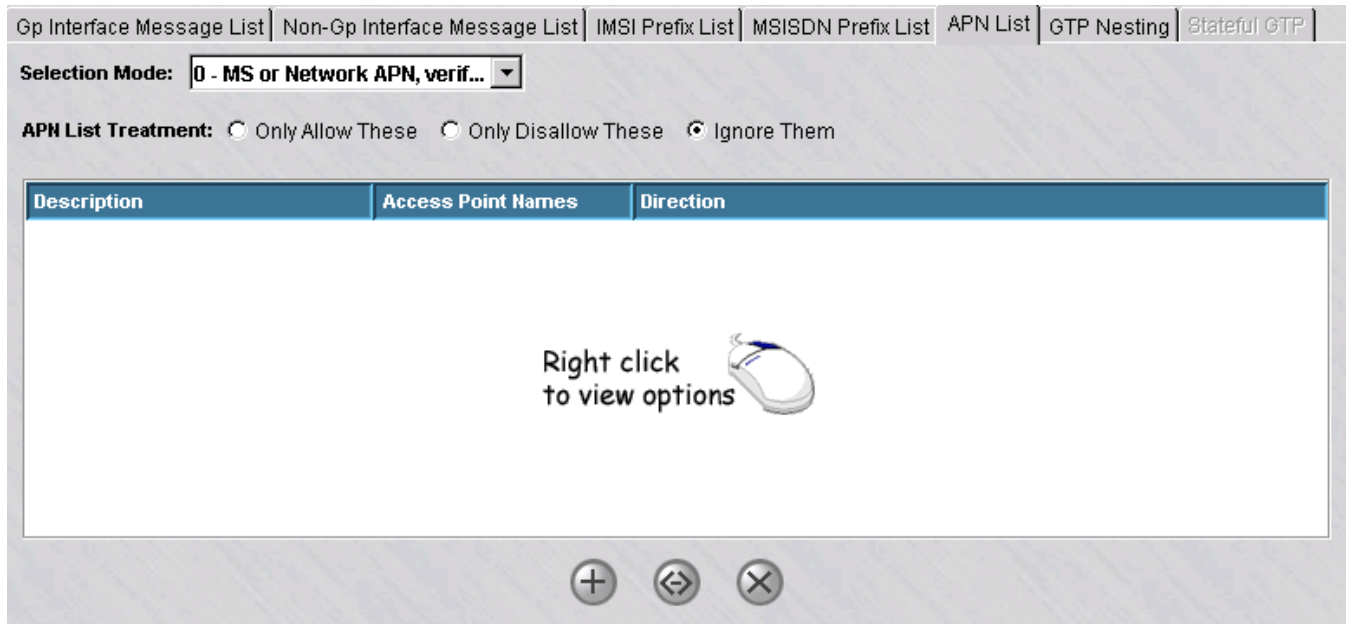
The following explains these fields:

- **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. 5-38) shows a sample screen):

Figure 5-18 GTP Application Filter Editor(APN List Tab)



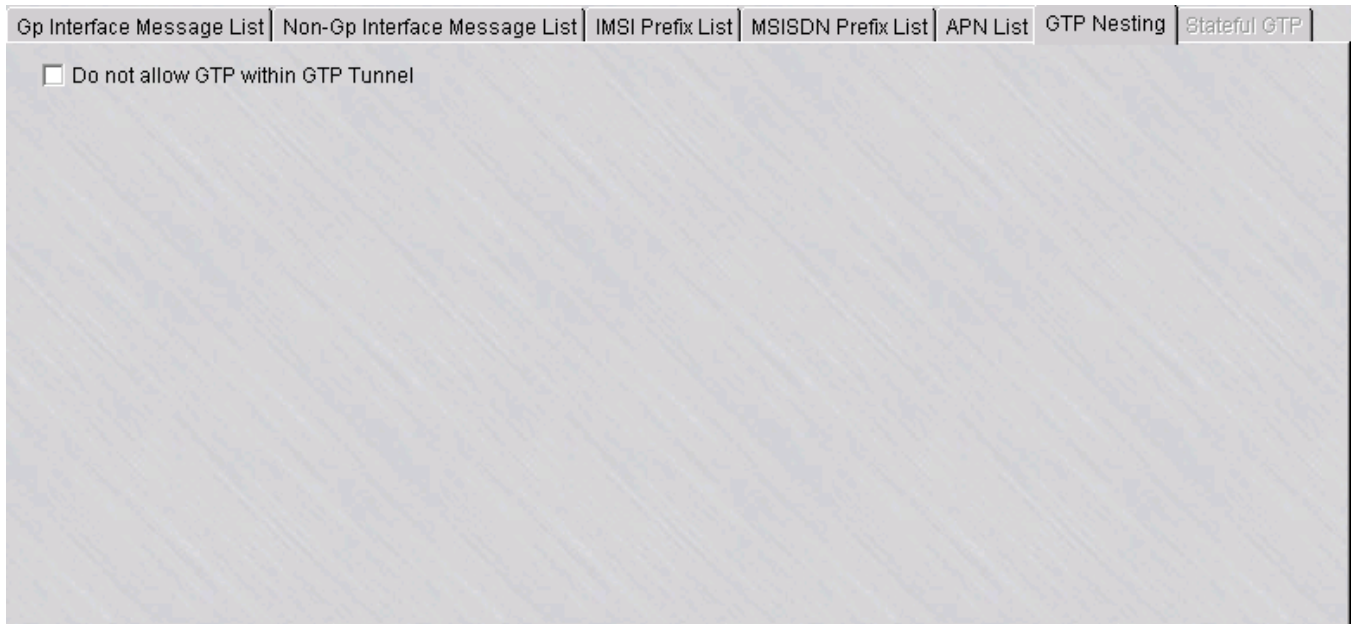
The following explains these fields:

- **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 *.lucent.com) or a fully qualified domain name (such as 127.0.0.1 or 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. 5-39) shows a sample screen):

Figure 5-19 GTP Application Filter Editor (GTP Nesting Tab)



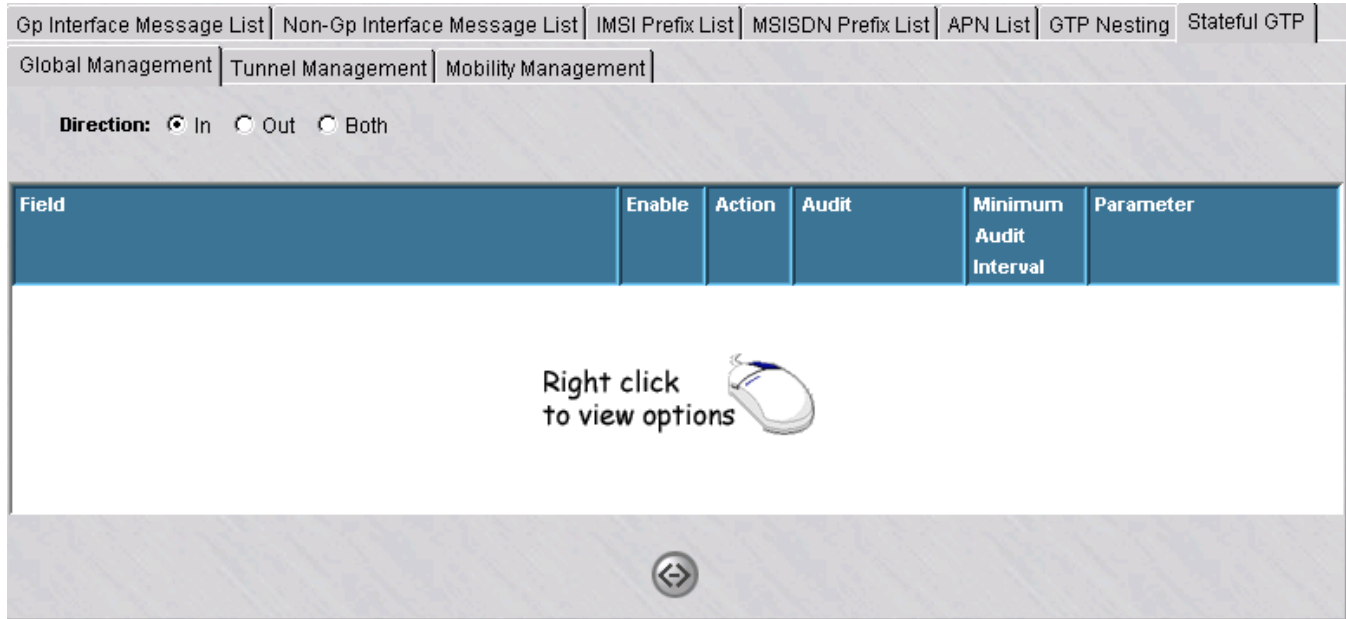
The following explains these fields:

- **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. 5-40) 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** From the **File** menu, select **Save and Close** to save the filter and close the Application Editor window.

.....
 END OF STEPS

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, “Lucent VPN Firewall Brick® Device Zone](#)

[Rulesets](#)". For additional information on service groups, see [Chapter 1, "Lucent VPN Firewall Brick® Device Zone Rulesets"](#). [Chapter 4, "Service Groups"](#).

END OF STEPS



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.

Figure 6-20 shows the H.323_VOIP Application Filter H.245 tab.

Figure 5-20 H.323_VOIP Application Filter (H.245 tab)

The following is an explanation of the "H323 VoIP" application filter parameters for 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.

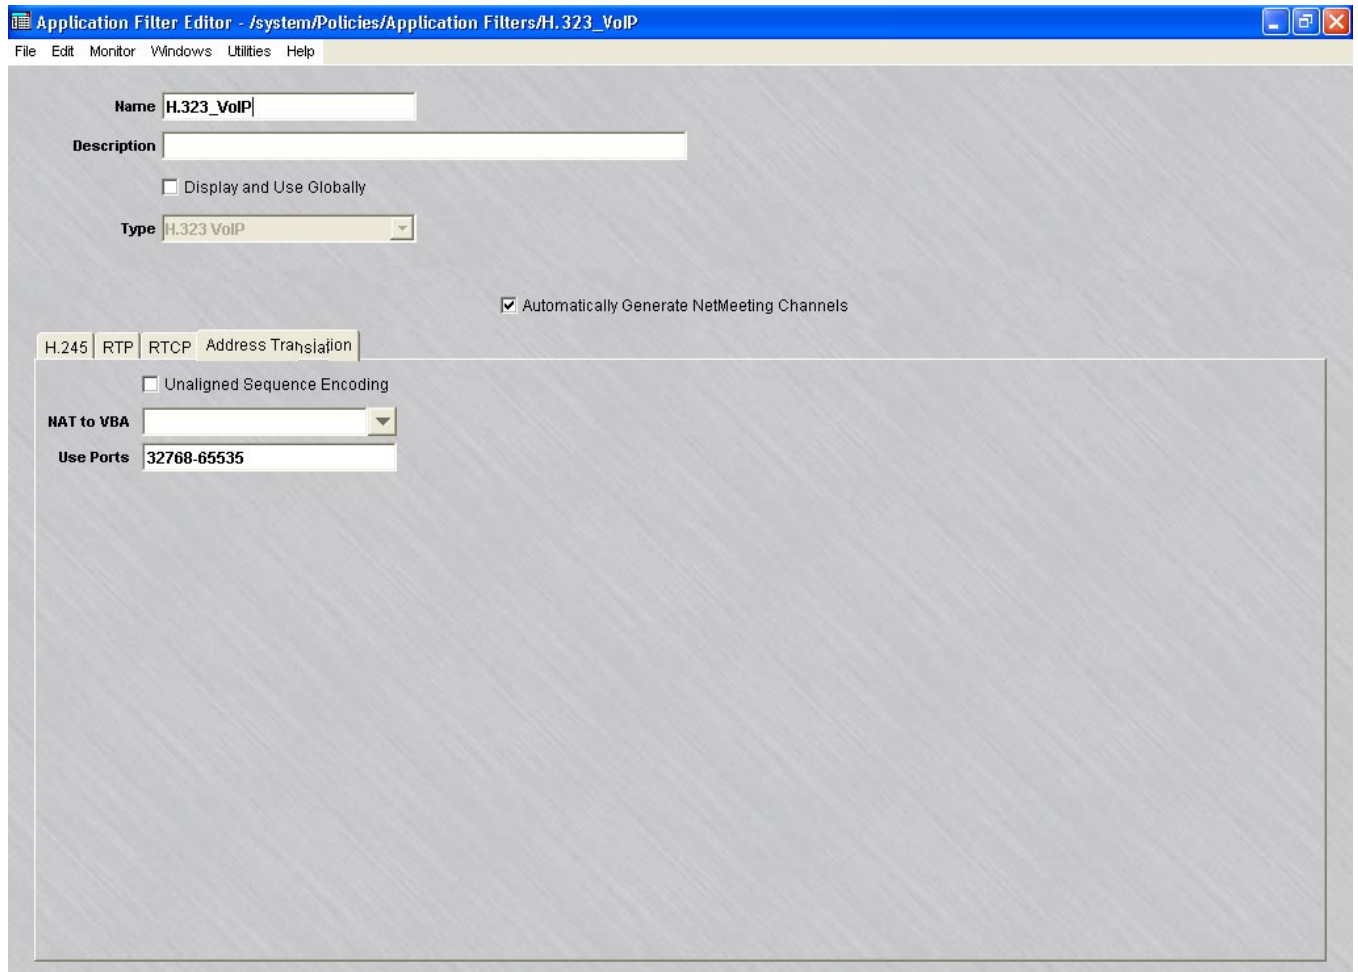
The Brick will NAT to the VBA (and PAT as needed) whenever a packet is processed that matched a rule that used one of the H.323 application filters and that application filter contained a "NAT to VBA" entry that contains the address in the packet.

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 5-21, “H.323_VOIP Application Filter (Address Translation Tab)” (p. 5-45) shows the H.323_VOIP Application Filter tab for the Port Address Translation (PAT) supports.

Figure 5-21 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.

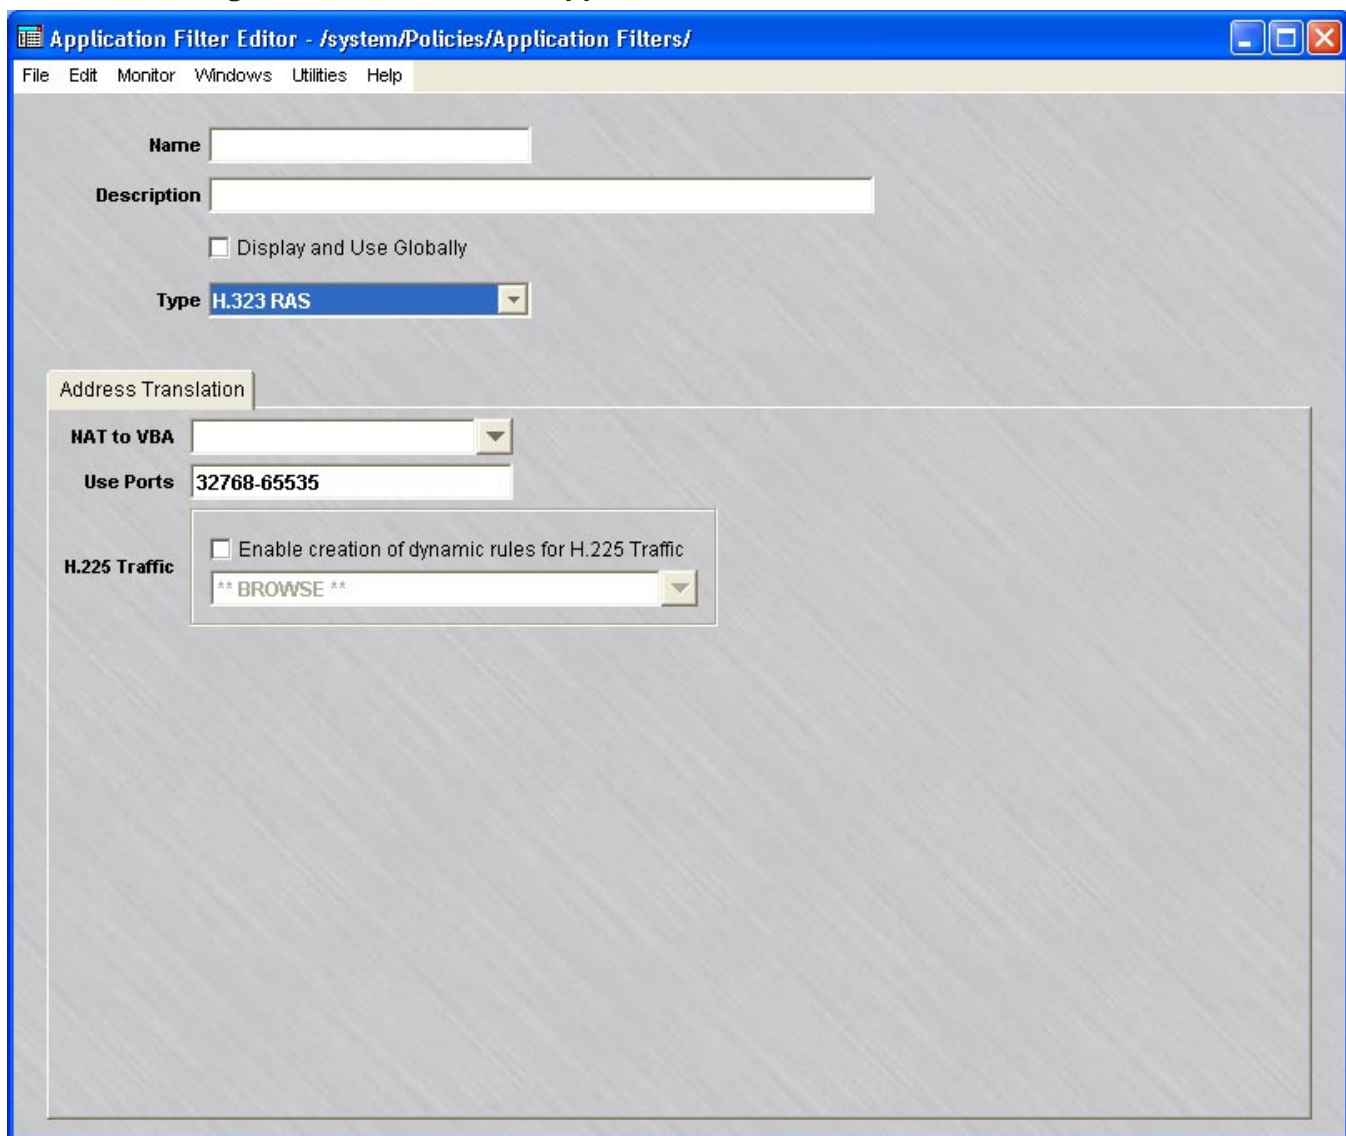
Important! The PAT feature is disabled when the NAT to VBA field is left blank.

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 5-22, "H.323_RAS Application Filter" (p. 5-46) shows the H.323_RAS Application Filter.

Figure 5-22 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 LSMS 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 LSMS 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. 5-92) 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 (Figure 5-21). 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 (Figure 5-22). 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 5-23, "Network Topology - Example 3" (p. 5-48)). 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 5-23 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.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 (i.e., dropped) by the Brick. Examples might include:

- *Strings that include extensions such as "mp3", "game", "exe", etc.*
- *Strings that contain a known virus.*

Configure an HTTP Application Filter

To configure an HTTP application filter, follow the steps below:

- 1 From the LSMS Navigator, open the Policies folder in the desired group and click **Application Filters**.

- In the right-hand column, right-click the mouse and select **New Application Filter** from the pop-up menu. The Application Filter Editor will appear. It is shown in [Figure 5-24, “Application Filter Editor”](#) (p. 5-51):

Figure 5-24 Application Filter Editor

- 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 LSMS 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. 5-92) 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 LSMS 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

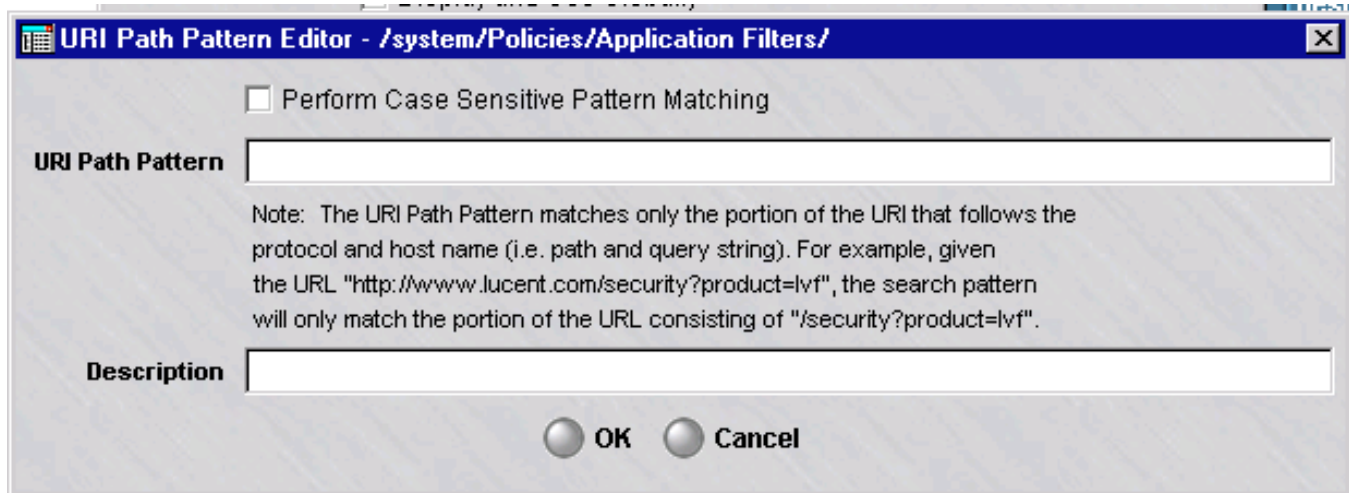
- **Commands**

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.

-
- 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 5-25, “URI Path Pattern Editor”](#) (p. 5-54).

Figure 5-25 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.

- 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 LSMS supports alternate encoding methods for the URI, including:

- %xx hex escaping
- 2 byte Unicode characters
- %U Microsoft encoding

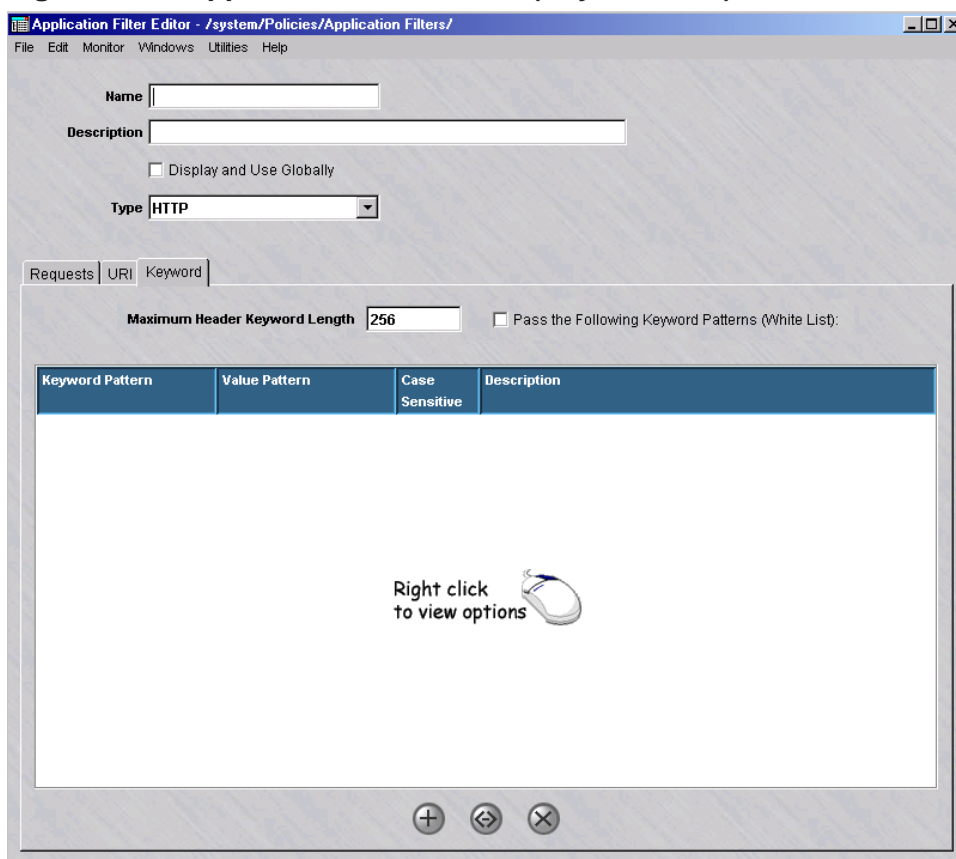
Regardless of encoding method, the Brick will fully decode the URI before attempting to perform pattern matching.

7 Click **Keyword** to display the **Keyword** tab.

Result

The Keyword tab of the Application Filter Editor is displayed (

Figure 5-26 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 a keyword may have in a request header.
 - 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.

END OF STEPS

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, “Lucent VPN Firewall Brick® Device Zone](#)

[Rulesets](#)". For additional information on service groups, see [Chapter 1, "Lucent VPN Firewall Brick® Device Zone Rulesets"](#). [Chapter 4, "Service Groups"](#).

END OF STEPS



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.

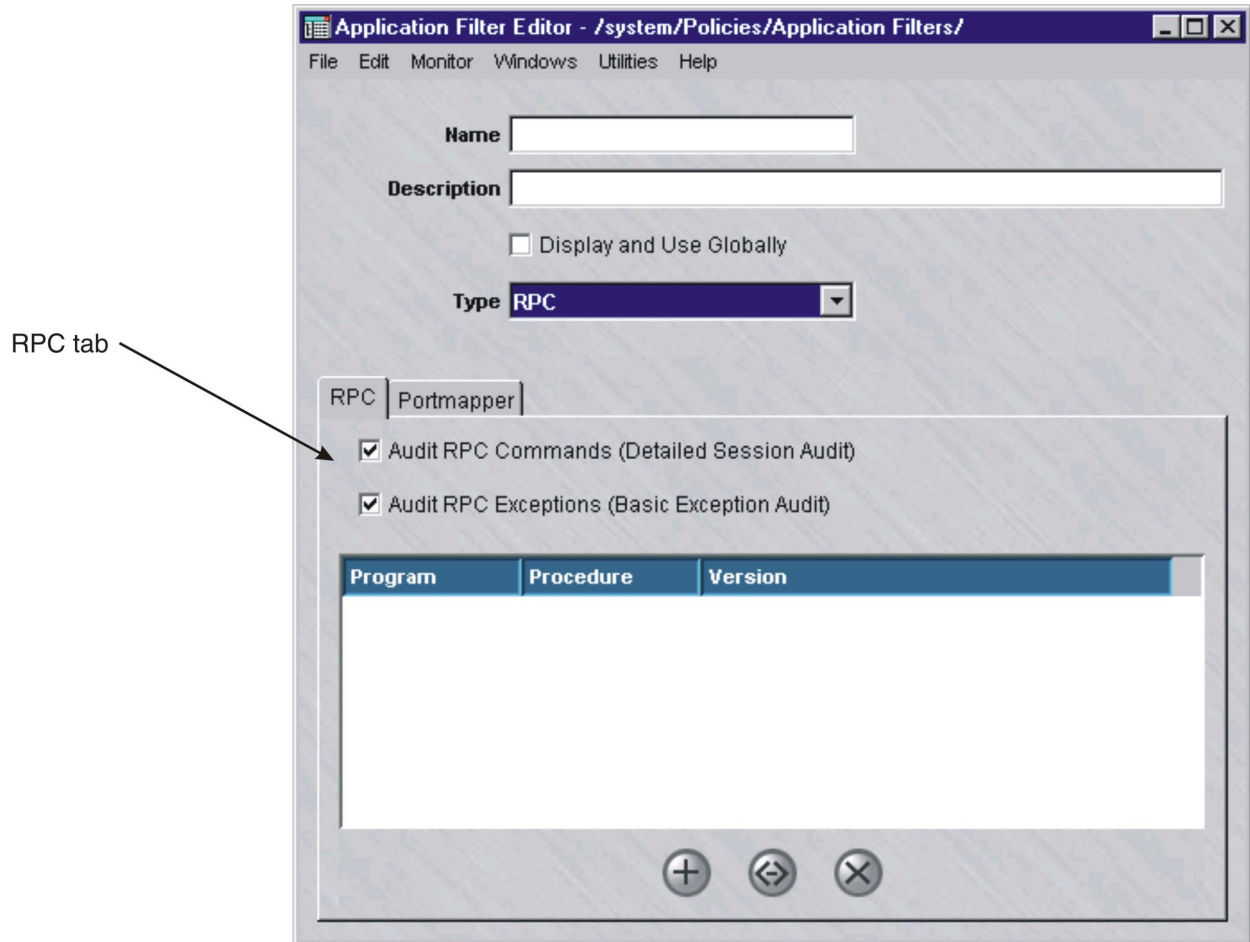
Configure an RPC Application Filter

To configure an RPC application filter, follow the steps below:

-
- 1 From the LSMS 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. You will see the RPC tab displayed, as shown in Figure 6-10.

Figure 5-27 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.

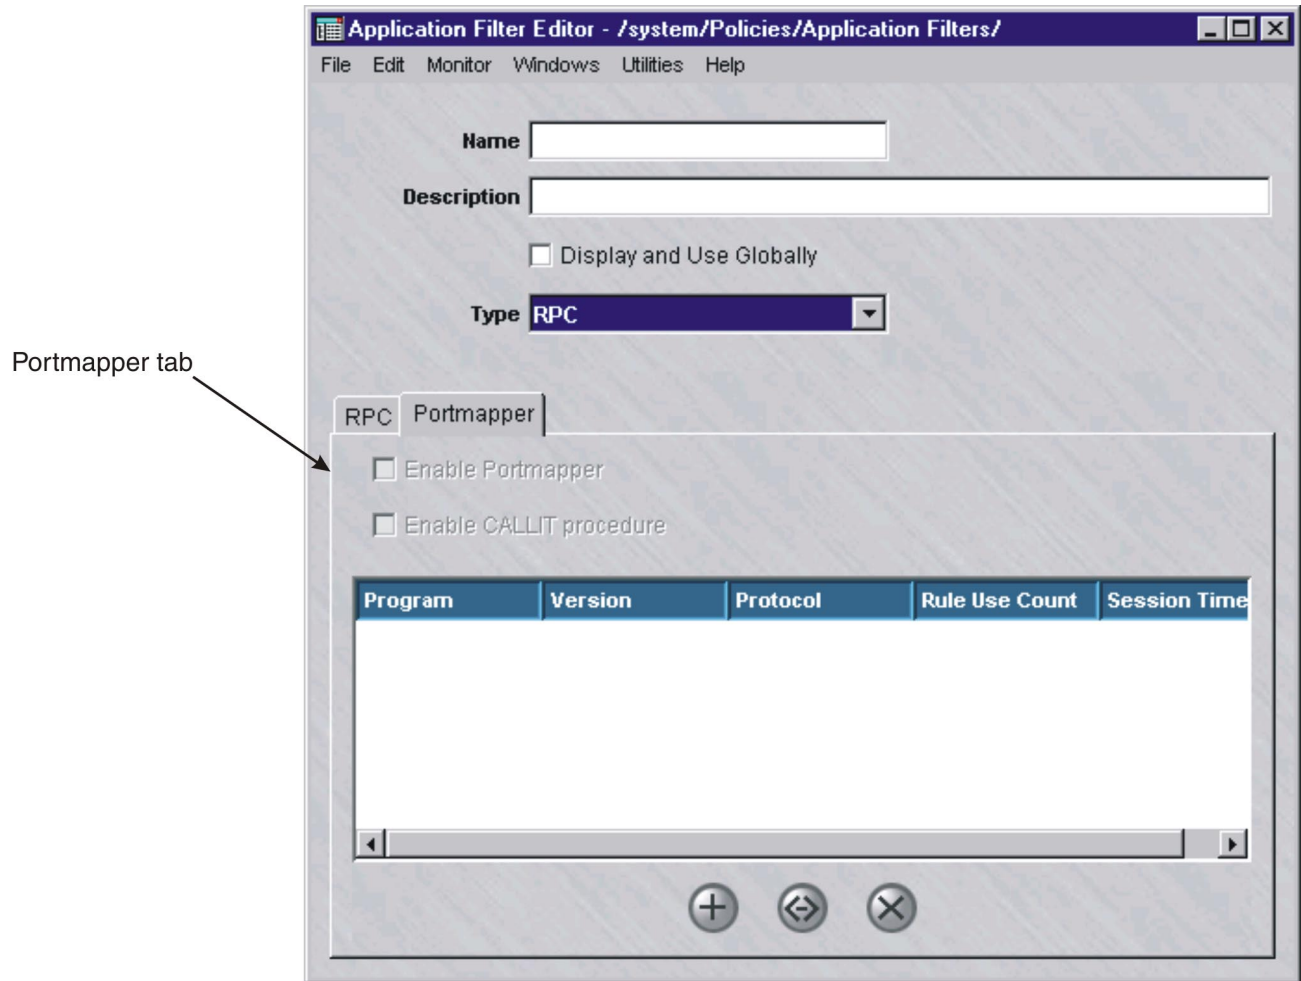
Important! If you are an LSMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the [“Global Application Filters”](#) (p. 5-92) section for an explanation of 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 5-28 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

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, "Lucent VPN Firewall Brick® Device Zone](#)

Rulesets". For additional information on service groups, see [Chapter 1, "Lucent VPN Firewall Brick® Device 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 signaling protocol used in applications such as Voice over IP (VoIP) and Instant Messaging services. SIP users have SIP ID's that function much like real-time email ID's. For instance, <sip:user1@example.com>. In a typical VoIP application, user1 registers her current "location", for example 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 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.

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.

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'ing (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'ing must be to/from the VBA. When NAT'ing 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 5-29, “SIP Proxy in the private network” (p. 5-67)). 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 5-30, “SIP Proxy on public side of the Brick” (p. 5-67))
- The Proxy in the DMZ (Figure 5-31, “SIP Proxy in the DMZ” (p. 5-68)). 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 LVE.

Figure 5-29 SIP Proxy in the private network

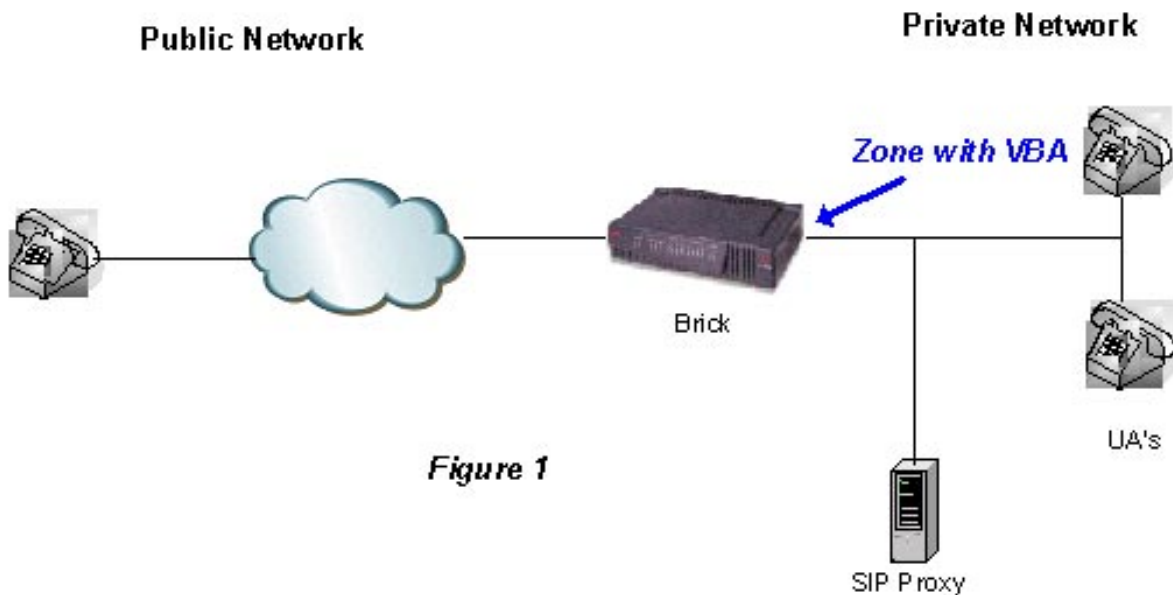


Figure 5-30 SIP Proxy on public side of the Brick





Figure 5-31 SIP Proxy in the DMZ



Create a SIP Application Filter

To configure a SIP application filter, follow the steps below:

- 1 From the LSMS 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 5-32, “SIP Application Filter Editor\(Options Default Tab\)”](#) (p. 5-69) will appear.

