

Lucent Technologies
Bell Labs Innovations



Lucent Security Management Server

Release 9.1

Administration Guide

260-100-017R9.1
Issue 1
August 2006

Copyright © 2006 Lucent Technologies. All Rights Reserved.

This material is protected by the copyright 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	xxii
Where to Find Technical Support	xxii
How to comment	xxiii

1 Getting Started

Overview	1-1
Log On and Off the LSMS Server or Compute Server	1-2
How to Use the Navigator Window	1-7
How to Operate the System	1-10
How to Organize the LSMS Interface	1-17
How to Apply Changes	1-20
Concurrency Control	1-22
Enable Concurrency Control	1-25
Force a Logout of an Administrator	1-27
Where to Begin	1-29

2	LSMS Redundancy	
	Overview	2-1
	LSMS Redundancy Concepts	2-2
	How Redundancy Works	2-6
	Redundant LSMS Monitoring	2-9
	How to Configure a Secondary LSMS or Compute Server	2-12
3	Configuring and Activating a Lucent VPN Firewall <i>Brick</i>[®] Device	
	Overview	3-1
	Before You Begin	3-2
	Configure a Brick on the LSMS	3-8
	Brick Device Failover	3-22
	To Set Up Brick Device Failover	3-26
	To Manually Initiate Failover	3-32
	To Activate a Brick	3-34
4	Configuring Lucent VPN Firewall <i>Brick</i>[®] Device Ports	
	Overview	4-1
	To Configure a Physical Port	4-2
	To Assign a Security Policy to a Port	4-8
	Static Routes	4-20
	Add a Static Route	4-22
	Modify a Static Route	4-25
	Activate or Deactivate a Route	4-26
	Delete a Static Route	4-27
	To Activate a Login Banner on the Brick Serial Port Console	4-28
5	Maintaining a Lucent VPN Firewall <i>Brick</i>[®] Device Configuration	
	Overview	5-1

To View a Brick Snapshot 5-3

To Modify a Brick 5-5

Apply Changes to a Brick 5-6

Delete a Brick 5-9

To Move a Brick Device 5-10

To Reboot a Brick Device 5-11

To Reboot a Brick Device via the LSMS 5-12

Refresh the MAC Table 5-14

ARP and MAC Handling in the Brick 5-16

Static MAC and ARP Assignments 5-18

Initiate a Ping or Traceroute from a Brick 5-20

Download Software to a Standalone Brick 5-22

Download Software to a Failover Brick 5-24

Download Software to Multiple Bricks 5-25

To Configure Intelligent Cache Management 5-28

6 Configuring VLANs on Lucent VPN Firewall Brick® Devices

Overview 6-1

What is a VLAN? 6-2

Why Build VLANs? 6-4

Forwarding Packets and VLAN Boundaries 6-5

Configure and Activate the Brick 6-6

Configure the Brick Physical Ports for VLAN-Tagged Traffic 6-7

Assign a Policy to the Ports 6-13

Associate a Network with a VLAN 6-17

What are VLAN Bridge Groups? 6-20

Enable the Brick to Support VLAN Bridge Groups 6-21

Configuring Bridging Between Specific VLANs 6-22

	Save and Apply the VLAN Configuration	6-23
7	Configuring Lucent VPN Firewall <i>Brick</i>[®] Device Partitions	
	Overview	7-1
	What are Brick Partitions?	7-3
	Configure Brick Partitions	7-4
	Use Static Routes with Partitions	7-6
	Allow Partitions to Intercommunicate with Static Routes	7-7
	Save and Apply the Brick Configuration	7-10
	Interpreting IP Addresses When Brick Partitions Are Configured	7-11
8	Creating Groups and Administrators	
	Overview	8-1
	What is a Group?	8-2
	To Create a Group	8-5
	To Maintain Groups	8-7
	What are LSMS and Group Administrators	8-9
	To Create Administrator Accounts	8-10
	To Assign Groups and Privileges	8-17
	To Maintain Administrator Accounts	8-21
	To Use the LSMS Messenger	8-25
9	Compute Servers	
	Overview	9-1
	What is a Compute Server?	9-2
	How to Configure a Compute Server	9-6
10	Remote Administration	
	Overview	10-1
	What is the Remote Navigator?	10-2

Install the Remote Navigator on Windows	10-3
Install the Remote Navigator on Solaris	10-7
Permitting Remote Administration on the LSMS	10-10
Create the Host Group	10-11
Create the Security Rules	10-12
Log in from a Remote Host	10-15
What Can a Remote Administrator Do?	10-17
11 Using the Configuration Assistant	
Overview	11-1
What is the Configuration Assistant?	11-3
Alarms	11-9
Detailed Policy Audit	11-11
Direct Paging	11-13
GUI and Status Monitor Parameters	11-18
Log Files	11-20
Log Transfer	11-24
Login Banner	11-27
LSMS Web Server	11-29
Reports	11-32
SNMP Agent	11-34
Software Download	11-37
TL1 Alarms	11-43
Tunable Parameters	11-45
User Authentication	11-48
Strong Passwords	11-51
VPN Debugging	11-53

12	Backing Up and Restoring Data	
	Overview	12-1
	Automatic Backup	12-2
	Manual Backup	12-3
	Scheduled Backups	12-6
	Restore Procedure	12-7
	Restore Scenarios on Redundant LSMS	12-10
	Other Restore Scenarios	12-11
13	Task Scheduler	
	Overview	13-1
	What is the Task Scheduler?	13-2
	Schedule Editor	13-3
14	Using the Status Monitor	
	Overview	14-1
	How to Access the Status Monitor	14-2
	How to Interpret the Status Monitor	14-3
	Status Overview Window	14-6
	Administrators Window	14-10
	Brick Status Windows	14-14
	Console Alarms Window	14-27
A	Administer a Lucent VPN Firewall <i>Brick</i>[®] Device Over the Internet from an Unregistered LSMS	
	Overview	A-1
	Background	A-2
	Configure the Brick	A-3
	Assign the Administrative Zone and Enter a VBA	A-4
	Add NAT Rules to the <i>administrativezone</i> Ruleset	A-5

	Activate the Remote Brick	A-8
B	Sizing Guidelines	
	Overview	B-1
	Sizing Tool	B-2
	Determine CPU Capacity	B-4
	Memory Utilization	B-6
	Disk Capacity for Log Files	B-7
	Disk Configuration	B-8
C	Changing the IP Address of the LSMS	
	Overview	C-1
	Procedure (Primary-only LSMS)	C-2
	Procedure (Primary/Secondary Redundant Pair)	C-4
	After the Update	C-6
D	Support for Non-IP Protocols	
	Overview	D-1
	Ethertype and DSAP Files	D-2
	Procedure for Passing Non-IP Packets	D-3
E	New Feature Setup	
	Overview	E-1
	What Do You Have Now?	E-2
	Using New Feature Setup	E-4

Index

List of tables

1	Getting Started	
1-1	Editing Buttons	1-14

List of figures

1	Getting Started	
1-1	Navigator Login Window	1-3
1-2	LSMS Remote Navigator Login Window	1-4
1-3	LSMS Navigator Locked Window	1-6
1-4	Navigator Window	1-7
1-5	Folders Panel	1-8
1-6	Contents Panel	1-9
1-7	Find IP Address search results screen	1-12
1-8	Move Column Header	1-16
1-9	Group (Folders and Subfolders)	1-17
1-10	Concurrency Control Editor	1-25
1-11	Administrators Status Window	1-27
2	LSMS Redundancy	
2-1	LSMS/Computer Servers Editor Window	2-12
2-2	Host Group Editor window (LSMS Host Group)	2-13
3	Configuring and Activating a Lucent VPN Firewall <i>Brick</i>[®] Device	
3-1	Brick Configuration	3-3
3-2	Firewall Configuration	3-5
3-3	LAN-LAN and Client Tunnels	3-6
3-4	Brick Editor (Brick Tab - Primary LSMS)	3-9

3-5	Brick Editor (Brick Tab - Redundant LSMS)	3-14
3-6	Brick Editor (Options Tab)	3-18
3-7	Example of Brick Device Failover Physical Topology	3-25
3-8	Brick Editor (Failover Tab, Brick Failover Enabled)	3-27
3-9	Ping Failover Editor	3-29
3-10	Make/Package Floppy Window	3-36
3-11	Make/Package Floppy Window (with Password Fields)	3-38
3-12	Browser Window	3-39
3-13	Make/Package Floppy Window	3-44
3-14	Make/Package Floppy Windows (continued)	3-45
3-15	Make/Package Floppy Windows (continued)	3-47

4 Configuring Lucent VPN Firewall *Brick*[®] Device Ports

4-1	Brick Editor(Physical Ports Tab)	4-3
4-2	Brick Ports Editor	4-4
4-3	Brick Editor (Policy Assignment Tab)	4-9
4-4	Brick Assignment Policy Tab (Basic)	4-10
4-5	Brick Policy Assignment Editor (Bandwidth Tab)	4-14
4-6	Brick Policy Assignment Editor (TOS Tab)	4-16
4-7	Policy Assignment Tab (Two Rulesets Assigned to Ether1)	4-17
4-8	Brick Editor (Static Routes Tab)	4-22
4-9	Brick Static Route Editor	4-23
4-10	Confirm Deletion Window	4-27

5 Maintaining a Lucent VPN Firewall *Brick*[®] Device Configuration

5-1	Brick Snapshot	5-4
5-2	Apply Brick Window	5-7
5-3	Confirmation Window	5-9

5-4	Warning Window for Rebooting Standalone Brick Device	5-12
5-5	Warning Window for Rebooting Brick Device in Failover Pair	5-13
5-6	Warning Window	5-14
5-7	Warning Windows (Standalone Brick)	5-22
5-8	Warning Window (Failover Brick)	5-24
5-9	Multiple Brick Operations Window	5-26
5-10	Brick Editor (Cache Management Tab)	5-30
5-11	Brick Cache Management Editor	5-31
5-12	Brick Cache Management Editor	5-32
5-13	Confirmation Window	5-33
6	Configuring VLANs on Lucent VPN Firewall <i>Brick</i>[®] Devices	
6-1	Flat Switched Network (no VLANs)	6-2
6-2	Switched Network (with two VLANs)	6-3
6-3	Switched Network (with VALNs and Brick on trunk)	6-3
6-4	Brick Editor (VLAN View)	6-7
6-5	Brick Ports Editor	6-9
6-6	Brick Editor (Policy Assignment Tab)	6-13
6-7	Brick Policy Assignment Editor	6-15
6-8	Brick Editor (VLAN/IP Assignment Tab)	6-17
6-9	Brick/VLAN IP Assignment Editor	6-18
7	Configuring Lucent VPN Firewall <i>Brick</i>[®] Device Partitions	
7-1	Brick Editor (Partition Tab)	7-4
7-2	Brick VLAN Partition Editor	7-5
7-3	Brick Static Route Editor	7-8
8	Creating Groups and Administrators	
8-1	Group Editor	8-5

8-2	GroupEditor Window	8-7
8-3	Administration Editor (Administrator Tab)	8-11
8-4	Administrator Editor (Authentication Tab)	8-13
8-5	Administration Editor (Authentication Service)	8-15
8-6	Group Administrator Privileges Window	8-18
8-7	Administrator Editor (Edit Mode)	8-21
8-8	Contents Panel (Administrators folder)	8-23
8-9	LSMS Messenger	8-25
8-10	Records of Sent Message	8-27
8-11	Messenger Envelope	8-28
8-12	LSMS Messenger (Message Sent)	8-28
9	Compute Servers	
9-1	LSMS Cluster Arrangement of Computer Servers and Redundant LSMS Pair	9-3
9-2	LSMS/Computer Servers Editor Window	9-6
9-3	Host Group Editor window (ComputeServers Host Group)	9-7
10	Remote Administration	
10-1	How to Install the Remove Navigator Window	10-5
10-2	How to Install the LSMS Remote Navigater Window (solaris)	10-9
10-3	LSMS Remote Navigator Login Window	10-15
11	Using the Configuration Assistant	
11-1	Configuration Assistant Window	11-5
11-2	Configuration Parameters Window	11-8
11-3	Alarms Parameter Group	11-9
11-4	Detailed Policy Audit Parameter Group	11-11
11-5	Direct paging parameter Group	11-13
11-6	FIPS Parameters Group	11-16

11-7	General Parameters	11-18
11-8	Log Files Parameters	11-21
11-9	Log Transfer Parameters	11-24
11-10	Same FTP Server	11-26
11-11	Login Banner Parameters	11-27
11-12	Login Banner Parameters	11-27
11-13	LSMS Web Server Parameters	11-30
11-14	Reports Parameters	11-32
11-15	SNMP Agent Parameters	11-35
11-16	Software Download Parameters	11-38
11-17	TLI Alarm Parameters	11-43
11-18	Tunable Parameters	11-46
11-19	User Authentication Parameters	11-49
11-20	Enable Strong Passwords Panel	11-52
11-21	VPN Debugging Parameters	11-53
13	Task Scheduler	
13-1	LSMS Schedule Editor Window (Initial View)	13-3
13-2	LSMS Schedule Editor Window (Initial View)	13-4
13-3	Edit Command Schedule Window	13-5
14	Using the Status Monitor	
14-1	Status Window Toolbar	14-4
14-2	Status Overview Window	14-6
14-3	Administrators Status Window	14-10
14-4	LSMS/LSCS Status Window	14-12
14-5	Brick Lists (All Bricks)	14-16
14-6	Single Brick Status Window	14-18

14-7	Select Summary Time and Date Window	14-21
14-8	Single Brick Ports Window	14-22
14-9	Single Brick Zones Window	14-24
14-10	Select Summary Time and Date Window (Brick Bandwidth Statistics)	14-25
14-11	Single Brick Zones Graphical Display	14-26
14-12	Console Alarms Window	14-27

D Support for Non-IP Protocols

D-1	Apply Brick Window	D-5
-----	--------------------------	-----

E New Feature Setup

E-1	New Feature Setup Window	E-3
E-2	New Option Key Window	E-5

About this information product

Purpose

Welcome to the *LSMS Administration Guide*. The purpose of this manual is to explain how to use Release 9.1 of the Lucent Security Management Server (LSMS) application.

Reason for reissue

Updated for R9.1 features.

Who Should Read this Book?

The *LSMS Administration Guide* is intended to be read by Network administrators who will be using the LSMS application to:

- Configure and install one or more Lucent Virtual Firewall *Brick*[®] devices so that they are communicating with the LSMS
- Configure Compute Servers as an alternative means of polling and collecting log information from Bricks.

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 Administration Guide* explains how to configure and activate a Brick and its associated interfaces. It describes the various functions of a Brick — firewall, tunnel endpoint, VLAN switch — and explains how to configure the Brick once its role has been identified.

This document also explains how to configure Compute Servers as an alternative means of collecting log information from Bricks managed by the LSMS.

The *LSMS Administration Guide* covers the following topics:

Chapter	Purpose
2. Getting Started	<p>This chapter explains how to log on and off the LSMS. It describes the Navigator window, which appears immediately after you log in, and explains how to use this window to access the system's functions.</p> <p>It also explains how to organize the groups, folders, and other objects that appear in the Navigator window.</p>
3. Configuring and Activating a Brick	<p>This chapter explains how to configure and activate a Brick device. It describes the questions you should ask before beginning the configuration, and it explains how to configure a Brick device from a standalone LSMS and a redundant LSMS. It also explains how to activate a Brick device from the LSMS host and from a remote host.</p>
4. Configuring Brick Ports	<p>This chapter explains how to configure the ports on a Brick device and assign security policies to the ports. It also explains how to create static routes and configure the intelligent cache management feature.</p>
5. Maintaining a Brick Configuration	<p>This chapter explains how to modify and delete a Brick configuration. It also explains how to move Brick devices, and how to reboot a Brick device and refresh its MAC table from the LSMS.</p> <p>Finally, it describes how to download a software upgrade to the Brick device from the LSMS.</p>
6. Configuring VLANs on Bricks	<p>This chapter explains how to configure a Brick to recognize, forward and filter VLAN-tagged frames.</p> <p>The Brick device now recognizes IEEE 802.1Q VLAN tagged Ethernet frames when received from the network, and can generate VLAN tagged frames.</p>
7. Configuring Brick Partitions	<p>This chapter explains the concept of a Brick partition, and describes how to configure a partition and how to use static routes with Brick partitions.</p>
8. Creating Groups and Administrators	<p>This chapter discusses the concept of a group, describes the <i>system</i> group (a special group provided with the LSMS), and explains how to create new groups.</p> <p>The chapter also describes the two types of Administrators (LSMS Administrators and Group Administrators), and how to create new Administrator accounts. It explains how to authenticate administrators with RADIUS and SecurID.</p>

Chapter	Purpose
9. Compute Servers	This chapter explains the concept of Compute Servers as an alternative means of collecting Brick log information and explains how to configure a Compute Server as a logging device in LSMS.
10. Remote Administration	This chapter explains how to administer the LSMS from a remote host. It explains how to install the LSMS Remote Navigator on the host (Windows and Solaris), and how to login to the LSMS using the Remote Navigator application. It also explains how to create the host group and security rules needed to allow the remote session to reach the LSMS through a Brick device.
11. Using the Configuration Assistant	This chapter explains how to use the Configuration Assistant to set a number of parameters that affect the system's operation and performance.
12. Backup Up and Restoring Data	This chapter explains how to backup the database and if necessary, restore this database. Configuration data that is managed by the LSMS, including policy, device, and VPN tunnel data, and stored in the database should be backed up on a regular basis.
13. Task Scheduler	This chapter explains how to use the LSMS Task Scheduler to schedule tasks such as database backups or log file transfers.
14. Using the Status Monitor	This chapter explains how to use the Status Monitor. The Status Monitor is a tool for monitoring the status of all Bricks, VPN tunnels, and LSMSs. It also shows all LSMS and Group Administrators currently logged into the LSMS you are logged into, and it displays all console alarm messages.