Figure 5-32 SIP Application Filter Editor(Options Default Tab)

Application Filter Editor - /system/Policies/Application Filters/sipDefault

File Edit Monitor Windows Utilities Help

Name

Description

Display and Use Globally

Type

Options Names and Other Addresses DNS Methods

Maximum Lengths

Entire Message

Message Header Part

Message Body Parts

Header Line

Request/Response Line

Keyword

SIP Audit Session Exception

Drop re-connect attempts after a violation (for this many seconds)

Media Transport UDP TCP

Media Maximum Streams in Session

Session Media VPN

Configure Inside and Outside of Zone Identically

Inside and Outside of Zone

Maximum Forward Count

Dynamic Ports

-
- 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.

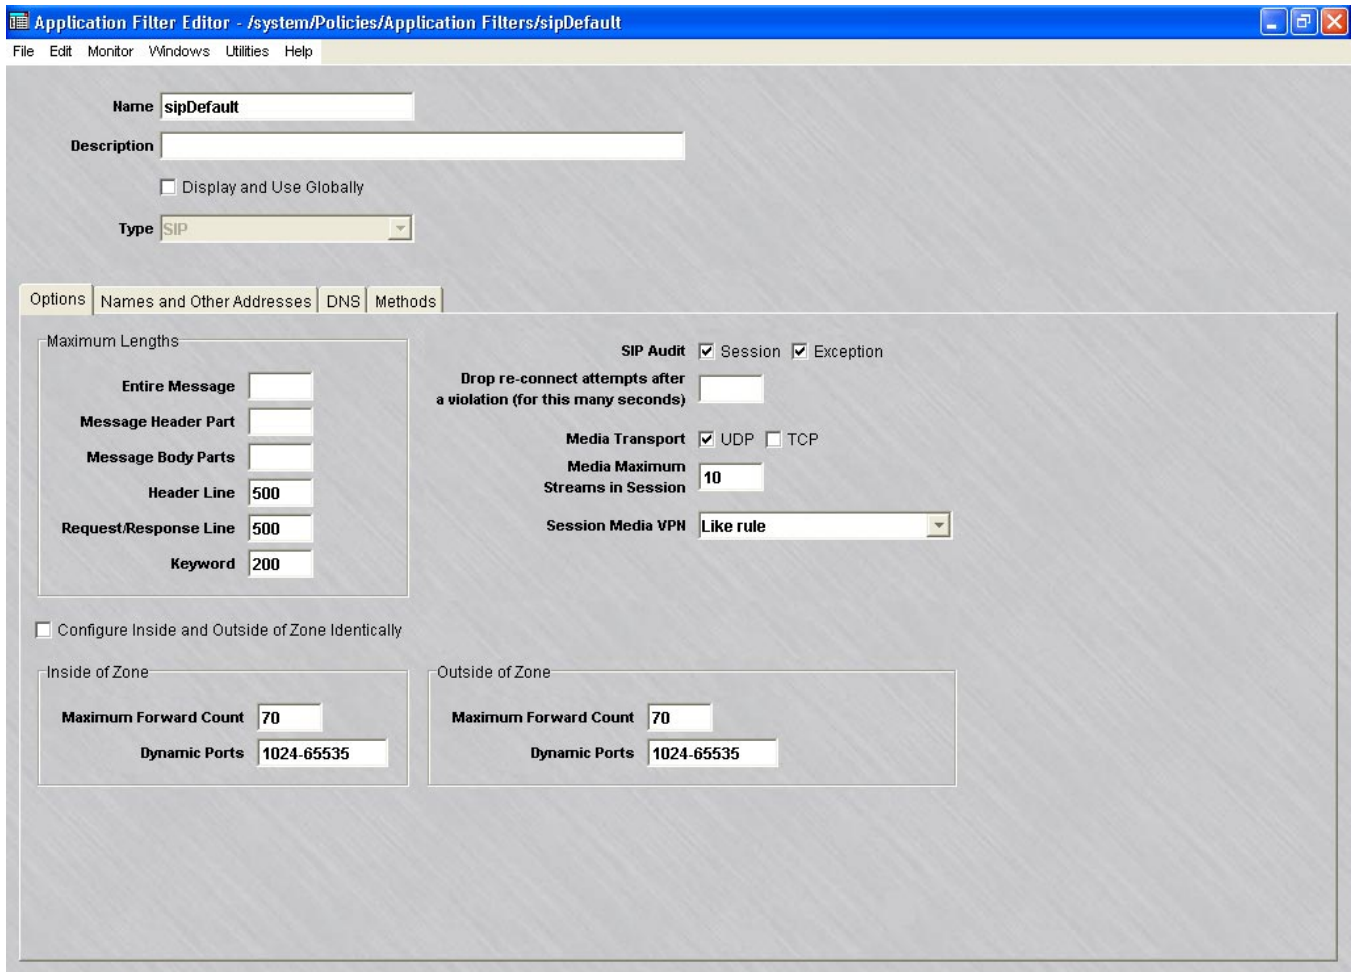
Important! If you are an LSMS Administrator, a checkbox entitled **Display and Use Globally** is displayed under the **Description** field. Click this checkbox if you want to make this a global application filter. Refer to the [“Global Application Filters”](#) (p. 5-92) section for an explanation of 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 (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 5-33, “SIP Application Filter Editor \(Options Tab\)- both zones”](#) (p. 5-71)).

- **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 5-33 SIP Application Filter Editor (Options Tab)- both zones



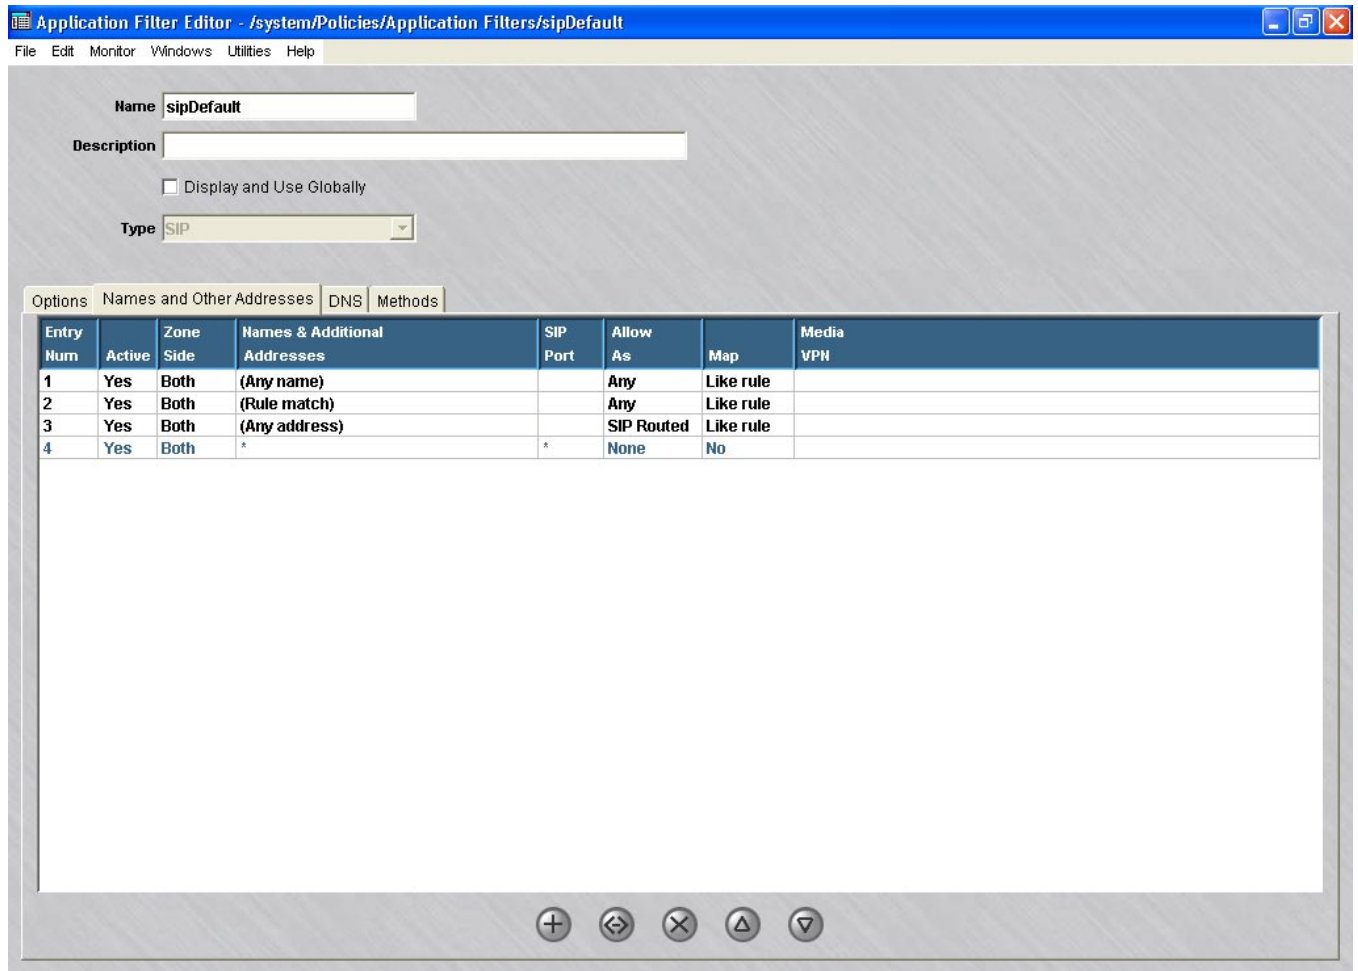
The following explains these fields:

- **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.

-
- 5 The **Names and Other Addresses** tab (see [Figure 5-34, "SIP Application Filter Editor \(Names and Other Addresses Tab\)"](#) (p. 5-73)) 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 5-34 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.

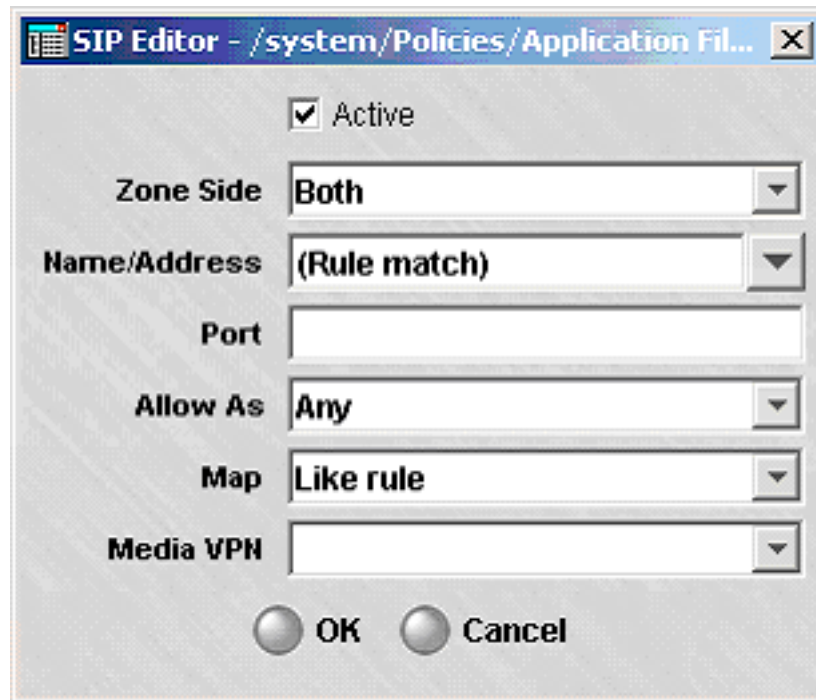
Additionally, this table can be used to ensure that certain domain names only originate from inside the zone.

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 5-35, “SIP Editor \(Names Entry Screen\)”](#) (p. 5-74)). 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 5-35 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 5-36, “SIP Application Filter Editor \(DNS Default Tab\)”](#) (p. 5-76)). 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 5-37, “SIP Application Filter Editor \(DNS Tab\)-both zones”](#) (p. 5-76)).

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 5-36 SIP Application Filter Editor (DNS Default Tab)

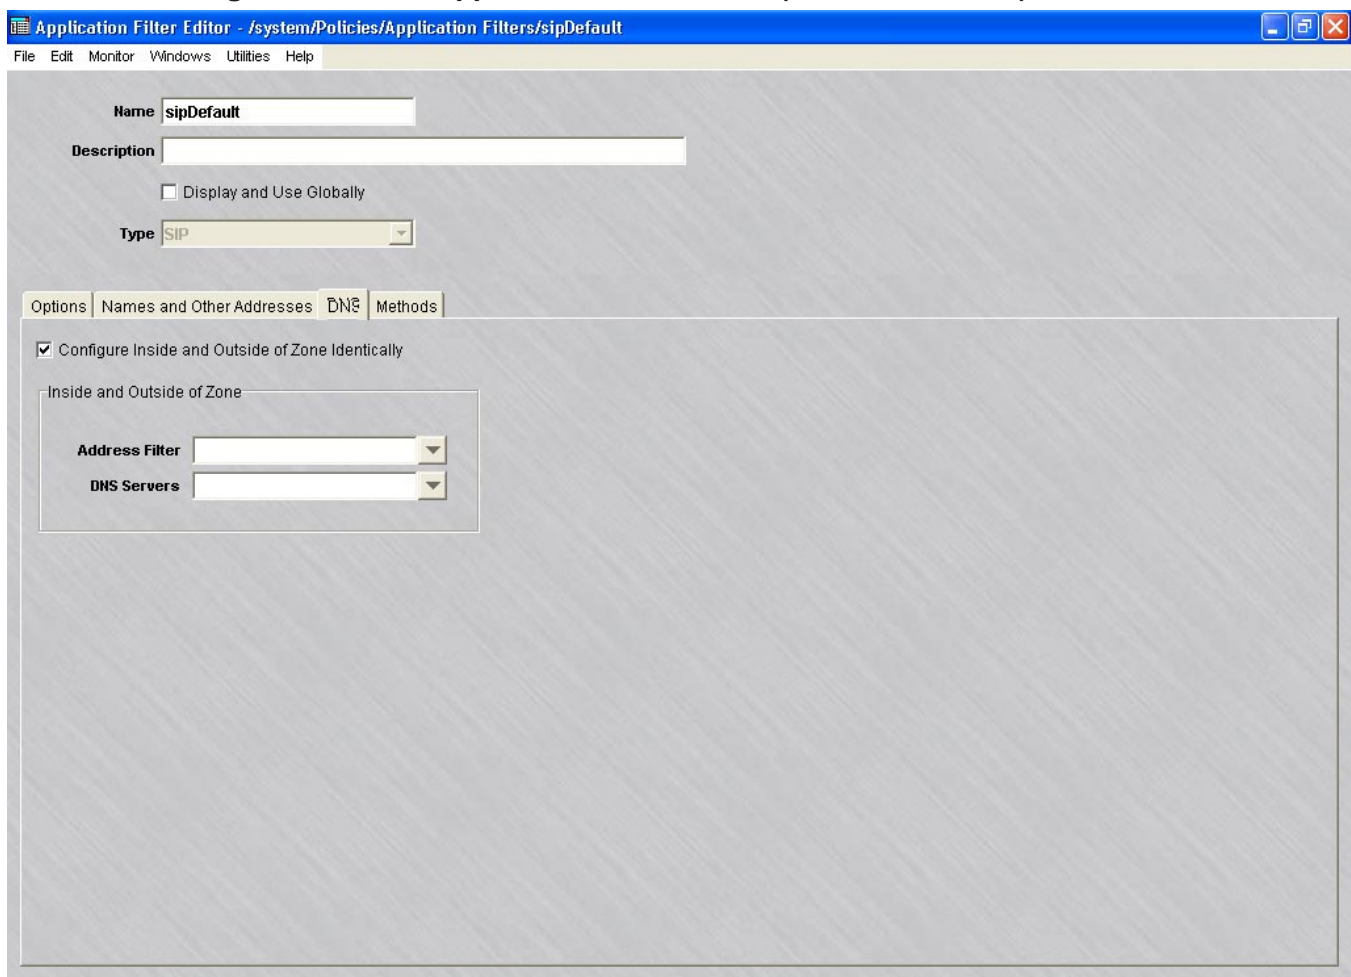
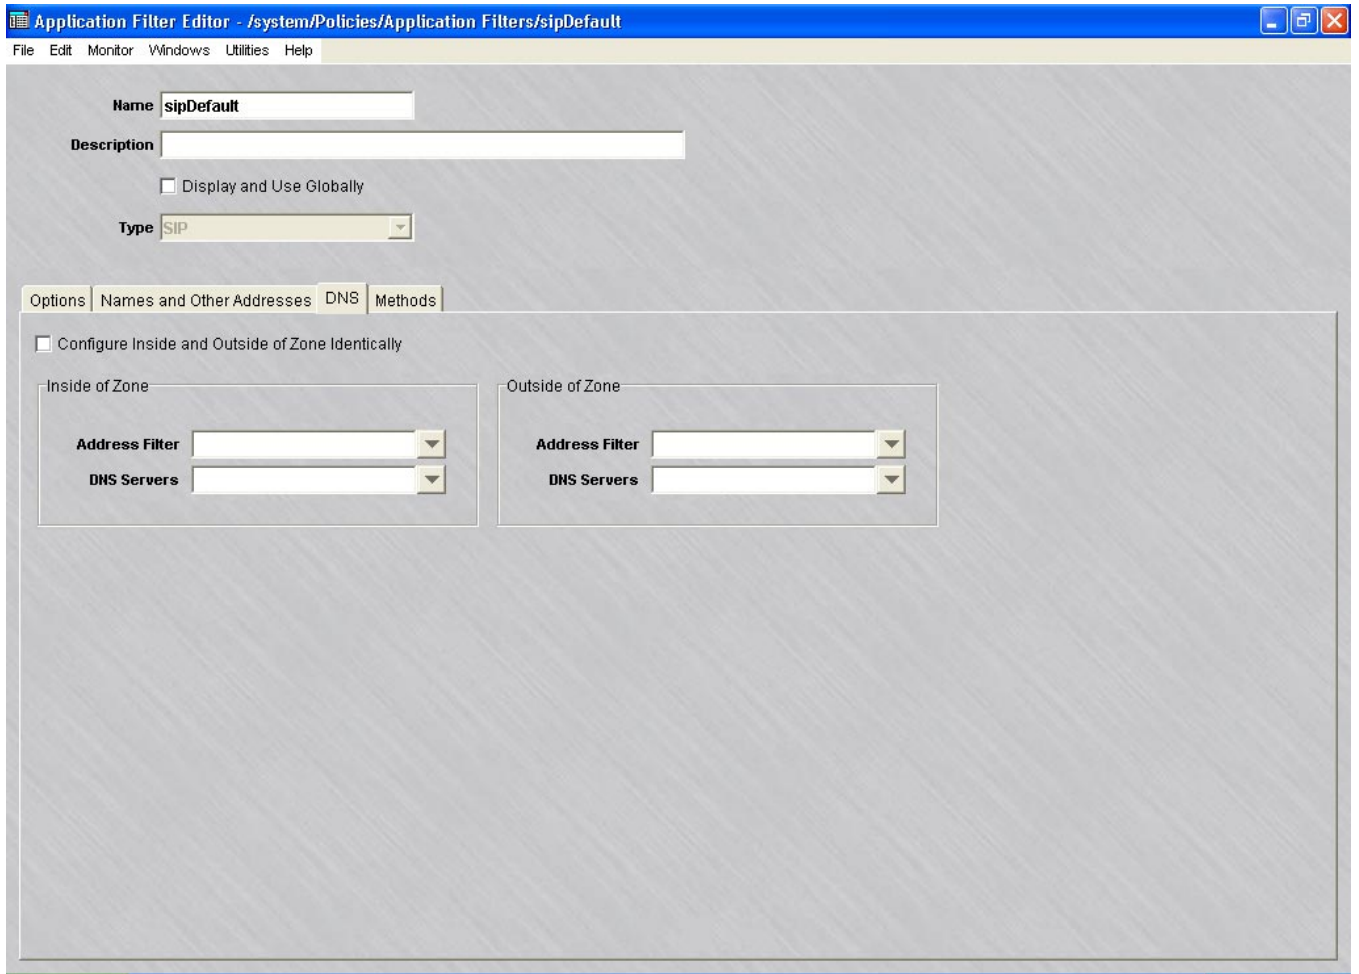


Figure 5-37 SIP Application Filter Editor (DNS Tab)-both zones



-
- 7 The **Methods** tab contains these options (see [Figure 5-38, “SIP Application Filter Editor \(Methods Default Tab\)”](#) (p. 5-78)):
- **Configure Inside and Outside of Zone Identically**
If checked (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 5-38, “SIP Application Filter Editor \(Methods Default Tab\)”](#) (p. 5-78)).

Figure 5-38 SIP Application Filter Editor (Methods Default Tab)

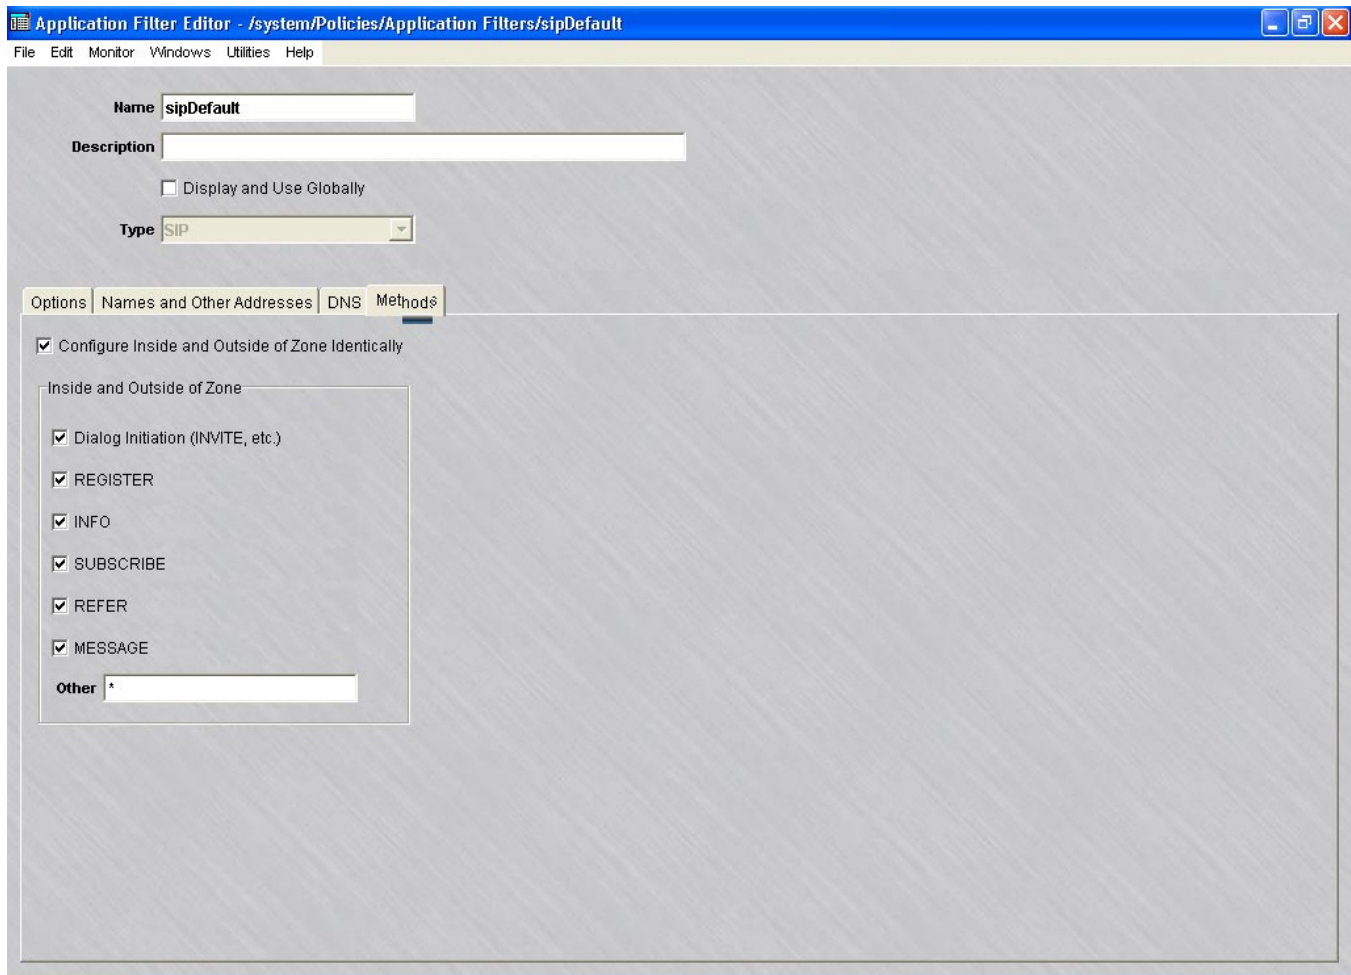
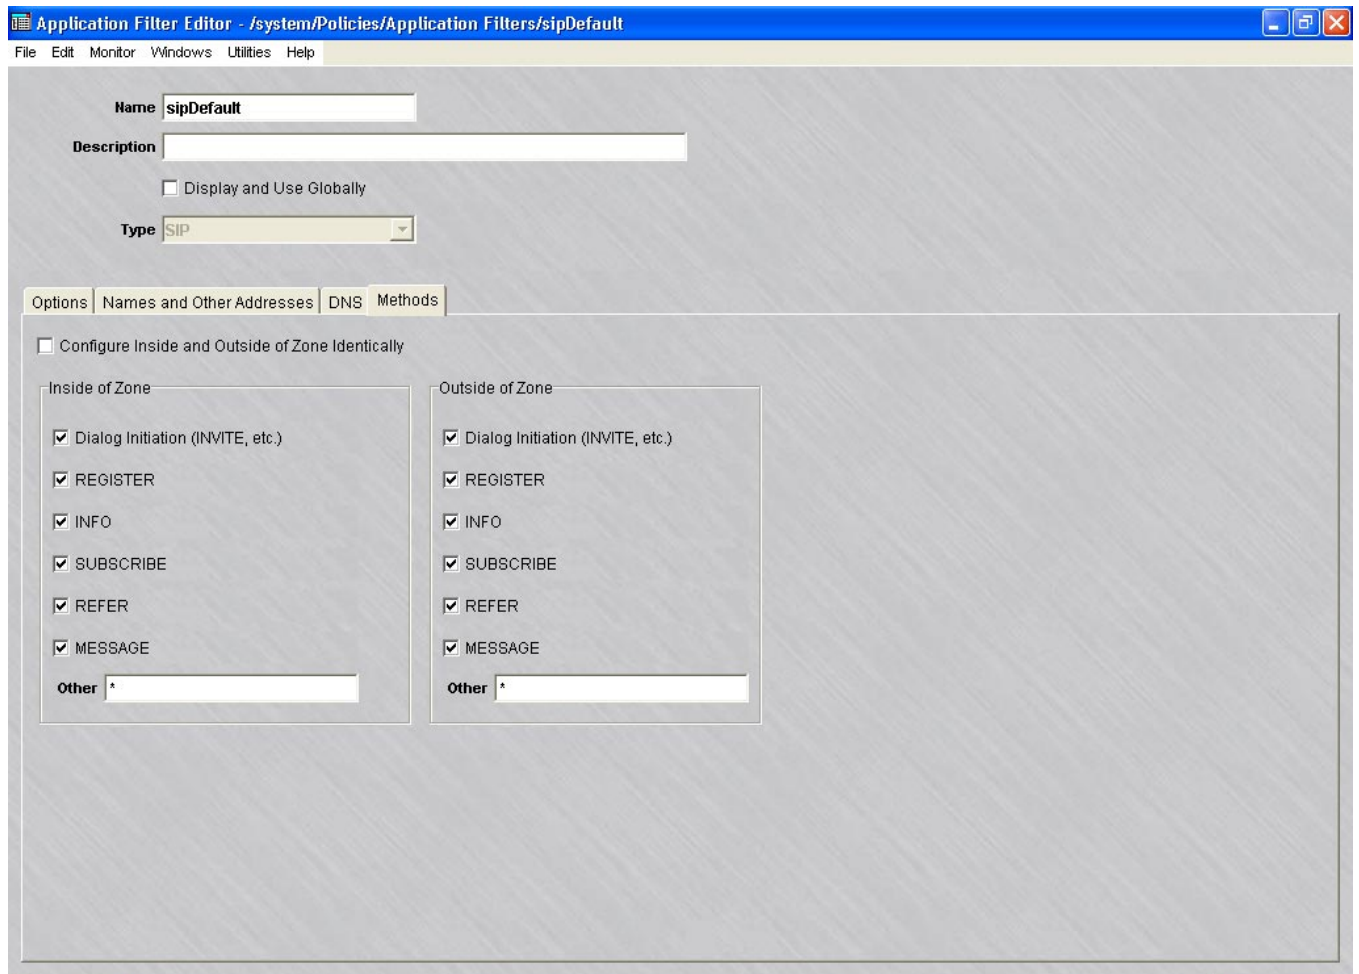


Figure 5-39 SIP Application Filter Editor (Methods Tab) - both zones



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 LSMS and, once installed, is created automatically for newly created groups. Similarly, an *smtpDefault* filter is created when an LSMS 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.

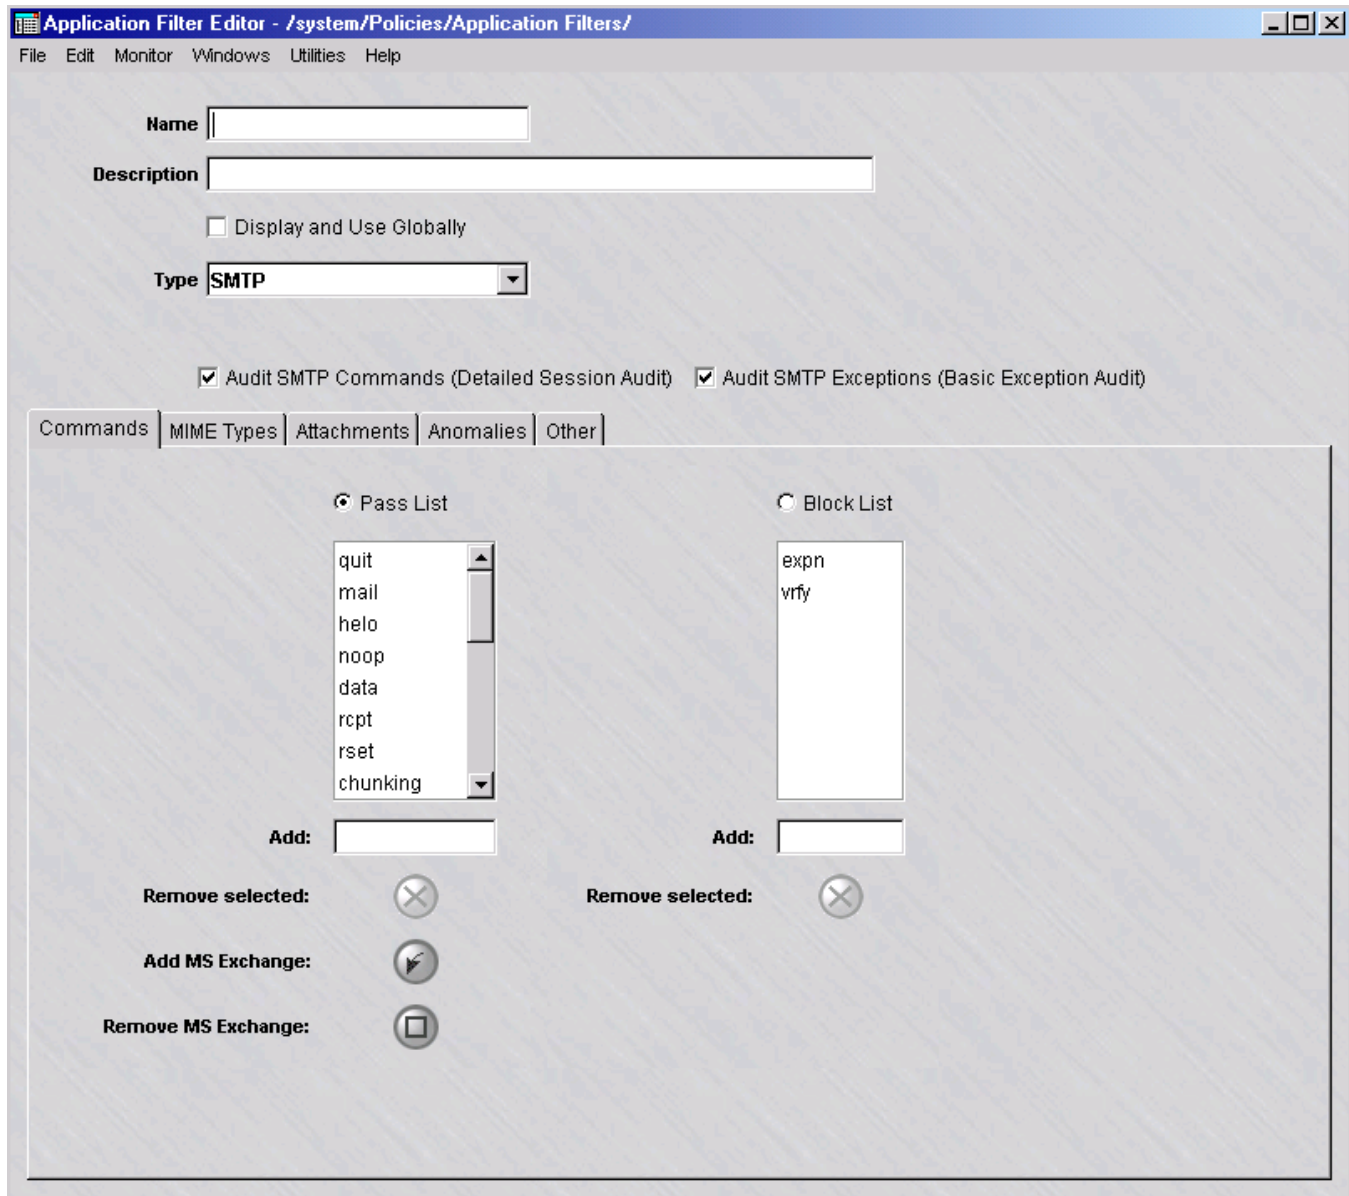
How to Configure an SMTP Application Filter

To configure an SMTP application filter, follow the steps below:

-
- 1 From the LSMS 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. “[How to Configure an SMTP Application Filter](#)” (p. 5-81) shows a sample of the screen.