The *LSMS Administration Guide* also contains five appendices that provide additional information about the system.

What is Not in this Book

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

- How to set up and manage security policies on one or more Bricks.
- How to create components of a security policy, such as host groups, service groups and dependency masks.
- How to set up network address translation (NAT).

- How to set up Brick to proxy incoming and outgoing HTTP, SMTP and FTP sessions.
- How to set up user authentication, using either a database residing on the LSMS, a RADIUS or SecurID server, or X.509 digital certificates.
- How to obtain and install X.509 digital certificates.
- How to set up LAN-LAN tunnels between Bricks.
- How to configure a Brick to serve as the endpoint of a client tunnel.

These and other topics are covered in the *LSMS Policy Guide*. Since these topics pertain primarily to the set up and administration of the software, we recommend that you read this *LSMS Administration Guide* — and perform all required hardware tasks — before you look at the *LSMS 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.
- The LSMS server platform operating system (Microsoft *Windows*[™] or Sun Microsystems *Solaris*[®]).

- 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 Getting Started

Overview

Purpose

This chapter explains how to log on and off the LSMS server or a Compute Server using the LSMS Navigator and LSMS Remote Navigator applications. These are utilities that have been provided with the LSMS for you to use to access the system.

This chapter also describes the Navigator window, which appears immediately after you log in. It explains how to use this window, and how to organize the groups, folders, and other objects that appear in the Navigator window. Finally, the chapter suggests where to turn to begin setting up your environment.

Contents

Log On and Off the LSMS Server or Compute Server	1-2
How to Use the Navigator Window	1-7
How to Operate the System	1-10
How to Organize the LSMS Interface	1-17
How to Apply Changes	1-20
Concurrency Control	1-22
Enable Concurrency Control	1-25
Force a Logout of an Administrator	1-27
Where to Begin	1-29



Log On and Off the LSMS Server or Compute Server

When to use

The computer running the LSMS application is called the LSMS *host*. You can log into the LSMS directly from the LSMS host, or you can log into the LSMS remotely from another host (Windows-based PC or laptop, or a Sun workstation) that has a LAN or Internet connection to the LSMS host.

To facilitate the collection of log data from one or more Lucent VPN Firewall *Brick*[®] devices being managed by the LSMS host, a set of up to ten Compute Servers can be configured and associated with a Primary or Secondary LSMS host through the LSMS GUI. For additional information about Compute Servers, refer to [Chapter 9, “Compute Servers”](#). Once it is added, you can log into a Compute Server (LSCS) directly or remotely from another host, using the LSMS Remote Navigator, in the same way that you would log into the LSMS host.

If you are logging in remotely, your login is secure because the ID and password you enter — as well as the entire remote administrative session — is protected by Triple DES encryption.

Important! To allow remote Administrator sessions (including your own) through the Brick, you first have to create two rules in the Administrative Zone or the NOC Gateway Zone, depending on how the Brick protecting the LSMS host is configured.

See *Chapter 9. Remote Administration* for instructions on creating the two rules required for remote access.

It is possible to install the LSMS Remote Navigator on the LSMS host so that you are running both the LSMS Navigator and the LSMS Remote Navigator on the same machine. The only reason to do this is if you intend to use the LSMS host to log into another LSMS or Compute Server remotely.

Log into an LSMS or Compute Server from the LSMS Host

The LSMS Navigator is installed automatically on the LSMS host when the LSMS application is installed. To log into an LSMS or Compute Server from the LSMS host using the LSMS Navigator, follow the steps below:

- 1 If the LSMS is running on Windows, click the **Start** menu and select:

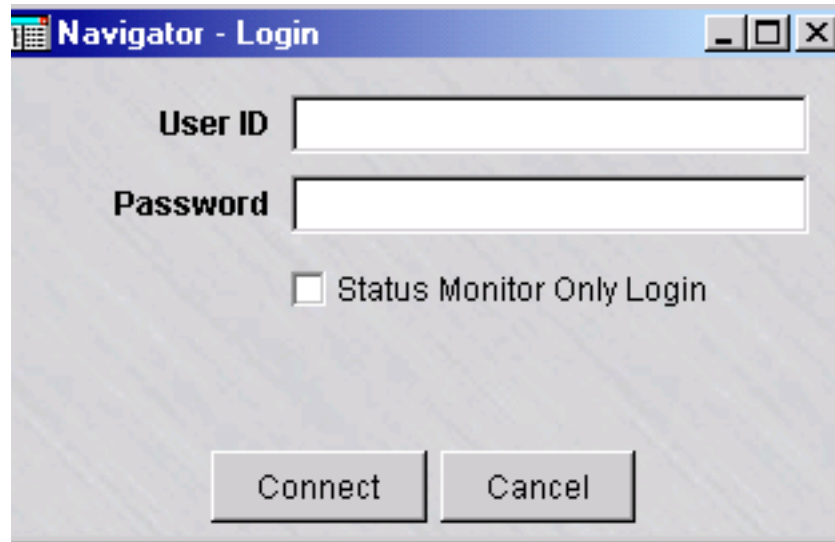
Programs ► Lucent Security Management Server ► LSMS Navigator

If the LSMS is running on Solaris, bring up the desktop menu by right-clicking on the display background and then click on LSMS Navigator or go to the installation root directory (*/opt/isms/lmf* if you used the defaults during installation) and enter:

```
./StartLSMSNavigator
```

from the command line. In either case, the login window shown in Figure 1-1 will appear.

Figure 1-1 Navigator Login Window



-
- 2 Enter your **Admin ID** and **Password**. The Admin ID and password are the ones that were created during LSMS installation, or those given to you by your LSMS Administrator.

If you want to access the Status Monitor of the LSMS or Compute Server, you can check **Status Monitor Only Login**. However, if you later decide you want to access the complete LSMS or Compute Server, you will have to exit the Status Monitor and log into the LSMS or Compute Server again. (See *Chapter 12. Using the Status Monitor* for a discussion of the Status Monitor.)

-
- 3 Click **Connect** or press **[Enter]**. A progress bar will appear and track the progress of the login. When you have successfully logged in, the Navigator window will appear (see "How to Use the Navigator Window" on page 3-5).

END OF STEPS

Log into an LSMS or Compute Server from a Remote Host

If you intend to log into the LSMS or a Compute Server remotely from another host — for example, a desktop at home or a laptop when traveling — you have to install

the LSMS Remote Navigator on that host. Instructions for installing the LSMS Remote Navigator are found in *Chapter 9. Remote Administration*.

To log in from a remote host using the LSMS Remote Navigator, follow the steps below:

- 1 If the remote host is running Windows, click the **Start** menu and select:


Programs ► Lucent Security Management Server ► LSMS Remote Navigator 9.0

If the remote host is running Solaris, go to the installation root directory (*/opt/isms/lmf* if you used the defaults during installation) and enter:

```
./StartLSMSNavigator
```

from the command line. In either case, the Remote Navigator window is displayed (Figure 1-2, “LSMS Remote Navigator Login Window” (p. 1-4)).

Figure 1-2 LSMS Remote Navigator Login Window



- 2 Enter your **Admin ID** and **Password**. The Admin ID and password are the ones that were created during LSMS installation, or those given to you by another administrator.

If you want to access the Status Monitor without logging into the rest of the LSMS, you can check **Status Monitor Only Login**. This is useful if:

- You have a special monitoring room with large screens, and you want to display the graphs to monitor the health of the system, or
- If you want to provide someone with the ability to monitor the system, but you do not want this person to be able to view or change the system's configuration.

3 Enter the URL of the LSMS or Compute Server. The URL is either:

`http://<IP_address>:<port_number>/LSMS`

— or —

`https://<IP_address>:<port_number>/LSMS`

where *<IP_address>* is the IP address of the LSMS or Compute Server and *<port_number>* is the port the web server is listening on. Ports 80 and 443 are the standard ports for HTTP and HTTPS, respectively. The port your web server is using was assigned during installation; if another port was entered, use it instead.

Each URL you enter will be placed in the drop-down list in the URL field, so that each time you enter this URL after the initial entry, you can simply select it from the drop-down list instead of typing it in. You can store multiple URLs in this list in the event that you need to log into more than one LSMS or Compute Server remotely.

4 Click **Connect**. When you have successfully logged in, the Navigator window will appear (see "How to Use the Navigator Window" on page 3-5).

END OF STEPS

Login to a "Locked" Navigator


Once you login to the LSMS Navigator or the LSMS Remote Navigator, your session will remain active until:

- You explicitly log out by clicking on **File** ► **Exit** from the menu bar
- OR -
- The LSMS Services are manually stopped.

However, after a period of inactivity, the LSMS Navigator will automatically lock itself. The Idle Time interval is set through the LSMS Configuration Assistant. For more information on this parameter, please refer to *Chapter 10. Using the Configuration Assistant*.

If you attempt to use a locked LSMS or LSCS Navigator, a window similar to the one shown in [Figure 1-3, “LSMS Navigator Locked Window”](#) (p. 1-6) (depending on whether you are locked out of an LSMS or LSCS) is displayed:

Figure 1-3 LSMS Navigator Locked Window



Click the **Unlock Navigator** button and enter the appropriate password for the administrator whose session is locked.

If you wish to login as a different administrator, simply click the **Logout** option and initiate a new session.

□

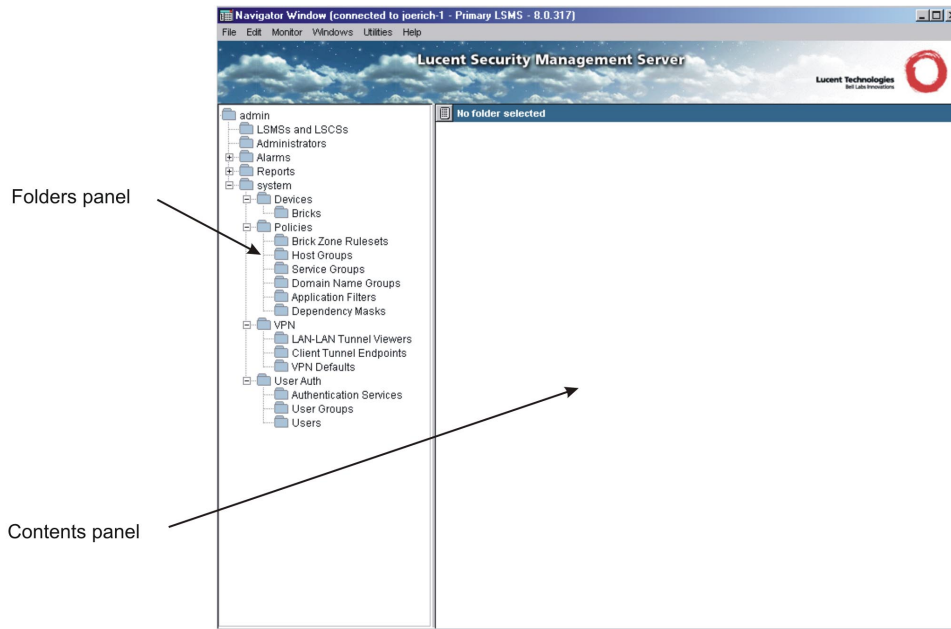
How to Use the Navigator Window

When to use

The window that appears immediately after login is called the **Navigator window**. This window is your doorway to the LSMS. Starting from this window, you will be able to centrally manage the Bricks in your network.

Figure 1-4 shows the Navigator window. The window work area is divided into two panels — a Folders panel on the left, and a Contents panel on the right.

Figure 1-4 Navigator Window



Folders Panel

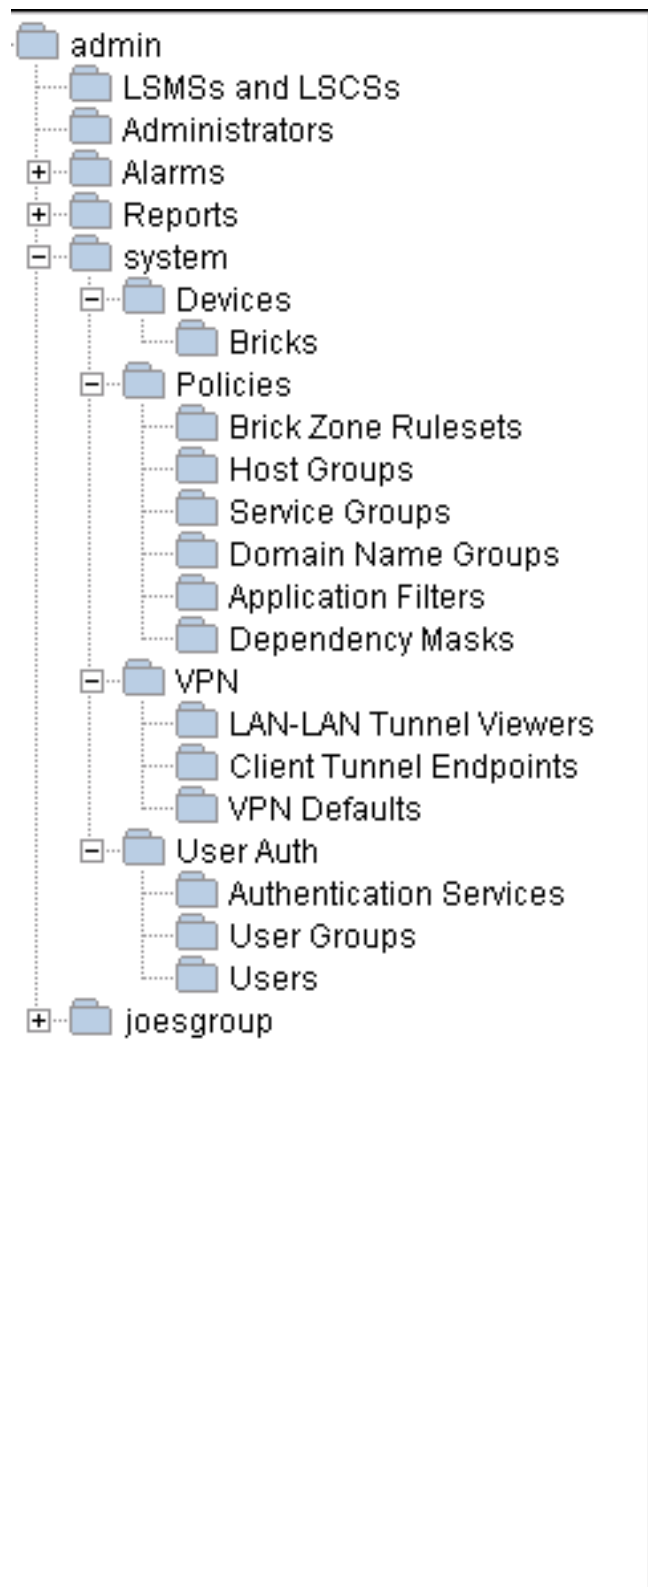
The Folders panel consists of a set of folders and subfolders organized into a hierarchical tree structure. Inside these folders, you will find the devices you are managing, the security policies and tunnel endpoints you have set up, and a variety of other system features and components.

Important! LSMS Administrators see all folders. The folders that Group Administrators see depends on their privileges. For example, Group Administrators with full policy privileges but no device privileges will see all policy folders, but no device folders.


This is explained in detail in *Chapter 8. Creating Groups and Administrators*.

Figure 1-5, “Folders Panel” (p. 1-8) shows a close-up of a Folders panel.

Figure 1-5 Folders Panel





Note that the folder at the very top of the tree structure is the Admin ID of the Administrator currently logged in. Each Administrator always sees his or her own Admin ID.

One level down on the tree structure are the Administrators, Alarms, Reports, and System folders. Alarms and reports are not part of any group, but are specific to each Administrator. The Alarms and Reports folders have subfolders under them. You can tell by the  to their left.

The System folder is a group. This group is provided with the LSMS application and automatically opens every time you log onto the LSMS. If you create additional groups, these will appear below the System group at the same level of the hierarchy. The new groups will automatically have the same subfolders that are originally found in the System group (shown in Figure 2-5). Groups you create will not expand automatically when you log on.

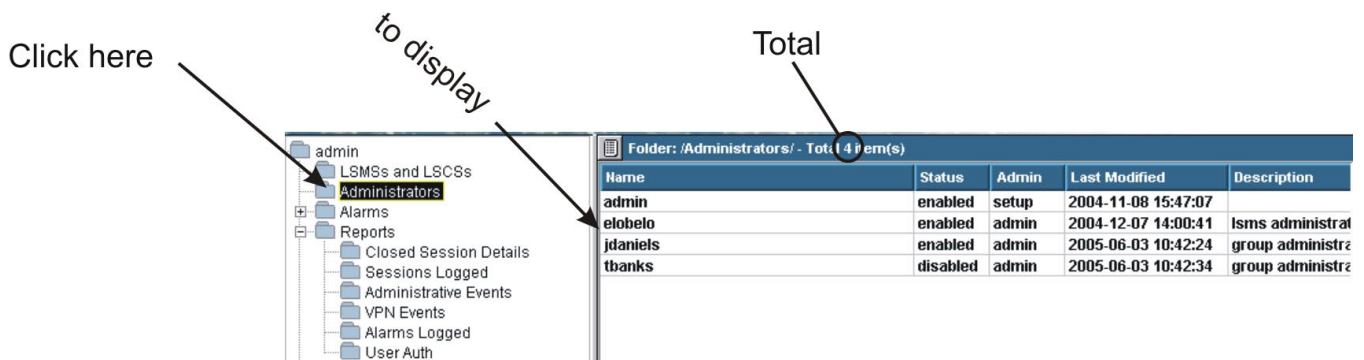
To open a folder that has subfolders, double-click the folder, or click  once.

Contents Panel

The Contents panel displays the contents of the lowest level of folders. You can tell that you have reached the lowest level when a folder no longer has  or has a  to its left.

To display the contents of a folder, click the folder once. Figure 1-6 shows the contents of an Administrators folder. You will note that there are four entries in the Contents panel, and the total is always shown above the entries. The entries are sometimes referred to as *leaves*.

Figure 1-6 Contents Panel



How to Operate the System

Overview

To perform actions using the LSMS, you have to manipulate the folders and their contents. There are several ways to do this:

- By using the menu bar at the top of each window.
- By using the mouse to select folders and leaves, and display a pop-up menu of actions.
- By using the buttons at the bottom of certain screens.

The method you choose to perform any given task does not matter. This redundancy has been deliberately built into the system to make it possible for you to perform any task from any screen in the manner that is easiest for you.

Menu Bar

A menu bar appears across the top left of every major LSMS screen. This allows you to perform almost all LSMS functions from anywhere in the system, without the need to exit the current window and change functions.

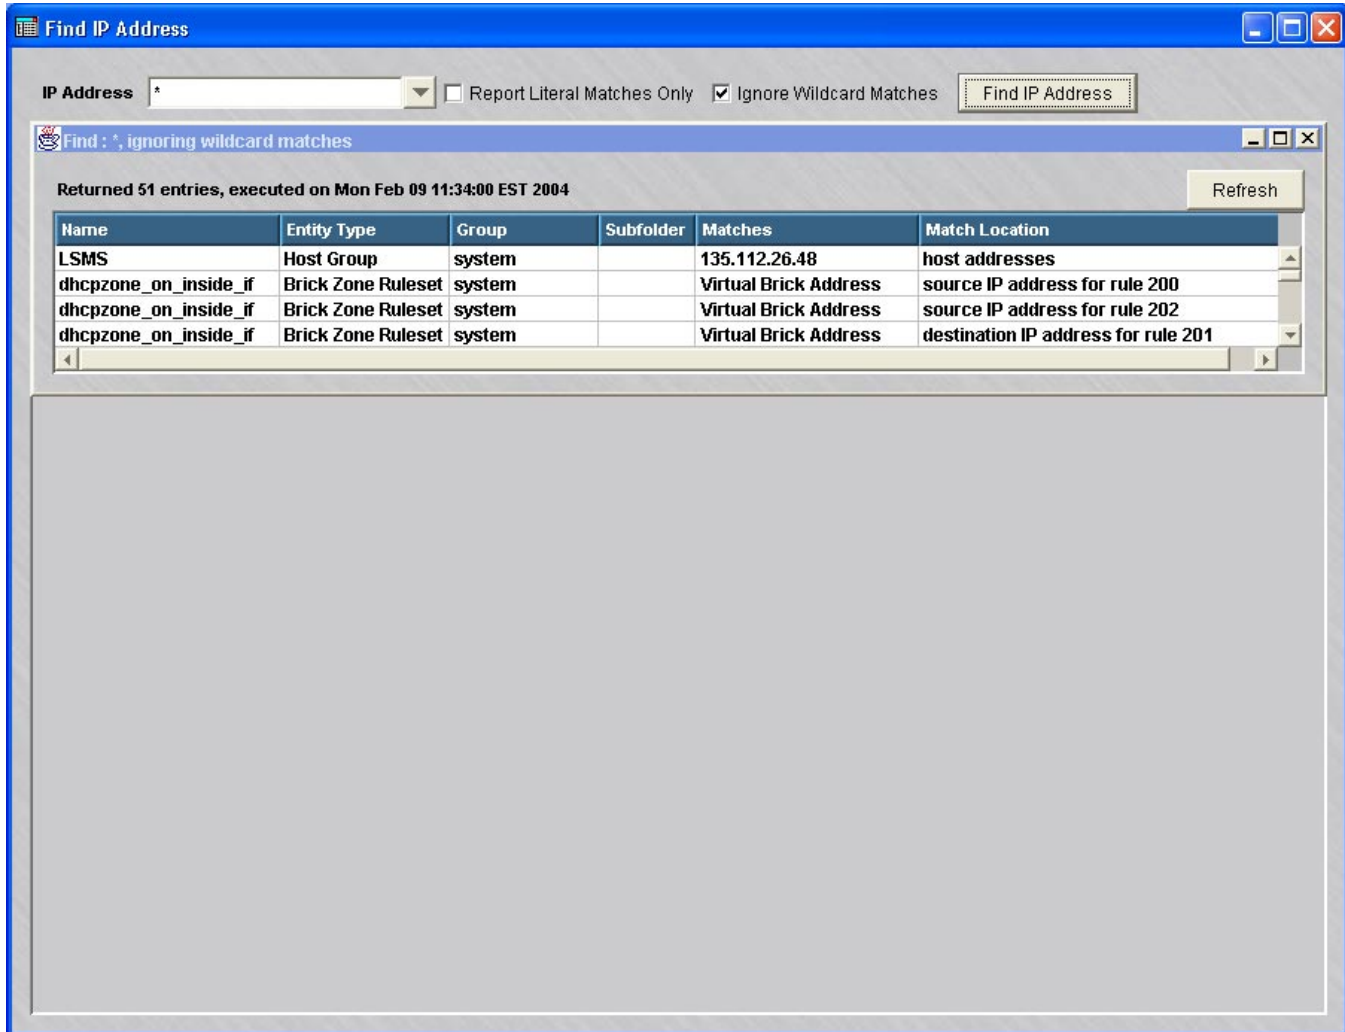
The menu bar contains the following menus and commands:

- File
 - Sub-menu items:
 - New Group - allows you to create, save, and apply a new group
 - New - allows you to create, save, and apply any of the following new entries:
 - Administrator
 - Alarm (Action, Trigger, TL1 Alarm)
 - Report Filter (Closed Session Details, Sessions Logged, Administrative Events, VPN Events, Alarms Logged, User Auth)
 - Device (Brick)
 - Policy (Brick Zone Ruleset, Host Group, Service Group, Domain Name Group, Application Filter, Dependency Mask)
 - VPN (LAN-LAN Tunnel, Client Tunnel Endpoint)
 - User Auth (Auth Service, User Group, User)
 - Group
 - Lock Navigator - locks the Navigator Window with password protection
 - Exit
- Edit
 - Sub-menu items:
 - Edit - allows you to edit the following entries:
 - Administrator
 - Alarm (Trigger, Action, TL1 Alarm)

- Report Filter (Closed Session Details, Sessions Logged, Administrative Events, Alarms Logged, User Auth)
 - Device (Brick)
 - Policy (Brick Zone Ruleset, Host Group, Service Group, Domain Name Group, Application Filter, Dependency Mask)
 - VPN (VPN Defaults, LAN-LAN Tunnel, Client Tunnel Endpoint)
 - User Auth (Auth Service, User Group, User)
 - Group
- Move - allows you to move any of the following entries:
 - Device (Brick)
 - Policy (Brick Zone Ruleset,, Host Group, Service Group, Domain Name Group, Application Filter, Dependency Mask)
 - User Auth (Auth Service, User Group, User)
 - Copy - allows you to copy any of the following entries:
 - Policy (Brick Zone Ruleset, Host Group, Service Group, Domain Name Group, Application Filter, Dependency Mask)
 - User Auth (Auth Service, User Group, User)
 - Delete - allows you to delete any of the following entries:
 - Administrator
 - Alarm (Trigger, Action, TL1 Alarm)
 - Report Filter (Closed Session Details, Sessions Logged, Administrative Events, VPN Events, Alarms Logged, User Auth)
 - Device (Brick)
 - Policy (Brick Zone Ruleset, Host Group, Service Group, Domain Name Group, Application Filter, Dependency Mask)
 - VPN (Client Tunnel Endpoint, LAN-LAN Tunnels)
 - User Auth (Auth Service, User Group, User)
 - Group
 - Find IP Address - a tool that allows you to find all the instances in an LSMS configuration where an IP address is used. It allows you to match a single address, a range of addresses, a subnet mask, a wildcard (*), or one of the special keywords available in the pulldown. A match is defined as any overlap between the IP address input by the admin and any IP address found in the system. Since * will match anything, there is an option, set by default, to ignore these matches so as to avoid having them clutter the output. Another option constrains matches to literal matches - no overlap matching is done when this option is set. Note that setting this option makes the Ignore Wildcard Matches option meaningless.

Matches are retrieved and displayed in a tabular form in a new window that floats in the inner area (see Figure 1-7). A series of different searches can be performed and displayed in separate windows while in the tool.

Figure 1-7 Find IP Address search results screen



Sub-menu items:

- Monitor - displays the LSMS Status Monitor and all console alarms that have been configured to display an on-screen message. The Status Monitor provides the following three views of your operation:

Sub-menu items:

- Status Overview - shows status of all Bricks, number of authenticated users, number of LAN-LAN tunnel endpoints, Brick packets, Brick sessions, Brick sessions by protocol type.
- Administrators - shows total number of administrators logged into an LSMS or Compute Server.

- LSMS/LSCS - shows status of each LSMS or LSCS (Compute Server).
- Brick Status - shows Brick status by category (all Bricks, monitored Bricks, lost Bricks, Bricks up, Bricks not up, Bricks by Parent Folder, Single Brick status, single Brick ports, single Brick bandwidth statistics).

If you are an LSMS Administrator, the Status Monitor shows the name of every LSMS and Group Administrator currently logged in and the status of all configured Bricks.

If you are a Group Administrator, the Status Monitor shows your name and the status of all Bricks over which you have view or full privilege. You will also see the names of all LSMS Administrators currently logged in, plus any other Group Administrators who have privileges for your group(s).

Sub-menu items:

- Console Alarms - accesses screen listing all console alarms with buttons to clear individual alarms or clear all alarms.

- Windows

Sub-menu items:

- Close All Open Windows - allows you to close all open window (except Navigator Window)
- Navigator Windows - lists all open windows and provides easy access to windows hidden behind other windows.

- Utilities - provides the listed utilities for the following devices:

Sub-menu items:

- Brick (Apply, Software Download, Make/Package Floppy, Initiate Failover, Reboot, Refresh MAC, Rehome To, View Brick Snapshot)
- Brick Zone Ruleset (Apply, View Policy History, View Policy Snapshot)
- Group (Apply)
- System Utilities (Configuration Assistant, Restart LSMS Services, LSMS Service Status, LSMS Log Viewer, LSMS Messenger)

- Help - displays the following online help options:

Sub-menu items:

- Contents
- Online Product Manuals
- Error Codes
- Subnet Mask Reference
- About...

The Brick Editor screen has the following additional menu item and commands:

- Brick Utilities (Software Download, Make/Package Floppy, Initiate Failover, Reboot, Refresh MAC, Rehome To, View Brick Snapshot)

Mouse Actions

You can use the right mouse button to perform most of the actions in the File and Edit menus, as well as many of the actions in the Utilities menu.

Right-click on a folder or subfolder in the Folders panel, and a menu will frequently pop up. If a menu does not pop up, no action can be taken on that particular folder, and you must look to a subfolder.

For example, if you right-click the System folder, a menu will pop-up and allow you to create a new group, edit an existing group, apply a group, or delete a group.

You can right-click in the Contents panel, or on an entry in the panel, to display a pop-up menu. You can also double-click an entry in the Contents panel, and if it is editable, it will appear in the appropriate editor to allow you to make any necessary changes.

You can also right-click in many of the other LSMS windows (those that have a viewing panel similar to the Contents panel). The procedures that are explained throughout this manual generally follow the right-click approach because this approach usually involves the least amount of mouse navigation.

Buttons

Certain LSMS windows provide special buttons that duplicate many of the menu and mouse functions. These include editing buttons, which appear on many windows and allow you to select an entry and perform an editing operation on it.

In addition, many of the windows provide special tunnel buttons that you can use to view or terminate existing tunnels, or configure new tunnels.

The buttons always appear unlabeled at the bottom of the window. To display a button's label, the button must be active. If it is not active, you have to select an entry from above to activate it. Once it is active, position the cursor on it to display its label.

The tables below show the editing and tunnel buttons, and explain what each one does.

Table 1-1 Editing Buttons















Button	Label	Purpose
	New	Creates a new entry.
	Duplicate	Duplicates the selected entry.
	Edit	Edits the selected entry.

Table 1-1 Editing Buttons (continued)

Button	Label	Purpose
	Delete	Deletes the selected entry.
	Up/Down	Moves the selected entry one row up or down with each click.
	Activate/ Deactivate	Activates and deactivates a rule in a Brick zone ruleset.

Tunnel Buttons

Button	Label	Purpose
	Client VPN	Displays the Client Tunnel Endpoint Editor. This window lets you configure a Brick to serve as the endpoint of a client tunnel.
	LAN-LAN VPN	Displays the LAN-LAN Tunnel Editor. This window lets you configure both endpoints of a LAN-LAN tunnel.
	Clear Tunnel Viewer	Removes all entries from the LAN-LAN tunnel viewer.
	View Folder Tunnels	Displays a Browse window that allows you to select a Brick folder. All tunnels using the devices in the folder will then be displayed in the LAN-LAN tunnel viewer.
	View Device Tunnels	Displays a Browse window that allows you to select a specific Brick. All tunnels using that Brick will then be displayed in the LAN-LAN tunnel viewer.
	Terminate All Sessions	Terminates all VPN client sessions.
	Terminate Session	Terminate the selected VPN client session.
	Refresh Tunnels	Refreshes the status of all the tunnels in a LAN-LAN tunnel viewer.

In addition, a Policy Snapshot button  appears on the Brick Interface Editor. It allows you to select a port and display a summary of the security policy associated with that port.


An Import Services Button  appears on the Service Group Editor. It allows you to import services from other services groups into a service group you are creating.

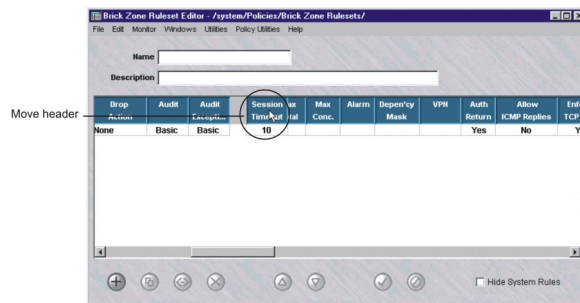
Table Columns

Many LSMS windows are formatted as tables, with columns and rows. In a number of cases, the tables contain more columns than can fit on one screen page, and so you must use the horizontal scrollbar at the bottom of the screen to display these columns.

For your convenience, you can rearrange the order of the columns to suit your needs. If a particular column contains information that is important, you may want to move it so that it displays when the screen first appears, without your having to resort to the scrollbar. You may also want to move certain columns next to each other because they provide information that is related.