Figure 5-40 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 LSMS 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. 5-92) 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 5-41, “SMTP Application Filter Editor (MIME Types Tab)” (p. 5-84) shows a sample screen):

Figure 5-41 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>

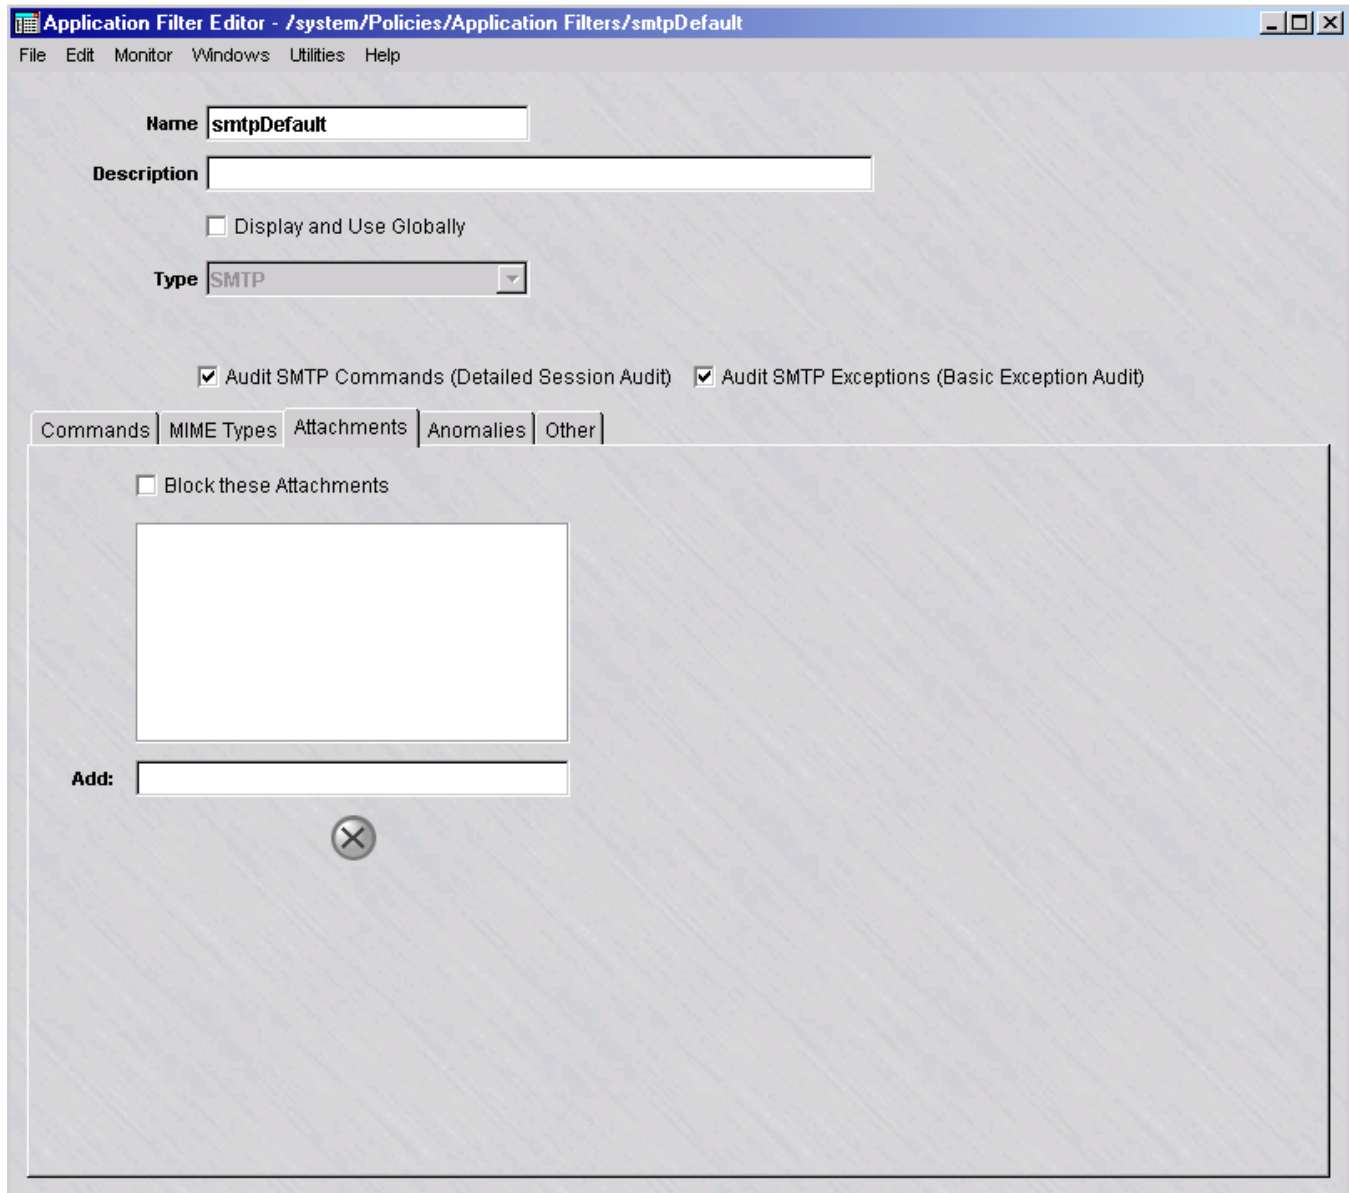
If the incoming mail contains a MIME type found in this list, then the entire MIME object (content) is removed and a message to this effect is inserted in the mail. An error message is sent back to the mail sender.

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 5-42, “SMTP Application Editor (Attachment Tab)” (p. 5-86) shows a sample screen):

Figure 5-42 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 5-43, “SMTP Application Filter Editor (Anomalies Tab)” (p. 5-87) shows a sample screen):

Figure 5-43 SMTP Application Filter Editor (Anomalies Tab)

The screenshot shows a window titled "Application Filter Editor - /system/Policies/Application Filters/smtpDefault". The window has a menu bar with "File", "Edit", "Monitor", "Windows", "Utilities", and "Help". The main area contains the following fields and controls:

- Name:** A text box containing "smtpDefault".
- Description:** An empty text box.
- Display and Use Globally
- Type:** A dropdown menu showing "SMTP".
- Audit SMTP Commands (Detailed Session Audit) Audit SMTP Exceptions (Basic Exception Audit)
- Navigation tabs: "Commands", "MIME Types", "Attachments", "Anomalies" (selected), "Other".
- Enable Protocol Anomaly Checking
- Maximum Domain Name Length:** Radio buttons for "64" and "255", with "255" selected.
- Maximum Message Length (MB):** An empty text box.

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**.

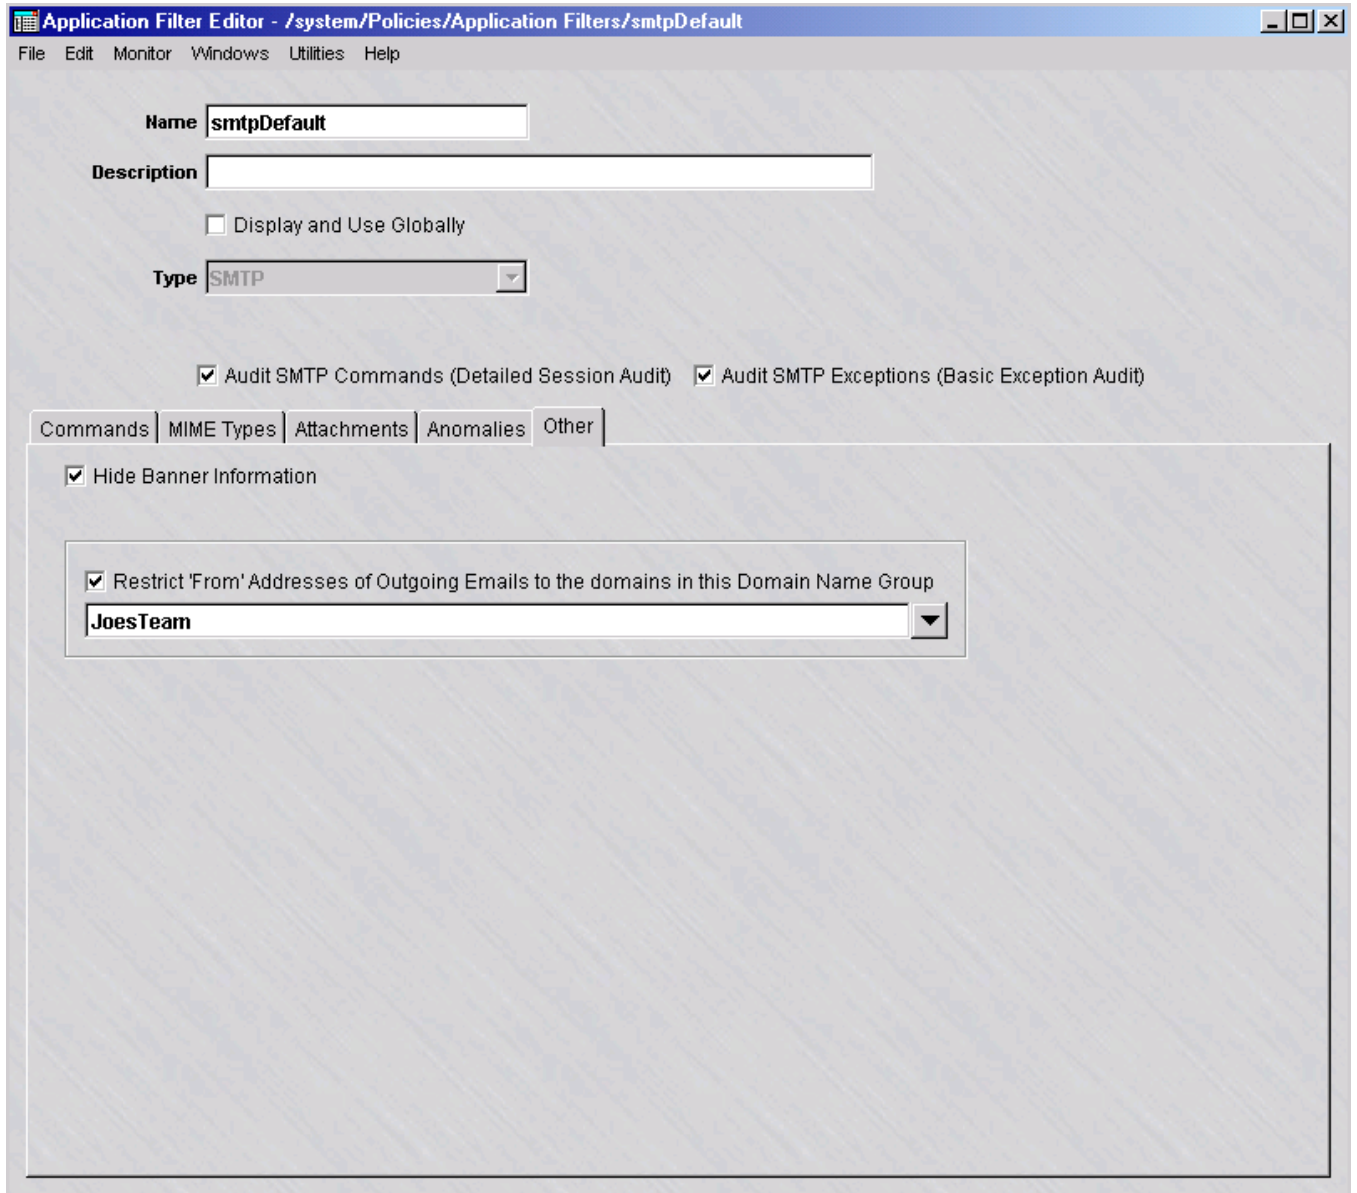
To thwart mail bomb attacks, a maximum message size may be specified. The allowed range is 0-9999. Setting the value to zero (0) means that the limitation is 64k. *If left blank, no limit is enforced.*

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 5-44, “SMTP Application Filter Editor (Other Tab)” (p. 5-89) shows a sample screen):

Figure 5-44 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

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, "Lucent VPN Firewall Brick® Device Zone](#)

Rulesets”. “Overview” (p. 1-1). For additional information on service groups, see Chapter 1, “Lucent VPN Firewall *Brick*[®] Device Zone Rulesets”. Chapter 4, “Service Groups”.

END OF STEPS



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 LSMS Administrators can create global application filters.

Create a Global Application Filter

If you are an LSMS 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.

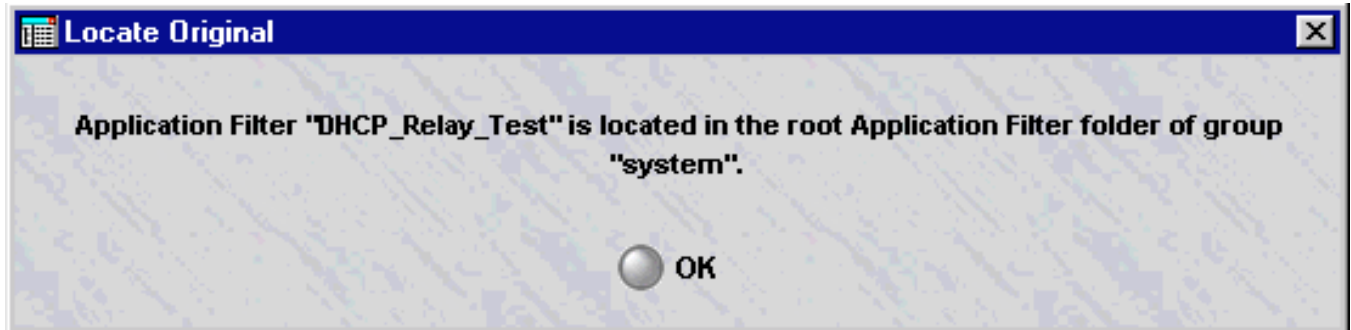
However, any global application filters created in a *different group* can be identified by a globe icon that appears to the left of the entry, as shown below.

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 5-45 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.

□

6 Network Address Translation

Overview

Purpose

This chapter explains how to set up Network Address Translation (NAT). NAT is a feature that enables the Lucent VPN Firewall *Brick*[®] device to map the source or destination addresses of inbound and outbound sessions to other addresses.

Contents

What is Network Address Translation?	6-2
How to Set Up Source Address Mapping	6-4
How to Set Up Destination Address Mapping	6-7
How to Set Up Destination Port Mapping	6-11
How to Perform Source Address Mapping with a Router	6-12
How to Perform Source Address Mapping without a Router	6-14
How to Perform Destination Address Mapping without a Router	6-16
Other Examples of a Brick Responding to ARPs	6-18



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.

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, for 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.



How 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 that you want to map the source addresses in the rule to, 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 (see [Figure 1-2, “Brick Zone Ruleset Editor”](#) (p. 1-10) in “Overview” (p. 1-1)).

- 4 Right-click in the Rule Viewer and select **New** from the pop-up menu. The Brick Zone Rule Editor will appear (see [Figure 1-5, “Brick Zone Rule Editor \(Basic Tab\)”](#) (p. 1-18) 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.
-

- 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 *sales_hosts*.

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 [“Overview” \(p. 1-1\)](#) 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**

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 IPSec Client users that need the "Local Presence" feature. Please refer to *Appendix A Local Presence* at the end of this manual for more details.

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



How 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. The Brick Zone Ruleset Editor will appear. See [Figure 1-2, “Brick Zone Ruleset Editor” \(p. 1-10\)](#) in [“Overview” \(p. 1-1\)](#).

- 3 Right-click an existing rule and select **New** from the pop-up menu. The Brick Zone Rule Editor will appear. See [Figure 1-5, “Brick Zone Rule Editor \(Basic Tab\)” \(p. 1-18\)](#) in [“Overview” \(p. 1-1\)](#).

-
- 4 Create an inbound rule. [Figure 6-2, “Local Access Rule” \(p. 6-8\)](#) 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 6-2 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.

-
- 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 6-2, “Local Access Rule” \(p. 6-8\)](#), 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 IPSec Client users that need the "Local Presence" feature. Please refer to *Appendix A Local Presence* at the end of this manual for more details.

END OF STEPS



How 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. See [Figure 1-2, “Brick Zone Ruleset Editor” \(p. 1-10\)](#) in “Overview” (p. 1-1).

- 3 Right-click an existing rule and select **New** from the pop-up menu. The Brick Zone Rule Editor will appear. See [Figure 1-5, “Brick Zone Rule Editor \(Basic Tab\)” \(p. 1-18\)](#) in “Overview” (p. 1-1).

- 4 Create an inbound rule (see [Figure 6-2, “Local Access Rule” \(p. 6-8\)](#)), 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




How 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



How to Perform Source Address Mapping without a Router

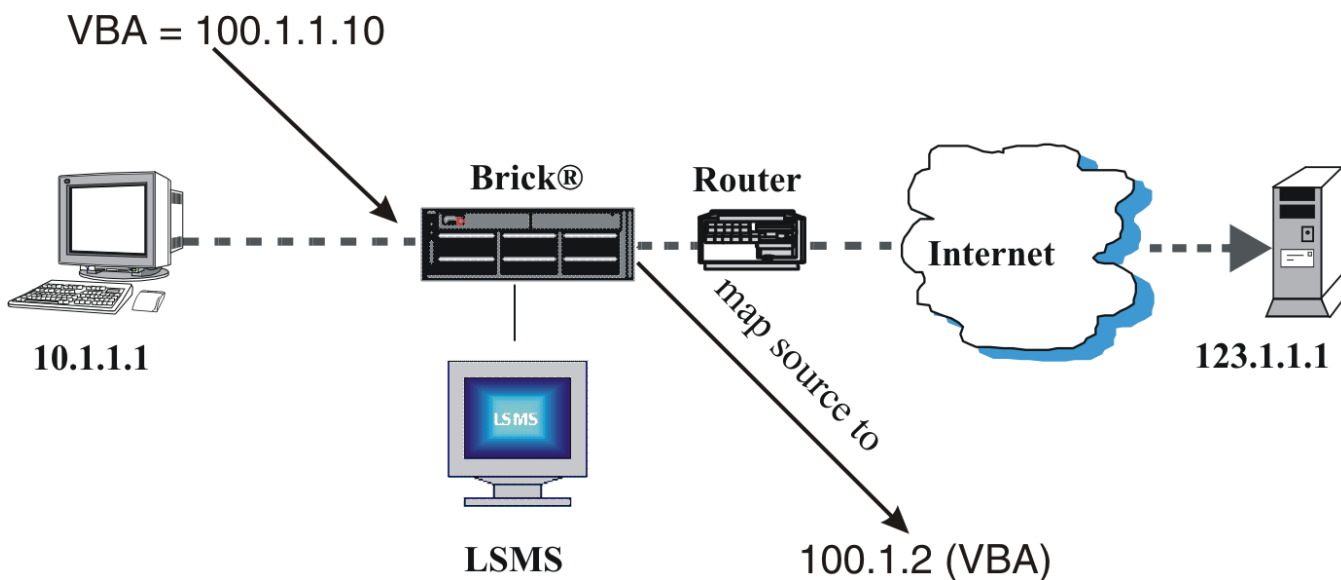
When to use

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

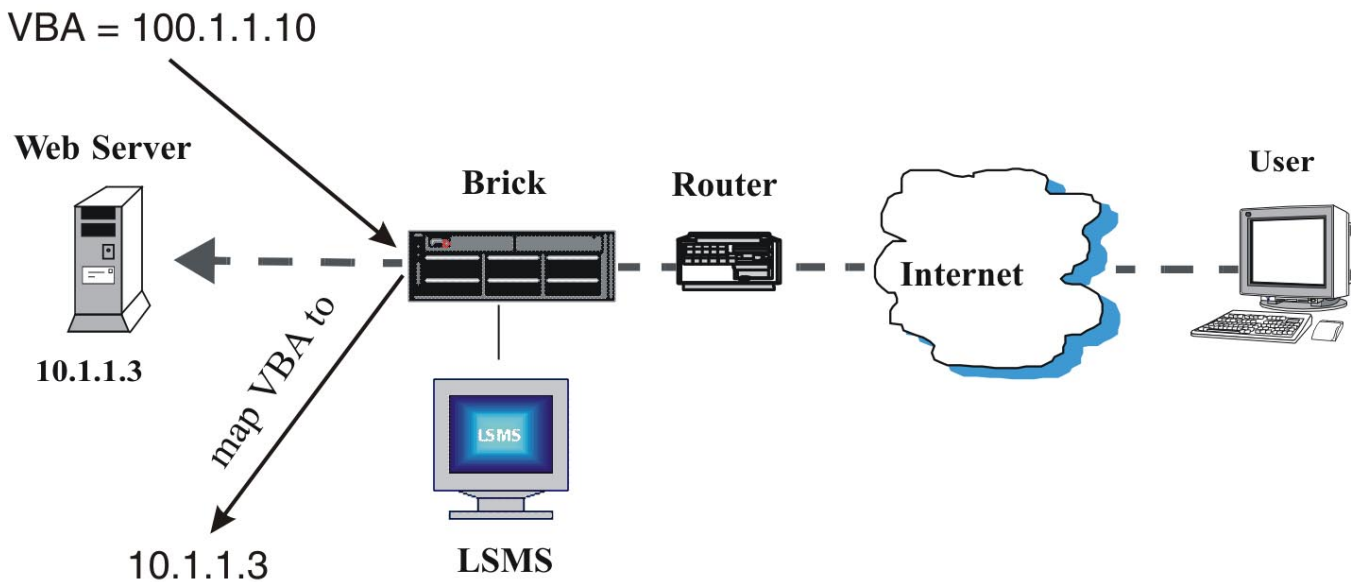


How to Perform Destination Address Mapping without a Router

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	10.1.1.10
Service	telnet
Action	Pass
Destination Address Mapping	Management Server



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, please refer to [Chapter 1, "Lucent VPN Firewall Brick® Device Zone Rulesets"](#).

-
- 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, please refer to *Chapter 4 Configuring Brick Ports* in the *LSMS 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
How to Set Up a Dependency Mask	7-3
How to Maintain a Dependency Mask	7-10
Example: RealAudio Session	7-16



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 Lucent VPN Firewall *Brick*[®] device 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.



How 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.

Display the Dependency Masks Editor

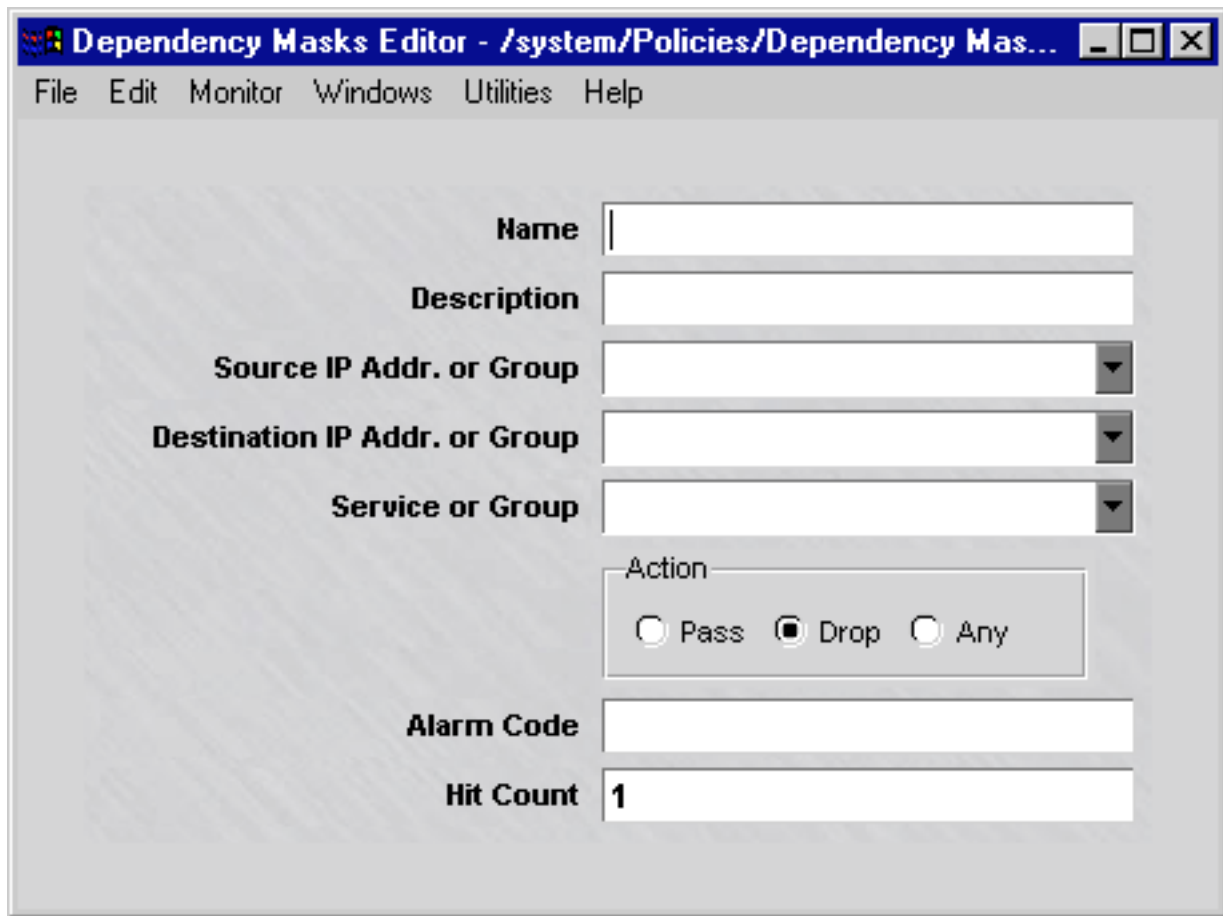
The easiest way to display the Dependency Masks Editor is to:

- 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. 7-4).

Figure 7-1 Dependency Masks Editor



END OF STEPS

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 [“View Dependency Masks”](#) (p. 7-10).

-
- 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 LSMS.
- *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. This 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.

You have the option of entering an alarm code in the dependency mask. If you do, the Brick will search the session cache for sessions matching that alarm code.

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

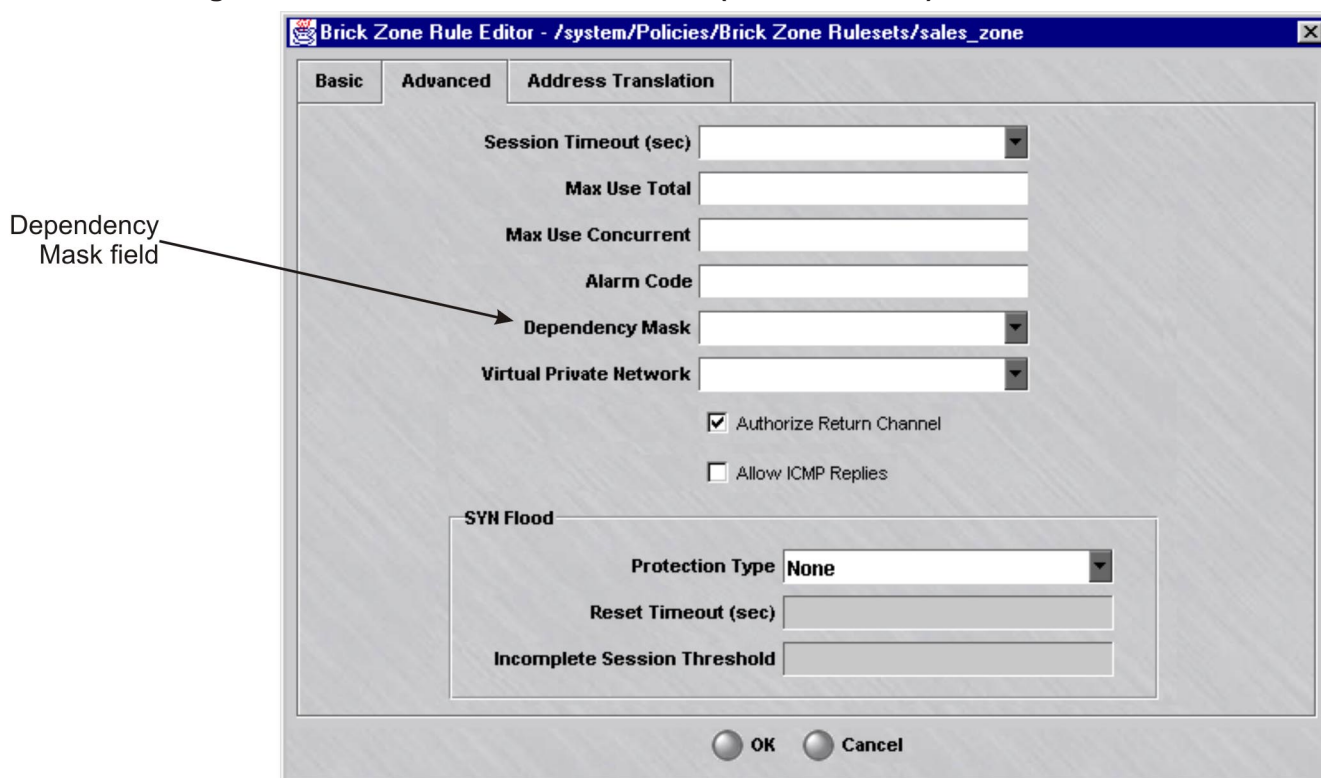
Associate the Dependency Mask with a Rule

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 “Overview” (p. 1-1). 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. It is shown in Figure 7-2, “Brick Zone Ruleset Editor (Advanced Tab)” (p. 7-8).

Figure 7-2 Brick Zone Ruleset Editor (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.

END OF STEPS



How to Maintain a Dependency Mask

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.

View Dependency Masks

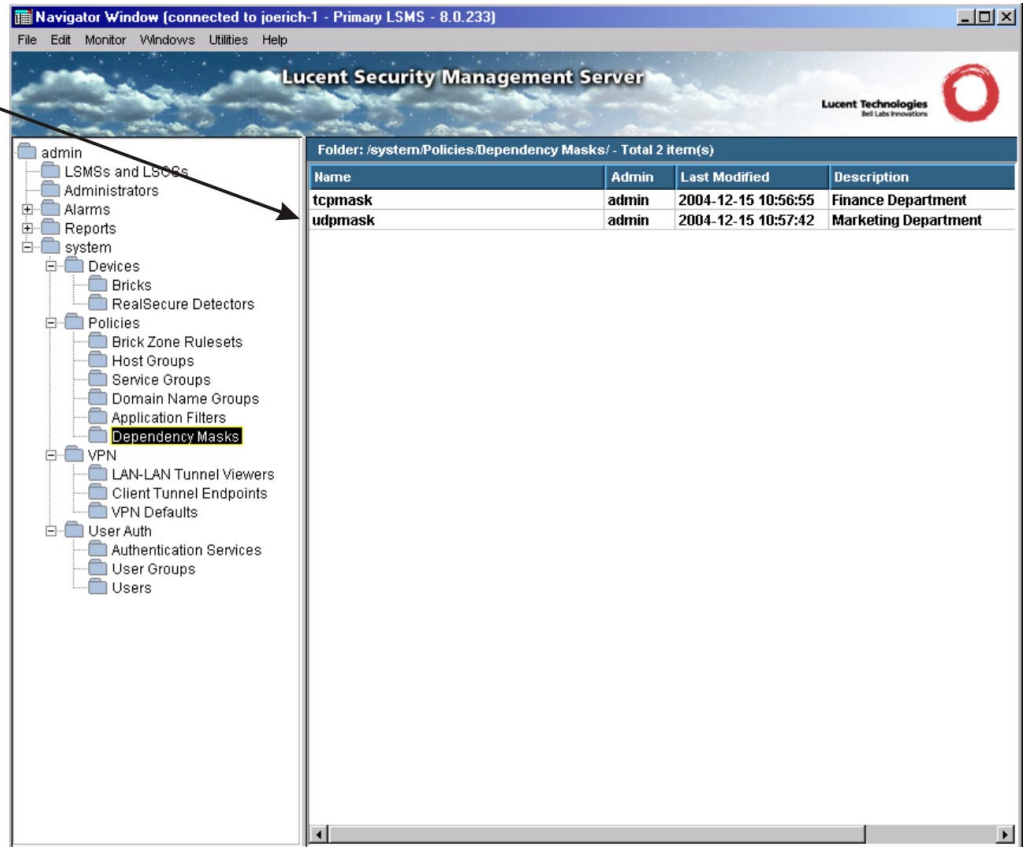
You must view the dependency masks before you can edit, move, copy, or delete a specific dependency mask. To view all the dependency masks that have been created to date in a particular group, follow the steps below:

-
- 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 7-3, “Navigator Window \(View Dependency Masks\)”](#) (p. 7-11)).

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 7-3 Navigator Window (View Dependency Masks)

Existing
Dependency
Masks



END OF STEPS

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. 7-12).

Figure 7-4 Dependency Masks Editor (Edit Mode)

The screenshot shows a window titled "Dependency Masks Editor - /system/Policies/Dependency Mas...". The menu bar includes "File", "Edit", "Monitor", "Windows", "Utilities", and "Help". The main area contains several fields for editing a dependency mask:

- Name:** sales_mask (greyed out)
- Description:** sales department
- Source IP Addr. or Group:** DESTINATION (dropdown menu)
- Destination IP Addr. or Group:** SOURCE (dropdown menu)
- Service or Group:** realaudio_control (dropdown menu)
- Action:** Pass Drop Any
- Alarm Code:** (empty text field)
- Hit Count:** 1

- 2 Change any information in the active fields.

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

END OF STEPS

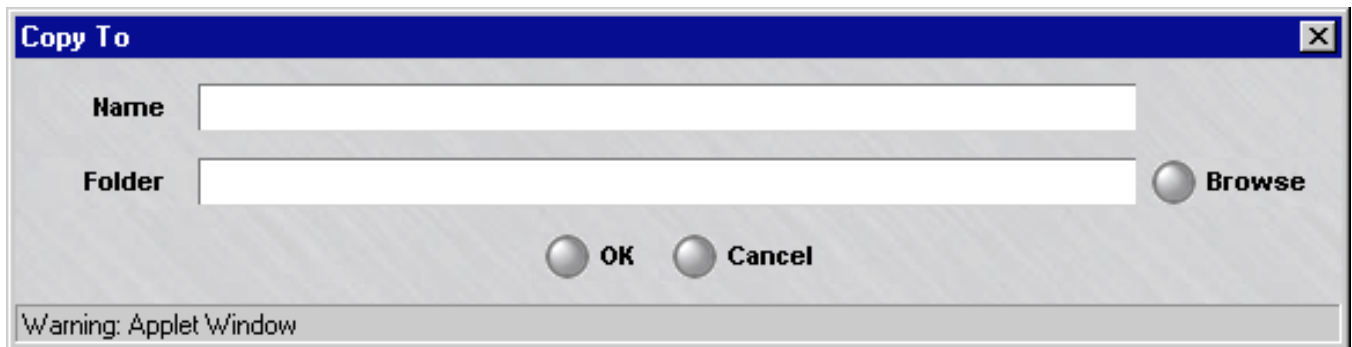
Copy a Dependency Mask

You can copy a dependency mask to a different folder in the same group, or to a folder in another group, provided you have *Edit* permission for that folder. However, you can only copy it to the folder labeled "Dependency Masks" or to a subfolder under that folder. If you try to copy it to another folder, you will get an error message.

To copy 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 copy and select **Copy** from the pop-up menu. A Copy To window will appear (see [Figure 3-8, "Copy To Window"](#) (p. 3-16)).