To move a column in a table, simply position the mouse cursor on the column header, left-click the mouse to grab the column header, and then drag the column header left or right until it is positioned where you want. [Figure 1-8, “Move Column Header”](#) (p. 1-16) shows a column header being dragged to the right.

Figure 1-8 Move Column Header



How to Organize the LSMS Interface

Overview

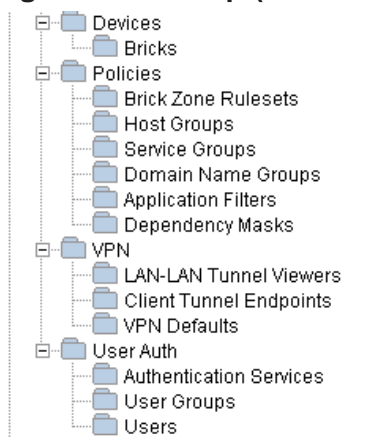
One of the first challenges facing an Administrator is to organize the various objects the LSMS will be managing. These objects include devices, policies, tunnels and user authentication components.

These objects are what you see when you look at the screen — and by organizing them precisely and logically, you can make the system that much easier to use, and your network security that much easier to manage.

Groups and Folders

The LSMS provides two mechanisms for organizing objects — groups and folders. A group is a collection of objects that are managed together. These objects are represented on the screen as nested folders and subfolders. Every group contains the folders and subfolders shown in [Figure 1-9, “Group \(Folders and Subfolders\)”](#) (p. 1-17). However, you must have the appropriate privileges (either view or full) to see all folders.

Figure 1-9 Group (Folders and Subfolders)



In addition to the folders and subfolders shown above, it is possible to create additional levels of subfolders. You can create subfolders under the Bricks subfolder, and also under every Policies subfolder except Dependency Masks.

System Group

One group is provided with the LSMS application. This group is called **System**, and whenever an LSMS administrator logs in, it is automatically opened and displayed on screen.

The **System** group is different from any other groups you may create in two important ways:

- Rulesets from other groups can be applied to Bricks in the **System** group. This is to permit a Managed Service Provider to maintain control over all configured Bricks by managing them in the **System** group, while at the same time allowing each customer's Group Administrator to control the customer's security policy. The **System** group is the only group that has the pre-configured rulesets needed to set up the management tunnel. These are the *nocgwzone* Brick zone ruleset and the *mgmt-tunnel* router tunnel ruleset.

Guidelines

The decision as to whether to place all your devices, policies, and tunnels in the **system** group, or whether to create additional groups, is at your discretion as an Administrator.

As a general rule, however, we recommend that only ISPs, and extremely large enterprises, create additional groups. ISPs in particular may find it useful to create groups for each customer, because each customer needs to be managed separately, and each customer may have a large internal organization.

It is important to keep in mind that — except for host groups, service groups and application filters, which can be given global status — objects such as devices and rulesets cannot be shared across groups. A tunnel can cross groups, but the host groups used in the tunnel definition cannot (unless they are made global). If, for example, you want to create a LAN-LAN tunnel from a device in one group to a device in another, you will not be able to take advantage of the ability of devices in the same group to easily share objects such as devices and rulesets.

Therefore, for most enterprises, we recommend using only the **System** group, and creating additional folders as needed. For example, if you are managing a large number of Bricks, you will probably want to create subfolders to organize the devices by location, department, and function. Similarly, if your security policies comprise many different rulesets, you can create subfolders to organize them logically.

An additional benefit of creating subfolders is that small subfolders generally display faster than large folders. This can be especially noticeable if you are logged onto the LSMS remotely.

Scenarios

The following are three possible scenarios for organizing the LSMS interface:

1. *All components in the System group*

As indicated above, this is the preferred scenario for most organizations. If additional hierarchy is required — if, for example, you want to organize your Bricks by department — folders and even subfolders can be created.

2. *Devices in System group, policies in other groups*

In this scenario, all *Bricks*[®] are configured in the Bricks folder of the System group, but the Brick and router rule sets and other policy components are in the Policies folders of other groups.

This scenario is possible because Group Administrators with full policy privilege in their group can automatically apply policies and tunnels to Bricks in the System group.

3. *All components in separate groups*

This scenario is primarily intended for ISPs, who need to keep the management of their individual customers' devices and policies completely separate.

Nonetheless, some ISPs may still prefer the second scenario, in which they maintain control over the Bricks — which are in the System group — while their customers maintain control over the policies and tunnels in their own groups.

LSMSs and Compute Servers (LSCSs)

A folder in the Folder panel called **LSMSs and LSCSs** is accessed to view, edit, and create the secondary LSMS(s) or Compute Server(s).



How to Apply Changes

Overview

Whenever an Administrator makes a change to information in the LSMS that affects a configured device, that change has to be applied to the affected Brick. Although the changes may be saved in the LSMS database, they will not take effect until they are applied to the device.

When to Apply

As a general rule, you have to apply any updated information that has been downloaded to a Brick. For example, Brick zone rulesets are assigned to Brick ports; if you assign a new ruleset to a port, that change has to be applied to the Brick.

Similarly, if you add a new rule to a ruleset, or change an existing rule, that change has to be applied to every Brick to which the ruleset has been assigned.

If you create a new Brick zone ruleset, it is applied automatically when you apply the ruleset assignment.

Similarly, if you add a new user account or create a new user group, you do not have to perform an apply. However, when you create a rule that uses the new user group, then you have to apply the ruleset.

What to Apply

There are five different types of apply actions that an Administrator with the appropriate privileges can perform. The type of apply you perform determines which devices are updated, and what information they are updated with.

When you perform the apply, the LSMS always displays a window that indicates the Bricks that will be updated by the apply. The window does not allow you to select specific devices. If you only want to update some of the devices shown, you will have to cancel the apply and perform a different type of apply.

For example, as the table below indicates, a group apply updates every device in the group. If you only want to update one device in the group (such as single Brick), then you would have to perform a Brick apply instead. The table below explains:

Apply	Result
Group	Updates each Brick in the group with all device, policy, and tunnel changes that apply to that device.
Brick	Updates the selected Brick with all device, policy, and tunnel changes that apply to it.

Apply	Result
Brick Zone Ruleset	Applies the changes in the selected Brick zone ruleset to each Brick the ruleset has been assigned to.
LAN-LAN Tunnel	Applies the changes in the LAN-LAN tunnel configuration to the <i>Bricks</i> [®] functioning as the tunnel endpoints. Also applies the zone rulesets of the tunnel endpoints.
Client Tunnel	Applies the changes in the client tunnel configuration to the Brick that is functioning as the client tunnel endpoint. Also applies the zone rulesets of the tunnel endpoints.

How to Apply

There are a number of ways to perform an apply. Regardless of which LSMS window is displayed, you can perform any apply, except a LAN-LAN or client tunnel, from the Utilities menu. In addition, there are other ways to perform each of the apply actions. The table below explains:

Apply	Do this...
Group	Right-click the appropriate group folder and select Apply from the pop-up menu.
Brick	Right-click the Brick in the Navigator window and select Apply from the pop-up menu — or — Select Save and Apply from the File menu in the Brick Editor after making changes.
Brick Zone 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 changes.
LAN-LAN Tunnel	Select Save and Apply from the File menu in the LAN-LAN Tunnel Editor after making changes.
Client Tunnel	Select Save and Apply from the File menu in the Client Tunnel Endpoint Editor after making changes.



Concurrency Control

Overview

LSMS is a carrier grade centralized management system capable of managing a large number of objects (Bricks, zone rulesets, Host groups, tunnels, service groups, and so forth). The group-based model allows the creation of multiple management domains, where each group contains a set of resources. Multiple administrators may have access and the desire to modify the same object simultaneously. This raises the potential for problems caused by simultaneous changes to an object by multiple administrators, or editing of an outdated instance within an object by an administrator.

The Concurrency Control feature prevents changes from being made to a managed object, such as a Brick or zone ruleset, by more than one administrator at a time. With Concurrency Control enabled, an object opened for **Edit** by an administrator is “locked out” to other administrators until the managing administrator completes and saves the changes.

Multiple administrators can open the same object in **View** mode, but only one administrator can open the object in **Edit** mode.

Edit mode

To make changes to an object, right-click on the name of the object in the Contents Panel and select **Edit** from the pop-up menu, or simply double-click on the object. The associated Editor window is opened, and the object is now in 'Edit' mode. Modifications can be made as needed.

If the Concurrency Control feature is enabled, the system prevents another administrator from opening the same object in 'Edit' mode and making changes at the same time. The operation is denied and a dialog box is displayed, informing the administrator that another administrator is currently editing the object. More detailed information, such as Admin Name, telephone number, and so forth are displayed if **Display Contact Information** is selected when enabling the Concurrency Control feature. Refer to the procedure [“Enable Concurrency Control”](#) (p. 1-25) for instructions.

View mode

To view the current configuration settings of an object, right-click on the name of the object in the Contents Panel and select **View**. The associated Editor window is opened in 'View' mode; the contents of the window display is shown in “view-only” mode and the parameter fields are greyed-out so no modifications can be made. An object can be opened in 'View' mode by multiple administrators at the same time

Brick console

As part of the Concurrency Control feature, LSMS prohibits more than one Brick console to be opened for a particular Brick at the same time. However, it is possible to open another Brick console by directly connecting to the serial port on the Brick.

Enabling concurrency control

By default, Concurrency Control is disabled. The feature is enabled via a dialog box that is accessed from the Utilities menu of the LSMS Navigator. Refer to the procedure [“Enable Concurrency Control”](#) (p. 1-25) for instructions.

Displaying contact information

An option can be enabled to display detailed contact information about the administrator who currently has an object “out for edit” and a means to send an instant message to that administrator.

If another administrator attempts to edit the object at the same time, and the **Display Contact Information** option is enabled, the following information about the administrator who has the object “out for edit”:

- Full Name
- Telephone #
- Pager #
- Name and IP address of the device on which it is opened (such as Navigator, Remote Navigator or Compute Server)
- Amount of time that the object has been out for edit

Important! The **Telephone#** and **Pager #** fields must be provisioned in the Administrator profile; otherwise, these fields will be blank.

The informational display also includes a button to send an instant message to communicate and coordinate activities with the other administrator.

Refer to the procedure [“Enable Concurrency Control”](#) (p. 1-25).

Lock status timeout

In a multi-LSMS/Compute Server environment, the amount of time needed to poll each device and determine the “out for edit” status of an object can be considerable. To avoid unnecessary lockouts, a timeout interval can be configured via the **Lock Status Timeout** field on the Concurrency Control Editor window. The default timeout interval value is **10** (seconds). When this time interval has elapsed, and the lockout information could be not retrieved in the time period, a warning message is issued to the LSMS GUI that information on whether an object is out for edit could not be obtained from every networked device and that there could be a possible concurrency violation.

For instructions on how to enable this option, refer to the procedure [“Enable Concurrency Control”](#) (p. 1-25).

Force logout

As part of the Concurrency Control feature, the system provides an option for an LSMS administrator to force the logout of any administrator from the LSMS, including another LSMS administrator. When the force logout feature is used, an AdminEventsLog record is created that shows the LSMS administrator who forced the logout of an administrator and the administrator who was logged out. This option can be useful, for example, to log out an administrator who has an object out for edit and has locked out other administrators from being able to access the same object for editing.

For instructions on how to force logout of an administrator, refer to [“Force a Logout of an Administrator”](#) (p. 1-27).



Enable Concurrency Control

When to use

Use this procedure to enable the Concurrency Control feature and, optionally, configure display contact information and the lockout timeout interval for this feature.

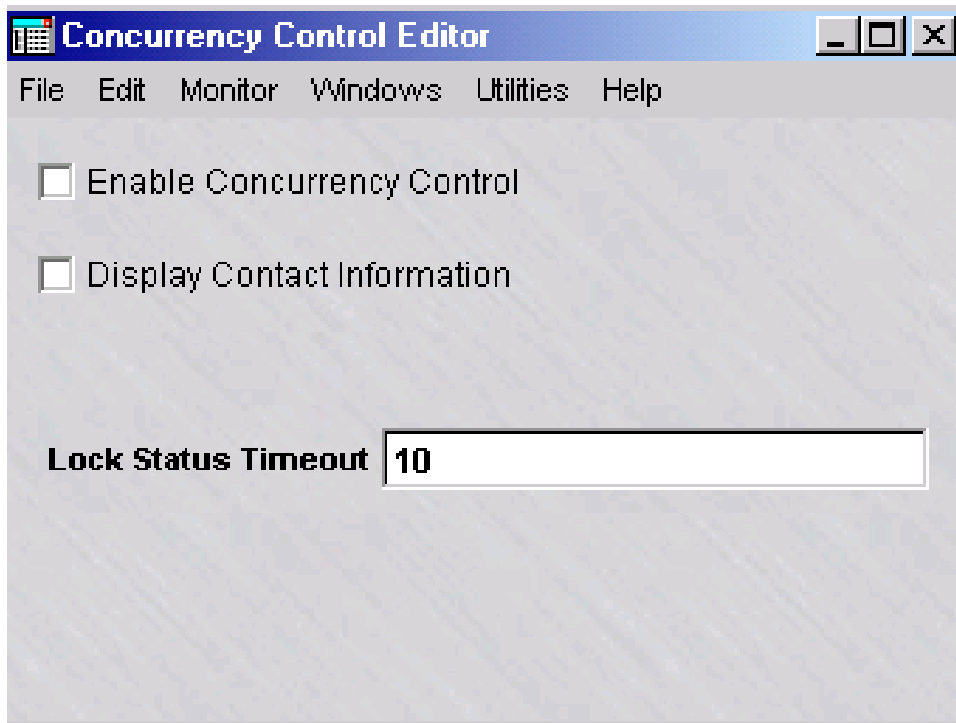
Task

Complete the following steps to enable the Concurrency Control feature and configure the related options, as necessary.

- 1 From the menu bar, select **Utilities > Edit LSMS Parameters**

Result The Concurrency Control Editor window is displayed ([Figure 1-10](#), “Concurrency Control Editor” (p. 1-25)).

Figure 1-10 Concurrency Control Editor



-
- 2 Configure the following options as needed:
- **Enable Concurrency Control**—click this checkbox to enable the Concurrency Control feature. The feature is disabled, by default.
 - **Display Contact Information**—click this checkbox to display detailed contact information about an administrator who has an object “out for edit” when the Concurrency Control feature is enabled.
 - **Lock Status Timeout**—change the lockout status timeout interval, as needed. The default value is **10** (seconds).

-
- 3 When you are finished configuring the settings, select **Save and Close** from the File menu on the Concurrency Control Editor window.

Result The settings are saved and the Concurrency Control Editor window is closed.

-
- 4 To just view (not edit) the current Concurrency Control feature settings, from the menu bar, select **Utilities > View LSMS Parameters**

Result The Concurrency Control Editor window is displayed in view-only mode. The fields are greyed out and cannot be changed.

END OF STEPS



Force a Logout of an Administrator

When to use

Use this procedure to force a logout of an administrator from the LSMS.

Important! You must be an LSMS administrator to log out another administrator.

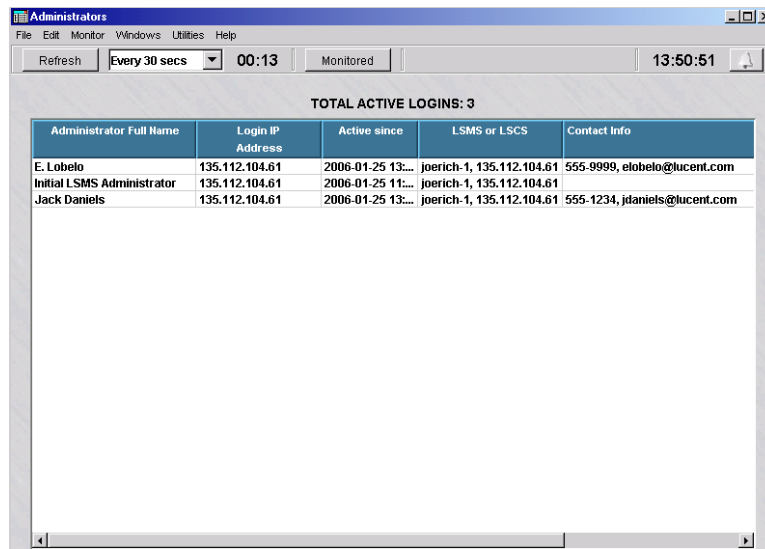
Task

Complete the following steps to force a logout of an administrator from the LSMS.

- 1 From the Menu bar, select **Monitor > Administrators**.

Result The Administrators Status window is displayed (Figure 1-11, “Administrators Status Window” (p. 1-27)).

Figure 1-11 Administrators Status Window



- 2 Right-click on the administrator to be logged out and select **Log Out** from the pop-up menu.

Result A confirmation dialog box is displayed, asking if you are sure that you want to log out the selected administrator.

- 3 Choose **Yes**.

Result The selected administrator is logged out of the LSMS, and is removed from the Administrators Status window.

If the logged out administrator attempts to access the LSMS GUI, a dialog box is displayed on that administrator's GUI instance with an error message *Session Terminated by LSMS*.

-
- 4 To close the dialog box, the logged out administrator needs to click **Ok**.

 - 5 From the LSMS administrator GUI instance, select **File > Close** to close the Administrators Status window.

END OF STEPS



Where to Begin

Guidelines

At this point, you should have a sufficient grasp of LSMS basics to begin... but where? The following suggests one standard approach, and indicates where in the documentation you can go to find instructions for each step:

1 *Configure and install the Brick protecting the LSMS host*

The first Brick you should install is the Brick that is protecting the LSMS. This Brick is usually connected directly to the LSMS.

See: Chapter 3. Configuring and Activating a Brick.

2 *Configure and install any additional Bricks .*

Deploy the other Bricks in your network. Develop rulesets to determine which traffic will be permitted through them, and which will be dropped. Assign these rulesets to ports on the Bricks.

See: Chapter 3. Configuring and Activating a Brick.

Chapter 2. Brick Zone Rulesets in the LSMS Policy Guide

3 *Configure any Bricks to be LAN-LAN or client tunnel endpoints*

Set up LAN-LAN and client tunnels. Provide tunnel endpoint addresses for the ports that will be used, and create the necessary ISAKMP/IPSec security associations.

See: Chapter 10. LAN-LAN Tunnels and Chapter 11. Client Tunnel Endpoints in the LSMS Policy Guide

4 *Set up user authentication*

If you will be setting up client tunnels, you have to authenticate the users. Decide whether to create a database on the LSMS, or use an external database such as RADIUS or SecurID. Create the required authentication services.

See: Chapter 8. User Authentication in the LSMS Policy Guide

END OF STEPS



2 LSMS Redundancy

Overview

Purpose

This chapter explains the concept of LSMS Redundancy and describes how to configure Primary and Secondary LSMSs in your operating environment.

Contents

LSMS Redundancy Concepts	2-2
How Redundancy Works	2-6
Redundant LSMS Monitoring	2-9
How to Configure a Secondary LSMS or Compute Server	2-12



LSMS Redundancy Concepts

Overview

To ensure the reliability, high availability, and data integrity of the network security environment, LSMS supports the concept of redundancy. In its basic form, two LSMSs can be installed and configured as a redundant pair. One LSMS takes over the management of Lucent VPN Firewall *Brick*[®] devices and associated security policies in the network in the event that the other LSMS fails for some reason.

In a basic redundant pair, one LSMS is installed and designated the Primary LSMS and the other LSMS is installed and designated the Secondary LSMS. Both LSMSs are active and share the same database. Each LSMS can be set up to manage its own set of Bricks, security policies, and tunnels. While configuration of redundant LSMSs is highly recommended, a single LSMS can be installed and configured as a Primary LSMS without a Secondary LSMS.

The common database is built on the Primary LSMS and replicated on the Secondary LSMS. This database is updated periodically over the network, and any user-initiated zone ruleset or interface parameter modification is shown immediately on either LSMS.

Heartbeat/keepalive messages are exchanged between the Primary LSMS and Secondary LSMS to establish connectivity. When connectivity is interrupted between redundant LSMSs, each LSMS keeps track of interim changes made in its own version of the database. When connectivity is restored, any interim changes made during the interruption in connectivity are reconciled in the common database.

LSMS redundancy in network design

The LSMS redundancy concept can be extended beyond the basic redundant LSMS pair to keep pace as the traffic volume and complexity of the network grows. A single Primary LSMS can be connected with up to three Secondary LSMSs, for added capacity and reliability, and to manage security for large-scale, multi-customer/multi-site networks, providing seamless redundancy and a single integrated view of the security policies and Bricks within the network to administrators. Each LSMS can handle up to 10,000 records per second, which represents the number of individual log requests retrieved from all managed Bricks, and up to 1000 managed Bricks, provided that the total logging rate from all managed Bricks does not exceed 10,000 records per second.

Compute servers

To further enhance the capacity and security management capabilities of the LSMS, one or more devices known as Lucent Secure Compute Servers (LSCSs), otherwise known as Compute Servers, can be linked to a Primary or Secondary LSMS to serve as log collection points for Brick log data, which frees up the computing resources of the LSMS for other activities. A Compute Server has all of the same basic firewall

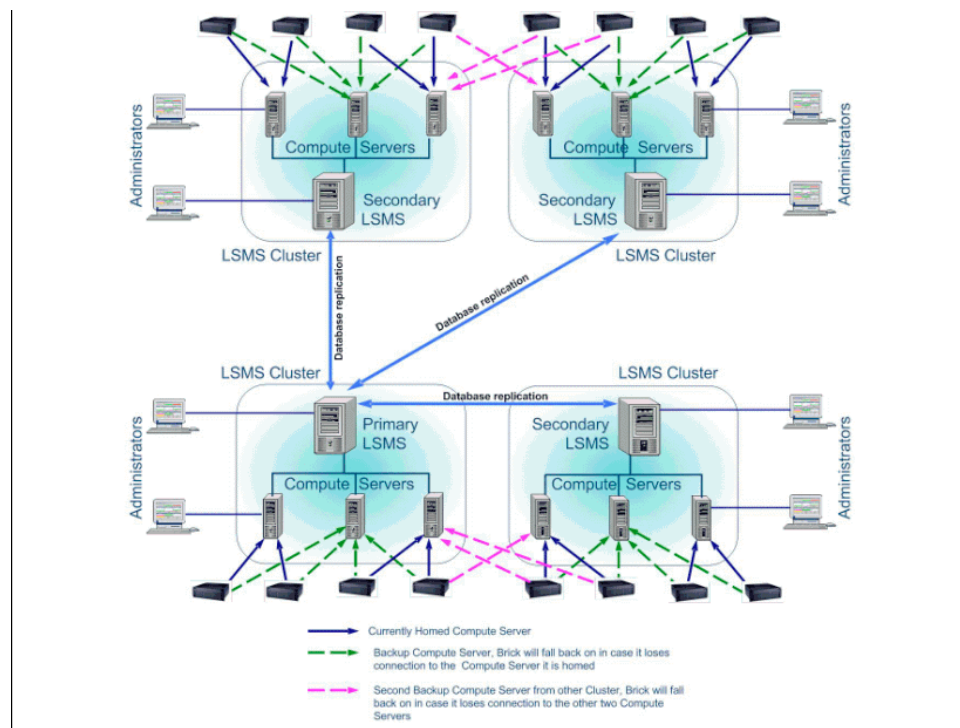
management functions of the an LSMS, except it does not maintain a copy of the LSMS database. A Brick can be accessed or configured from an LSMS, Compute Server, or the Brick console.

The LSMS and its related Compute Servers can be configured as a unit or *cluster* of up to five Compute Servers, providing an additional level of redundancy for communications with a Brick.

For additional information about Compute Servers, refer to [Chapter 9, “Compute Servers”](#).

Figure: scalability of the LSMS network design

A basic redundant LSMS pair (Primary/Secondary LSMS) can be installed and configured to ensure uninterrupted, secure traffic between the managed Bricks and protected devices. Additional Secondary LSMSs and Compute Servers can be added later as network security requirements grow. The following figure illustrates how to optimize the use of redundant LSMSs and Compute Servers for large scale, multi-site networks.



Installation

When you install a Primary or Secondary LSMS, the license key entered while running the installation program designates the LSMS being installed as either a Primary or Secondary LSMS.

Instructions for installing a Primary or Secondary LSMS are provided in the *LSMS Installation Guide*.

If you are installing a Primary LSMS, the capability to manage additional Bricks or IPSec users, or optional new features can be added to the basic installation by using the New Feature Setup function described in [Appendix E, “New Feature Setup”](#). Any Secondary LSMS installed and associated with the Primary LSMS automatically receives the feature option keys from the Primary LSMS.

If you are installing a Secondary LSMS, the following general steps should be followed:

1. Install the Primary LSMS, using the procedure for a *Windows™* or *Solaris®* server platform provided in the *LSMS Installation Guide*.
2. Log into the Primary LSMS and add the Secondary LSMS to the database, associating it with the Primary LSMS, using the procedure [“How to Configure a Secondary LSMS or Compute Server”](#) (p. 2-12).
3. Install the Secondary LSMS, using the procedure for a Windows or Solaris server platform provided in the *LSMS Installation Guide*.

While installing the Secondary LSMS, you must enter the same Secondary LSMS Name that was used while adding the Secondary LSMS on the Primary LSMS. The installation program also prompts for the IP address of the Primary LSMS so the Secondary LSMS can communicate with the Primary LSMS.

The above steps also apply for a Compute Server; Compute Servers are not supported on a Solaris server platform. To configure a Compute Server, follow the procedure [“How to Configure a Compute Server”](#) (p. 9-6).

It is possible to install a Primary LSMS and defer the installation of one or more Secondary LSMSs to a later date. If this is done, however, all Bricks managed by the Primary LSMS will require a manual update (save, apply, and reboot) when the Secondary LSMS(s) is installed.

Set LSMS/LSCS priority

Once the Primary LSMS, Secondary LSMS(s), and Compute Server(s) have been installed and configured, the order, or priority in which Secondary LSMSs or Compute Servers take over management of a Brick in the event that the main (Priority 1) LSMS loses connection with the Brick or reboots. The order and priority in which the LSMSs/LSCSs take over is configured in the Home LSMS/LSCS Priority Table of the Brick Editor.

Entries in the Priority Table can be edited, deleted, or rearranged.

The LSMS from which the Brick is configured automatically becomes the Priority 1 LSMS. This is the management device to which the Brick is initially *homed*. A Brick

always attempts to home to its Priority 1 LSMS after rebooting or after LSMS services have been restarted.

If the Secondary LSMS in a redundant pair has been configured on the Primary LSMS and installed, its name and IP address appear automatically as the Priority 2 LSMS in the Home LSMS/LSCS Priority Table, unless you manually change the Priority 2 entry. Up to five LSMSs or Compute Servers can be entered in the priority list.

If a Brick is currently homed to the first Compute Server in the list, and the connection to that Compute Server is lost, the Brick re-homes to the next available Compute Server or Primary/Secondary LSMS in the priority list.

A Brick can be configured from any LSMS (Primary, Secondary, or Compute Server). The LSMS where the Brick is created is automatically configured as the Priority 1 LSMS for the Brick. You can add to or reorder the priority list for a Brick as needed using the Home LSMS/LSCS Priority Table in the Brick Editor.

By default, the Brick automatically rehomes to the Priority 2 device if communications is lost with the Priority 1 device. The **LSMS Rehome Options** box on the Brick Editor determines what happens when the Brick re-establishes contact with the Priority 1 device.

If Compute Servers are used for logging purposes, it is recommended that one Compute Server be specified as the main logging server, a second Compute Server from the same LSMS cluster be designated on the priority list as the first alternate, and another Compute Server from a different LSMS cluster be designated as the second alternate.

Details about provisioning the rehome priority list and related parameters for a Brick are provided in the procedure [“Basic Configuration - LSMS Redundancy”](#) (p. 3-14).

□

How Redundancy Works

Overview

LSMS redundancy requires database synchronization and a heartbeat between the Primary and Secondary LSMS(s).

Database Synchronization

The members of the LSMS redundancy configuration synchronize their data with each other at the time that data is modified. The synchronization is initiated within five seconds of the time the change is made.

Heartbeat

In a redundant LSMS configuration, each LSMS transmits *heartbeat* messages to the other across the network in both directions. Heartbeat messages indicate connectivity and the availability of each LSMS to the other.

If a heartbeat request fails to complete within a specified period of time, the request is considered a time-out or missed heartbeat. The default is to transmit one heartbeat every second; if five heartbeats are missed, then failover occurs and an alarm is automatically generated.

Once connectivity is restored between the LSMSs, the LSMS that stayed active immediately copies any database changes that occurred during the failure to the LSMS that has come back on-line.

Load Sharing

An LSMS and LSCS, in combination, share the load for Brick log data. Hence, the terms LSMS and LSCS are synonymous in the context of the following discussion on load balancing.

LSMS redundancy supports any degree of load sharing. Here, for example, are two possible scenarios:

- Home all the Bricks to one LSMS, and use the other LSMS for administrative duties
- Home half of the Bricks to one LSMS and half to the other

In the first scenario, one LSMS is the workhorse, attending to all the tasks required for normal firewall and VPN operation, such as keeping logs.

The other LSMS is used as a Brick maintenance platform from which administrators can provision new Bricks, apply policy or brick changes, add users, user groups, host groups and service groups, and perform all other tasks that require human intervention.

There are trade-offs with such an approach, however. If all Bricks are homed to one LSMS, Brick provisioning tasks, performed from the other LSMS, will be fast. However, if the LSMS with all the Bricks homed to it fails, logging for all Bricks switches to the Priority 2 LSMS.

While the load may be distributed across the two LSMSs in various ways, there are more factors to consider than simply the number of Bricks in the network. The speed of the connections between the two LSMS and many LSCSs, and between the Bricks and the LSMS they are "homed" to plays a role as well.

Rehoming

A Brick rehomes only when it loses connection with the LSMS to which it is currently homed, or when an Administrator manually rehomes it. This is true even if the priority order of the LSMSs is changed.

For example, suppose we have a situation in which a Brick is connected to a redundant pair of LSMSs. LSMS_A is its Priority 1 LSMS and LSMS_B is its Priority 2 LSMS.

If an Administrator changes the priority order, so that LSMS_B is now the Priority 1 LSMS, this will *not* cause the Brick to rehome. The Brick stays homed to its Priority 1 LSMS, even though this LSMS is now its Priority 2 LSMS. The new priority order is applied only when the Brick connection to LSMS_A is broken, or if the Administrator manually rehomes the Brick.

When you change the priority order, and then save and apply the change, all the apply does is specify the priority order that will be used the next time the Brick needs to rehome, either because of a loss of connectivity with the LSMS to which it is currently homed or a Brick reboot.

If you manually rehome a Brick to its priority 2 LSMS, it will stay homed to the priority 2 LSMS until the connection is lost. Once the connection is lost, it will try to home to its Priority 1 LSMS. Therefore, if you intend to permanently rehome a Brick, you must change its priority order as well.

Important! When a policy is applied by an Administrator, the LSMS to which the administrator is logged in will try to apply the policy to all affected Bricks, whether they are homed to that LSMS, or to the other member of the redundant pair. If the LSMS fails to apply to a device that is homed to the other LSMS, the LSMS will display the following error message:

```
Could not apply policy to brick <brickname>, please try to apply from the other LSMS.
```

The quickest way to determine which LSMS a brick is homed to is to check the LSMS Status Monitor. You can also issue the `display LSMS` command from the a local or remote brick host. For details of brick host operation, see the *Introduction*

to the Brick Command Line Interface chapter in the LSMS Tools and Troubleshooting Guide.



Redundant LSMS Monitoring

Overview

You can use the Status Monitor and the Log Viewer to monitor the status of redundant of LSMSs/LSCSs and their managed Bricks.

Status Monitor

The Status Monitor provides a number of tools that you can use to monitor the status of redundant LSMSs/LSCSs and their managed Bricks. For a more detailed discussion of the Status Monitor, refer to [Chapter 14, “Using the Status Monitor”](#).

Administrators & LSMS window

The Administrators & LSMS window provides status information about the LSMS you are currently logged into and all of the other LSMSs/LSCSs in your redundant network. Compute Servers are indented and grouped with their associated Primary or Secondary LSMS. This window indicates whether each LSMS is up or down, and shows the name, IP address and the number of Bricks assigned (both homed and not homed) to each LSMS/LSCS.

To display the Administrators & LSMS window, open the Monitor Menu and select **Administrator & LSMS**.

Status windows

The Status Monitor provides a number of windows that show Brick status. There is a Status Overview window, as well as a number of individual Brick Status windows that you can use to view different subsets of bricks (all Bricks, up Bricks, lost Bricks, and so forth). This information is accessible from the Monitor menu.

If a Brick is not homed to the LSMS to which you are currently logged into, the only up-to-date information that will be provided in any status window is the Brick operational status (UP/LOST) and the LSMS to which the Brick is currently or previously homed. The status of all managed Bricks across all LSMSs/LSCSs is shared and can be displayed on the Status Monitor of the LSMS that you are currently logged into.