Figure 7-5 Copy To Window



-
- 2 In the **Name** field, enter the name you want to give the copy. If you are copying the dependency mask to the same group, you must assign the copy a new name.
 - 3 In the **Folder** field, click **Browse** and select the Dependency Masks folder you want to copy this dependency mask to.
 - 4 Click **OK** to copy the dependency mask and dismiss the Copy To window. You will be returned to the Navigator window.

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.

To move 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 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.

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

Delete a Dependency Mask

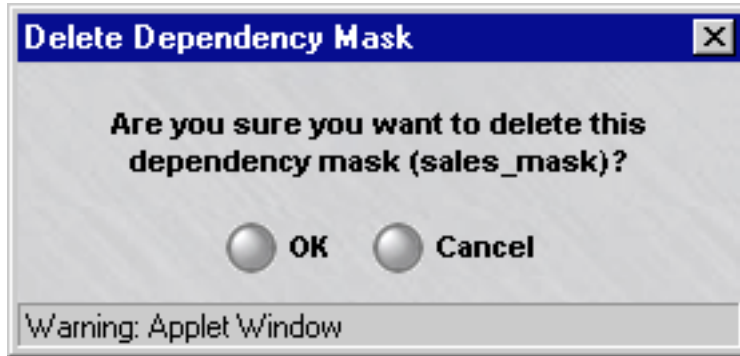
Before deleting a dependency mask, make sure that it is not currently used in a rule. If you attempt to delete a dependency mask used in a rule, you will get an error message.

To delete 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 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. 3-18) will appear.

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” \(p. 7-16\)](#) 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

The service group *realaudio_control* is included with the LSMS application. The protocol is TCP and the destination port is 7070, which is the standard RealAudio port.

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.

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

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 LSMS 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.
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 Proxies

Overview

Purpose

This chapter explains the LSMS proxy feature. This feature allows you to create rules that send HTTP, SMTP and FTP sessions to a proxy host running the Lucent Proxy Agent application.

To set up this feature on the LSMS, you have to put an entry in the Proxy Table for each Lucent VPN Firewall *Brick*[®] device that will be sending sessions to the Lucent Proxy Agent. You also have to assign a ruleset to the port on each Brick that is connected to the proxy host. This chapter explains how to do this.

This chapter also explains how to set up proxy load sharing, so the workload can be distributed to multiple proxy hosts.

Contents

How the Proxy Feature Works	8-2
How to Make an Entry in the Proxy Table	8-4
How to Maintain the Proxy Table	8-9
How to Assign a Ruleset to the Brick Port	8-11
How to Set Up Proxy Load Sharing	8-18



How the Proxy Feature Works

Overview

A Brick can transparently “reflect,” or redirect, sessions matching a particular rule to a host running the Lucent Proxy Agent for further inspection and processing. The Brick knows where the Lucent Proxy Agent is located by looking at the information in its Proxy Table.

Proxy Table

Each Brick maintains a Proxy Table, which consists of the following :

- Zone
- Service (protocol/destination port/source port)
- Proxy host
- Proxy port
- Reflection type (single | dual)
- Encryption flag
- Encryption key

See [Table 8-1, “Proxy Table without Load Sharing” \(p. 8-19\)](#) and [Table 8-2, “Proxy Table with Load Sharing” \(p. 8-19\)](#) at the end of the chapter.

When a packet matches a rule that has *Proxy* or *VPN Proxy* as its action, the Brick uses the service associated with that packet to match an entry in the Proxy Table. If an entry is found, the Brick will reflect the packet to the proxy host and port named in that entry. Otherwise, the Brick will generate an error message.

Lucent Proxy Agent

The LSMS 8.0 CD-ROM contains the Lucent Proxy Agent software. The Lucent Proxy Agent works with TrendMicro’s Interscan VirusWall (ISVW) and 8e6 Technologies’ X-Stop suite of software (XServer application) to protect your mail servers and web servers.

The software does the following:

- Scans all email sent by SMTP for known viruses, and blocks undesired SMTP commands
— and —
- Scans all HTTP sessions between a client and web server for known viruses, blocks user access to unauthorized websites, blocks undesired HTTP commands, and blocks hostile mobile code, such as Java applets and ActiveX controls.

If you want to create rules that send HTTP, SMTP or FTP sessions to the Lucent Proxy Agent for processing (and referred to as *proxy rules*), you have to install the software

on a host, configure the software, and then connect the host to the Brick that will be sending it proxied sessions.

For instructions on installing and configuring the Lucent Proxy Agent application, see the *Lucent Proxy Agent Installation and User Guide*. This manual is part of the LSMS documentation set.

Important! As of Release 8.0, the Lucent Proxy Agent comes with the LSMS application software. However, any proxy software that supports the Lucent Brick Reflection Protocol will work.

For additional information, see the technical note "Lucent Brick Reflection Protocol, Specifications and Application Examples," which can be found at this website:

www.lucent.com/security



How to Make an Entry in the Proxy Table

Proxy table entries

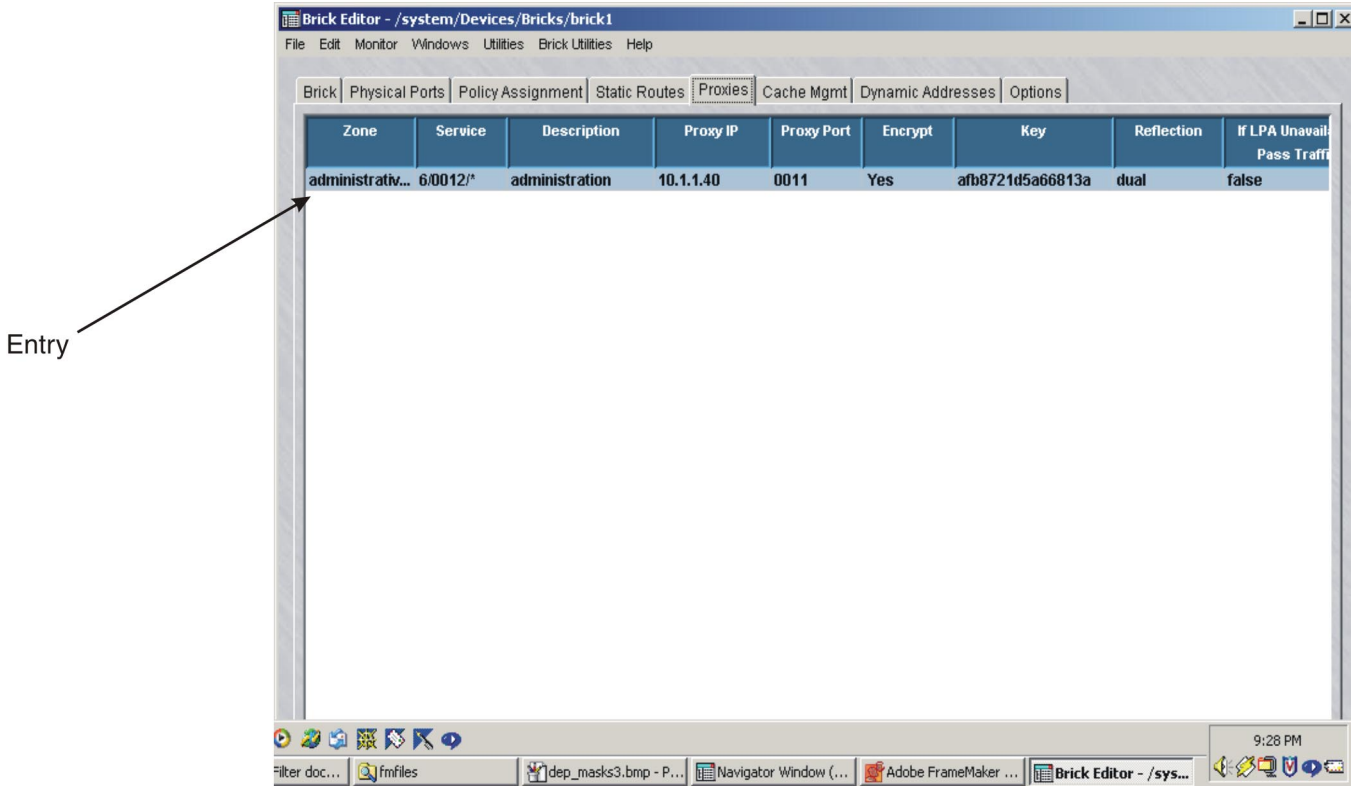
Every Brick that will be reflecting sessions to the proxy hosts must have an entry in its Proxy Table for each service (HTTP, SMTP, FTP) it will be proxying. Hence, if a Brick will be sending only HTTP sessions to the proxy hosts, only one entry in the table is required for each Brick.

To put an entry in a Brick Proxy Table, follow the steps below:

-
- 1 With the Navigator window displayed, open the appropriate group folder, and then open the Devices folder.
-
- 2 Click **Bricks** to display all configured Bricks in the Navigator window.

- 3 Double-click the Brick to display the Brick Editor, and then click **Proxies** to display the Proxies tab. A typical Proxies tab, with one entry, is shown in Figure 8-1, “Brick Editor (Proxies Tab)” (p. 8-5).

Figure 8-1 Brick Editor (Proxies Tab)



-
- 4 Right-click in the **Proxy** panel and select **New** from the pop-up menu. The Brick Proxy Editor will appear. This is the window you use to put an entry in the Proxy Table. It is shown in [Figure 8-2, “Brick Proxy Editor”](#) (p. 8-6).

Figure 8-2 Brick Proxy Editor

The screenshot shows the 'Brick Proxy Editor' dialog box with the following configuration:

- Zone:** client-zone (with a 'Browse' button)
- Protocol:** TCP
- Description:** SMTP
- Proxy Host IP Addr:** 10.10.10.112
- Proxy Port:** 2501
- Encrypt Reflection Channel?
- Key:** 1111111111111111
- Reflection Type:** dual
- Pass Traffic if LPA is Unavailable

Buttons at the bottom: OK, Cancel

-
- 5 In the **Zone** field, click **Browse** and select the Brick zone ruleset from the Browse window that appears. This is the ruleset that contains the proxy rules.
-
- 6 In the **Protocol** field, the default is TCP. Since HTTP, SMTP and FTP are TCP applications, leave the default in place.
-
- 7 In the **Dest Port** and **Src Port** fields, enter the destination ports and source ports. The destination port is usually 80 for HTTP, 25 for SMTP and 21 for FTP. The destination port is required, but the source port is optional, and you can leave the default asterisk in place.

Important! When you create rules with *Proxy* or *VPN Proxy* as the action, the service in the rules has to match the protocol, destination port, and source port you enter in [Step 6](#) and [Step 7](#).

-
- 8** In the **Description** field, you can enter an optional description of this entry. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers, and certain special characters.
-
- 9** In the **Proxy Host IP Addr** field, enter the IP address of the proxy host. This IP address is automatically added to a host group created by the LSMS called *Proxy_hosts*. This host group is used in rules found in the *proxyzone* ruleset. See “[proxyzone Rules](#)” (p. 8-13) below.
-
- 10** In the **Proxy Port** field, enter the port on which the Lucent Proxy Agent will be listening. This must be a port that is not in use. Usually a port in the 10-20,000 range works well.
- The port you enter here must be the same port that you enter as the local port when configuring the Lucent Proxy Agent. See the *Lucent Proxy Agent Installation and User Guide* for details.
-
- 11** By default, the reflection channel is authenticated and encrypted, and a key is provided, which you can change, if necessary. This key must match the key that is entered when configuring the Lucent Proxy Agent. See the chapters cited above in the *Lucent Proxy Agent Installation and User Guide*.
- To disable the key, uncheck the **Encrypt Reflection Channel** checkbox. The **Key** field will be greyed out. If your proxy host software does not support the encryption portion of the Lucent Brick Reflection Protocol, you must do this.
-
- 12** By default, the reflection type is *Dual*. To change this, select **Single** from the drop-down list in the **Reflection Type** field. If you make this change you must also deactivate rule #232 in the *proxyzone* ruleset, or the ruleset you are using in its place, and activate rule #233. See “[proxyzone Ruleset](#)” (p. 8-11) below.
- Important!** From a security standpoint, dual reflection is preferred over single. Single is less secure because the proxy host connects directly to the server without going through the Brick.
- However, single reflection is faster, and it is required to perform proxy forwarding on HTTP sessions. See the *Lucent Proxy Agent Installation and User Guide* for additional details.

.....

13 If the **Pass traffic if LPA is unavailable** checkbox is checked, packets will be passed through the Brick even if the Brick is unable to forward proxied packets to the Lucent Proxy Agent; for example, when the Brick cannot communicate with the LPA.

.....

14 Click **OK** to dismiss the Brick Proxy Editor and return to the Proxy Table.

.....

15 Display the File menu and select **Save**.

Important! An entry is automatically added to the Proxy Table for each zone that will be performing user authentication. By default, these rules are hidden. To display them, unclick the **Hide System Proxy Entries** checkbox.

.....

END OF STEPS

.....



How to Maintain the Proxy Table

When to use

Administrators are responsible for keeping the Proxy Table accurate and up to date. You can modify, duplicate or delete any entry in the table, and you can reorder the entries in the table.

Modify an Entry

To modify an entry in the Proxy Table, follow the steps below:

- 1 With the Proxies tab of the Brick Editor displayed (see [Figure 8-1, “Brick Editor \(Proxies Tab\)”](#) (p. 8-5)), double-click the entry you want to modify. The Brick Proxy Editor ([Figure 8-2, “Brick Proxy Editor”](#) (p. 8-6)) will appear with the entry populated in its fields.
- 2 Change any of the information in the entry, as necessary.
- 3 Click **OK** to dismiss the Brick Proxy Editor and return to the Proxies tab of the Brick Editor.
- 4 Display the File menu and select **Save**.

END OF STEPS

Duplicate an Entry

To duplicate an entry in the Proxy Table, follow the steps below:

- 1 With the Proxies tab of the Brick Editor displayed (see [Figure 8-1, “Brick Editor \(Proxies Tab\)”](#) (p. 8-5)), select the entry you want to duplicate and click the **Duplicate** button. The Brick Proxy Editor ([Figure 8-2, “Brick Proxy Editor”](#) (p. 8-6)) will appear with the entry populated in its fields.
- 2 Change any of the information in the entry, as necessary. If you do not change at least one field, you will have two identical entries in the Proxy Table.

-
- 3 Click **OK** to dismiss the Brick Proxy Editor and return to the Proxies tab of the Brick Editor. The duplicate you just created and edited will appear in the table.
-

- 4 Display the File menu and select **Save**.
-

END OF STEPS

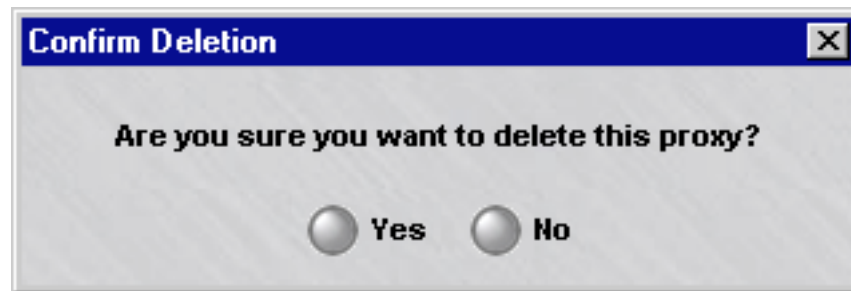
Delete an Entry

To delete an entry from the Proxy Table, follow the steps below:

- 1 With the Proxies tab of the Brick Editor displayed (see [Figure 8-1, “Brick Editor \(Proxies Tab\)”](#) (p. 8-5)), select the entry you want to delete and click the **Delete** button.

A Confirmation window similar to the one shown in [Figure 8-3, “Confirmation Window \(Proxies\)”](#) (p. 8-10) will appear.

Figure 8-3 Confirmation Window (Proxies)



-
- 2 Click **Yes** to confirm the deletion and dismiss the pop-up window. The entry will be deleted, and you will be returned to the Proxies tab of the Brick Editor.
-

END OF STEPS



How to Assign a Ruleset to the Brick Port

When to use

When an incoming or outgoing session matches a proxy rule, a Brick will send the session to the proxy host. The proxy host, using the two third-party scanning engines, will perform virus scanning, command blocking and content filtering, and then either:

- Send the session to its final destination (single reflection)
or
- Send the session back to the Brick, which then sends the session to its final destination (dual reflection).

A ruleset is required to protect the proxy hosts from attack, while at the same time permitting the required Brick-proxy host communication.

proxyzone Ruleset

A pre-configured ruleset has been provided with the LSMS specifically for the purpose of enabling Brick-proxy host communication. The ruleset is called *proxyzone*, and it has been pre-configured with all the required rules.

Figure 8-4, “proxyzone Ruleset” (p. 8-12) shows the *proxyzone* ruleset. You can use this ruleset, or you can create one of your own, as long as it includes the appropriate rules.

Figure 8-4 proxyzone Ruleset

The screenshot shows the Brick Zone Ruleset Editor window for the 'proxyzone' ruleset. The window title is 'Brick Zone Ruleset Editor - /system/Policies/Brick Zone Rulesets/proxyzone'. The menu bar includes File, Edit, Monitor, Windows, Utilities, Policy Utilities, and Help. The 'Name' field contains 'proxyzone' and the 'Description' field contains 'Zone protecting the proxies'. Below these fields is a table of rules.

Rule Number	Active	Direction	Source	Destination	Service	VLAN ID	Action
230	Yes	In	Proxy_hosts	*	*	*	Drop
231	Yes	In	Bricks_VBA	Proxy_hosts	Proxy_Listening...	*	Pass
232	Yes	Out	Proxy_hosts	Bricks_VBA	*	*	Pass
233	No	Out	Proxy_hosts	*	*	*	Pass
234	Yes	Out	Proxy_hosts	bricks	Reflection_OOB...	*	Pass
65535	Yes	Both	*	*	*	*	Drop

At the bottom of the window, there are several control buttons: a plus sign (+), a square with a plus sign, a square with a minus sign, a square with an X, a triangle pointing up, a triangle pointing down, a checkmark, and a circle with a slash. To the right of these buttons is a checkbox labeled 'Hide System Rules'.

proxyzone Rules

The table below explains the purpose of each rule found in the *proxyzone* ruleset.

Rule #	Purpose
230	<p>Drops any incoming sessions whose source has the same IP address as a host protected by <i>proxyzone</i>. This prevents an attacker from mounting a spoofing attack by pretending to be a proxy host.</p> <p>The source is a host group called <i>Proxy_hosts</i> that was created automatically by the LSMS. It contains the IP address of every proxy host. It obtains these addresses from the Proxy Table.</p>
231	<p>Allows sessions reflected by the Brick to connect to the proxy hosts.</p> <p>The source is a host group called <i>Bricks_VBA</i> that was created automatically by the LSMS. It contains the VBAs of every zone connected to the Brick. Sessions reflected by the Brick will have the zone's VBA as their source. The destination is the host group <i>Proxy_hosts</i>.</p> <p>The service is a service group called <i>Proxy_Listening_ports</i> that was created automatically by the LSMS. It contains the ports on which the proxy hosts are listening. The ports were taken from the Proxy Table.</p>
232	<p>Allows proxied sessions using dual-reflection to be reflected to the Brick.</p> <p>The source and destination are the host groups <i>Proxy_hosts</i> and <i>Bricks_VBA</i>, respectively. Sessions reflected back to the Brick will have a proxy host as the source and the zone's VBA as the destination.</p>
233	<p>Allows proxied sessions using single-reflection to be reflected to the Brick.</p> <p>Since dual-reflection is the default, this rule is inactive. If you chose Single as the reflection type in the Proxy Table, you must disable rule #232 and enable this rule.</p>

Rule #	Purpose
234	<p>Allows proxy hosts to query a Brick about reflected sessions. It passes all traffic from the proxy hosts to the Brick, using the protocol and ports set aside for reflection.</p> <p>The source is the host group <i>Proxy_hosts</i>, and the destination is a host group called <i>Bricks</i> that was created automatically by the LSMS. It contains the IP addresses of every Brick connected to the Proxy Zone.</p> <p>The service is the service group <i>Reflection_OOB_services</i> that was created automatically by the LSMS to support reflection (UDP/1024/*).</p>

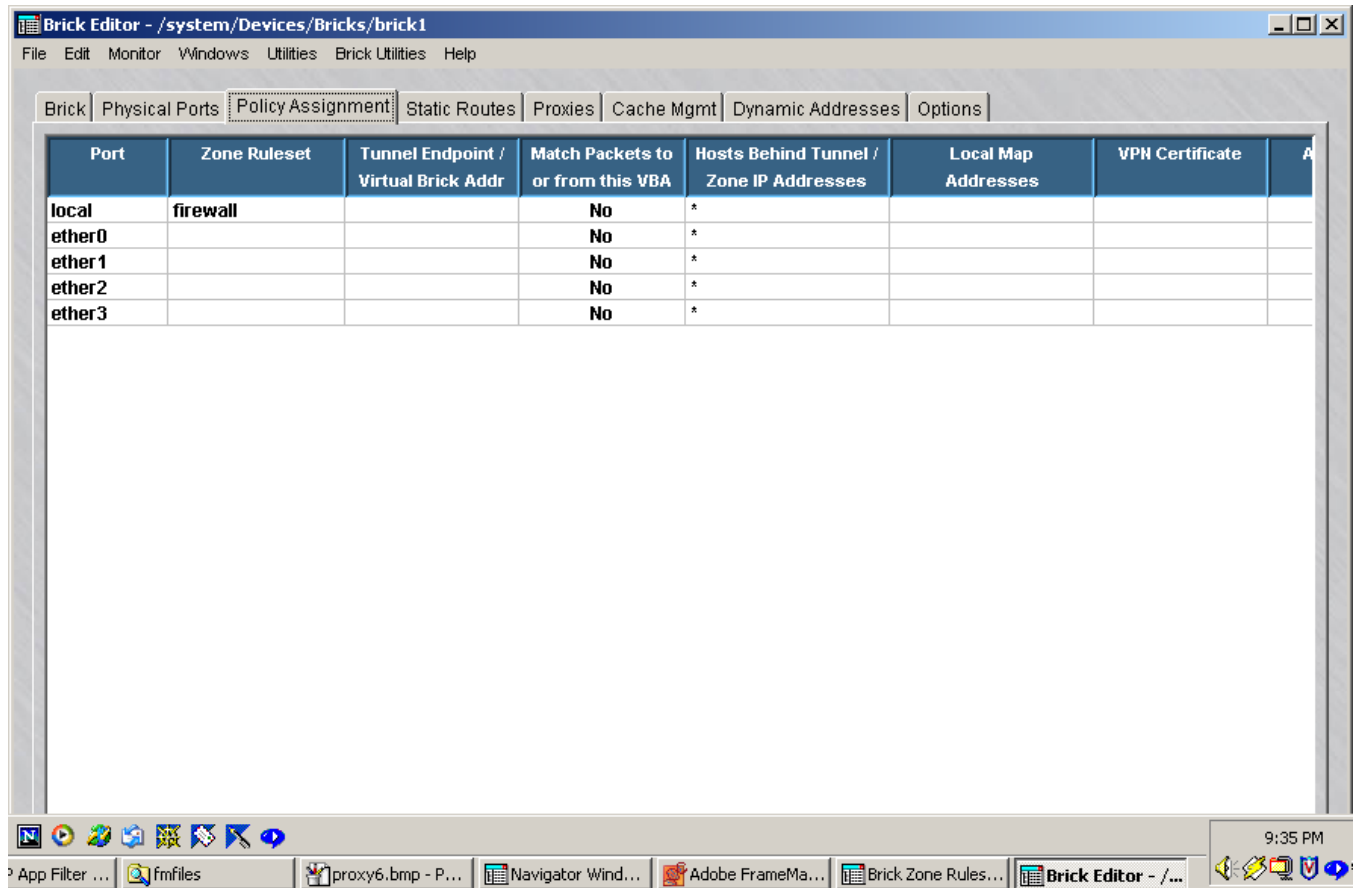
Assign the Ruleset to a Port

If you will be using the *proxyzone* ruleset, you have to assign it to the port of every Brick that will be communicating with the proxy host. If you are using a ruleset you created, assign that ruleset instead. To assign the ruleset to a port, follow these steps:

-
- 1 With the Navigator window displayed, open the appropriate group folder, then open the Devices folder.
 - 2 Click **Bricks** to display all configured Bricks in the Navigator window.
 - 3 Double-click the Brick you want to assign the ruleset to. The Brick Editor will appear.

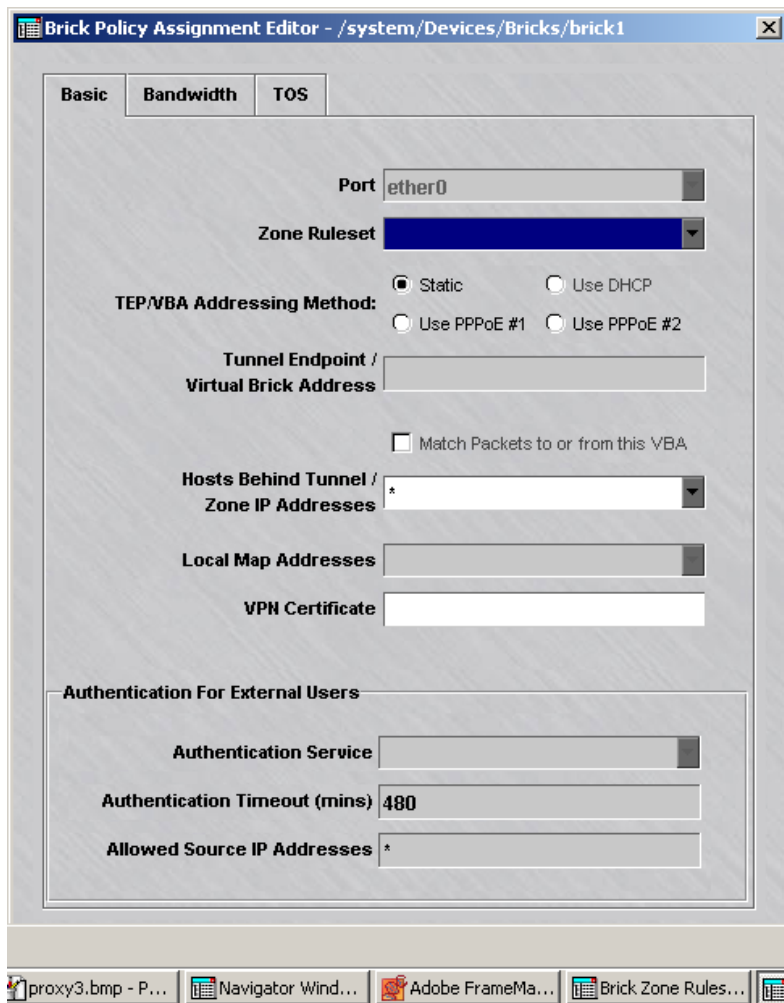
- 4 Click **Policy Assignment** to display the Policy Assignment tab. It is shown in [Figure 8-5, “Policy Assignment Tab”](#) (p. 8-15).

Figure 8-5 Policy Assignment Tab



-
- 5 Double-click the port that will be used to transmit to the proxy hosts. The Brick Policy Assignment Editor will appear. It is shown in [Figure 8-6, “Brick Policy Assignment Editor”](#) (p. 8-16).

Figure 8-6 Brick Policy Assignment Editor



-
- 6 In the **Zone Ruleset** field, display the drop-down list and select a ruleset, or select **Browse** and use the Browse window to select a ruleset.
 - 7 In the **Hosts Behind Tunnel/Zone IP Addresses** field, enter the IP addresses of the proxy hosts. The easiest way to do this is to select the *Proxy_hosts* host group from the drop-down list. This is the host group that was created by the LSMS and automatically populated with the IP addresses of all proxy hosts, taken from the Proxy Table.
-

There are several other ways to enter this information:

- Display the drop-down list and select the asterisk. If the port is dedicated to the proxy hosts, the asterisk is acceptable.
- Type the IP addresses directly into the field. You can enter a single IP address, a range of IP addresses, in the format 10.1.1.1-10.1.1.10, or an IP address with subnet mask (for example, 123.123.123.123/24).
- Display the drop-down list, select **Browse**, and then select a host group you created for this purpose. See [“Overview” \(p. 2-1\)](#).

-
- 8** None of the other fields are applicable to this ruleset, so click **OK** to dismiss the Brick Policy Assignment Editor and return to the Policy Assignment tab.

-
- 9** Display the File menu and select **Save**.

END OF STEPS



How to Set Up Proxy Load Sharing

When to use

The LSMS allows Administrators to set up proxy load sharing, so that the workload can be shared among multiple proxy hosts. Without this feature, for example, all HTTP traffic in a given zone on a particular Brick, if proxied, would have to go through the same proxy host.

For some customers and network loads, this is not desirable. In addition, if that single proxy host goes down, or is unreachable from the Brick for whatever reason, no proxied traffic will pass.

Round-Robin Distribution

Proxy load sharing distributes the load on a round-robin basis to any number of configured proxy hosts. To set up load sharing, an Administrator has to put multiple entries in the Proxy Table for a given combination of Brick, zone, and service. This allows the Brick to redirect traffic to more than one IP address/port for this combination. The multiple proxy entries are referred to as a *load sharing group*.

You can put duplicate entries in a load sharing group. If you do this, the Brick will include this proxy host twice (or as many times as you put the entry in the Proxy Table) in the load sharing round robin. This means this proxy host will receive more traffic than proxy hosts with a single entry in the load sharing group.

Heartbeat

Every proxy host sends heartbeats using the reflection out-of-band channel to the Brick to which it is homed. The Brick uses these heartbeats to determine which, if any, proxy hosts are available in a load sharing group.

If the Brick stops receiving heartbeats from a proxy host in a load sharing group, it no longer uses that host for proxying. When the Brick begins receiving heartbeats again, it resumes sending proxied traffic to the host.

Third-party proxy software does not support the heartbeat. Therefore, any proxy hosts using software other than the Lucent Proxy Agent cannot be included in a load sharing group.

Examples

[Table 8-1, “Proxy Table without Load Sharing” \(p. 8-19\)](#) shows a typical Proxy Table *without* proxy load sharing. The table has two proxy host addresses (10.1.1.1 and

10.1.1.2), but only one for each zone/service combination. Hence, no load sharing is taking place.

Table 8-1 Proxy Table without Load Sharing

Zone	Service	Proxy Host	Proxy Port	Reflection Type	Encrypt. Flag	Encrypt.Key
Inside	6/80/*	10.1.1.1	8080	Dual	Yes	0123456789abcdef
Inside	6/25/*	10.1.1.2	8025	Single	No	abc-def0123456789
Out-side	6/80/*	10.1.1.1	8090	Dual	Yes	0123456789abcdef

Table 8-2, “Proxy Table with Load Sharing” (p. 8-19) shows the same Proxy Table after load sharing has been introduced. As you can see, for certain zone/service combinations (Inside/HTTP;Outside/HTTP), there is a second entry in the table. Hence, load sharing will take place.

Note that the *Reflection Type* and *Encryption Flag* in the second entry have to match the first, but the *Encrypt. Key* can be different.

Table 8-2 Proxy Table with Load Sharing

Zone	Service	Proxy Host	Proxy Port	Reflection Type	Encrypt. Flag	Encrypt.Key
Inside	6/80/*	10.1.1.1	8080	Dual	Yes	0123456789abcdef
Inside	6/80/*	12.1.1.1	8081	Dual	Yes	fedcba9876543210
Inside	6/25/*	10.1.1.2	8025	Single	No	abcdef0123456789
Out-side	6/80/*	10.1.1.1	8090	Dual	Yes	0123456789abcdef
Out-side	6/80/*	12.1.1.1	8091	Dual	Yes	fedcba9876543210

□

9 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 a Lucent VPN Firewall *Brick*[®] device 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?	9-2
How Authentication Works	9-4
Before Setting Up User Authentication	9-8
To Set Up Local Password Authentication	9-11
How to Set Up RADIUS and SecurID Authentication	9-19
How to Set Up VPN Certificate Authentication	9-29



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 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 Lucent IPsec Client application is provided on a separate CD-ROM that comes with the LSMS. For instructions on installing and using the client application, see the *Lucent IPsec Client Installation and User Guide*.

Types of Authentication

The LSMS supports four types of authentication:

- **Local Password**

The simplest way to perform authentication is to set up your own user database on the LSMS. 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 LSMS will verify the ID and password based on information in the database (see [Figure 9-1, "User Editor" \(p. 9-12\)](#)).

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 LSMS, and have these users authenticated by RADIUS, you can do that, too. In either case, the LSMS will contact the RADIUS server after the user has requested authentication and supplied an ID and password (see [Figure 9-1, "User Editor" \(p. 9-12\)](#)).

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 LSMS, and have these users authenticated by SecurID, you can do that, too. In either case, the LSMS will contact the ACE/Server after the user has requested authentication and supplied an ID and password (see [Figure 9-1, "User Editor" \(p. 9-12\)](#)).

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.

LSMS 8.0 supports digital certificates. These can be obtained from certification authorities such as Entrust.

You can only use a VPN certificate to authenticate client users.

□

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 LSMS, initiating the authentication process.