Log viewer

You can use the Log Viewer to view log records. To launch the Log Viewer on a Windows platform, open the Start menu and select:

Programs ► Lucent Security Management Server ► LSMS LogViewer

On Solaris, make the installation directory (usually */opt/isms/lmf*) the present working directory, and enter the following command from the Solaris command line:

```
./LogViewer
```

For details on the operation and use of the LSMS Log Viewer see *Chapter 3. LSMS Log Viewer* in the *LSMS Reports, Alarms, and Logs Guide*.

The various logs maintained by the LSMSs/LSCSs can be merged by setting an option in the Reports Wizard to run a Bricks log report across all of the LSMSs/LSCSs and merge the results.

The LSMS/LSCS to which a Brick is homed keeps the log records for that Brick. Thus, if a Brick that is homed to its Priority 1 LSMS loses connectivity and rehomes to its priority 2 LSMS, the 1 LSMS Administrative Events Log shows this as a **Brick Lost** event, and the Priority 2 LSMS Administrative Events log shows it as a **Brick Up** event. Any other activity that occurs while the Brick is homed to the Priority 2 LSMS is logged and remains in the Priority 2 LSMS logs.

For a rehomed Brick, no data will be found in reports and logs of its priority 1 LSMS. After rehoming, the Brick log data resides in the LSMS logs which it has rehomed.

Under normal circumstances, the primary LSMS initiates a refresh to the Secondary LSMS every five minutes, and messages about these periodic refreshes appear in the Administrative Events log on both LSMS.

When a change is made on the Primary LSMS, the Secondary LSMS is immediately sent a message to synchronize the database with the Priority 1 LSMS. This is logged in the Administrative Events log as log type **126, Refresh Status**, with the status **Initiated**, as shown below:

```
126:i:scheduler:013123::REFRESH INITIATED
```

On the Secondary LSMS, whenever a database synchronization (refresh) succeeds, an event **126** is logged with status **Successful** in the Administrative Events log, as shown below:

```
126:i:scheduler:013125::REFRESH SUCCESSFUL
```

If a database synchronization fails on the Secondary LSMS, an event **126** is logged with status **Failed** in the Administrative Events Log, as shown below:

```
126:i:scheduler:013126::REFRESH FAILED
```

When either primary LSMS or secondary LSMS loses connectivity with its peer an event **125, LSMS Status**, is logged in the Administrative Events Log indicating the loss, as shown below:

```
125:i:scheduler:015913::LSMS_TW0:LOST::
```

After the Primary or Secondary LSMS re-connects with its peer, an event **125, LSMS Status**, is logged in the Administrative Events Log indicating the LSMS that was

re-contacted. The version of the software on the peer is also a part of this log event. It is shown below:

```
125:i:scheduler:020115::LSMS_TWO:CONTACTED:7.2.185:
```

During the time the Primary and Secondary LSMS have different versions of the software, the database synchronization will be disabled, and when any change to the database is made, a LSMS error **N9005** will be logged in the Administrative Events Log. The output message will be similar to the one shown below:

```
N9005 - WARNING - Secondary LSMS (version 9.0.184) has a different
version than the Primary LSMS (version 9.0.185) and needs upgrading.
Please Upgrade the Secondary LSMS to version 9.0.185.
```

This event will be logged in the Administrative Events log each time a change is made on that LSMS.

Whenever there is a change to the database and either LSMS is unable to initiate a database synchronization, it will log an appropriate error message.

On the primary LSMS, **N9007** will be logged to the Administrative Event Log with an appropriate **Reason** and on the secondary, **N9006** will be logged to the Administrative Event Log. This is shown below:

```
N9007 - WARNING - Changes made may not be visible on the Secondary LSMS
(redundantNT_2). Reason - Could not connect to the Secondary LSMS.
```

□

How to Configure a Secondary LSMS or Compute Server

When to use

Use this task to configure a Secondary LSMS or Compute Server.

Before you begin

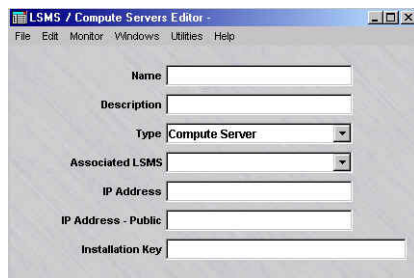
Before you begin this task, obtain the installation key of the type of LSMS or Compute Server that you are configuring. For information about installation keys, refer to the *LSMS Installation Guide*.

Task

To configure a new Secondary LSMS or Compute Server, follow the steps below:

- 1 With the Navigator window displayed, right-click the **LSMS and LSCSs** folder and select **New LSMS and Compute Servers** from the pop-up menu. The LSMS / Compute Servers Editor window is displayed ([Figure 2-1, “LSMS/Computer Servers Editor Window”](#) (p. 2-12) shows a sample window).

Figure 2-1 LSMS/Computer Servers Editor Window



- 2 Enter values in the following fields:
 - **Name** - The name of the Secondary LSMS or Compute Server, 1-45 characters.
 - **Description** - A textual description of the LSMS or Compute Server (Bricks supported, customer(s) supported, and so forth).
 - **Type** - If a Secondary LSMS is being added, click the down-arrow next to this field and select **Secondary LSMS**.
If a Compute Server is being added, click the down-arrow next to this field and select **Compute Server**.
 - **Associated LSMS** - For the Secondary LSMS or Compute Server being added, click the down-arrow next to this field and select the associated LSMS.

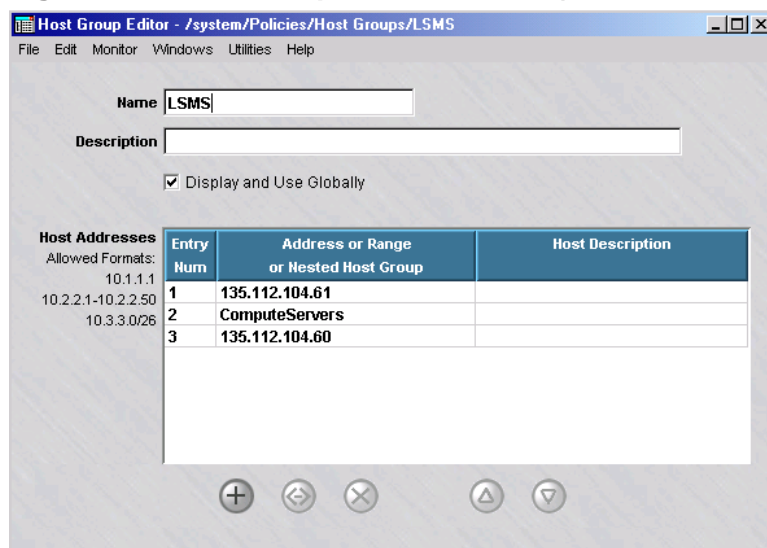
- **IP Address** - The real IP address of the LSMS or Compute Server. On private networks, this is the real IP address of the LSMS/LSCS that can be mapped to a virtual address using NAT.
- **IP Address - Public** - This is the Virtual Brick Address (VBA) of the Brick protecting the LSMS/LSCS used for NAT. This field is optional.
- **Installation Key** - The license key used to install the LSMS.

3 From the File menu, select **Save and Close**.

The new LSMS or Compute Server is configured. If it is a Secondary LSMS, an entry for the new LSMS appears in the LSMS Host Group. If it is a Compute Server, an entry is added to the ComputeServers Host Group.

Figure 2-2, “Host Group Editor window (LSMS Host Group)” (p. 2-13) shows a sample Host Group Editor window.

Figure 2-2 Host Group Editor window (LSMS Host Group)



END OF STEPS



3 Configuring and Activating a Lucent VPN Firewall *Brick*[®] Device

Overview

Purpose

This chapter explains how to configure and activate a Brick. The process is the same regardless of the model of Brick you are using. The process consists of the following activities:

- Using the LSMS, create an instance of the Brick and enter the required configuration parameters, such as Brick name and IP address
- Create a floppy disk or USB drive containing the configuration information and use the disk to activate the Brick device, package the configuration files for remote floppy/USB drive creation using an external floppy drive or USB drive connected to the USB port of the Brick device (for certain models), or use the floppyless bootstrap method that copies a "boot image" to the Brick serial port

Contents

Before You Begin	3-2
Configure a Brick on the LSMS	3-8
Brick Device Failover	3-22
To Set Up Brick Device Failover	3-26
To Manually Initiate Failover	3-32
To Activate a Brick	3-34



Before You Begin

Overview

All Bricks have to go through the same configuration and activation process to become operational. However, before you begin the process of configuring and activating a Brick, there are a number of questions you need to ask yourself about this Brick and the purpose it serves in the network.

Will the Brick be Operating as a Bridge or Router?

Important! Do not attach two or more interfaces that are configured with the same VLAN/subnet to the same switched network unless that network is running spanning tree protocol.

A Brick is essentially a bridging device. It passively "listens" on all its ports, in promiscuous mode. This means it accepts any traffic it "hears," regardless of the destination MAC or IP address. One important advantage of having the Brick operate as a pure bridge is that it does not interact with surrounding network equipment directly, so no connected device has to be reconfigured when a Brick is added to a network.

During the configuration process, you will be asked to provide the Brick with an IP address and subnet mask, which will then be automatically assigned to each of the Brick physical ports (see "How to Configure a Brick on the LSMS" on page 3-7). If you make no changes to the addresses, the Brick will bridge all traffic on all ports.

If you change the IP address of one or more of the physical ports — for example, if you assign one IP address to two of the Brick ports, and then assign a second IP address to the other two ports — packets traveling between two ports with the same IP address will be bridged., while packets traveling between two ports with different IP addresses will be routed from one subnet to the other.

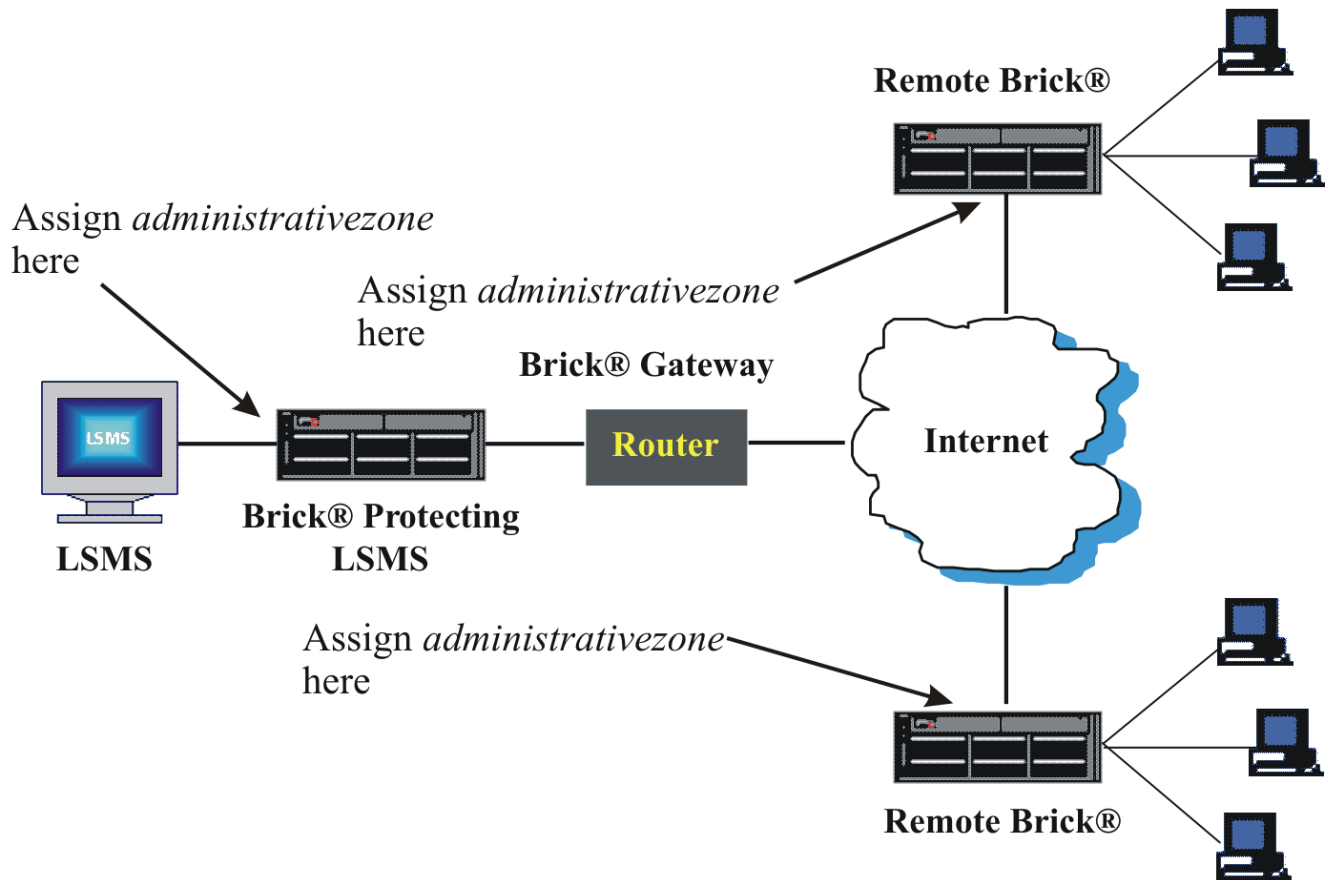
The Brick does not participate in dynamic routing protocols. It can be configured to pass packets to other subnets that are not directly connected to it by creating static routes (see "How to Configure and Maintain Static Routes" in *Chapter 4. Configuring Brick Ports*).

Will the Brick be Directly Connected to the LSMS?

We strongly recommended that you connect at least one Brick directly to the LSMS host to prevent an attacker from gaining access to the LSMS. Figure 3-1 shows a

configuration in which one Brick is connected directly to the LSMS, and two other Bricks are connected to the LSMS remotely.

Figure 3-1 Brick Configuration



If the Brick being configured protects the LSMS, you should assign the pre-configured Brick zone ruleset *administrativezone* to the port connected to the LSMS host (see "How to Assign a Security Policy to a Port" on page 4-5 for instructions). This is a ruleset provided with the LSMS software specifically for this purpose. It contains rules that drop all traffic to the LSMS except from the Bricks it is managing. (See *Appendix B. Pre-Configured Brick Zone Rulesets* in the *LSMS Policy Guide* for a detailed description of the *administrativezone* ruleset.)

If the Brick being configured is connected to the LSMS remotely, you should still apply *administrativezone* to the port connecting it to the LSMS. This is usually the port connected to the router that has been designated the Brick gateway during the configuration process.

When assigning *administrativezone* to a port, you also have to indicate the IP addresses connected to the port that will be protected by this ruleset (the ruleset along with the IP addresses is referred to as a *zone*). It is best if the LSMS is the only host connected to the Brick on its port. If this is the case, enter an asterisk in the **Zone IP Addresses** field when assigning *administrativezone* to the port. If other hosts are connected, enter the IP address of the LSMS host instead (see page 3-7).

This configuration will prevent an administrator from creating a policy on that port that disallows access from the Brick to the LSMS host, thereby rendering that Brick unmanageable. (If this does accidentally happen, simply change the policy on the LSMS and reboot the Brick to allow it to recollect its policy from the LSMS.)

Will the Brick be Functioning as a Firewall?

If this Brick will be used as a firewall protecting internal LANs from attack via the Internet, you have to configure the Brick appropriately. [Figure 3-2, “Firewall Configuration” \(p. 3-5\)](#) shows a Brick configured as a firewall.