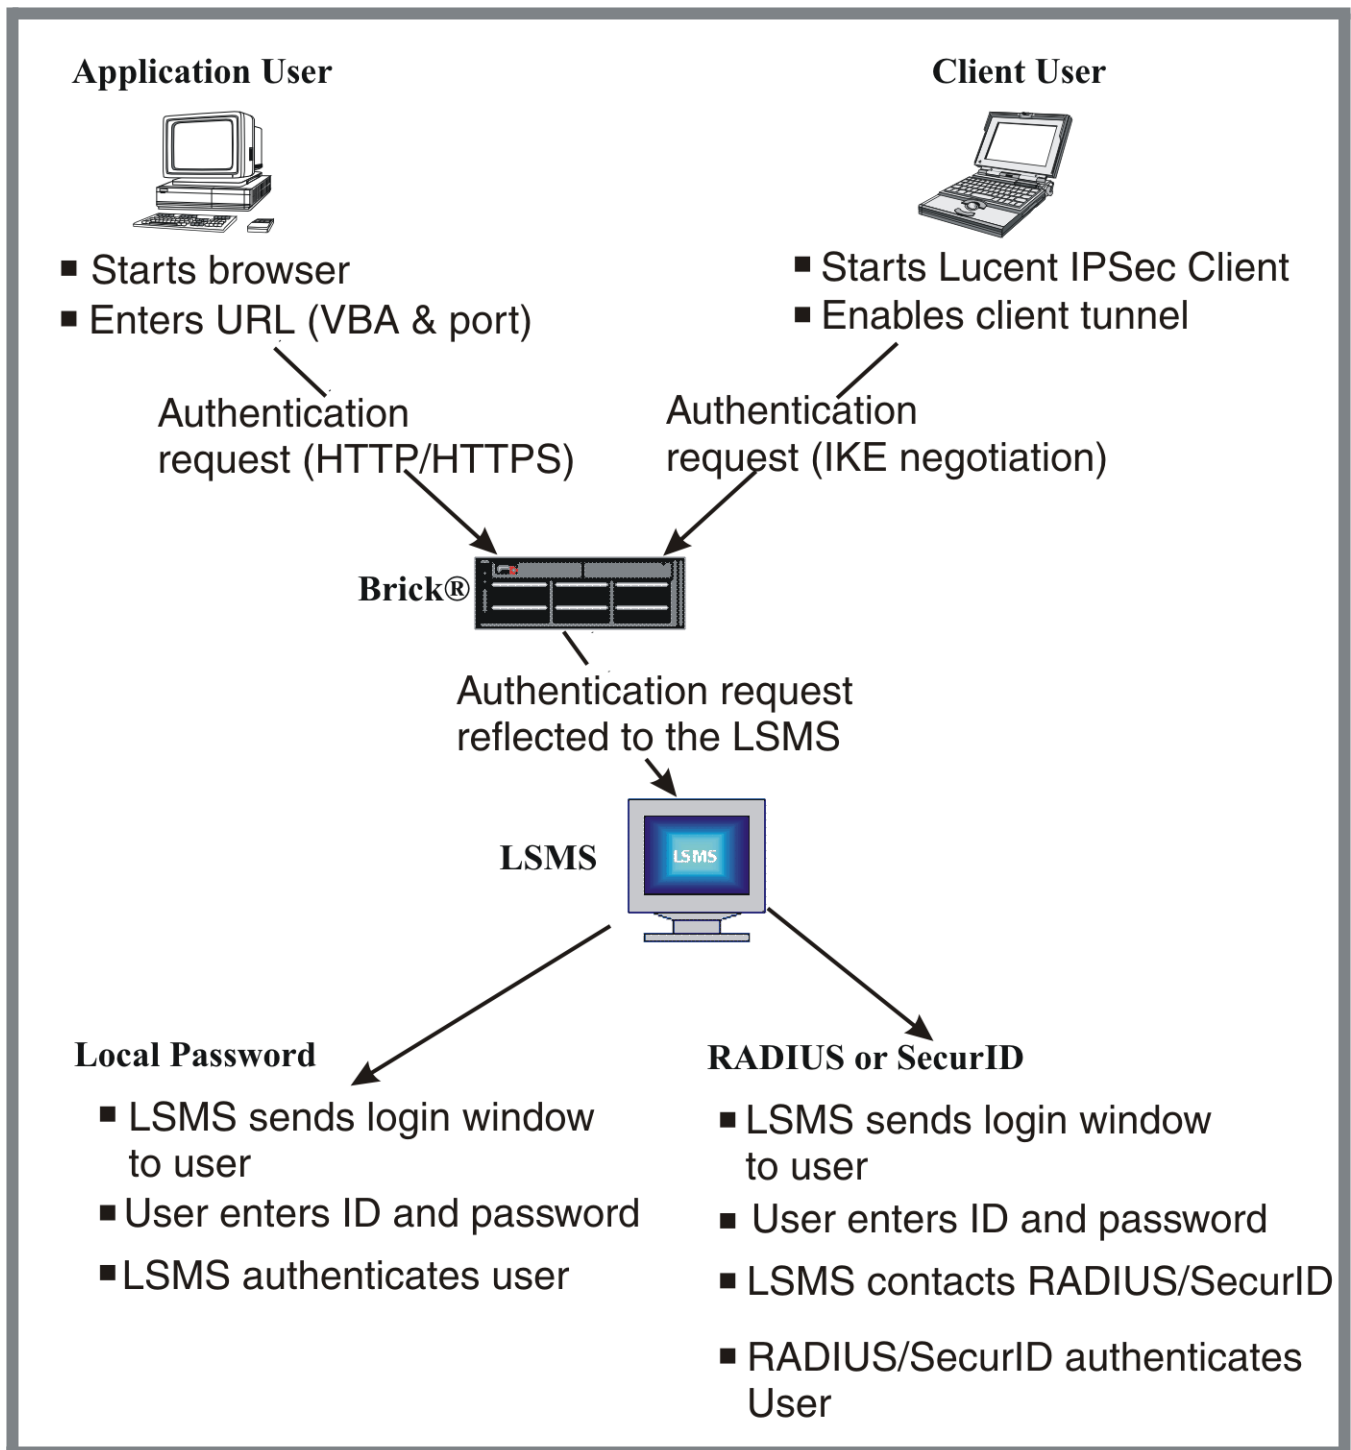
Login Procedure

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 LSMS 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 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 LSMS application includes an authentication server that processes incoming authentication requests from application users only.

During installation of the LSMS 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. See *Chapter 10. Using the Configuration Assistant in the LSMS Administration Guide* for instructions.

If the authentication server is HTTPS, you will need a digital certificate. See [Chapter 10, “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 LSMS to listen for authentication requests (see [Figure 9-1, “User Editor”](#) (p. 9-12)). 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 LSMS, 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, "Lucent VPN Firewall Brick® Device 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 LSMS to communicate with the RADIUS or SecurID ACE/Server. You also must add several entries to the *services* file on the LSMS.

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 LSMS.

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 LSMS, you have to add a rule to the NOC Gateway or Administrative Zone to permit the LSMS 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 LSMS 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 LSMS.

- 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 LSMS 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 LSMS. 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 LSMS 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 LSMS:

<LSMS_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.

To add the appropriate entries to the *services* file, follow the steps below:

- 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 LSMS 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 LSMS database, and are verified by the LSMS.

If you choose to set up **Local Password** authentication, you do not have to create an authentication service, one has been provided with the LSMS 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 (the default) via the Configuration Assistant, password restrictions that comply with Sarbanes-Oxley (SOX) requirements would be applied for local password creation and authentication. In this case, the 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

For more information about setting the strong password (SOX compliance) option via the Configuration Assistant, refer to the *Using the Configuration Assistant* chapter in the *LSMS Administration Guide*.

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 9-1, “User Editor”](#) (p. 9-12).

Figure 9-1 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 controls:

- User ID**: A text input field.
- Enable User**: A checked checkbox.
- Full Name**: A text input field.
- Description**: A wide 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 to its 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 9-2, "Password Panel" \(p. 9-13\)](#).

Figure 9-2 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” \(p. 12-1\)](#).

- 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 9-2, “Password Panel” \(p. 9-13\)](#)).

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.

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

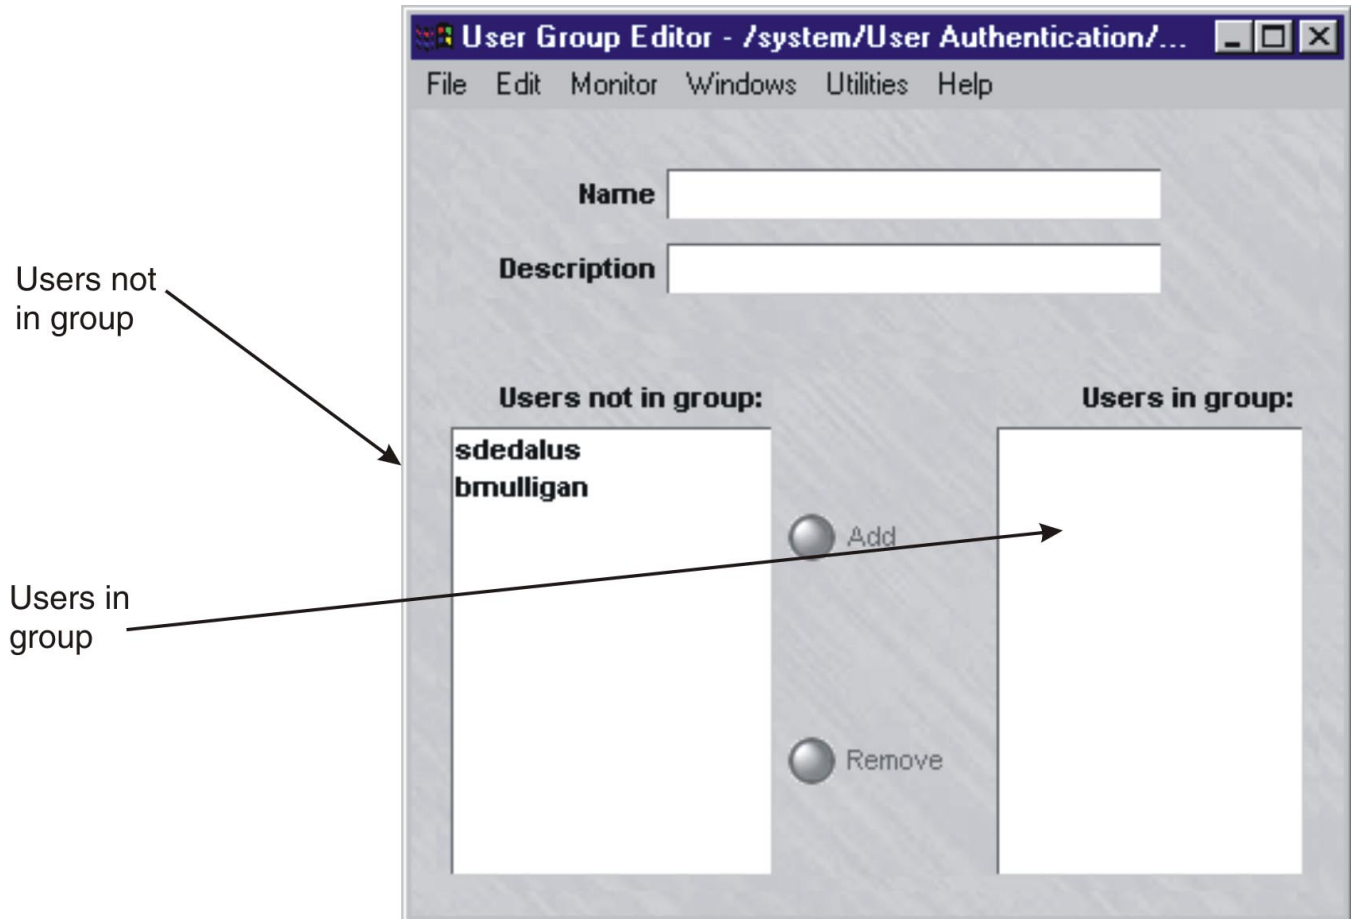
Create a User Group

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 9-3](#), “User Group Editor” (p. 9-16).

Figure 9-3 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

LSMS-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 LSMS automatically creates four user groups, which are described in [Table 9-1, “LSMS-defined user groups” \(p. 9-17\)](#). 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 9-1 LSMS-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.

Table 9-1 LSMS-defined user groups (continued)

User group	Description
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>



How 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 LSMS 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 LSMS, and the rest on a given external server.

- **LSMS Host**

If the database will reside on the LSMS 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. 9-11)). 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 (“Create a User Group” (p. 9-15)).

- **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 LSMS-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 LSMS database. However, to avoid replicating your entire external database on the LSMS, 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 LSMS will try the default authentication service for that zone.

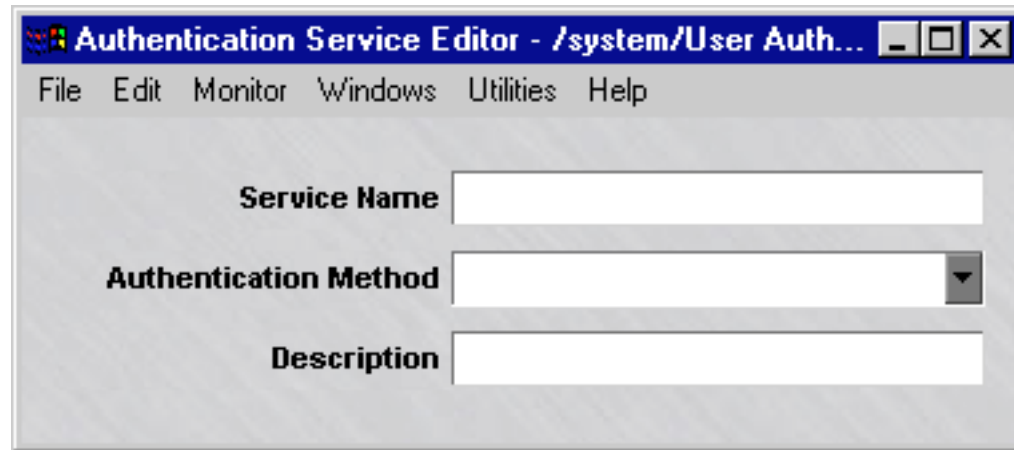
Create a RADIUS Authentication Service

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. The Authentication Service Editor will appear. It is shown in “Rule” (p. 6-13).

Figure 9-4 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.

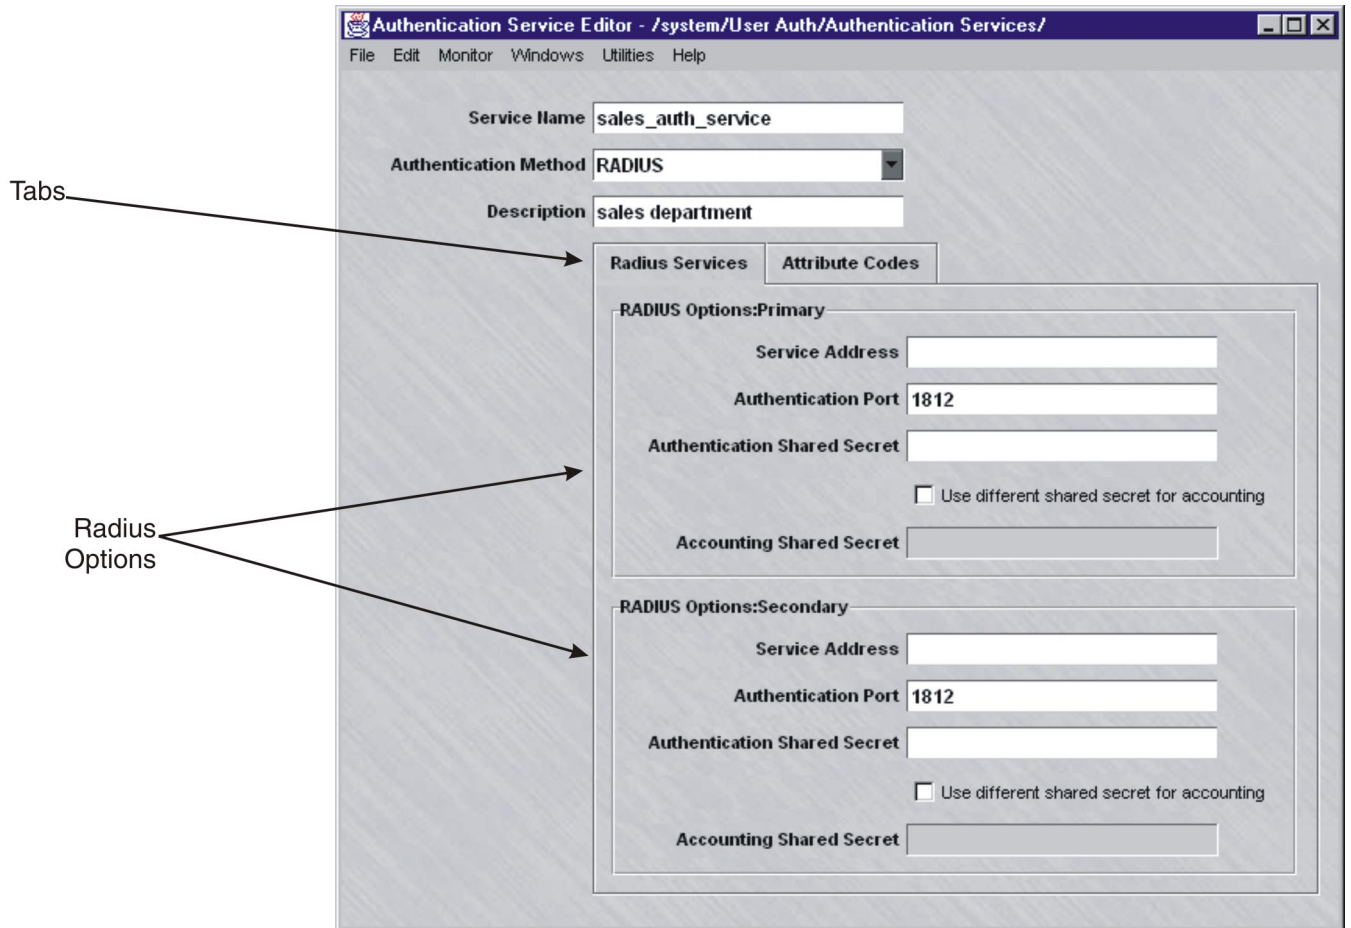
-
- 3 In the **Description** field, you can enter an optional description of this authentication service. The description can contain up to 80 characters. It can consist of upper and lower case letters, numbers, and certain special characters.

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 two tabs (**Radius Services** and **Attribute Codes**).

The Radius Services tab, which includes two **RADIUS Options** boxes, is the tab that is displayed first (see [Figure 9-5, “Authentication Service Editor \(RADIUS Service Tab\)”](#) (p. 9-21)).

Figure 9-5 Authentication Service Editor (RADIUS Service Tab)



5 In the **RADIUS Options: Primary** box, enter the following information:

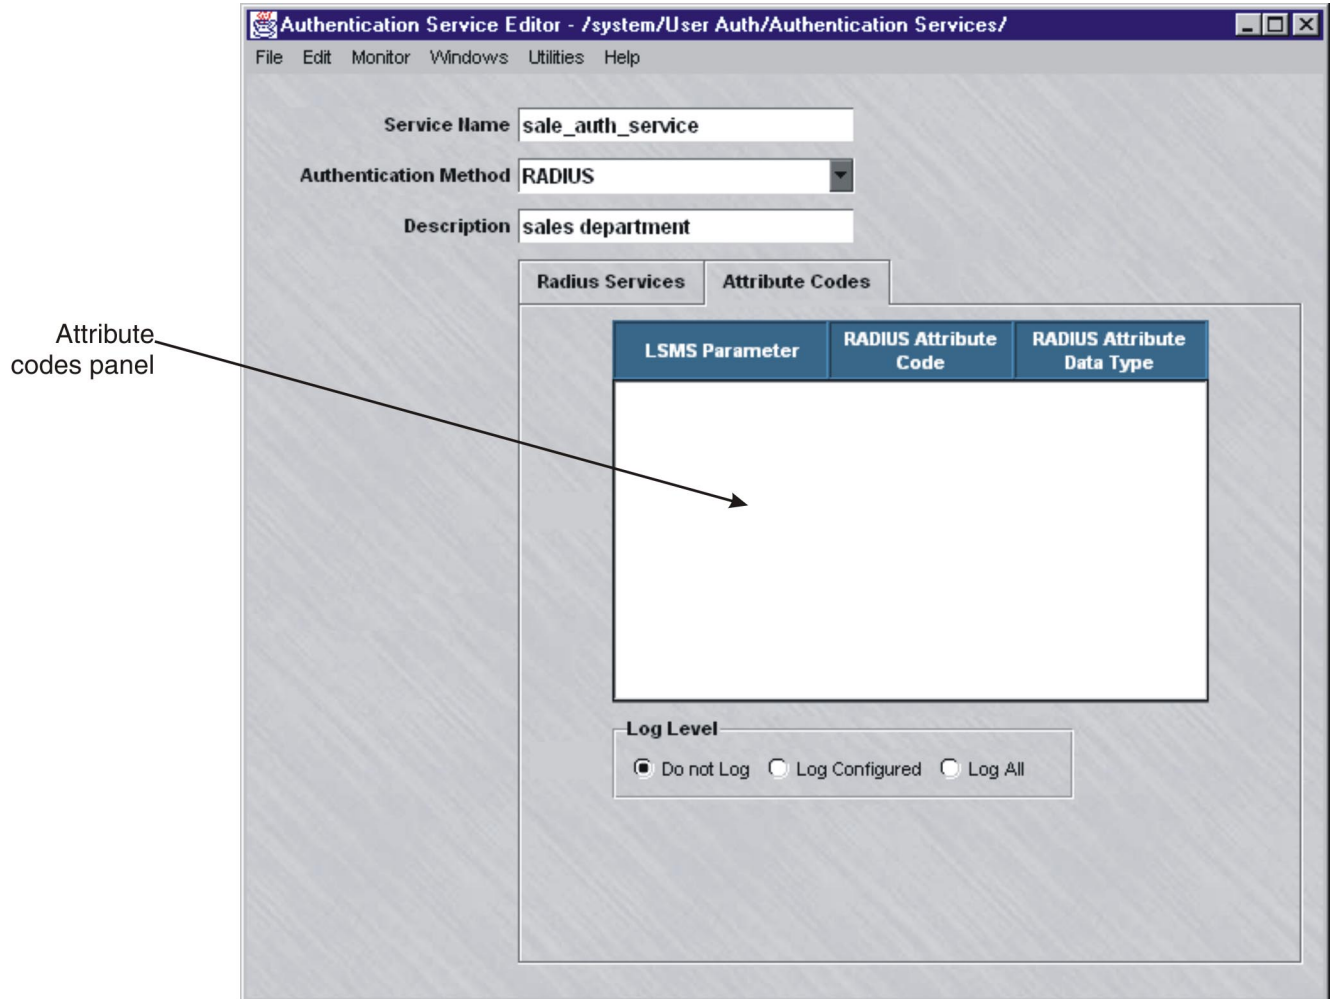
Field	Explanation
Service Address	The IP address of the RADIUS server
Authentication Port	The number of the port on which the RADIUS authentication service is listening. The default is 1812, but you can change this if necessary.

Field	Explanation
Authentication Shared Secret	<p>The Authentication Shared Secret is the string used to encrypt communication between the LSMS and the RADIUS server.</p> <p>The same value must be entered on both servers. The value is case-sensitive.</p>
Accounting Shared Secret	<p>Typically, the Accounting Shared Secret is the same as the Authentication Shared Secret.</p> <p>However, if the RADIUS server is configured with a different Accounting Shared Secret, you can set the shared secret for the accounting port here.</p> <p>To activate this field, click the box labeled Use different shared secret for accounting.</p>

-
- 6 If your environment includes a secondary RADIUS server, you can enter the same information for this server in the **RADIUS Options: Secondary** box. If the primary RADIUS server cannot be reached, the LSMS will try to reach the secondary RADIUS server.
-
- 7 If you want the LSMS to use certain information contained in the RADIUS protocol messages that the RADIUS server sends to the LSMS when authenticating users, click

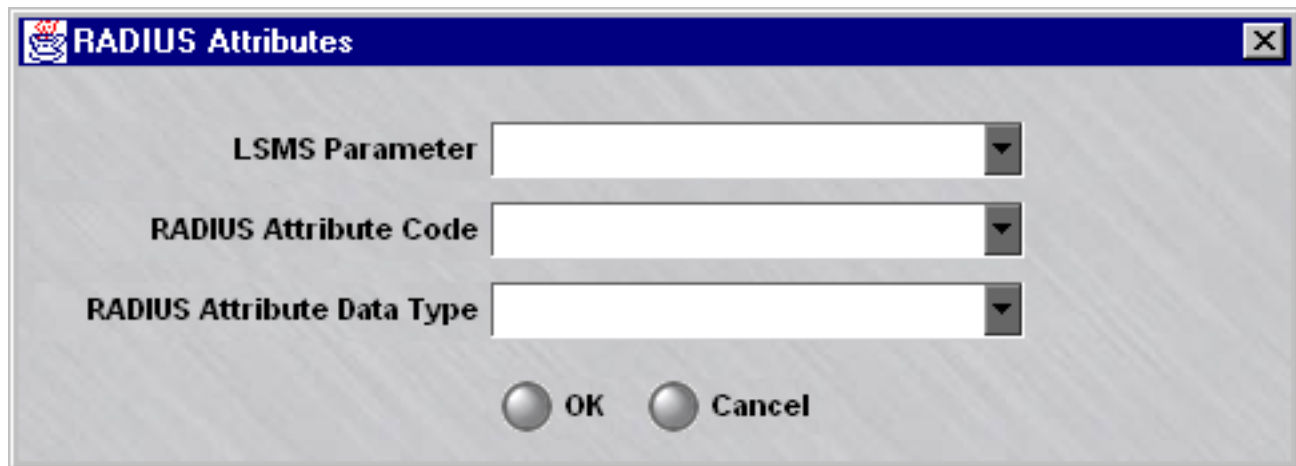
Attribute Codes to display the Attribute Codes tab. It is shown in [Figure 9-6](#), “Authentication Service Editor (Attribute Codes Tab)” (p. 9-23).

Figure 9-6 Authentication Service Editor (Attribute Codes Tab)



- Right-click in the Attribute Codes panel and select **NEW** from the pop-up menu. The RADIUS Attributes window will appear. It is shown in [Figure 9-7, “RADIUS Attributes Window”](#) (p. 9-24).

Figure 9-7 RADIUS Attributes Window



- In LSMS 8.0, it is possible to configure the LSMS to use some of the information contained in the RADIUS protocol messages that the RADIUS server sends to the LSMS when authenticating users.

In the LSMS **Parameter** field, select the LSMS parameter that corresponds to the RADIUS attribute code you want to configure on the RADIUS server. Then, select the appropriate RADIUS attribute code and data type.

This is explained in greater detail in [Appendix B, “Pre-Configured Lucent VPN Firewall Brick® Device Zone Rulesets”](#) and [Appendix D, “RADIUS Attributes”](#).

- 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

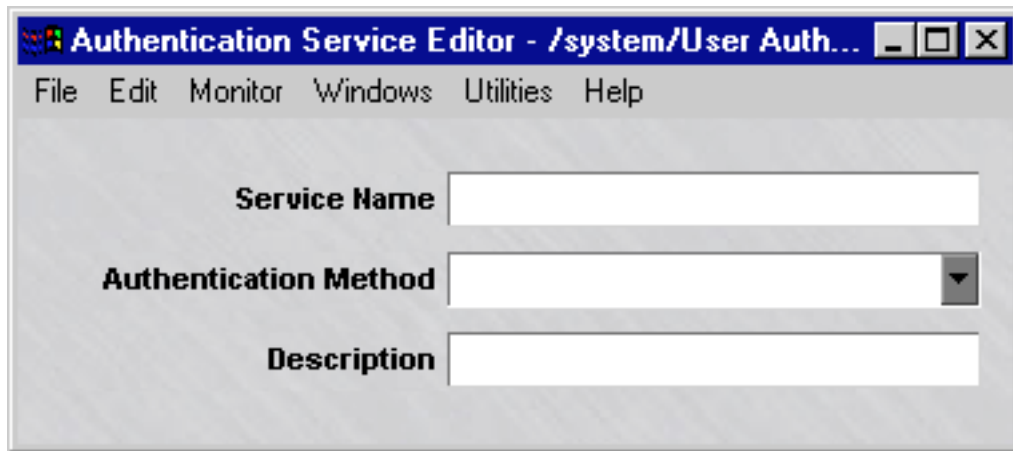
END OF STEPS

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 9-8, "Authentication Service Editor"](#) (p. 9-25).

Figure 9-8 Authentication Service Editor

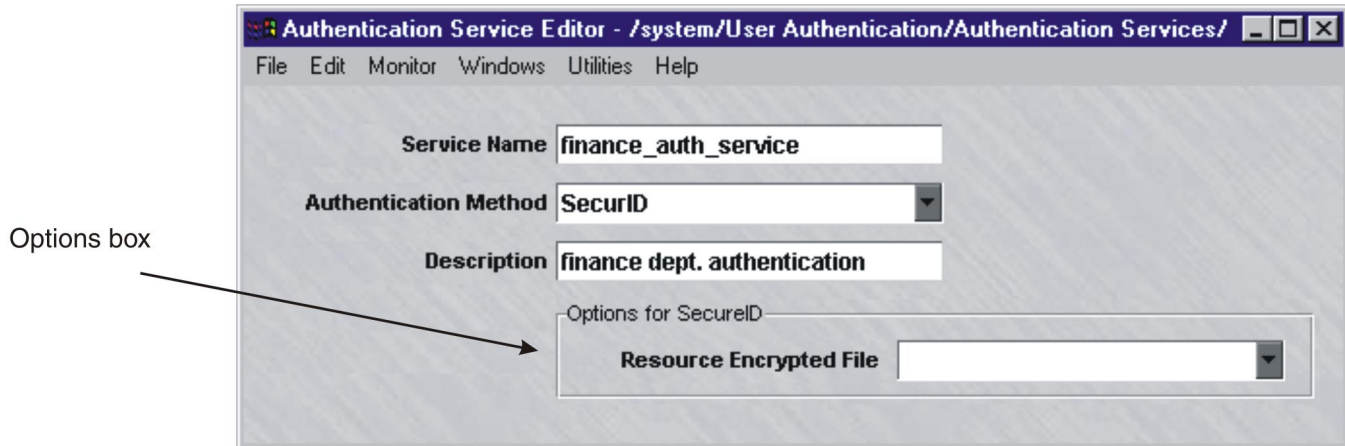


- 2 In the **Service Name** field, enter a unique name to identify this authentication service. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers, and certain special characters.
- 3 In the **Description** field, you can enter an optional description of 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 9-9, “SecurID Options”](#) (p. 9-26).

Figure 9-9 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 LSMS 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 LSMS:

<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 LSMS 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



How to Set Up VPN Certificate Authentication

When to use

You can use a digital certificate to authenticate VPN client users only (digital certificate authentication is not available for external users). To do this, you first have to obtain a digital certificate from a certification authority (CA) and install it on the LSMS host. The procedure for doing this is described in [Chapter 10, “Digital Certificates”](#).

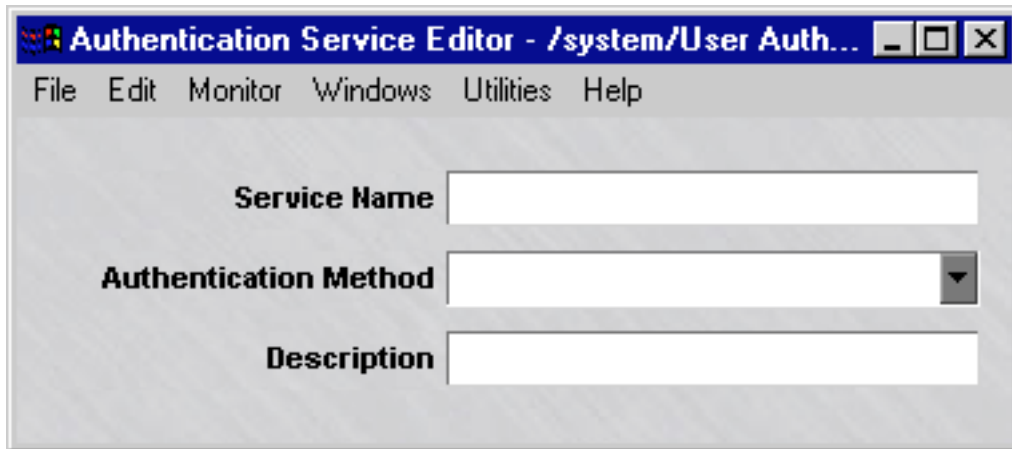
Once you have obtained and installed the certificate, you have to create a VPN certificate authentication service, and then you have to assign it to a client tunnel endpoint.

The following explains how to create the authentication service. The procedure for entering the authentication service is the same as for RADIUS or SecurID ([“Create a SecurID Authentication Service” \(p. 9-25\)](#)).

Task

- 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” \(p. 9-11\)](#).

Figure 9-10 Authentication Service Editor



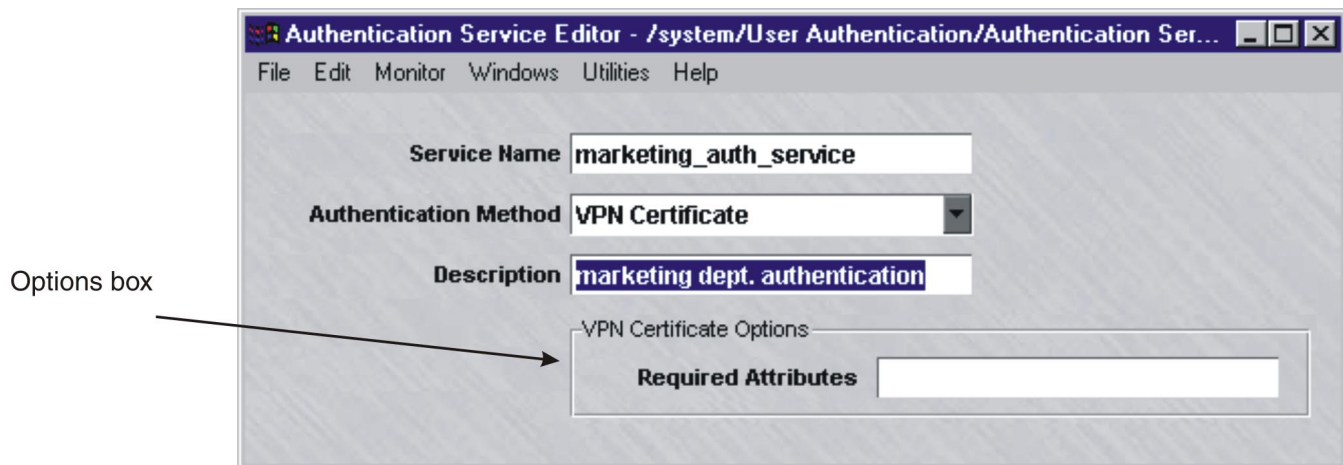
- 2 In the **Service Name** field, enter a unique name to identify this authentication service. The name can contain up to 44 characters. It can consist of upper and lower case letters, numbers, and certain special characters.

-
- 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 9-11, “VPN Certificate Options”](#) (p. 9-30).

Figure 9-11 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 9-12, "Attributes" \(p. 9-31\)](#) shows an example of attributes entered in the **Required Attributes** field.

Figure 9-12 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

.....



10 Digital Certificates

Overview

Purpose

This chapter explains how to obtain and install digital certificates. Digital certificates are used:

- To configure a secure (HTTPS) LSMS 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?	10-2
What is the Certificate Manager?	10-3
How to Start the Certificate Manager	10-5
How to Obtain and Install a Server Certificate	10-7
How to Obtain and Import a VPN Certificate	10-12
Group Assignment for VPN Certificates	10-19
Pending Certificate Signing Requests	10-23
Server Versus VPN Certificate Signing Request	10-24



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.

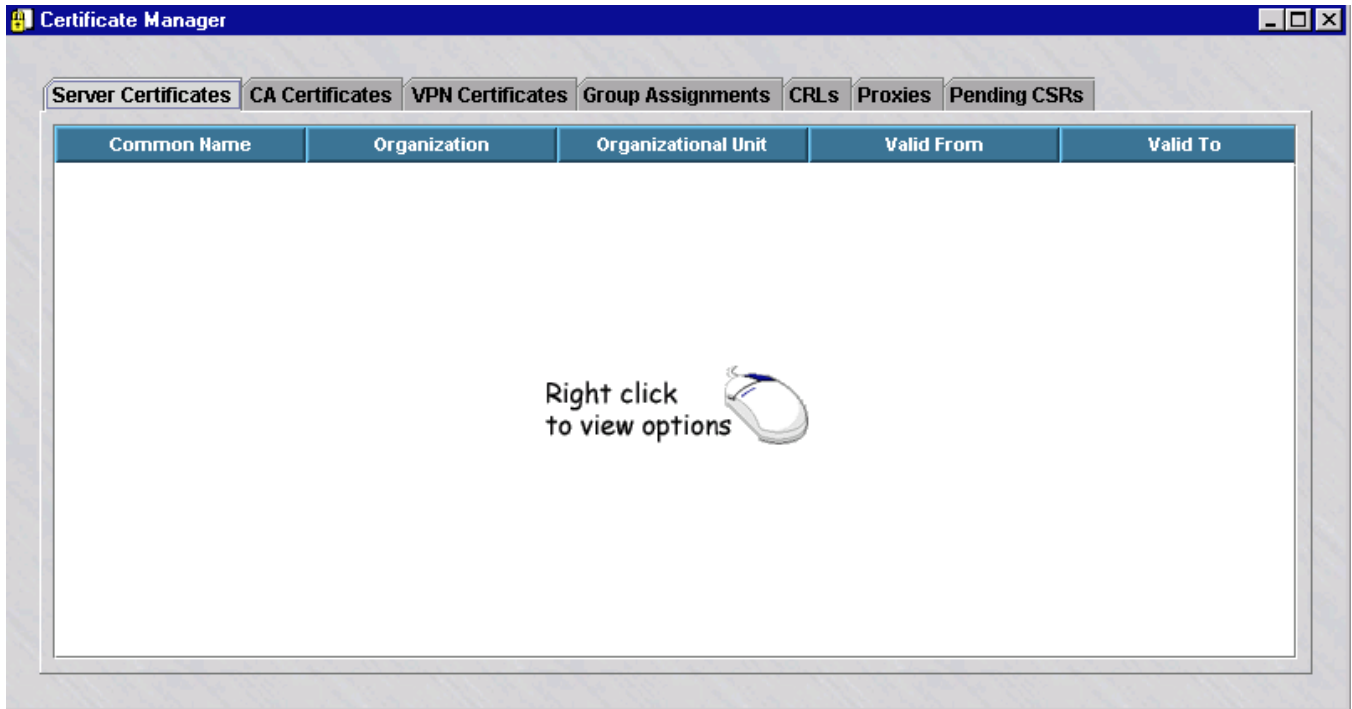


What is the Certificate Manager?