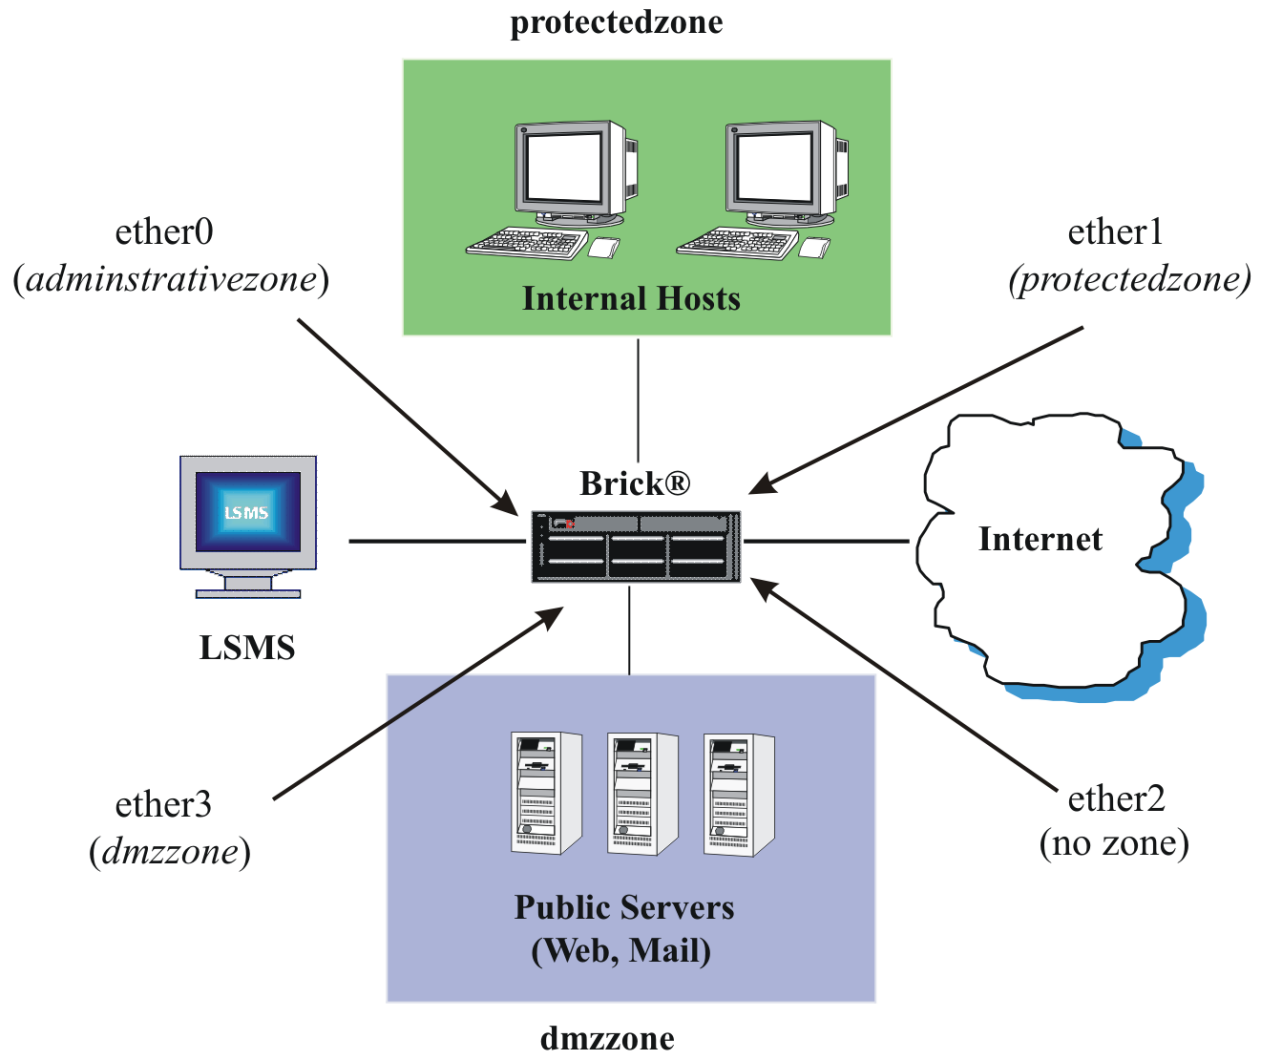
In this example, the Brick is deployed between the Internet and an internal LAN consisting of two zones (Brick zone rulesets). All communication between the Internet hosts in the two zones is monitored by the *protectedzone* and *dmzzone* rulesets that were assigned to ethers1 and 3.

Since this Brick is also protecting the LSMS, the *administrativezone* ruleset was applied to ether0, the port connecting the Brick and LSMS. No ruleset was applied to ether2, since it is connected to the Internet.

Deploying the Brick at the port to the Internet protects the internal network from external intrusion and attack. Attacks between hosts in the internal network can also be mitigated by connecting them to separate ports on the Brick. This ensures that all

communication between these hosts must pass through the Brick so that their traffic is also scrutinized by the rulesets on the Brick.

Figure 3-2 Firewall Configuration



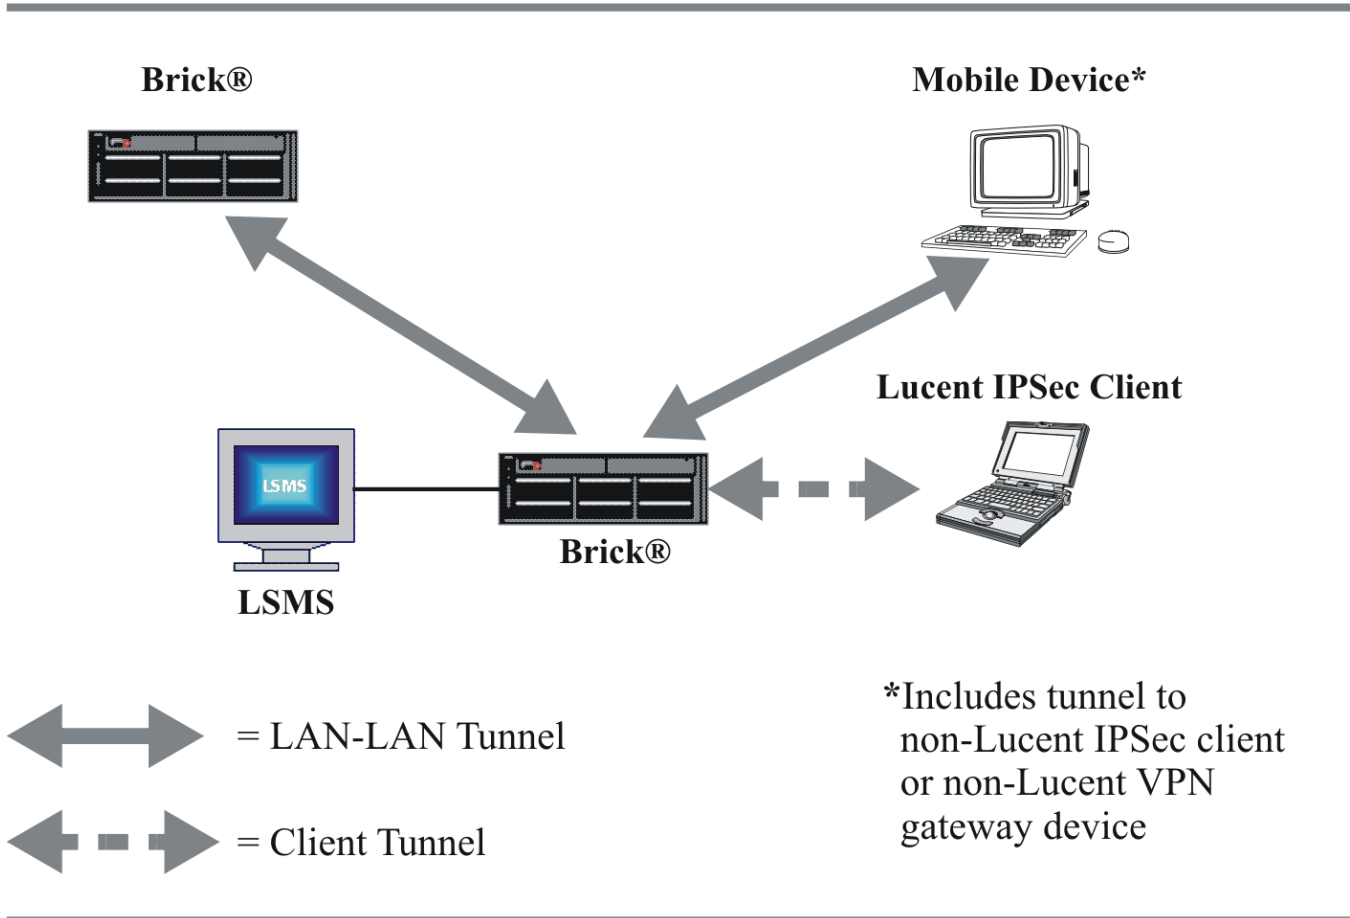
Will the Brick be Functioning as a Tunnel Endpoint?

If the ports on this Brick will be terminating LAN-LAN or client tunnels, you have to assign tunnel endpoint addresses to the ports when configuring the Brick.

[Figure 3-3, “LAN-LAN and Client Tunnels” \(p. 3-6\)](#) shows the different kinds of tunnels that can terminate on a Brick port. LAN-LAN tunnels include tunnels from the Brick to another Brick or to an unmanaged device. A client tunnel is a tunnel from a host running the Lucent IPsec client to the Brick.

Important! The tunnel endpoint address is also the Virtual Brick Address (VBA). The VBA can be used when performing network address translation (see *the Network Address Translation* in the *LSMS Policy Guide*).

Figure 3-3 LAN-LAN and Client Tunnels



Will the Brick Need a "Dynamic" IP Address?

In some environments, such as a home office or branch office, a service provider may only grant users a single dynamic IP address, which may change from time to time. Also, in a DSL network environment, dynamic IP addresses are assigned by DSL modems using PPPoE. A Brick can operate in such settings in either of three different modes: with a static, private IP address in which the router performs a Network Address Translation function, or in which the Brick itself obtains the dynamic address via either DHCP or PPPoE.

When the Brick is provisioned with a static IP address, there are any number of low-end routers on the market that provide a DHCP client operating on their "outside address". These devices have the ability to source address translate any outbound TCP

or UDP connection so that their public addresses may be shared by inside "private" sources. In addition, the external device must support mapping inbound connections to specific private addresses. Typically, this configuration would involve a Model 50 Brick with a small office router such as the Linksys Etherfast router. The router performs Network Address Translation (NAT) on the inbound and outbound packets from the Brick while the LSMS learns about any IP address change for the Brick.

When the Brick is provisioned as a DHCP client or as PPPoE, the Brick communicates with the DHCP server itself or the DSL modem itself and obtains a dynamic IP address, netmask, and gateway from them. In these cases, any outside router does not typically do NAT. The Brick may do NAT or VPN for the network "behind" it, using this public address.



Configure a Brick on the LSMS

When to use

Once you have addressed all of the preceding questions, you are ready to begin configuring the Brick. To configure a Brick, you have to display the Brick Editor and enter the configuration information requested. This information is then saved in the LSMS database.

The basic information required to configure a Brick has to be entered in the Brick tab of the Brick Editor. The information differs somewhat, depending on whether you are configuring the Brick from a Primary LSMS or from an LSMS that is part of a redundant pair.

The section below explains how to configure a Brick device from a Primary LSMS, and the section that follows explains how to configure the Brick device from an LSMS that is part of a redundant pair.

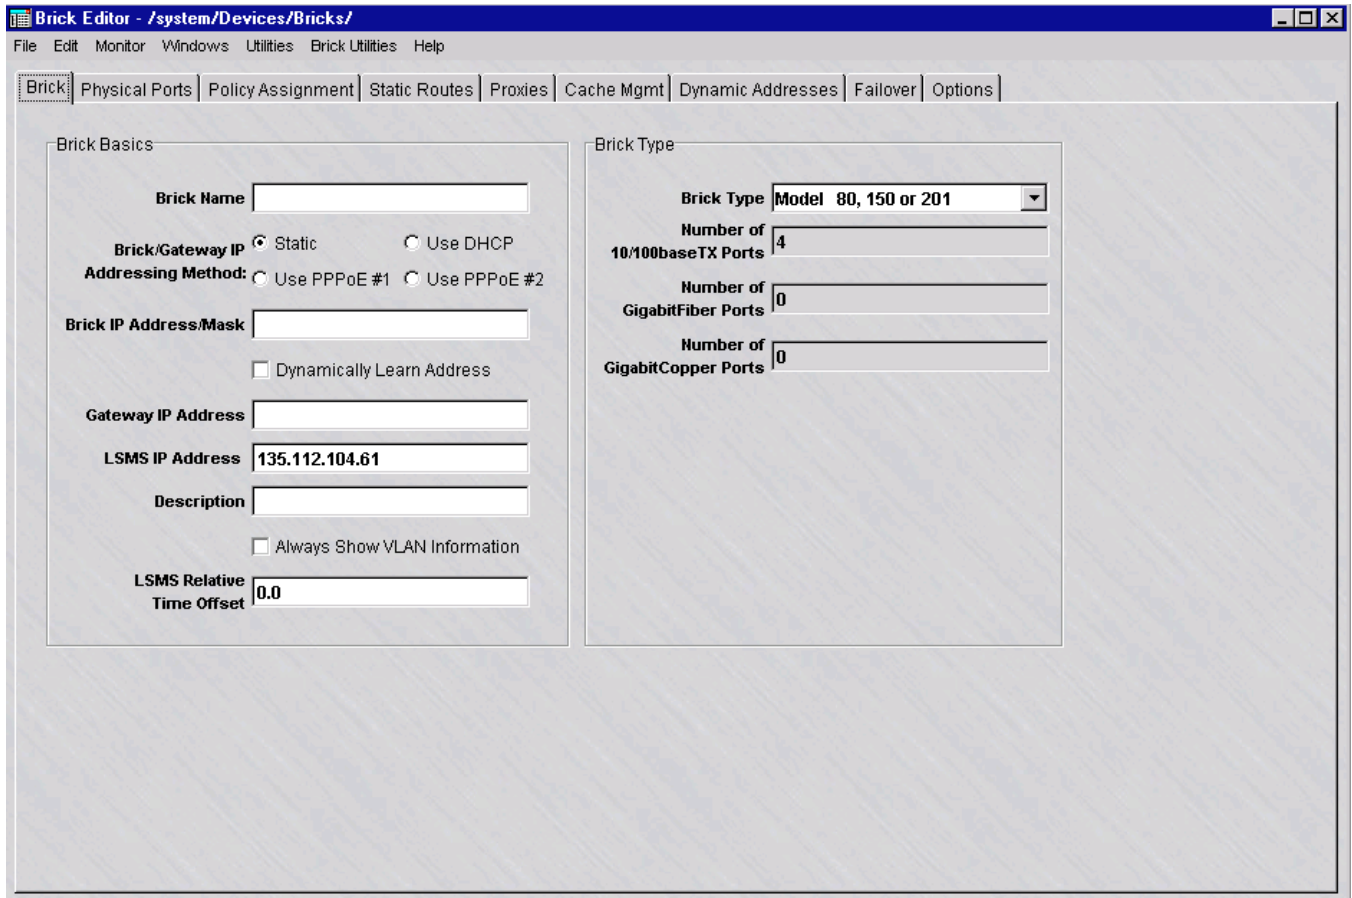
Basic Configuration - Primary LSMS

If your LSMS has not been specifically configured to be part of a redundant pair, it is a Primary LSMS. Complete the following steps to configure a Brick using a Primary LSMS.

- 1 Open the folder of the group in which you want to put the Brick and open the **Devices** folder.

- Right-click the **Bricks** folder and select **New Brick** from the pop-up menu. The Brick Editor (Brick Tab) is displayed (see [Figure 3-4, “Brick Editor \(Brick Tab - Primary LSMS\)”](#) (p. 3-9)).

Figure 3-4 Brick Editor (Brick Tab - Primary LSMS)



- In the **Brick Name** field, enter the name to be used to identify this Brick. This field is required. The name of the Brick must be unique and can contain 1 to 80 lowercase alphanumeric characters.
- To assign a static IP address to the Brick and a static gateway that the Brick will use to talk to the LSMS, click the checkbox that says "Static." In the **Brick IP Address/Mask** field, enter both the IP address and subnet mask of the Brick.

Alternatively, you may dynamically assign these addresses. This can be accomplished by either using DHCP or PPPoE. Your service provider can tell you if DHCP is

required. If your network uses DSL lines installed on DSL modems that require the CPE side equipment to talk with the modem, then PPPoE is most likely required.

Click the checkbox **Using DHCP** to enable DHCP. Click the checkbox **Using PPPoE#1** or **Using PPPoE#2** to enable Point-to-Point Over Ethernet #1 (PPPoE#1) or PPPoE #2. Checking any of these boxes causes the **Dynamically Learn Address** box to be checked and grayed out and the Gateway field to be grayed out.

PPPoE #1 and #2 stand for two different PPPoE sessions that are configured on the Dynamic Addresses Brick screen in support of the following configurations:

- A single port on the Brick connecting to a single PPPoE modem with a single PPPoE session (#1 or #2) active.
- A single port on the Brick connecting to a single PPPoE modem with two PPPoE sessions (#1 and #2) active.
- Two ports on the Brick connecting to a different PPPoE device.
- A single port on the Brick connecting to two PPPoE devices simultaneously.

You may click the **Dynamically Learn Address** checkbox if:

- The Brick itself does not acquire its management address dynamically.
- The Brick is in a home office / branch office setting.
- The service provider grants only a single IP address which is updated periodically.
- There is a device somewhere between the Brick and the LSMS that performs a Network Address Translation (NAT) function on the Brick address using an address that may change over time. This is often (but not always) a low end router immediately in front of the Brick.
- The router has been specifically configured to allow the following inbound connections from the LSMS to the Brick:
 - TCP 910 (administrative channel and asynchronous commands)
 - UDP 1024 (reflection channel for LSMS based user authentication proxy)
 - UDP 9014 (proxy IKE from LSMS VPN gateway controller service)

If the **Dynamically Learn Address** checkbox is checked, the LSMS will not attempt to contact the Brick until it has made first contact with the LSMS.