Definition

Certificate Manager is the name given to the LSMS utility that is used to obtain and import digital certificates. An example is shown in [Figure 10-1, “Certificate Manager”](#) (p. 10-3)

Figure 10-1 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 LSMS 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 LSMS 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 LSMS 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 LSMS yet).

All digital certificates that have been obtained and installed will be shown in the Certificates Panel.



How to Start the Certificate Manager

Overview

The Certificate Manager is run locally on the LSMS host. If the LSMS is installed on the Solaris operating system, you can also access the Certificate Manager remotely using a remote host that supports X Windows.

The following explains how to start the Certificate Manager on a Windows LSMS, on a Solaris LSMS, and from a remote host that supports X Windows.

Start the Certificate Manager on Windows

- 1 To start the Certificate Manager on a Windows host, click **Start** on the Windows taskbar and select:

Programs ► Lucent Security Management Server ► Utilities ► Certificate Manager

END OF STEPS

Start the Certificate Manager on Solaris

To start the Certificate Manager on a Solaris workstation, follow the steps below:

- 1 Open an X-term window and enter:

```
<LSMS_install_directory>/certmgr &
```

where *<LSMS_install_directory>* is the directory in which the LSMS application was installed (*/opt/isms/lmf* if you used the default directory during the installation).

This makes the installation root directory the Present Working Directory.

- 2 Enter:

```
./certmgr
```

to display the Certificate Manager window (see [Figure 10-1, “Certificate Manager” \(p. 10-3\)](#)).

END OF STEPS

Start the Certificate Manager on a Remote Host

If the LSMS is installed on the Solaris operating system, you can start the Certificate Manager and display it on a remote host that supports X windows. The following explains how to do this:

-
- 1 On the remote host, open an X-term window and enter

```
xhost + <LSMS_IP_address>
```

This will allow the LSMS to open a Certificate Manager client window on the remote host's X server.

-
- 2 Telnet to the LSMS and enter

```
DISPLAY=<Remote_Host_IP_Address>:0.0  
export DISPLAY  
<LSMS_install_directory>/certmgr &
```

This will open the Certificate Manager window on the remote host.

END OF STEPS



How to Obtain and Install a Server Certificate

Overview

Server certificates are used when the LSMS 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 LSMS
- Obtain and import the Certificate Authority's CA certificate(s) on the LSMS

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
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
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 LSMS web server, this must resolve to the IP address of the LSMS 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

-
- 4** When you have finished entering this information, click **OK**.

A private/public key pair will now be generated. The public key will be wrapped in the CSR and written to the file specified in Step 3.

.....

END OF STEPS

.....

Obtain the Server Certificate

Once the CSR has been created and saved in a file, 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. 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 sever certificate. See ["Import the Server Certificate"](#) (p. 10-10) 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 LSMS. 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 LSMS, 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 LSMS.*

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

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 LSMS. 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 LSMS. 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 LSMS software and are already imported on the LSMS. 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
.....



How to Obtain and Import a VPN Certificate

Overview

You can use X.509 digital certificates to authenticate 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 LSMS
- Obtain and import the Certificate Authority's CA certificate(s) on the LSMS

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**.

A private/public key pair will now be generated. The public key will be wrapped in the CSR and written to the file specified in Step 4.

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:

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.

END OF STEPS

Import the VPN Certificate

Once the CA issues your VPN certificate, you have to import it on the LSMS. 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 LSMS. If you have VPN certificates in Entrust EPF format, follow the procedure in the [“How to Convert an Entrust EPF File to PKCS12”](#) (p. 10-17) 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 LSMS, 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 LSMS.*

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 LSMS. 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 LSMS. 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 LSMS software and are already imported on the LSMS. 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 LSMS CD to convert the files to PKCS12 format.

Copy the directory *Lucent-EPF-Conversion* and its contents from the LSMS 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 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 LSMS 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.

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



Group Assignment for VPN Certificates

Overview

VPN certificates must be assigned to an LSMS 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 LSMS
- 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 LSMS. 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 LSMS automatically retrieve the updated CRL for you. When CRL Update is configured, the LSMS 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.

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

Configure Proxies for CRL Updates

If you have configured Automatic CRL Updates, you may have to configure a proxy to allow the LSMS to contact the Certificate Authority web site. Each LSMS (primary, secondary, and compute server) will perform CRL updates once a day for the Bricks that are homed to that LSMS. If there is an HTTP Proxy server between the LSMS 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 LSMS 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 LSMS 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 LSMS 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.



11 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?	11-2
Set LAN-LAN Tunnel Defaults	11-5
Set Up a LAN-LAN Tunnel	11-12
Set Up a LAN-LAN Tunnel with UDP Encapsulation	11-23
Redundant LAN-LAN Tunnels	11-25
Maintain LAN-LAN Tunnels	11-26
How to Set Up Service Level Agreements	11-30



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

Lucent Virtual Firewall *Brick*[®] devices and unmanaged devices can all serve as the endpoint of a LAN-LAN tunnel. The drawing in [Figure 11-1, “Types of LAN-LAN Tunnel Endpoints”](#) (p. 11-2) illustrates the different types of endpoints.

Figure 11-1 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. See *Chapter 4. Configuring Lucent VPN Firewall Brick[®] Device Ports* in the *LSMS 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 LSMS. It could be a Brick that is managed by another LSMS, another Group Administrator, another vendor router or firewall, or a non-Lucent IPSec client application.

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. LSMS 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 SAs is the Internet Key Exchange (IKE) process.
Default parameters for automatic key exchange are provided with the LSMS 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 Security Associations (SAs), which contain the encryption/authentication types and keys, manually at both endpoints. This has to be done extremely carefully, 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.

IKE on the Brick

In previous releases, the LSMS performed proxy IKE for Client and LAN-LAN VPN tunnels. With the IKE on the Brick feature, introduced in LSMS V8.0, the IKE process of establishing Security Associations (SAs), containing authentication and encryption information, at both endpoints of the tunnel is handled locally on the Brick instead of through the LSMS. Moving the IKE negotiation process to each Brick makes the tunnels setup process quicker and more reliable, and frees up computing resources in the LSMS.

Internet key exchange version 2 (IKEv2)

The LSMS and managed Bricks, as of Release 9.1, 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 some Brick models (50, 150, and 1100A). This greatly increases the number of tunnels per second that can be established on these devices as well as reducing the overall load on the Bricks.

LAN-LAN tunnel defaults

Before you set up a LAN-LAN tunnel, you should examine the LAN-LAN tunnel defaults provided with the LSMS application to determine whether or not any of them need to be modified.

The LSMS 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-Lucent device.

□

Set LAN-LAN Tunnel Defaults

When to use

Use this task to configure or modify the default parameter settings for LAN-LAN tunnels.

Task

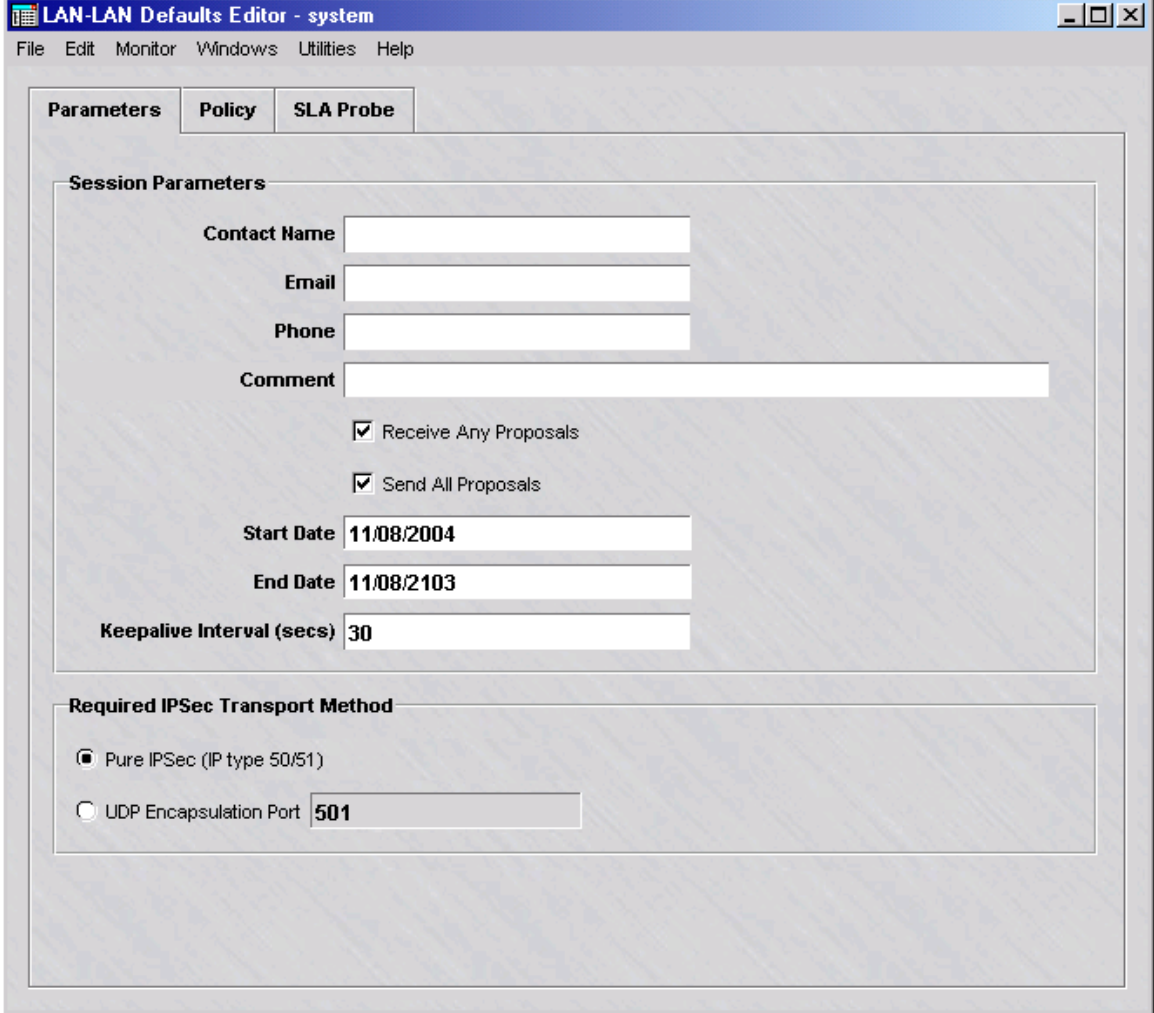
Complete the following step to display and/or modify defaults for LAN-LAN tunnels:

- 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 11-2, “LAN-LAN Defaults Editor \(Parameters Tab\)”](#) (p. 11-6).

Figure 11-2 LAN-LAN Defaults Editor (Parameters Tab)



LAN-LAN Defaults Editor - system

File Edit Monitor Windows Utilities Help

Parameters Policy SLA Probe

Session Parameters

Contact Name

Email

Phone

Comment

Receive Any Proposals

Send All Proposals

Start Date

End Date

Keepalive Interval (secs)

Required IPSec Transport Method

Pure IPSec (IP type 50/51)

UDP Encapsulation Port

- 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 E-mail Phone Comment	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.
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 “Set Up a LAN-LAN Tunnel with UDP Encapsulation” (p. 11-23) 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 11-3, “LAN-LAN Defaults Editor (Policy Tab)” (p. 11-8)

The Policy tab contains the IKE SA Proposal and IPSec Proposal defaults.

Figure 11-3 LAN-LAN Defaults Editor (Policy Tab)

LAN-LAN Defaults Editor - system

File Edit Monitor Windows Utilities Help

Parameters Policy SLA Probe

IKE SA Proposal

Diffie-Hellman Group Group 1

Encryption Type TRIPLE DES CBC

Auth Type HMAC SHA1

SA Lifetime (sec) 86400

IPSec SA Proposal

Protocol ESP-50

Encryption Type TRIPLE DES CBC

Auth Type HMAC SHA1

SA Lifetime (sec) 14400

SA Lifetime (Kbytes) 1000000

Enable Perfect Forward Secrecy

Enable LZS Compression

- 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	The default is Diffie-Hellman Group 5 . This exchange key method can be changed to Group 1 , Group 2 , or Group 14 . Group 2 provides more security than Group 1 , but rekey time may be better using Group 1 .
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 .
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 , and AES CBC 256 .
Auth Type	The default is HMAC SHA1 , but you can change this to HMAC MD5 .

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 LSMS. 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



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 Tunnels** 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. 11-12) or [“Task, Method 2: from the Devices folder”](#) (p. 11-12).

Result The LAN-LAN Tunnel Editor is displayed, showing the 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 **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 6 .
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.

6

If...	Then...
IKEv1 was chosen as the VPN Type in Step 5	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.
IKEv1 was chosen as the VPN Type in Step 5	The Initiator field is displayed. 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. If a TEP does not have a fixed IP address, it must be the initiator, to allow LSMS to automatically select the address.

-
- 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...
Preshared Key is chosen	<p>A default preshared key is randomly generated by the system and hidden by asterisks (*) in the Preshared Key field. 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.

-
- Click the **Endpoint 1** tab of the LAN-LAN Tunnel Editor.

Result The Endpoint 1 tab is displayed.

-
- 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. 11-12) 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. 11-12) was used to display the LAN-LAN Tunnel Editor, this field is already populated.

- 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 Lucent VPN Firewall Brick Device Ports* chapter in the *LSMS 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.</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 LSMS, such as a TEP on a non-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-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 LSMS). 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>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>

If...	Then...
IP Address was chosen in Step 2	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.
If Name was chosen in Step 2	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.

4

If...	Then...
Preshared Key was chosen as the Authentication Method on the Tunnel tab	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 Domain Name or Email Address is chosen, enter the appropriate value in the ID field.

If...	Then...
<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 [“Set LAN-LAN Tunnel Defaults”](#) (p. 11-5).

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 [“Set LAN-LAN Tunnel Defaults”](#) (p. 11-5).

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 [“How to Set Up Service Level Agreements”](#) (p. 11-30) for instructions on how to set up Service Level Agreements and configure the fields for the SLA Probe tab for a tunnel.

END OF STEPS

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



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 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 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 LSMS 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 LSMS 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. 11-13) 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 LSMS, 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.



Maintain 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.

View Existing Tunnels

The LSMS provides these six viewers that you can use to view existing LAN-LAN tunnels. The table below indicates what each viewer shows.

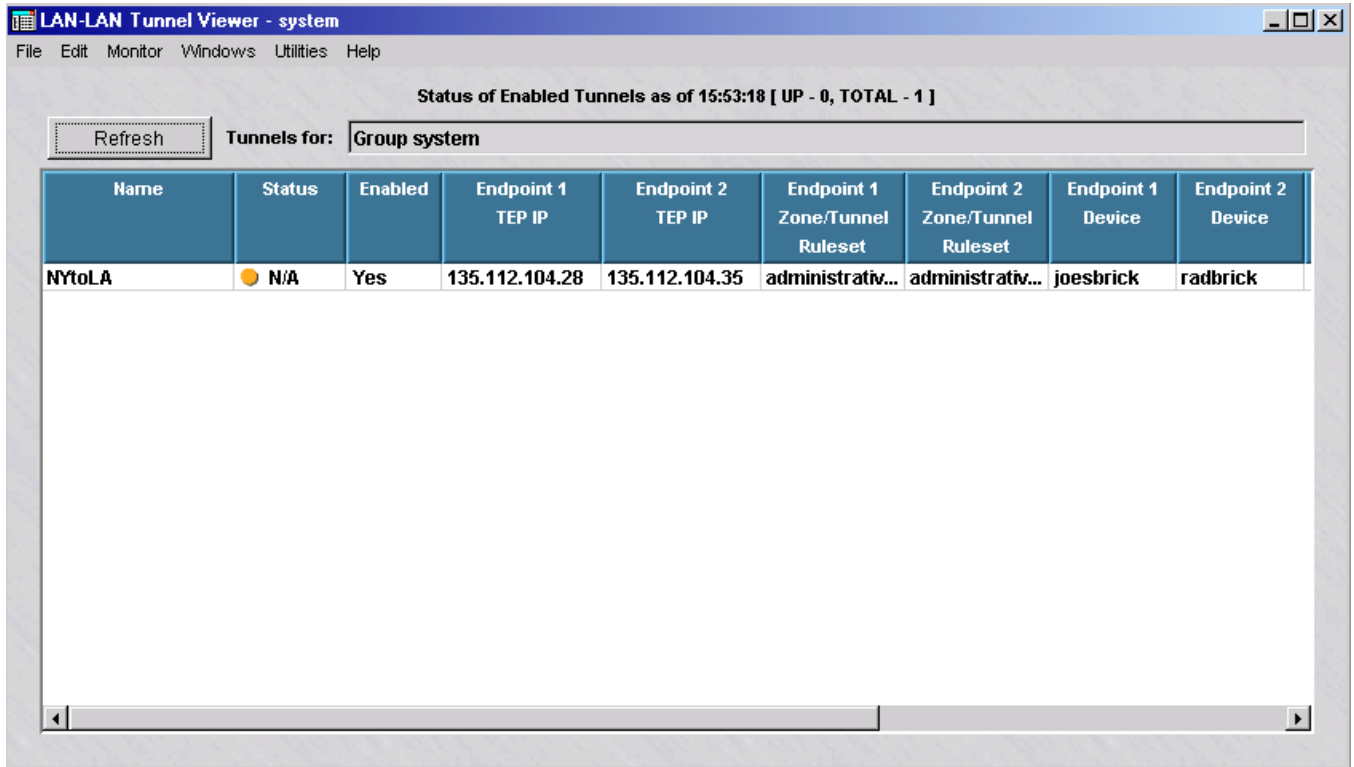
Viewer	Shows . . .
All Brick Tunnels	All tunnels terminating on Bricks in the current group
All Group Tunnels	All tunnels terminating on devices (Bricks) in the current group
Selected Device Tunnels	All tunnels terminating on the device you specify
Selected Folder Tunnels	All tunnels terminating on devices 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 11-4, “All Group Tunnels Viewer” (p. 11-27) shows an example of the All Group Tunnels Viewer with one tunnel displayed.

Figure 11-4 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 LSMS 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 “Set LAN-LAN Tunnel Defaults” (p. 11-5)).

To refresh the status, click the **Refresh** button.

Modify a Tunnel

To modify the configuration of an existing tunnel, follow the steps below:

- 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 “ [How to Set Up Service Level Agreements](#)” (p. 11-30)).
-

- 3 Display the File menu and select **Save and Apply**.

END OF STEPS

Delete a Tunnel

To delete the configuration of one or more existing tunnels, follow the steps below:

- 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

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.



How 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

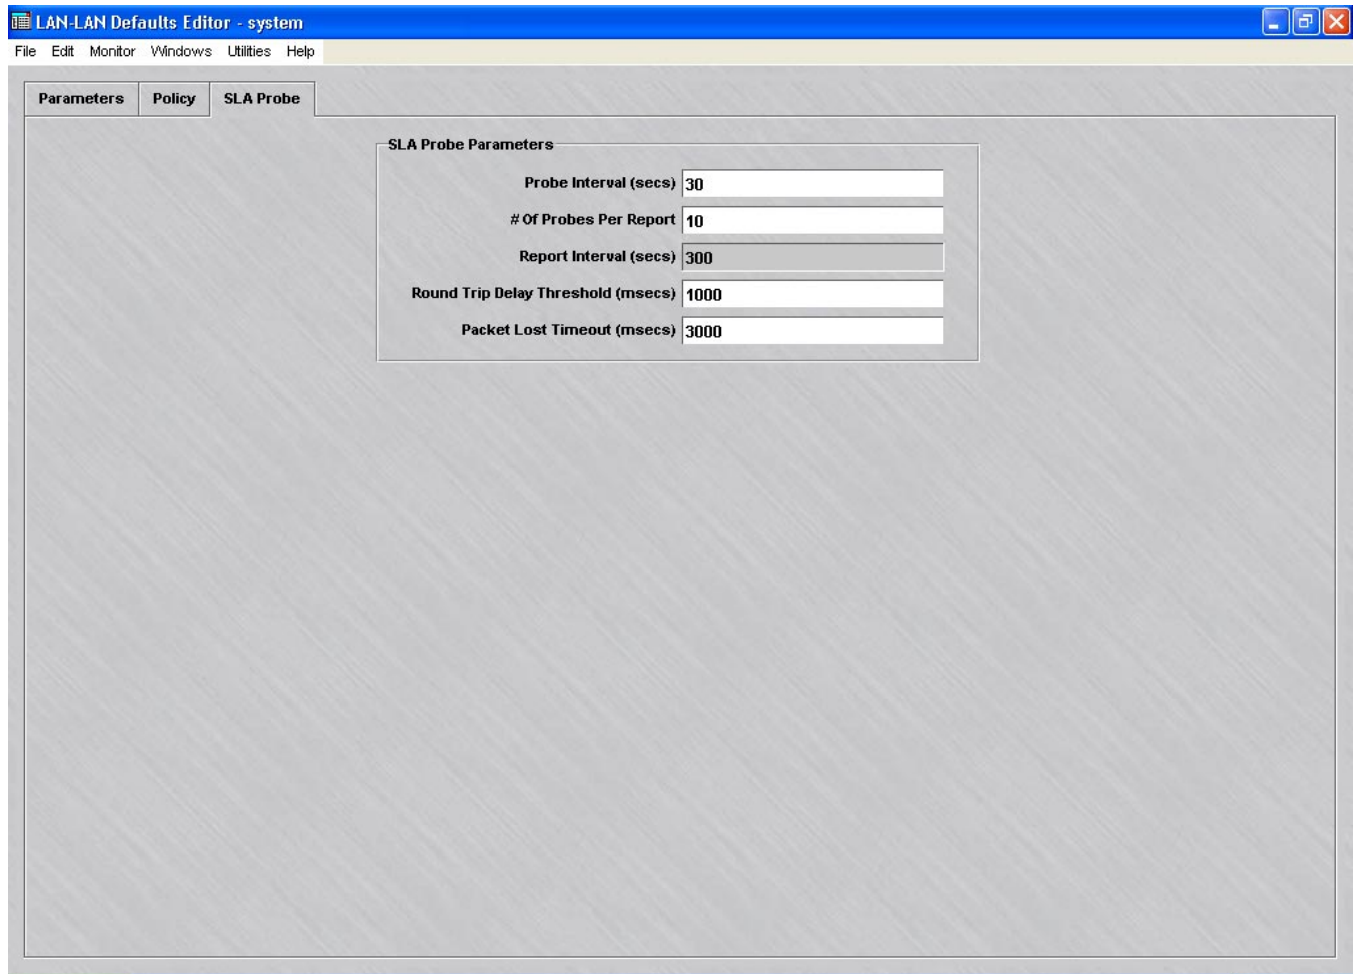
Any or all of these may be in place for a given customer of a Service Provider. SLAs are particularly well suited to work with the QoS/bandwidth management feature. This feature creates limits and guarantees on session and packet throughput, and can report when those limits and guarantees are not met.

Set the Default SLA Parameters

The LSMS 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 11-5, “Default SLA Parameters”](#) (p. 11-31) 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 11-5 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

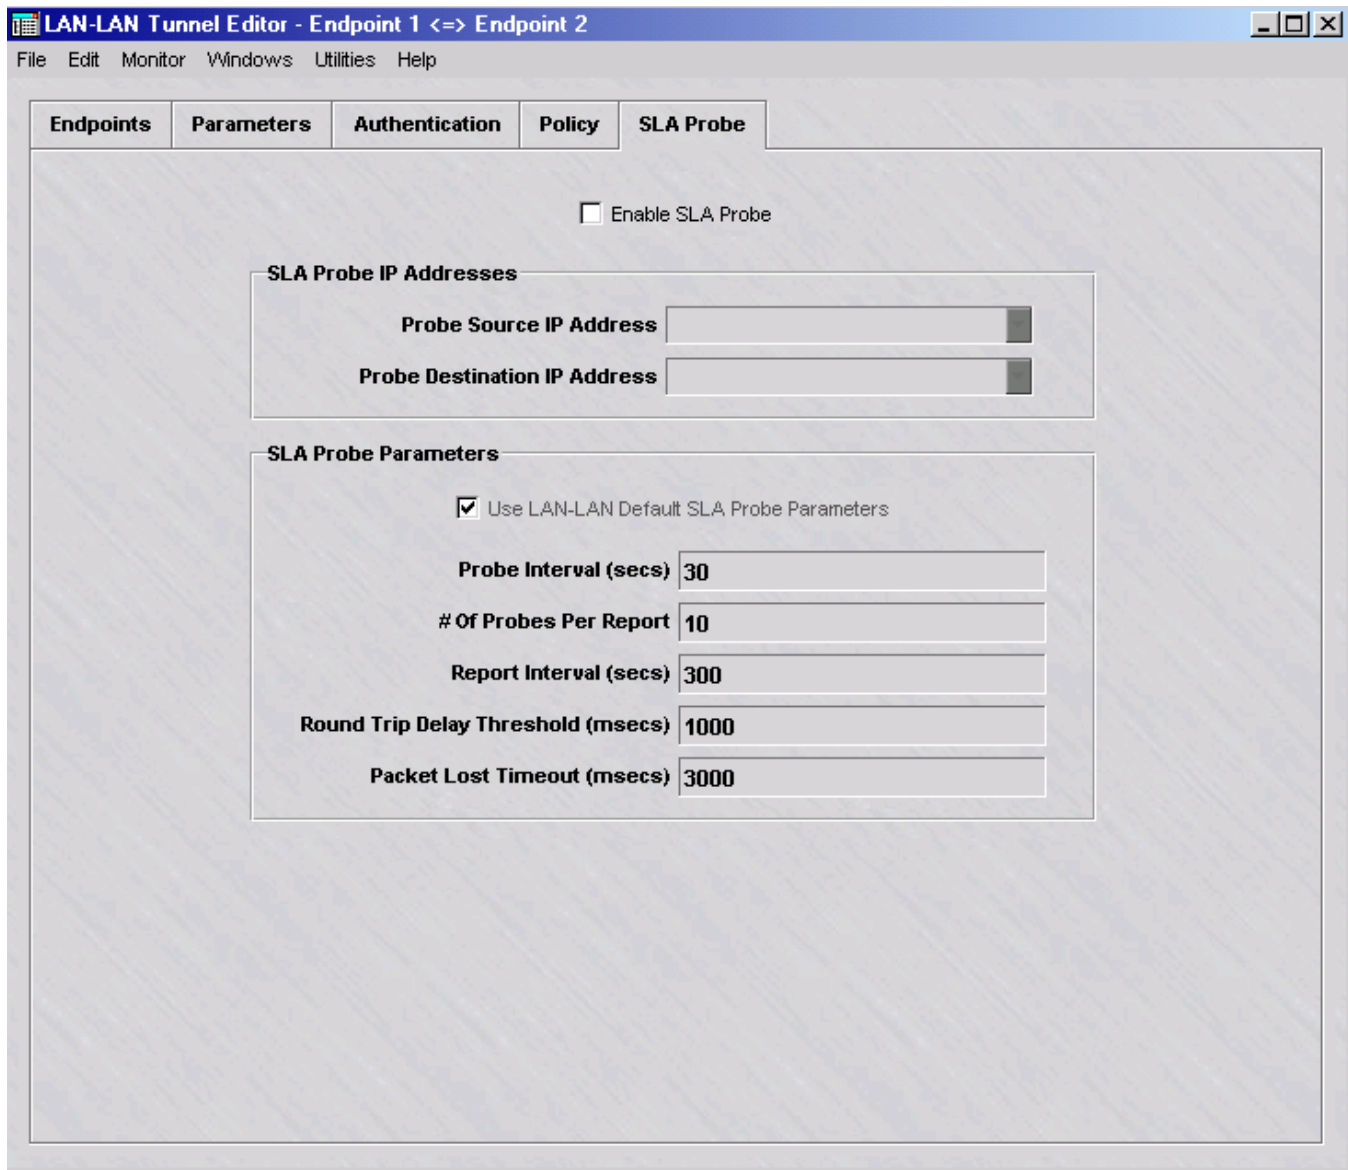
When setting up a LAN to LAN tunnel, you can use the default SLA parameters, or you can change the parameters for this particular tunnel.

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 11-6, “LAN-LAN Tunnel Editor (SLA Probes tab)” (p. 11-32)).

Figure 11-6 LAN-LAN Tunnel Editor (SLA Probes tab)



-
- 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 [“View Existing Tunnels”](#) (p. 11-26) 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 11-4, “All Group Tunnels Viewer”](#) (p. 11-27)).

Figure 11-7, “Tunnel Round Trip Delay Statistics” (p. 11-34) shows a typical tunnel round trip delay statistics display, and Figure 11-8, “Group Round Trip Delay Statistics” (p. 11-34) shows a group round trip delay statistics display.

Figure 11-7 Tunnel Round Trip Delay Statistics

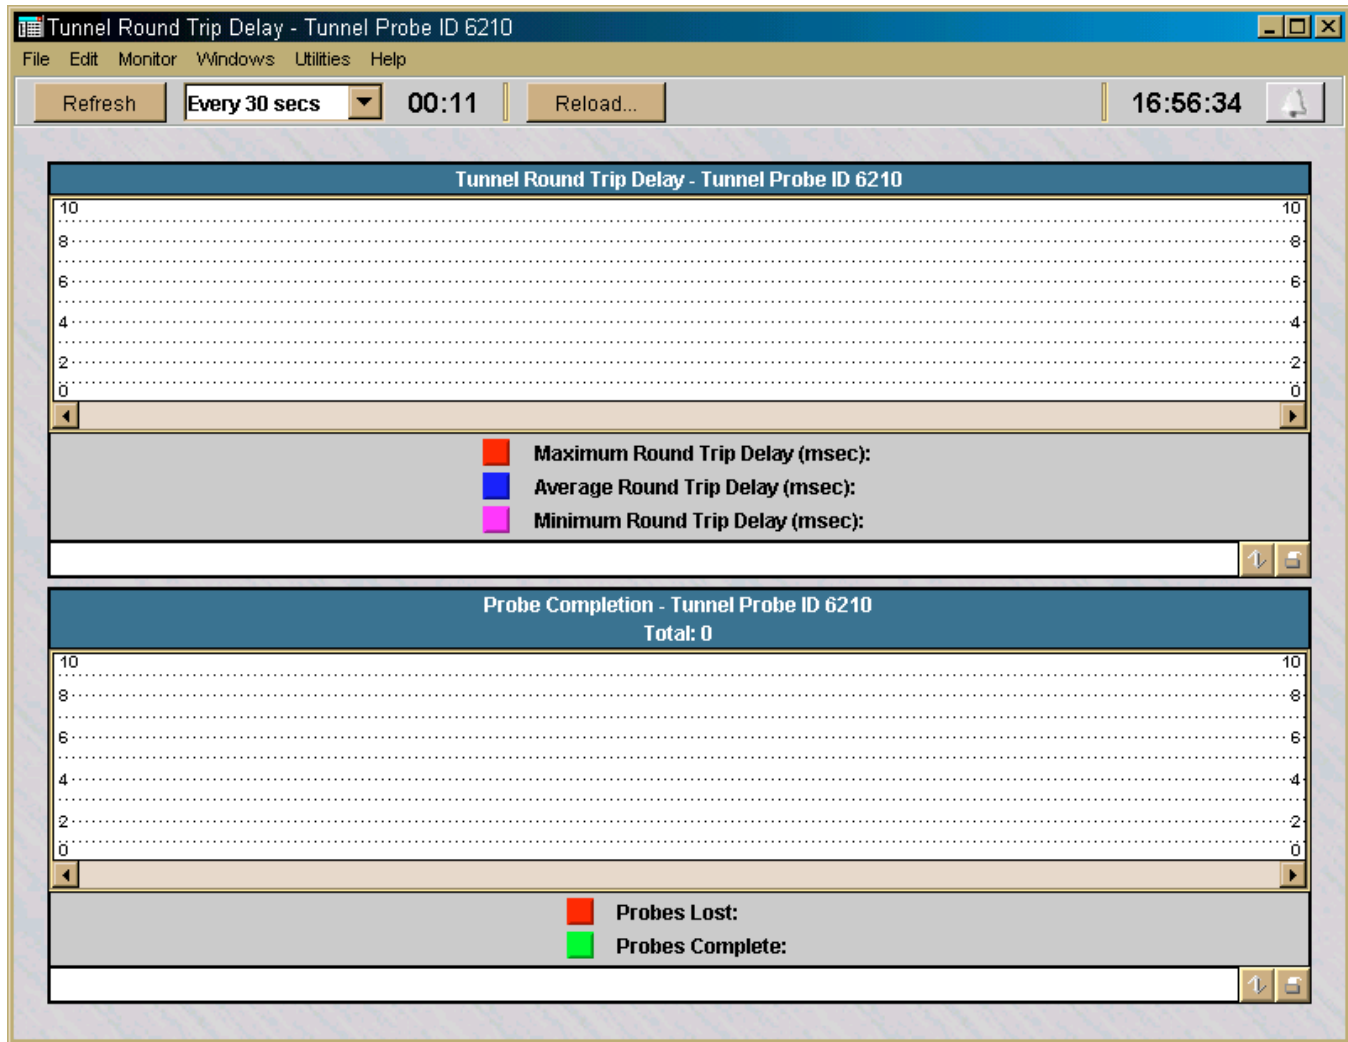
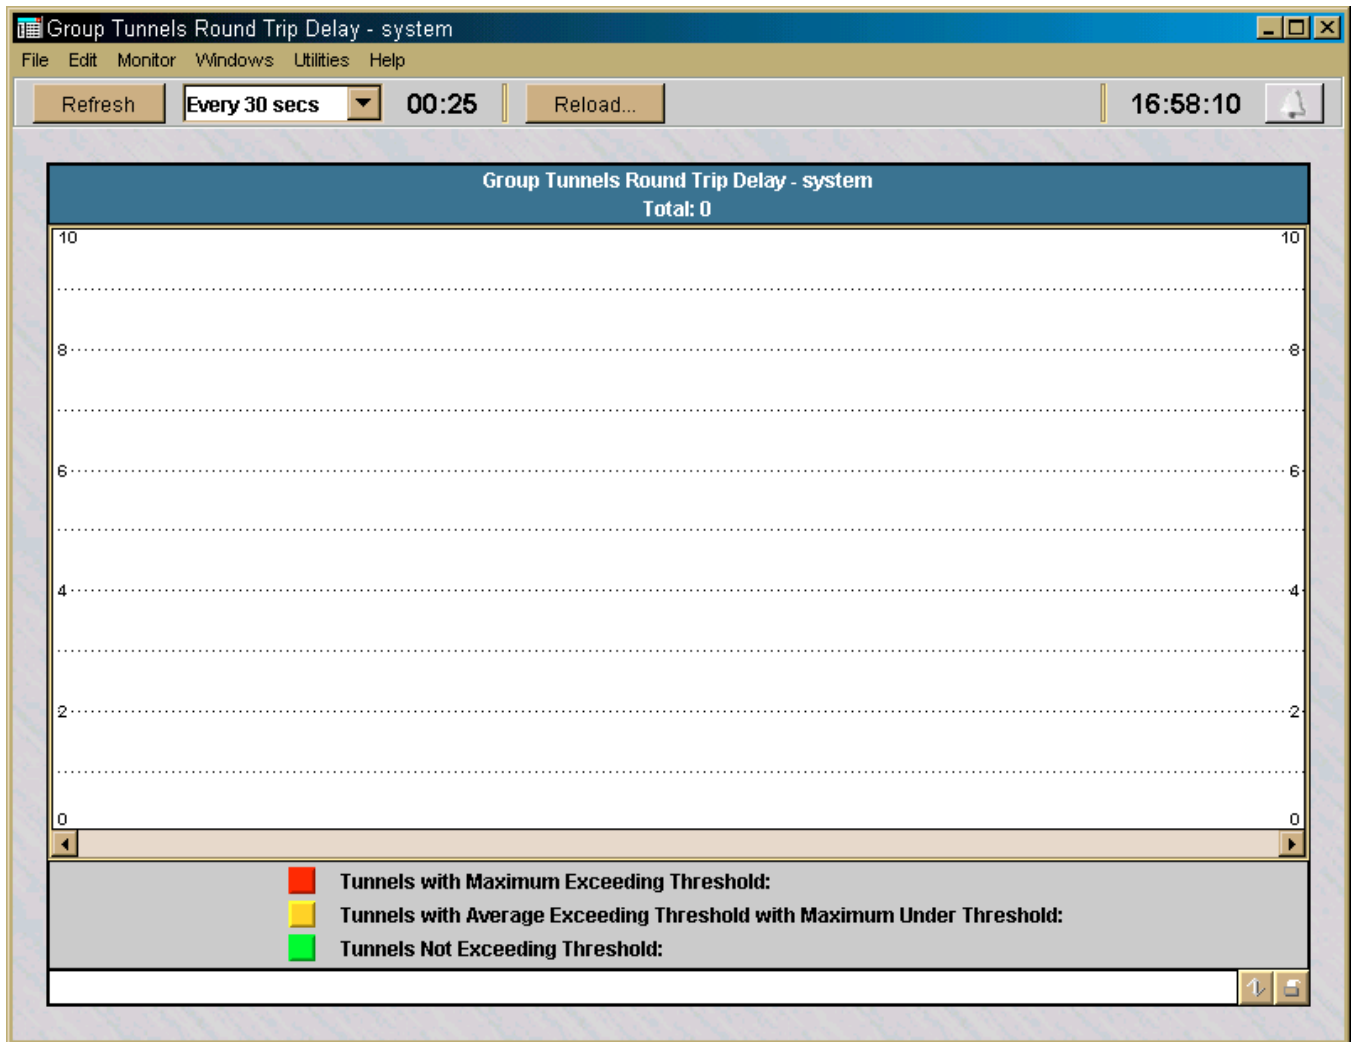


Figure 11-8 Group Round Trip Delay Statistics



□

12 Client Tunnel Endpoints

Overview

Purpose

This chapter explains how to set up a client tunnel endpoint so that users of the Lucent IPsec Client application (or a compatible IPsec client application) can establish a tunnel between their hosts and a Lucent VPN Firewall *Brick*® device. 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?	12-2
To Set the Client Tunnel Defaults	12-6
To Set Up a Client Tunnel Endpoint	12-22
To Set Up a Client Tunnel with UDP Encapsulation	12-29
What to Do Next	12-31
Maintain Client Tunnel Endpoints	12-37
Create a Message for Lucent IPsec Client Users	12-43



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 Lucent IPsec Client and a port on a Brick.

Tunnel Endpoints

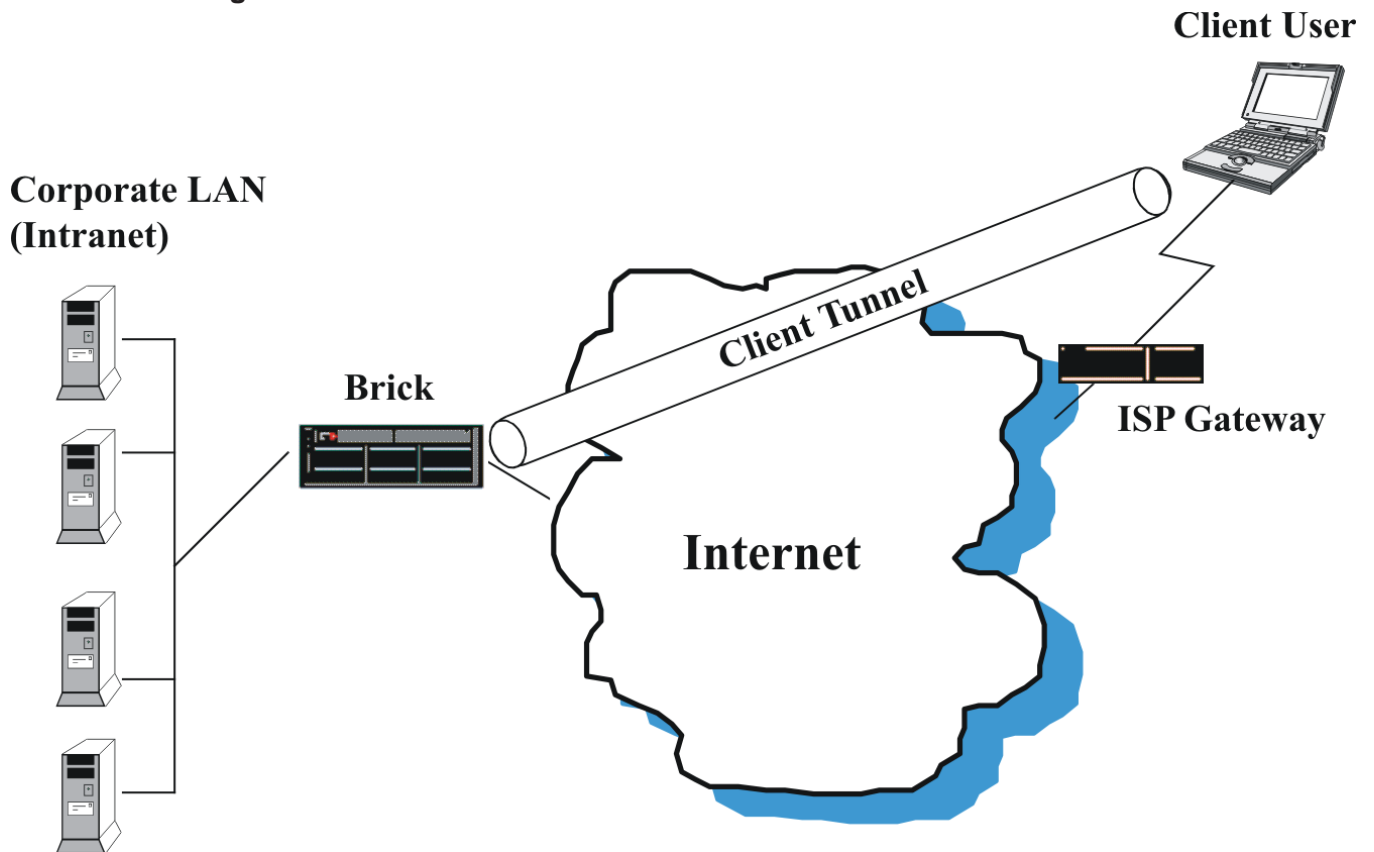
The client user endpoint is represented by the IP address of the host running the Lucent IPsec Client. This can be either a fixed IP address or an IP address temporarily assigned by an Internet Service Provider (ISP).

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. See the *Configuring Lucent VPN Firewall Brick[®] Device Ports* in the *LSMS Administration Guide*.

Scenarios

The diagram in [Figure 12-1, “Client Tunnels”](#) (p. 12-3) illustrates the following scenario.

Figure 12-1 Client Tunnels



In [Figure 12-1, “Client Tunnels”](#) (p. 12-3), the endpoints of the tunnel are a mobile user laptop computer running the 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 LSMS 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.

IKE hardware acceleration

Some IKE (v1 and v2) operations are now executed with hardware acceleration on some Brick device models (50, 150, and 700 Max). This greatly increases the number of tunnels per second that can be established on these devices as well as reducing the overall load on the Brick devices.

Packet data gateway (PDG) accounting

The Packet Data Gateway (PDG) accounting feature can be enabled for a client tunnel endpoint. When this feature is enabled, the LSMS 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.

Client tunnel defaults

Before you set up a client tunnel, you should examine the client tunnel defaults provided with the LSMS application to determine whether or not any of them need to be modified.

The LSMS 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 LSMS application. See [Chapter 9, “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 12-2, “Client Defaults Editor \(Parameters Tab\)”](#) (p. 12-7).

Figure 12-2 Client Defaults Editor (Parameters Tab)

Client Defaults Editor - system

File Edit Monitor Windows Utilities Help

Parameters | IKEv1 Gateway | IKEv2 Gateway | Remote Client ID | Policy

Session Parameters

Keepalive Interval (secs) 300

Idle Timeout (mins) 30

Receive Any Proposals

Primary DNS 0.0.0.0

Secondary DNS 0.0.0.0

Primary WINS 0.0.0.0

Secondary WINS 0.0.0.0

Client Firewall Pass if Client Initiated

Disable Built-in Firewall if IPSec Client version is 6.0.1+

Allow client to save password

Allowed IPSec Transport Methods

Pure IPSec (IP type 50/51)

UDP Encapsulation Port(s) 501

NAT Traversal Port 4500

- 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 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 that receives the message generates a message back to the client.</p> <p>The default is 300 seconds. You can change this to any 1-second interval between 10 and 172800 seconds.</p> <p>Enter a value of zero (0) to disable keepalive/heartbeat messages.</p> <p>If the client does not receive after three heartbeat intervals (2 missed, plus one more), it disables the tunnel.</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>
Primary/Secondary WINS/DNS	<p>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.</p>

Field	Explanation
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 Lucent IPsec Client, if the two settings are different.</p>
Disable Built-in Firewall if IPsec Client version is 6.0.1+	<p>Placing a check in this checkbox will disable the built-in firewall. This feature only works if using IPsec Client V6.0.1 or higher. Note: Disabling the firewall also removes the "Allow Multi-session Protocol" feature from the client's GUI.</p>
Allow Client to Save Password	<p>The Allow Client to Save Password checkbox is a convenience for the 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. 12-29) for a more detailed discussion of UDP encapsulation.</p> <p>If NAT Traversal Port or UDP Encapsulation Port(s) is selected, LSMS 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, LSMS 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](#).

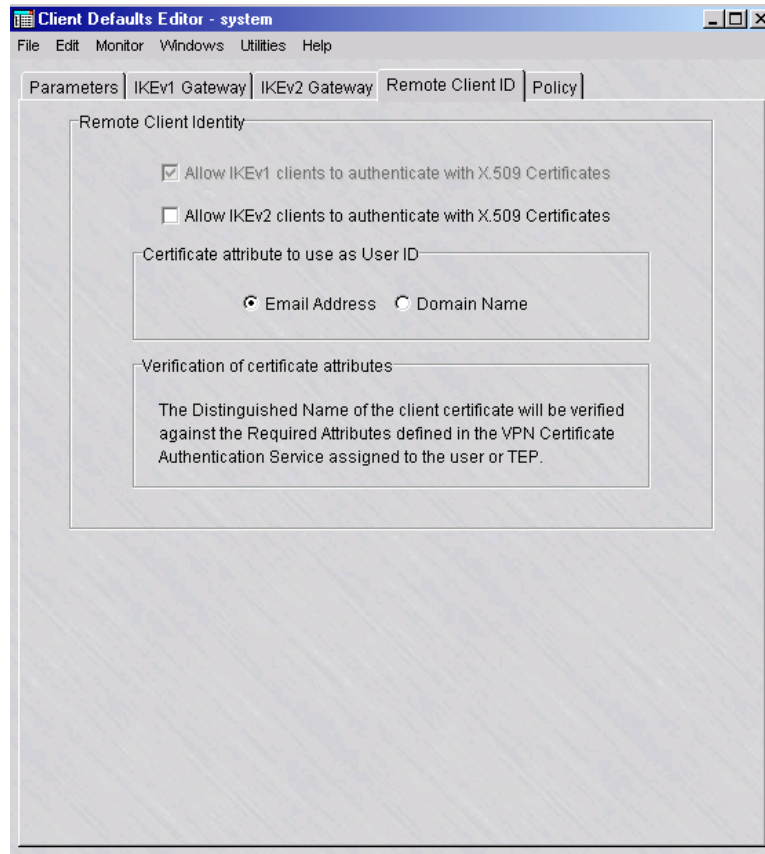
7	If	Then
	<p>Preshared Key was selected on the IKEv1 Gateway tab in Step 6</p>	<p>Use the default Preshared Key and Group ID assigned by the system, or change these field values.</p> <p>Click the Generate Key button to generate a random 20-character Preshared Key. A default group key is generated randomly 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-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 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 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 12-3, “Client Defaults Editor (Remote Client ID Tab)” (p. 12-14)).

Figure 12-3 Client Defaults Editor (Remote Client ID Tab)



-
- 9** 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**.

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.

- 10** Click **Policy** to display the Policy tab.

Result The Policy Tab is displayed (Figure 12-4, “Client Defaults Editor (Policy Tab)” (p. 12-16)).

Figure 12-4 Client Defaults Editor (Policy Tab)

The screenshot shows the 'Client Defaults Editor - system' window with the 'Policy' tab selected. The window contains three main sections:

- IKE SA Proposal:**
 - Diffie-Hellman Group: Group 2
 - Encryption Type: TRIPLE DES CBC
 - Auth Type: HMAC SHA1
- IPsec SA Proposal:**
 - Protocol: ESP-50
 - Encryption Type: TRIPLE DES CBC
 - Auth Type: HMAC SHA1
 - Enable Perfect Forward Secrecy
 - Enable LZS Compression
- IPsec SA Lifetimes:**

	IKEv1	IKEv2
SA Lifetime (secs)	14400	14400
SA Lifetime (Kbytes)	5000000	5000000

The Policy tab contains the IKE SA Proposal and IPsec SA Proposal defaults. 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.

The IPsec SA Proposal is used to negotiate IPsec SA parameters (such as Encryption and Auth Type) and to set up matching IPsec SAs in the peers.

- 11 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	The choices are: Group 1 , Group 2 , Group 5 , and Group 14 . The default is Group 2 . Group 2 provides more security than Group 1 , but rekey time may be better using Group 1 . If the IKE on the Brick option is enabled, Group 5 can be chosen, which provides more security than Group 1 and Group 2 .
Encryption Type	The options are TRIPLE DES CBC (default), DES CBC , AES CBC 128 , AES CBC 192 , and AES CBC 256 .
Auth Type	The default is HMAC SHA1 , but you can change this to HMAC MD5 .

- 12 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 options are TRIPLE DES CBC (default), DES CBC , AES CBC 128 , AES CBC 192 , and AES CBC 256 .
Auth Type	The default is HMAC SHA1 , but you can change this to HMAC MD5 .

Field	Explanation
<p>SA Lifetime (secs and Kbytes) (IKEv1 and IKEv2 client TEP settings)</p>	<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>
<p>Enable Perfect Forward Secrecy</p>	<p>By default, this checkbox is unchecked. If it is checked, a second Diffie-Hillman key exchange will take place during processing. This can improve security, but it also can impact re-keying performance.</p>

Field	Explanation
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 Lucent IPsec Client.</p>

.....
 END OF STEPS

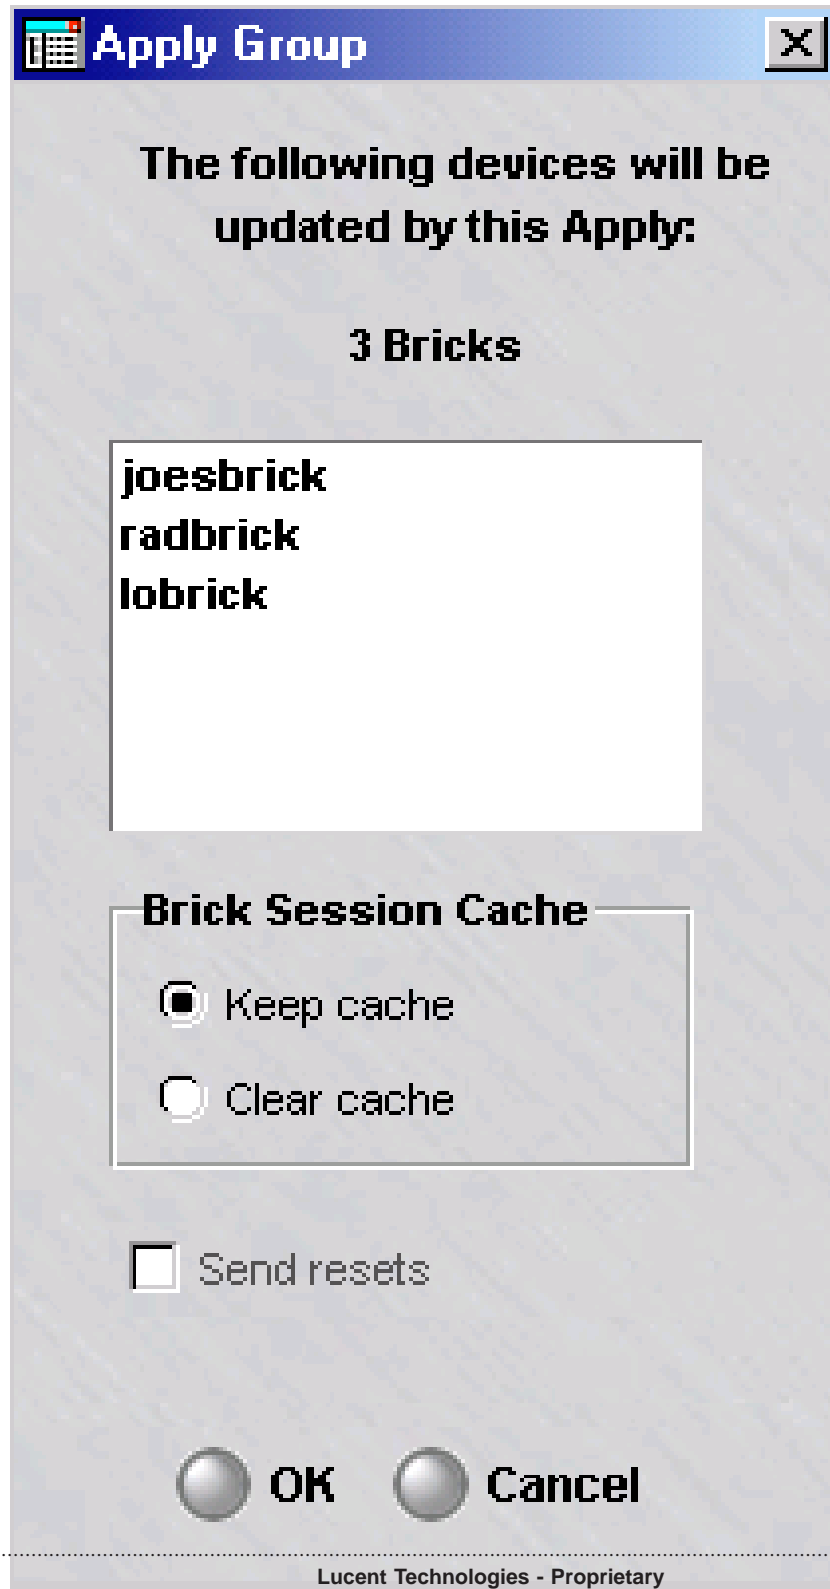
Task: save and apply the tunnel defaults

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 12-5, “Apply Group Window” (p. 12-20))

Figure 12-5 Apply Group Window



-
- 2 Click **OK** to save and apply the client tunnel defaults to the affected Brick(s).

Result The new policies are compiled and downloaded to the affected Brick(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 9, “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 10, “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

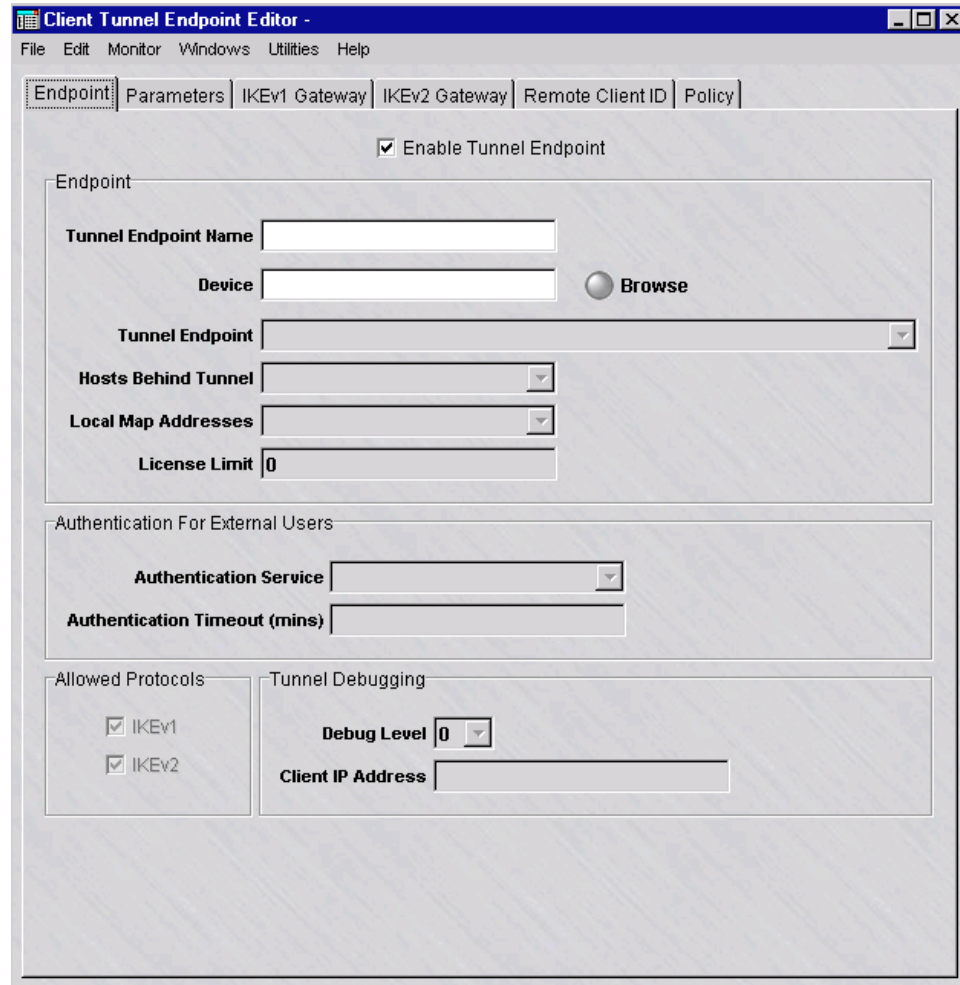
Complete the following steps to display the Client Tunnel Endpoint Editor from the VPN folder.

- 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 12-6, “Client Tunnel Endpoint Editor \(Endpoint Tab\)”](#) (p. 12-23).

Figure 12-6 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 12-6, “Client Tunnel Endpoint Editor \(Endpoint Tab\)”](#) (p. 12-23)) 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. 12-22) or [“Task, Method 2: from the Devices folder”](#) (p. 12-24).

Result The Client Tunnel Endpoint Editor is displayed, showing the Endpoint tab.

If [“Task, Method 1: from the VPN folder”](#) (p. 12-22) 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. 12-24) 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 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. See *Chapter 4. Configuring Lucent VPN Firewall Brick® Device Ports* in the *LSMS 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 LSMS application is Local Password. If you are using RADIUS, SecurID or VPN Certificate, you have to create the authentication service yourself. See [Chapter 9, “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 debugging level. When tunnel debugging is enabled, the Brick writes debug messages to the VPN log, identified by zone, tunnel endpoint name, and remote TEP (client IP). The available choices 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.

END OF STEPS

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. 12-28).

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. 12-6) 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: complete the PDG Accounting tab (optional)

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 set an option to prohibit traffic on a client tunnel endpoint until accounting has been started, and to configure the interim update intervals, in time period and traffic volume. The fields on this tab are optional.

- 1 Click **PDG Accounting** to display the PDG Accounting tab on the Client Tunnel Endpoint Editor.

Result The PDG Accounting tab is displayed.

- 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 **Accounting Interval Minutes** 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.
-

- 5 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.

-
- 6 In the **MCC-MNC** field, enter the 5-6 digit user ID to be used to authenticate the user, which is compared with the Mobile Country Code (MCC) and Mobile Network Code (MNC) extracted from the International Mobile Subscriber Identity (IMSI) code sent by the RADIUS server during the user authentication process.

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 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 IPSEC Client Licenses on Tunnel Endpoints” \(p. 12-31\)](#) for instructions on how to allocate an 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.

END OF STEPS



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 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 a Lucent proprietary mechanism and only works with Bricks and 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** checkbox 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** checkbox 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 LSMS 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 LSMS 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 checkbox 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)** checkbox 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**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



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 IPSEC Client Licenses on Tunnel Endpoints”](#) (p. 12-31).

Allocating 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 IPSEC Client License limit.

The LSMS software installation key and feature keys determine how many client licenses are available on your LSMS instance. You may distribute these licenses among the Groups on your LSMS instance, then define user license limits for each Client TEP, up to the limit allocated for that Group.

Task

To allocate IPsec Client licenses:

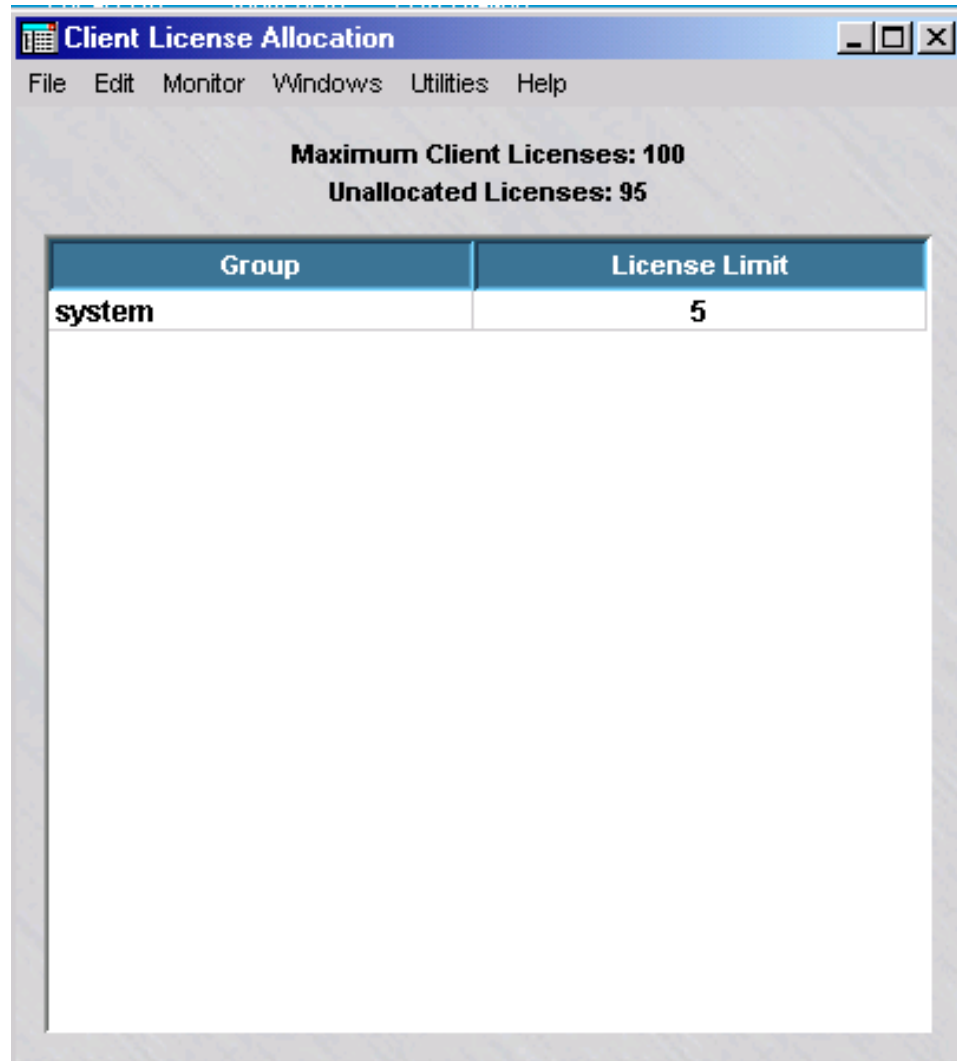
- 1 Log in using the LSMS Navigator.

- 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 12-7, “Client License Allocation Window” (p. 12-32) shows a sample window).

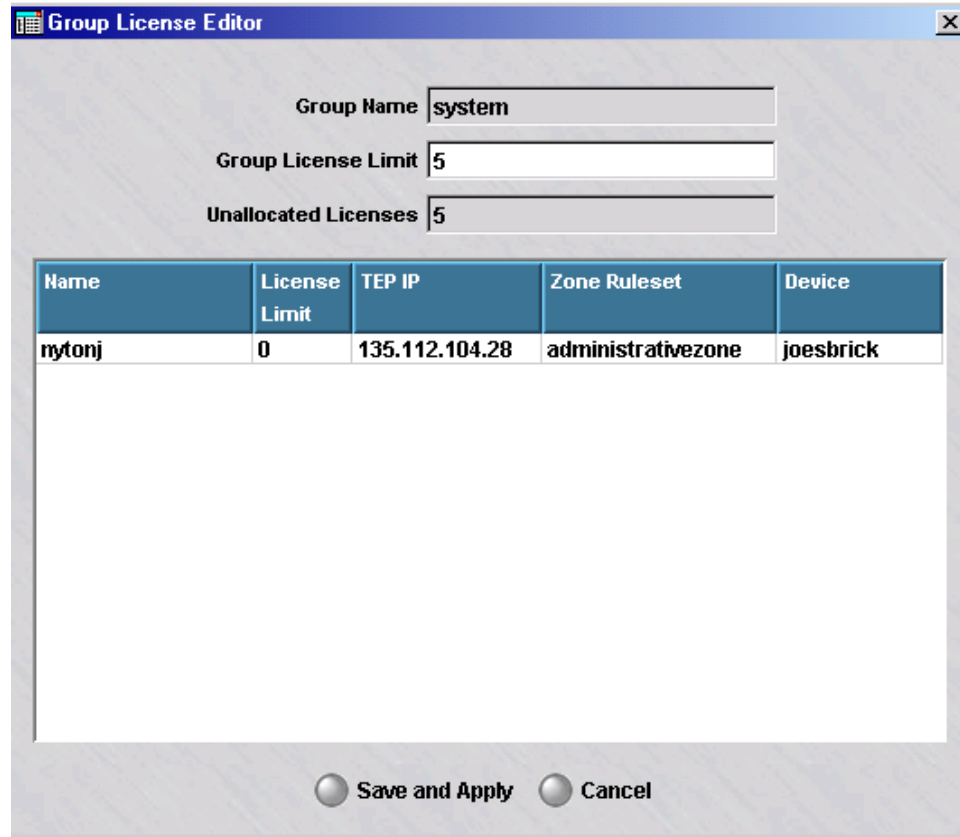
Figure 12-7 Client License Allocation Window



-
- 4 Right-click on a Group and select **Edit** or double-click on the Group.

The Group License Editor window is displayed (Figure 12-8, “Group License Editor Window” (p. 12-33) shows a sample window).

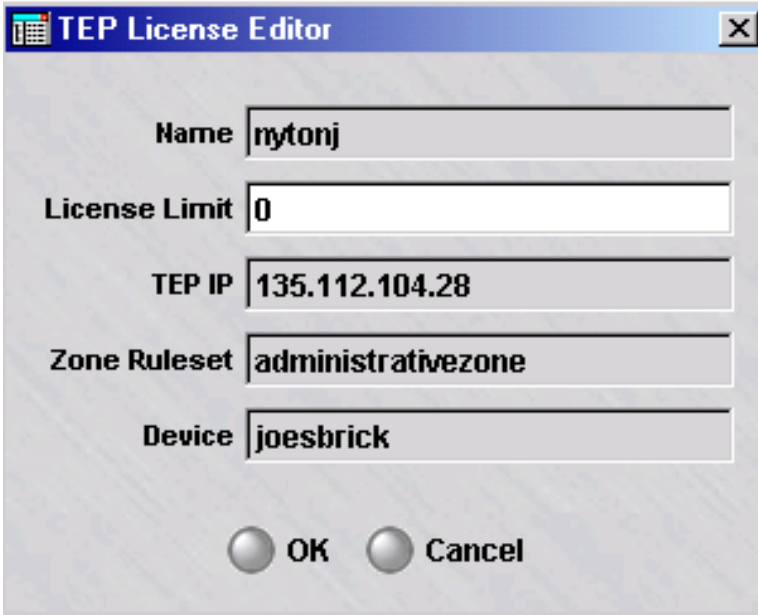
Figure 12-8 Group License Editor Window



-
- 5 Right-click on each TEP in the list and select **Edit**.

A Tunnel Editor Endpoint window is displayed (Figure 12-9, “Tunnel Endpoint Editor Window” (p. 12-34) shows a sample window).

Figure 12-9 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 text input fields, each with a label to its left: "Name" (value: nytonj), "License Limit" (value: 0), "TEP IP" (value: 135.112.104.28), "Zone Ruleset" (value: administrativezone), and "Device" (value: 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. LSMS Release 9.1 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 LSMS 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 LSMS application.
- **RADIUS and SecurID**
If you are using RADIUS or SecurID, you have to add a rule to the zone protecting the LSMS to allow the LSMS 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 10, “Digital Certificates”](#) in the *LSMS Policy Guide*).

For a detailed explanation of user authentication and instructions for creating the required rules, see [Chapter 9, “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.

Field	Explanation
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 “ LSMS-defined user groups ” (p. 9-17) in Chapter 9, “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.



Maintain Client Tunnel Endpoints

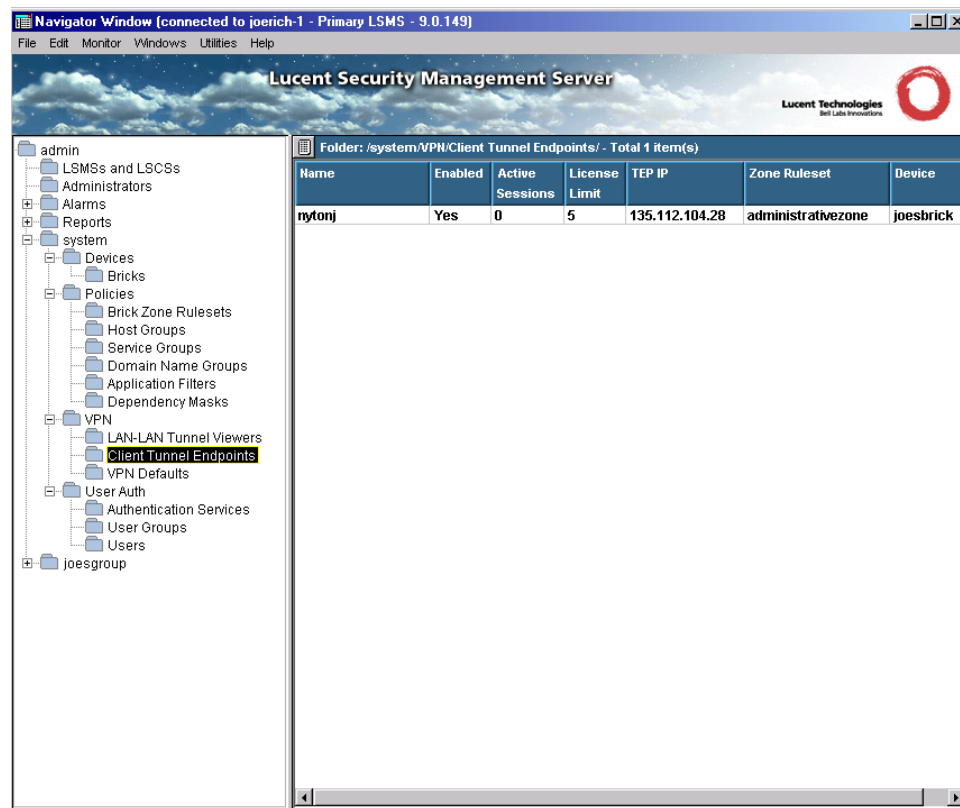
Overview

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.

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 12-10, “Client Tunnel Endpoints Window”](#) (p. 12-37) shows this display.

Figure 12-10 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)

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 12-6, “Client Tunnel Endpoint Editor \(Endpoint Tab\)”](#) (p. 12-23)) by checking or unchecking the **Enable Tunnel Endpoint** checkbox.

Display Active Sessions

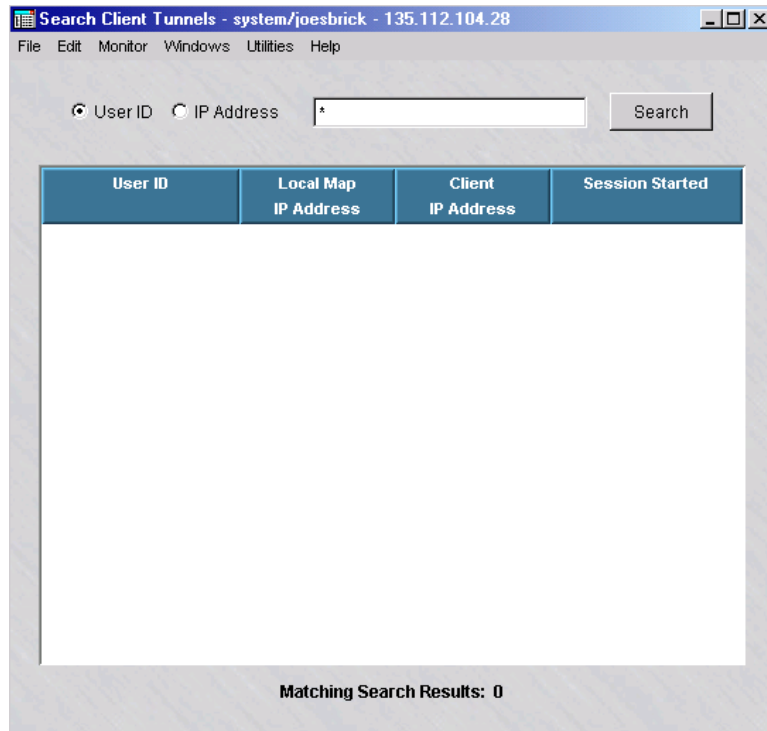
Use this procedure to search for and display active client tunnel sessions, based on the search criteria that you specify. You can search for and display all active client tunnel sessions by the originator’s user ID, IP address, or the local address to which the client is mapped.

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 12-11, “Search Client Tunnels Window” (p. 12-39)

Figure 12-11 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 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

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

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 12-6, “Client Tunnel Endpoint Editor \(Endpoint Tab\)”](#) (p. 12-23)) 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


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 12-12, “Confirmation Window \(Client Tunnel Endpoint\)”](#) (p. 12-42) is displayed.

Figure 12-12 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



Create a Message for Lucent IPsec Client Users

When to use

Use this task to create a message that all 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, etc.

Task

To create a message:

- 1 Using a standard text editor such as 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 LSMS software (*users* on the Windows platform, *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

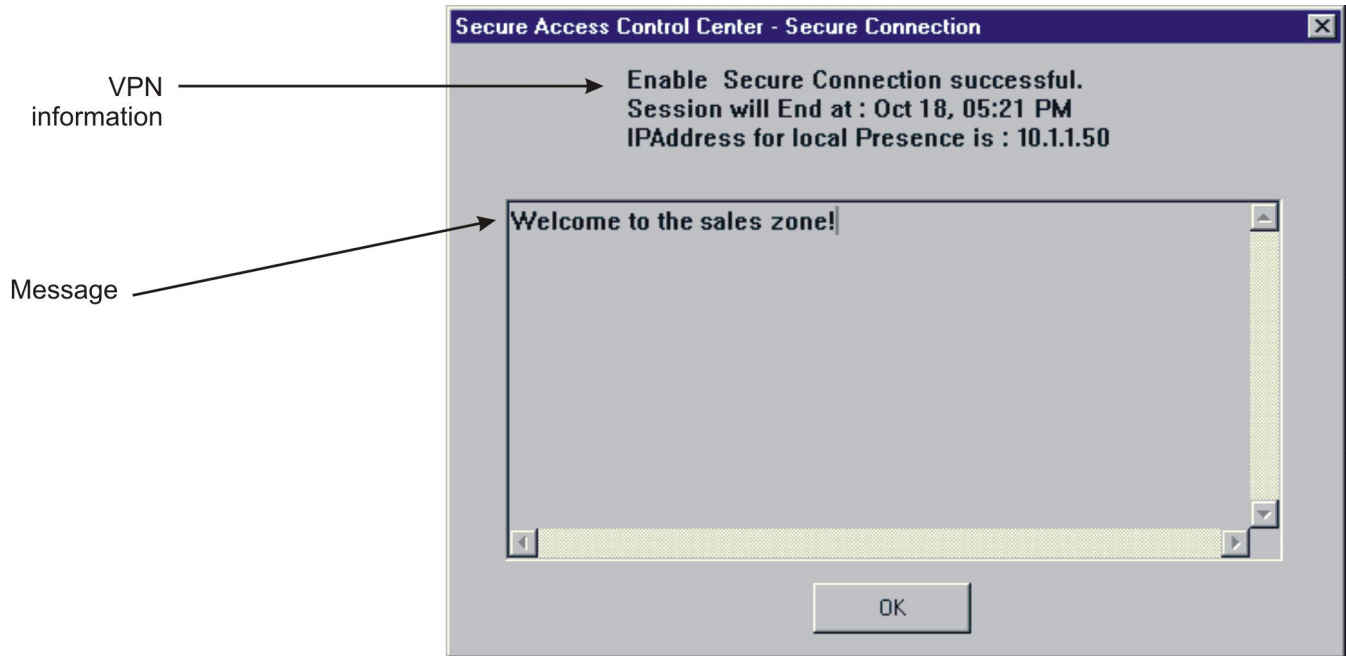
Sample Message

Figure 12-13, “Client Message” (p. 12-44) shows a typical message that a 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 12-13 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 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 Lucent VPN Firewall *Brick*® device 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 Bricks that pass the local presence packets. If some of your user community is running older versions of the Lucent IPsec Client, additional NAT rule changes must be made also.

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 Lucent IPsec Client) specify local presence mapping.

The host group **Active_VPN_Users** is created automatically by the LSMS. When a VPN client is authenticated, the LSMS 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 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.)

Because the Brick depends on the VPN Gateway Controller (VGC) process on the LSMS to fill in the host group, and because the VGC doesn't fill in local presence addresses for LAN-LAN, you can't use local presence for LAN-LAN tunnels.

□

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 Lucent IPSec Client package is in use by your remote users that access your network via a VPN tunnel.

1. Scenario 1 - All of your Windows NT and Windows 2000 client users are running at least IPSec Client version 3.0.xxx or later. All of your Windows 95 and Windows 98 client users are running at least version 3.1.xxx or later. For most users of LSMS 8.0, this is the expected scenario.
2. Scenario 2 - Some of your Win NT and Win 2K client are using a version earlier than 3.0.xxx. -and / or - some of your Win 95 and Win 98 users are using a version earlier than 3.1.xxx.

Subscribers to LSMS Software Support can easily download updated versions of the Lucent IPSec Client from the "private zone" under www.lucent.com/security



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.

(See *Chapter 3 Configuring and Activating a Brick* in the *LSMS Administration Guide* for a more detailed discussion of port assignments.)

Task

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. See [Figure 2-5, “Host Group Entry Window \(Nested Host Groups\)”](#) (p. 2-13) in *Chapter 3. Configuring and Activating a Lucent VPN Firewall Brick® Device* in the *LSMS 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. See in *Chapter 4. Configuring Lucent VPN Firewall Brick® Device Ports* in the *LSMS 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 *Chapter 12 Client Tunnel Endpoints* earlier in this manual.

END OF STEPS



Local Presence - Scenario 2

Setting up local presence

To give users of older versions of the Lucent IPsec Client an IP address on the local LAN, you must make use of the LSMS'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 *Chapter 6 Network Address Translation* earlier in this guide.

Set Up NAT Rule Changes for Older 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.

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. (See *Chapter 4. Configuring Lucent VPN Firewall Brick® Device Ports* in the *LSMS 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 *Chapter 3. Configuring and Activating a Lucent VPN Firewall Brick® Device* in the *LSMS 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 12, “Client Tunnel Endpoints”](#).

END OF STEPS



Appendix B: Pre-Configured Lucent VPN Firewall *Brick*[®] Device 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 LSMS 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 LSMS, and so forth.

The **administrativezone** ruleset is usually assigned to the port on the Brick that is connected to the LSMS. It drops all traffic to the LSMS 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 LSMS 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 Lucent Proxy Agent application, and is described in [Chapter 8, "Proxies"](#). The **vpnzone** ruleset is used to set up LAN-LAN tunnels, and is described in [Chapter 11, "LAN-LAN Tunnels"](#).

Contents

administrativezone	B-3
dhcpzone_on_inside_if	B-8

dhcpzone_on_internet_if	B-14
firewall	B-20
nocgwzone	B-27



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 LSMS to send audit data to the LSMS and request the LSMS 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 LSMS, 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 LSMS, 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 LSMS. 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>LSMS</p> <p>This is the LSMS host group that contains the IP addresses of all LSMSs.</p>
Service	<p>brick_to_SMS_Services</p> <p>This is a service group that was created automatically to facilitate Brick to LSMS communication. It consists of the following protocol/ports:</p> <p>TCP/9000-9001/*</p> <p>UDP/9014/*</p> <p>TCP/900/*</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 LSMS to download a policy and send configuration information to any Brick connected to the LSMS. 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 LSMS, which is inside the Administrative Zone, to reach the bricks, which are outside the zone.</p>
Source	<p>LSMS</p> <p>This is the LSMS host group that contains the IP addresses of all LSMSs.</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 LSMS, 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 LSMS.</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 LSMS 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>
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>
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>

Field	Explanation
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 LSMS 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 LSMS, 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 LSMS, 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 LSMS. 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.
Destination	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
Service	TCP/*/910 The rule applies only to Brick traffic with a source port of 910.

Field	Explanation
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.



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 LSMS 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 LSMS, 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	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
Service	brick_to_SMS_Services This is a service group that was created automatically to facilitate Brick to LSMS communication. It consists of the following protocol/ports: TCP/9000-9001/* UDP/9014/* TCP/900/*
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 LSMS 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 LSMS, which is inside the zone.

Field	Explanation
Source	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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_to_SMS_Services This is a service group that was created automatically to facilitate Brick- to-LSMS communication. It consists of the following protocol/ports: TCP/9000-9001/* UDP/9014/* TCP/900/*
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 LSMS 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 LSMS, 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	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS, 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 LSMS, 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 LSMS 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 LSMS, 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	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
Service	brick_to_SMS_Services This is a service group that was created automatically to facilitate Brick-to-LSMS communication. It consists of the following protocol/ports: TCP/9000-9001/* UDP/9014/* TCP/900/*
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 LSMS 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 LSMS, which is inside the Administrative Zone, to reach the bricks, which are outside the zone.

Field	Explanation
Source	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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_to_SMS_Services This is a service group that was created automatically to facilitate Brick-to-LSMS communication. It consists of the following protocol/ports: TCP/9000-9001/* UDP/9014/* TCP/900/*
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 LSMS 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 LSMS, 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	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS, 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 LSMS, 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 LSMS. 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 LSMS, which is outside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.

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 LSMS. The following explains how the rule works:

Field	Explanation
Direction	In To Zone The purpose of the rule is to allow traffic from the LSMS, which is outside the Firewall Zone, to reach the Brick, which is inside the zone.
Source	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS-to-Brick communication. It consists of the following protocol/ports: TCP/910/* ICMP/8/0 UDP/1024/* UDP/9014/*
Action	Pass The rule permits traffic matching the above fields, which can only be coming from the LSMS, 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 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 LSMS to send traffic to the LSMS. 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 LSMS, 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 LSMS. It was automatically created when the NOC Gateway Zone was created.
Destination	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.

Field	Explanation
Service	brick_to_SMS_Services This is a service group that was created automatically to facilitate Brick to LSMS communication. It consists of the following protocol/ports: TCP/9000-9003/* UDP/9014/* TCP/900/*
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 LSMS to send traffic to any Brick connected to the LSMS. The following explains how the rule works:

Field	Explanation
Direction	Out Of Zone The purpose of the rule is to allow traffic from the LSMS, which is inside the NOC Gateway Zone, to reach the Bricks, which are outside the zone.
Source	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.

Field	Explanation
Destination	bricks This is a host group that contains the IP addresses of all the Bricks that are being managed by this LSMS.
Service	brick_from_SMS_Services This is a service group that was created automatically to facilitate LSMS-to-Brick communication. It consists of the following protocol/ports: TCP/910/* ICMP/8/0 UDP/1024/* UDP/9014/*
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 LSMS, which is inside the zone.

Field	Explanation
Source	bricks This is a host group that contains the IP addresses of all the Bricks that are being managed by this LSMS. It was automatically created when the NOC Gateway Zone was created.
Destination	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS. 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 LSMS, which is inside the zone.

Field	Explanation
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS.
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 LSMS. 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 LSMS, which is inside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS.
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 LSMS. 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 LSMS, which is inside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS.
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 LSMS. 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 LSMS, which is inside the zone.
Source	Wildcard asterisk (*) The wildcard asterisk represents all existing and future Bricks.
Destination	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS.
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 LSMS 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 LSMS, which is inside the NOC Gateway Zone, to reach the Bricks, which are outside the zone.
Source	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS 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 LSMS 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 LSMS, which is inside the NOC Gateway Zone, to reach the Bricks, which are outside the zone.
Source	LSMS This is the LSMS host group that contains the IP addresses of all LSMSs.
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 LSMS 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 LSMS.
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 Lucent VPN Firewall *Brick*® device 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 (i.e., 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 LSMS. 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 LSMS 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 *LSMS 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 LSMS to communicate with a database server that provides user authentication, access control and accounting functionality.

It is possible to configure the LSMS to use some of the information contained in the RADIUS protocol messages that the RADIUS server sends to the LSMS when authenticating users. To do this, you have to assign RADIUS attribute codes to certain LSMS parameters, and then you have to indicate the RADIUS attribute's data type.

The following lists the LSMS 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
IKEv1 Preshared Key	D-5
Timeout Information	D-6
Local Presence IP Information	D-8
Primary WINS/Secondary DNS	D-11
Sample Configuration	D-13



User Group Information

Overview

An administrator can configure the LSMS to process user group information contained in the RADIUS protocol message.

How to Configure the LSMS and RADIUS Server

To configure the LSMS 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 LSMS, create a user group and assign that user group to one or more rules in one or more zones assigned to one or more Lucent VPN Firewall *Brick*® devices. 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 LSMS user group can be empty.

The procedure for creating an LSMS user group is described in “[Overview](#)” (p. 9-1) (see “[Create a User Group](#)” (p. 9-15)).

- 3 Create a RADIUS authentication service. On the **Attribute Codes** tab of the Authentication Service Editor, set the LSMS 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
- 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

The procedure for creating a RADIUS authentication service is described in “Overview” (p. 9-1) (see “How to Set Up RADIUS and SecurID Authentication” (p. 9-19)).

END OF STEPS

How the Feature Works

Once the RADIUS server and LSMS have been properly configured, the user group feature will work with both application users and 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 LSMS 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 LSMS will parse the string or text accompanying each attribute as a UTF-8 string. The LSMS will perform a character-by-character match on all user groups configured in the LSMS group to which this user is a member.
3. If a user group is found in the LSMS group that matches the string described above, the LSMS will apply the IP address associated with the end user performing the authentication to the LSMS 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 LSMS.)
5. Regardless of whether or not a user group is not found in the LSMS group that matches the string described above, the LSMS 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 LSMS 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 LSMS 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



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 LSMS to process user timeout and idle timeout information contained in the RADIUS protocol message.

How to Configure the LSMS and RADIUS Server

To configure the LSMS 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 LSMS 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” \(p. 9-1\)](#) (see [“How to Set Up RADIUS and SecurID Authentication” \(p. 9-19\)](#)).

END OF STEPS

How the Feature Works

Once the RADIUS server and LSMS have been properly configured, the User Timeout feature applies to both application users and 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 LSMS 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 LSMS 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 LSMS 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 LSMS.

□

Local Presence IP Information

Overview

An administrator can configure the LSMS 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 Lucent IPsec Client users who are enabling a client tunnel.

How to Configure the LSMS and RADIUS Server

To configure the LSMS 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 LSMS 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” \(p. 9-1\)](#) (see [“How to Set Up RADIUS and SecurID Authentication” \(p. 9-19\)](#)).

END OF STEPS

How the Feature Works

Once the RADIUS server and LSMS have been properly configured, the local presence IP feature will work with 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 LSMS 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 LSMS 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 LSMS 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 LSMS 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 LSMS finds a local IP address in the RADIUS message, the LSMS 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 LSMS for that tunnel.
5. If the LSMS determines that the local IP address in the RADIUS message has already been allocated to another tunnel, the LSMS 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 LSMS.

6. If the LSMS 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 LSMS will allow that client to authenticate. However, the LSMS will log a warning message indicating that a local IP address assigned by RADIUS conflicts with the configured Local Map Pool. The LSMS 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 LSMS to process any WINS/DNS information contained in the RADIUS protocol message.

How to Configure the LSMS and RADIUS Server

To configure the LSMS 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 LSMS Parameter to **Primary WINS** 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
- 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. 9-1) (see “How to Set Up RADIUS and SecurID Authentication” (p. 9-19)).

END OF STEPS



Sample Configuration

RADIUS attribute codes

The following is an example of a typical LSMS-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 LSMS sends to the RADIUS server for authentication.

□

Index

A Accounting shared secret

- Accounting shared secret, [9-21](#)
- ACE/Server see [SecurID, 9-19](#)
- Activate a rule, [1-57](#)
- administrativezone ruleset
 - administrativezone, [B-3](#)
- Application users, [9-2, 9-4](#)
 - Login procedure, [9-4](#)
- Apply
 - Brick zone ruleset, [1-50](#)
 - LAN-LAN tunnel, [11-21](#)
- Assign to an interface
 - Brick zone ruleset, [1-11](#)
- Authenticatio service, [12-21](#)
- Authentication port, [9-21](#)
- Authentication see [User authentication, 9-1](#)
- Authentication service, [9-19, 9-25, 9-27](#)
 - RADIUS, [9-19](#)
 - SecurID, [9-25](#)
 - VPN certificate, [9-27](#)
- Authentication shared secret, [9-21, 9-21](#)
- Authentication timeout, [9-6, 9-27](#)
- Authorize return channel, [1-33](#)
- Automated key exchange, [11-3](#)

B Brick, [11-2, 12-2](#)

- LAN-LAN tunnel endpoint, [11-2](#)
- Brick devices
 - supported, [xxi](#)
- Brick Proxy Editor, [8-4](#)
- Brick session cache see [Session cache, 1-5](#)
- Brick Zone Ruleset Editor, [1-10](#)
- Brick zone rulesets, [1-1, 1-57](#)
 - add rules-based routing IP addresses, [1-45](#)
 - Advanced features, [1-30](#)
 - Apply, [1-50](#)
 - Assign to an interface, [1-11](#)
 - Brick Zone Ruleset Editor, [1-10](#)
 - Copy, [1-49](#)
 - Create, [1-11](#)
 - customized zones, [1-4](#)
 - Definition, [1-2](#)
 - Delete, [1-50](#)
 - Move, [1-49](#)
 - Packet filtering, [1-2](#)
 - rules-based routing and, [1-4](#)
 - View, [1-47](#)

C CA see [Certification Authority, 10-1](#)

- Certificate manager, 10-7
 - Create certificate signing request, 10-7
 - Start on a remote host, 10-6
 - Start on Solaris, 10-5
 - Start on Windows NT, 10-5
 - Certificate signing request
 - Create, 10-7, 10-7
 - Submit to certificate authority, 10-9
 - Change session timeout, 1-31
 - Client Defaults Editor, 12-6
 - Parameters tab, 12-6
 - Client firewall, 12-9
 - Client tunnel endpoint
 - set up with UDP encapsulation, 12-29
 - Client tunnel endpoints
 - maintain, 12-37
 - set up, 12-22
 - Client tunnels, 12-1
 - authentication service, 12-5
 - Authentication type, 12-17
 - Client Defaults Editor, 12-6
 - Client firewall, 12-9
 - defaults, 12-4
 - display active sessions, 12-38
 - Encryption type, 12-17
 - Host behind tunnel, 12-21
 - host groups and, 12-5
 - Keep alive interval, 12-8
 - Save password, 12-9
 - Scenarios, 12-2
 - Set the defaults, 12-6
 - Client users, 9-2, 9-4
 - Login procedure, 9-4
 - Configuration Assistant
 - user authentication parameters, 9-11
 - Copy
 - Brick zone ruleset, 1-49
 - Dependency mask, 7-13
 - Host group, 2-18, 3-15
 - Service group, 4-18
 - Create, 9-11, 9-11
 - Brick zone ruleset, 1-11
 - Certificate signing request, 10-7
 - Dependency mask, 7-5
 - Host group, 2-5, 3-4
 - RADIUS authentication service, 9-19
 - Rule, 1-17
 - SecurID authentication service, 9-25
 - Service group, 4-4
 - User account, 9-11
 - User group, 9-15
 - CSR see Certificate Signing Request, 10-7
-

D Delete

- Brick zone ruleset, 1-50
- Dependency mask, 7-14
- Host group, 2-19, 3-17
- LAN-LAN tunnel, 11-28
- Proxy Table entry, 8-10
- Rule, 1-56
- Service group, 4-19
- Dependency masks, 7-1, 7-18
 - Associate the with a rule, 7-7
- Copy, 7-13
- Create, 7-5
- Definition, 7-2
- Delete, 7-14

- Dependency Masks Editor, [7-3](#)
- Modify, [7-12](#)
- Move, [7-14](#)
- RealAudio, [7-16](#)
- View, [7-10](#)
- Dependency Masks Editor, [7-3](#)
- Destination address mapping, [6-5](#), [6-16](#), [A-7](#)
 - Direct, [6-9](#)
 - Pool, [6-9](#)
- Destination port mapping, [6-11](#)
- Deactivate a rule, [1-57](#)
- Diffie-Hellman group, [11-9](#)
- Digital Certificates, [10-1](#)
- Digital certificates, [10-24](#)
- Duplicate
 - Proxy Table entry, [8-9](#)
 - Rule, [1-55](#)

E Enable compression, [12-19](#)

- Enable LAN-LAN tunnel;Disable LAN-LAN tunnel;LAN-LAN tunnels
 - Enable, [11-28](#)
- Enable perfect forward secrecy, [11-9](#), [12-18](#)
- Endpoints see Tunnel Endpoints, [11-2](#)

F firewall ruleset

- firewall, [B-6](#)

FTP traffic

- content filtering for, [1-4](#)

H Host group

- client tunnels and, [12-5](#)
- Host groups, [2-1](#), [2-22](#), [3-1](#)
 - Benefits, [2-3](#)

- Copy, [2-18](#), [3-15](#)
- Create, [2-5](#), [3-4](#)
- Definition, [2-2](#), [3-2](#)
- Delete, [2-19](#), [3-17](#)
- Host Groups Editor, [2-4](#)
- Modify, [2-17](#), [3-15](#)
- Move, [2-19](#), [3-16](#)
- Provided with the LSMS application, [2-9](#), [3-9](#), [4-10](#), [5-93](#)
- Uses, [2-2](#), [3-2](#)
- View, [2-15](#), [3-13](#)
- Host Groups Editor, [2-4](#)
- Hosts behind tunnel, [12-21](#)
 - Client tunnel, [12-21](#)
- HTTP traffic
 - content filtering for, [1-4](#)

I IKE SA proposal

- Encryption type, [12-17](#)

IKE see Internet Key Exchange, [11-3](#)

IKEv1

- See: Internet Key Exchange Version 1 (IKEv1)
 - client tunnel endpoint defaults, [12-18](#)

IKEv2

- See: Internet Key Exchange Version 2 (IKEv2)
 - client tunnel endpoint defaults, [12-18](#)

Internet Key Exchange, [11-3](#)

Internet Key Exchange Version 1 (IKEv1), [11-3](#), [D-5](#)

Internet Key Exchange Version 2 (IKEv2), [11-3](#), [12-3](#)

IPSec proposal, [11-9](#)

- Authentication type, [11-9](#), [12-17](#)
- Enable compression, [12-19](#)
- Enable compression;Enable compression, [11-10](#)

- Enable perfect forward secrecy, [12-18](#)
 - Enable perforward secrecy, [11-9](#)
 - Encryption type, [11-9](#), [12-17](#)
 - Protocol, [11-9](#), [12-17](#)
 - SA lifetime (kbytes), [11-9](#)
 - SA lifetime (Kbytes), [12-18](#)
 - SA lifetime (secs), [12-18](#)
 - SA lifetime(secs), [11-9](#)
 - IPSec SA proposal, [12-17](#)
 - Authentication type, [12-17](#)
 - ISAKMP proposal, [12-17](#)
 - Diffie-Hellman group, [12-17](#)
-
- K** Keepalive interval, [12-8](#)
-
- L** LAN-LAN Defaults Editor, [11-5](#)
 - Parameters tab, [11-5](#)
 - LAN-LAN Tunnel Editor
 - Main tab, [11-13](#)
 - LAN-LAN Tunnel EditorLAN-LAN tunnels
 - LAN-LAN Tunnel Editor, [11-13](#)
 - LAN-LAN tunnels, [11-1](#), [11-13](#), [11-33](#)
 - Automated key exchange, [11-3](#)
 - Brick, [11-2](#)
 - defaults, [11-4](#)
 - Delete, [11-28](#)
 - Disable, [11-28](#)
 - Endpoints, [11-2](#)
 - IPSec transport method, [11-7](#)
 - keepalive interval, [11-7](#)
 - LAN-LAN Defaults Editor, [11-5](#)
 - maintain, [11-26](#)
 - Modify, [11-27](#)
 - Receive any proposals, [11-7](#)
 - Save and apply, [11-21](#)
 - Set defaults, [11-5](#)
 - set up, [11-12](#)
 - set up with UDP encapsulation, [11-23](#)
 - View, [11-26](#)
 - Local password, [9-3](#), [9-11](#), [9-11](#), [9-15](#)
 - Create user account, [9-11](#), [9-11](#)
 - Create user group, [9-15](#)
 - How to set up, [9-11](#)
 - Local Presence, [A-1](#)
 - Local presence
 - Destination address mapping, [A-7](#)
 - Source address mapping, [A-7](#)
 - Log-On Data Corporation (XServer application), [8-2](#)
 - Lucent Brick Reflection Protocol, [8-2](#)
 - Lucent IPSec
 - create message for users, [12-43](#)
 - Lucent Netcare Professional Services, [xxii](#)
 - Lucent Proxy Agent, [8-2](#)
-
- M** Max use concurrent, [1-32](#)
 - Max use total, [1-32](#)
 - Modify
 - Dependency mask, [7-12](#)
 - Host group, [2-17](#), [3-15](#)
 - LAN-LAN tunnel, [11-27](#)
 - Proxy Table entry, [8-9](#)
 - Rule, [1-55](#)
 - Move
 - Brick zone ruleset, [1-49](#)
 - Dependency mask, [7-14](#)
 - Host group, [2-19](#), [3-16](#)

- Service group, [4-19](#)
-
- N** NAT see Network Address Translation, [6-1](#)
- Network address translation, [6-1](#), [6-16](#)
 - Destination address mapping, [6-5](#), [6-16](#)
 - Destination port mapping, [6-11](#)
 - Source address mapping, [6-14](#)
 - With a router, [6-12](#)
 - Without a router, [6-14](#), [6-16](#)
 - nocgwzone ruleset
 - nocgwzone, [B-27](#)
-
- P** Packet Data Gateway (PDG) accounting, [12-4](#)
- Packet filtering, [1-2](#)
 - Pool, [6-6](#), [6-9](#)
 - Pre-configured brick zone rulesets, [B-1](#)
 - Pre-configured Brick zone rulesets, [B-3](#), [B-6](#), [B-27](#)
 - Pre-configured brick zone rulesets, [B-27](#)
 - Pre-configured Brick zone rulesets, [B-38](#)
 - Proxies, [8-1](#)
 - Brick Proxy Editor, [8-4](#)
 - Load sharing, [8-18](#)
 - Log-On Data Corporation (XServer application), [8-2](#)
 - Lucent Brick Reflection Protocol, [8-2](#)
 - Lucent Proxy Agent, [8-2](#)
 - Proxy Table, [8-2](#)
 - proxyzone ruleset, [8-11](#), [8-12](#), [8-14](#)
 - Reflection, [8-4](#)
 - TrendMicro (Interscan VirusWall), [8-2](#)
 - Proxy load sharing, [8-18](#)
 - Examples, [8-18](#)
 - Heartbeat, [8-18](#)
 - Round-robin distribution, [8-18](#)
 - Proxy servers
 - rules-based routing and, [1-4](#)
 - Proxy Table, [8-2](#), [8-4](#), [8-9](#), [8-9](#), [8-10](#), [9-6](#)
 - Delete an entry, [8-10](#)
 - Destination port, [8-4](#), [9-6](#)
 - Duplicate an entry, [8-9](#)
 - Examples, [8-18](#)
 - Make an entry, [8-4](#)
 - Modify an entry, [8-9](#)
 - Proxy port, [8-4](#), [9-6](#)
 - Source port, [8-4](#)
 - proxyzone ruleset, [8-11](#), [8-12](#), [8-14](#)
 - Assign to an interface, [8-14](#)
 - Rules, [8-12](#)
-
- R** RADIUS, [9-3](#), [9-19](#), [9-19](#), [9-21](#), [9-21](#), [9-21](#), [9-21](#), [9-21](#)
- attribute codes, [D-1](#)
 - Authentication port, [9-21](#)
 - Authentication service, [9-19](#)
 - Required rule, [9-8](#)
 - RealAudio, [7-16](#)
 - Receive any proposals, [11-7](#)
 - Reflection, [8-4](#)
 - Renumber
 - Rules, [1-57](#)
 - Required attributes
 - Required attributes, [9-29](#)
 - Resource encrypted file, [9-9](#), [9-25](#)
 - Routers, [12-2](#)
 - Rules
 - Activate, [1-57](#)
 - Associate alarm, [1-32](#)
 - Brick-specific, [1-4](#)

- Create, [1-17](#)
- Deactivate, [1-57](#)
- Delete, [1-56](#)
- Duplicate, [1-55](#)
- Modify, [1-55](#)
- Numbering, [1-6](#)
- Renumber, [1-57](#)
- Types, [1-6](#)
- View, [1-53](#)
- Rules-based routing, [1-4](#)
 - add IP addresses for, [1-45](#)

S SA lifetime (Kbytes), [11-10](#), [12-18](#)

- SA lifetime (secs), [11-9](#), [11-10](#), [12-18](#), [12-18](#)
- SecurID, [9-3](#), [9-9](#), [9-25](#), [9-25](#), [9-25](#)
 - Authentication service, [9-25](#)
 - Required, [9-8](#)
 - Resource encrypted file, [9-9](#), [9-25](#)
- Security rules see Rules, [1-17](#)
- Service Group Editor, [4-3](#)
- Service groups, [4-1](#), [4-21](#)
 - Benefits, [4-2](#)
 - Copy, [4-18](#)
 - Create, [4-4](#)
 - Definition, [4-2](#)
 - Delete, [4-19](#)
 - Move, [4-19](#)
 - Provided with LSMS application, [4-21](#)
 - Service Group Editor, [4-3](#)
 - Uses, [4-2](#)
 - View, [4-14](#)
- Services file
 - Services file, [9-10](#)
- Session timeout, [1-31](#)

- SMTP traffic
 - content filtering for, [1-4](#)
- Source address mapping, [6-14](#), [A-7](#)
 - Direct, [6-6](#)
 - Pool, [6-6](#)
- Strong password (SOX compliance), [9-11](#)
- SYN flood protection, [1-34](#)

T Technical support, [xxii](#)

- TrendMicro (Interscan VirusWall), [8-2](#)
- Tunnel endpoints, [11-2](#), [11-2](#), [12-2](#), [12-2](#)
 - Brick, [11-2](#), [12-2](#)
 - Router, [12-2](#)
 - Types, [11-2](#)

Tunnels

- see Client Tunnels, [11-1](#)
- see LAN-LAN Tunnels, [11-1](#)

U UDP encapsulation

- set up client tunnel endpoint with, [12-29](#)
- set up LAN-LAN tunnel with, [11-23](#)

UMA

- See: Unlicensed Mobile Access (UMA)

Unlicensed Mobile Access (UMA), [D-5](#)

User account, [9-11](#), [9-11](#), [9-11](#)

- Create, [9-11](#)

User authentication, [9-1](#), [9-10](#), [9-29](#)

- Application users, [9-2](#), [9-4](#)
- Authentication process, [9-5](#)
- Authentication server, [9-5](#), [9-6](#)
- Authentication timeout, [9-6](#), [9-27](#)
- Client users, [9-2](#), [9-4](#)
- How it works, [9-4](#)
- Local password, [9-3](#), [9-11](#)

- Proxy Table, [9-6](#)
- RADIUS, [9-3](#)
- SecurID, [9-3](#)
- Types of authentication, [9-3](#)
- Types of users, [9-2](#)
- VPN certificate, [9-3](#)
- User groups, [9-15](#)
 - Create, [9-15](#)
 - LSMS-defined, [9-17](#)
- Users, [9-2](#), [9-2](#)
 - Application users, [9-2](#)
 - Client users, [9-2](#)

V View

- Brick zone rulesets, [1-47](#)
- Dependency masks, [7-10](#)
- Host groups, [2-15](#), [3-13](#)
- LAN-LAN tunnels, [11-26](#)
- Rules, [1-53](#)
- Service groups, [4-14](#)
- Virtual brick address, [6-5](#), [6-9](#), [6-14](#)
- VPN certificate, [9-3](#), [9-27](#), [9-29](#), [9-29](#)
 - Authentication service, [9-27](#)
 - How to set up, [9-29](#)
 - VPN certificate, [9-29](#)