When a device NATs a Brick address on the way to the LSMS, that Brick will not support IPsec Client tunnels, but will support LAN-LAN tunnels with UDP encapsulation. The Brick will support IPsec client tunnels; but may not be effective because, typically, there is no way to obtain the current client address.

The LSMS automatically assigns the address and mask to each of the Brick physical ports. You can view this on the Physical Ports tab of the Brick Editor (see Figure 4-1 in *Chapter 4. Configuring Brick Ports*). If the ports are not changed, the Brick functions as a "pure" bridge on all ports.

To change the IP address and subnet mask for any of the ports, you have to display the Brick Ports Editor (see Figure 4-2 in *Chapter 4. Configuring Brick Ports*). If you change one or more ports, the Brick routes packets on those ports. (It may also be necessary to add static routes to reach other subnets; this is explained in "How to Configure and Maintain Static Routes" in *Chapter 4. Configuring Brick Ports*).

Once the configuration is saved, the mask will be dropped from this field. It still can be seen by displaying the Brick Ports Editor.

-
- 5** In the **Gateway IP Address** field, enter the IP address of the device (usually a router) that will serve as the "default route" for this Brick. This is where the Brick will send all traffic not destined for a local subnet.

If the LSMS is on a different subnet than the Brick, you have to enter the IP address of the router that the LSMS is using to communicate with the Brick in this field. This is the only route that the Brick can use to contact the LSMS when the Brick first boots. The gateway address must fall within the IP range of one of the Brick ports.

If you enter a conflicting route in the Static Routes Table, it may override this **Gateway IP Address** field (see Note below).

Important! The default gateway on the LSMS should be set to the address of whatever device is the "next hop." If the LSMS is directly connected to the Brick, use the Brick IP address for the gateway. If there is a router between the LSMS and the Brick, use the router address for the gateway on the Brick.

If the LSMS host is running Windows, make sure the *Default Gateway* field is set properly in the TCP/IP protocol properties.

If the LSMS host is running Solaris, make sure the */etc/defaultrouter* setting is set properly.

-
- 6** In the **Description** field, enter a brief description of this Brick. The description is optional. It can contain up to 80 characters (letters, numbers and certain special characters).

Important! *ALWAYS SHOW VLAN INFORMATION*

If you will be configuring the Brick to recognize, forward and filter VLAN traffic, you must check the **Always Show VLAN Information** checkbox (see Figure 3-5). This will add two new tabs (*VLAN/IP Assignment and Partitions*) to the row of tabs in the Brick Editor, and it will add a number of VLAN-related columns to the table in the Physical Ports tab. It will also add several VLAN fields onto the Brick Ports Editor and Policy Assignment Editor.

Once you check this checkbox and save the configuration, you will not be able to return to the pre-VLAN view. The VLAN/IP Assignment tab and the changes to the Physical Ports tab will remain in effect, regardless of whether you actually make use of the VLAN feature or not.

If you plan to use the VLAN feature, turn to *Chapter 6. Configuring VLANs on Bricks* for an explanation of how to configure the Brick's physical ports and assign policies to the ports to handle VLAN traffic.

7 The **LSMS Relative Time Offset** field, is only used if:

- The Brick is in a different time zone than the LSMS
-AND-
- You need to add time and day restrictions to the rules.

Enter the time offset value in hours, with a "+" or "-" to indicate whether the time zone for the Brick is "ahead" or "behind" the time on the LSMS.

For example, if the Brick is in Los Angeles and the LSMS is in New York, the Brick is three hours behind the LSMS. The offset value is "-3.0". Similarly, if a Brick is in a time zone 8 1/2 hours ahead of the LSMS, the offset is "+8.5".

The LSMS and Brick synchronize their times once an hour.

For additional details, see "How to Add Time and Day Restrictions to a Rule" in *Chapter 2. Brick Zone Rulesets in the LSMS Policy Guide*.

8 In the **Brick Type** field, select the model of the Brick you are configuring from the drop-down list. The options are:

- Model 20 or 50
- Model 80, 150 or 201
- Model 300
- Model 350
- Model 500
- Model 1000 (3/4/0)
- Model 1000 (5/4/0)
- Model 1000 (7/2/0)
- Model 1000 (9/2/0)
- Model 1100 (7/0/13)
- Model 1100 (7/4/1)
- Model 1100 (7/6/1)

After you enter the model type, the LSMS fills in the number of ports. The Model 20 and Model 50 have three 10/100baseTX Ethernet ports. The Models 80, 150, and 201 each have four such ports, the Model 300 eight, and the Model 500 14 (plus one GigabitFiber port).

The Model 1000 comes in these four configurations:

- Three 10/100baseTX ports and four GigabitFiber ports
- Five 10/100baseTX ports and four GigabitFiber ports
- Seven 10/100baseTX ports and two GigabitFiber ports
- Seven 10/100baseTX and two GigabitFiber ports (also includes an encryption accelerator card)
- Nine 10/100baseTX ports and two GigabitFiber ports

The Model 1100 comes in these three configurations:

- Seven 10/100baseTX ports and thirteen GigabitCopper ports
- Seven 10/100baseTX ports, four GigabitFiber ports, and one GigabitCopper ports
- Seven 10/100baseTX ports, six GigabitFiber ports, and one GigabitCopper ports

See the *User's Guide* for the respective Brick model for complete details about the hardware configuration.

-
- 9 Display the File menu and select **Save**. You have just entered the basic information necessary to activate a Brick. However, there are certain optional configuration parameters you should consider before activating the Brick. See "Configuration Options" on page 3-18 for an explanation of these parameters.

Important! If you open the Monitor menu and select **Status Overview** to display the Status Monitor, the Brick being activated appears in the Brick Status graph as *LOST*. Once the activation process is complete, the status will change to *UP*, indicating the Brick and the LSMS are communicating.

In previous releases of the LSMS, if you selected **Console Alarms** from the Monitor Menu, you would see an alarm indicating the LSMS cannot contact the Brick. This alarm is no longer generated (nor is an alarm generated the first time a Brick is contacted after the LSMS services are started).

LOST and CONTACTED alarms will be generated after a transition from CONTACTED to LOST or from CONTACTED to *LOST* back to *CONTACTED* (as they are now).

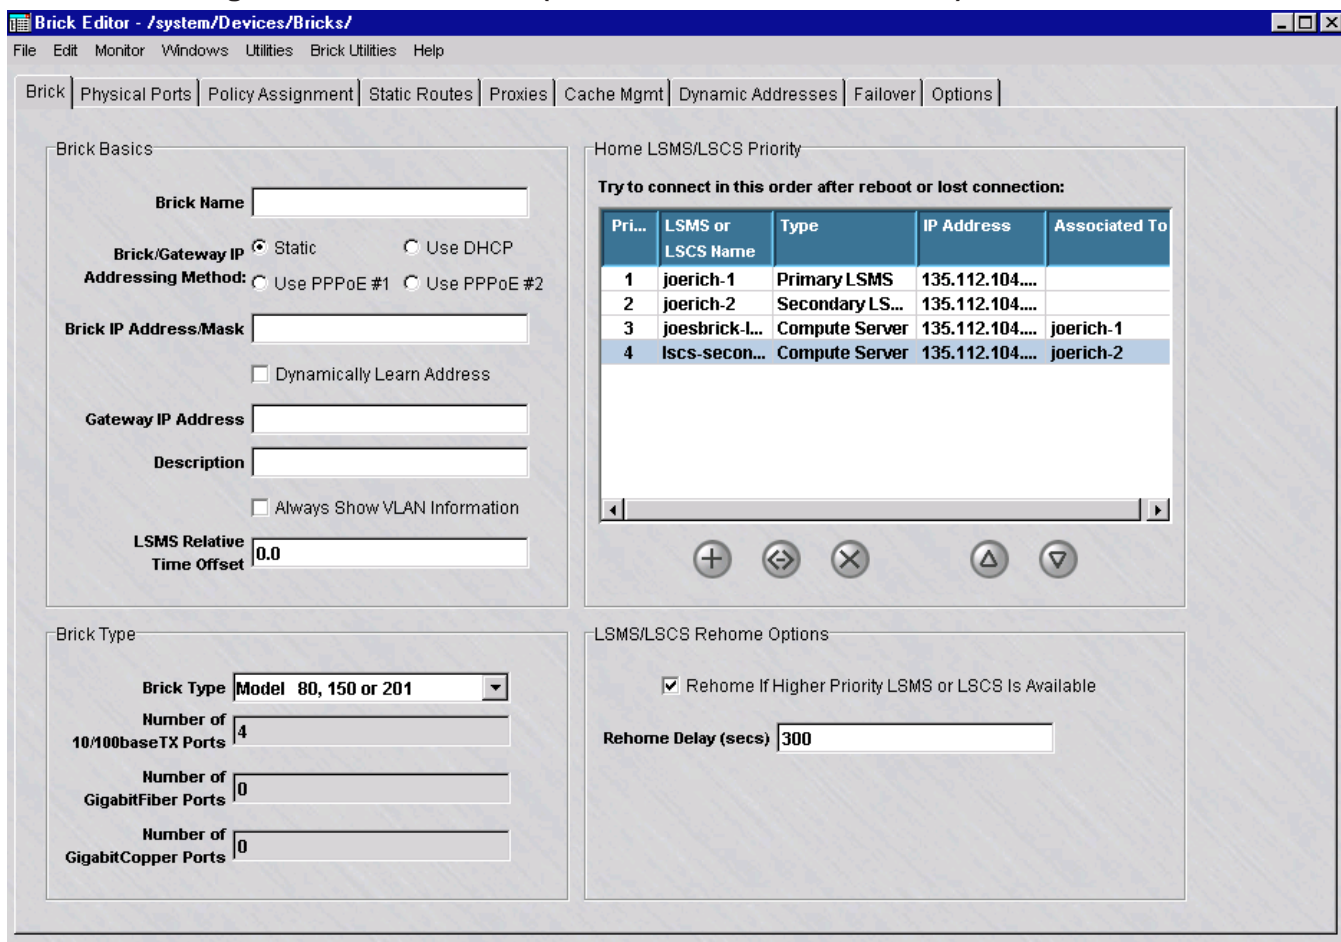
END OF STEPS

Basic Configuration - LSMS Redundancy

Use the following steps to configure the LSMS Redundancy options for a Brick.

- 1 A redundant pair of LSMSs consists of two active LSMS. Every Brick managed by this redundant pair has one LSMS designated its *priority 1* LSMS and the other its *priority 2* LSMS. When a Brick has to contact its LSMS (for example, to perform logging or user authentication), it attempts to contact the *priority 1* LSMS first. If it cannot contact the *priority 1* LSMS, it tries to contact the *priority 2* LSMS. Any Brick can be managed by either LSMS at any time. Up to five LSMSs or Compute Servers can be designated in the priority list for taking over management of a Brick. Regardless of which LSMS a given Brick is currently homed to, it can be managed by either LSMS. If the LSMS you are using is part of a redundant pair, the Brick tab of the Brick Editor has a different appearance than the Brick tab for a Primary LSMS. [Figure 3-5, “Brick Editor \(Brick Tab - Redundant LSMS\)”](#) (p. 3-14) shows the Brick tab for a redundant LSMS.

Figure 3-5 Brick Editor (Brick Tab - Redundant LSMS)



As this figure indicates, the fields for entering the Brick name, Brick IP address, and optional gateway address and description, as well as the model type, are the same as those for a Primary LSMS.

However, instead of an LSMS IP Address field, this tab has an **Home LSMS/LSCS Priority** panel and a **LSMS/LSCS Rehome Options** panel. The following explains how to use these two panels:

- **Home LSMS/LSCS Priority**

The LSMS that you are logged into when you create a Brick is automatically considered the Brick *priority 1* LSMS, and the Brick is said to be *homed* to this LSMS. You can list the homing priority of each associated device for the Brick in this list, including the Secondary LSMS (if applicable) and Compute Servers in the cluster, up to five entries.

Since Compute Servers are dependent on their associated LSMS, it is recommended that all of the LSMSs/LSCSs assigned to the Brick not be in the same cluster.

A Brick will always try to home to its *priority 1* LSMS after rebooting or after LSMS services have been restarted.

If a single Compute Server has been configured, a Brick can be set up to home to this server if connectivity is lost between the Primary or Secondary LSMS. The name of the Compute Server and its priority is displayed in this panel. If more than one Compute Server has been configured and associated with the managing LSMS, the name and rehome priority of each one is displayed in the panel.

The **Home LSMS/LSCS Priority** panel gives the name and IP address of each LSMS or Compute Server in the cluster. It also indicates which LSMS is the *priority 1* LSMS and which is the *priority 2* LSMS. If a single or cluster of Compute Servers has been configured with the managing LSMS, the name and priority of each one is displayed. From this panel, you can do all of the following:

- Change the priority of the LSMSs or Compute Server(s). Select the server (LSMS or Compute Server) and click the **Down** button to lower its priority of connectivity with the Brick, or select server and click the **Up** button to raise its priority of connectivity with the Brick. The LSMS priority can even be changed while you are creating the Brick.
- Change the name and/or IP address of one of the LSMSs or Compute Server(s). Select the server and click the **Edit** button. This is helpful if you are using an LSMS or Compute Server with a private unregistered IP address.

- Add a new LSMS or Compute Server. Click the **New** button. Only two LSMS can be entered in this panel. If you converted a pre-V8.0 standalone LSMS installation to a Primary LSMS in a redundant installation, you would need to add the secondary LSMS to each previously configured Brick.
- Remove an LSMS or Compute Server from a Brick configuration. Select the server and click the **Delete** button. You can use the delete function if you remove one of the LSMSs or Compute Server(s) and add a new one.
- **LSMS Rehome Options**

If a Brick is currently connected to a lower priority LSMS or Compute Server, and connectivity is lost, the Brick automatically rehomes to a higher priority LSMS or LSCS — provided the checkbox labeled **Rehome If Higher Priority LSMS or LSCS Is Available** remains checked (the default). If you uncheck it, the Brick will remain connected to its priority 2 LSMS/LSCS until it loses contact with this server, is rebooted, or the LSMS services are restarted.

The **Rehome Delay** field determines the delay, in number of seconds, before the Brick rehomes when a higher priority LSMS/LSCS becomes available. The default is 300 seconds. After this time period has elapsed, the Brick repeatedly attempts to rehome to the higher priority LSMS/LSCS.

When you are finished, display the File menu and select **Save**.

Important! When a Brick fails over to its Priority 2 LSMS, all log records are switched to that LSMS as long as the Brick is homed there.

For this reason, it is a good idea to leave the **Rehome If Higher Priority LSMS Is Available** checkbox checked.

END OF STEPS

Dynamic Addresses

Complete the following steps for dynamic addressing.

- 1 The checkbox called "Treat the two PPPoE sessions as a redundant pair" means that wherever a PPPoE (#1 or #2) session is used in the Brick configuration, the one actually used is the first session that becomes available. You don't really have two different sessions going on at once, but one that has an alternative link.

When this box is checked, the use of PPPoE #1 or PPPoE #2 is equivalent everywhere except in the VLAN assignment and interface assignment. For example, using PPPoE #1 as a VBA and PPPoE #2 as the target of a static route means the same thing. The Brick will use the address/route associated with the link that is currently selected as active.

-
- 2 If you are acquiring any Brick address via DHCP, you may need to change some of the following settings:
 1. **Allow DHCP responses from these servers** - This allows you to select an IP address or host group list of IP addresses that the Brick will allow to serve its requests. By default, the Brick will accept a reply from any IP address.
 2. **Allow DHCP addresses in range** - This allows you to restrict the addresses that the Brick will accept for its own addresses. It prevents someone from spoofing the DHCP reply and, for example, assigning an address that really belongs to another interface. By default, the Brick will accept any IP assignment.
 3. **Broadcast Discover/Unicast Request(direct to server)** - Typically, DHCP operates in broadcast mode and this checkbox should be left on broadcast. However, in a sensitive environment where the DHCP request address should not be broadcast all over the network, you may click Unicast Request in which case the Brick will sequentially ask every IP address in the **Allow DHCP responses from these servers** field (which means there had better not be very many!).
-
- 3 Configuration for setting up the PPPoE #1 and #2 sessions is displayed along the bottom. All fields are optional except the two Keep-Alive fields. If a User ID is entered, either a Password or CHAP Key is required. The MAC Address identifies which physical device to use, and must be entered as six pairs of hex characters delimited by colons. The CHAP Key can be entered as either text or hex characters.

END OF STEPS

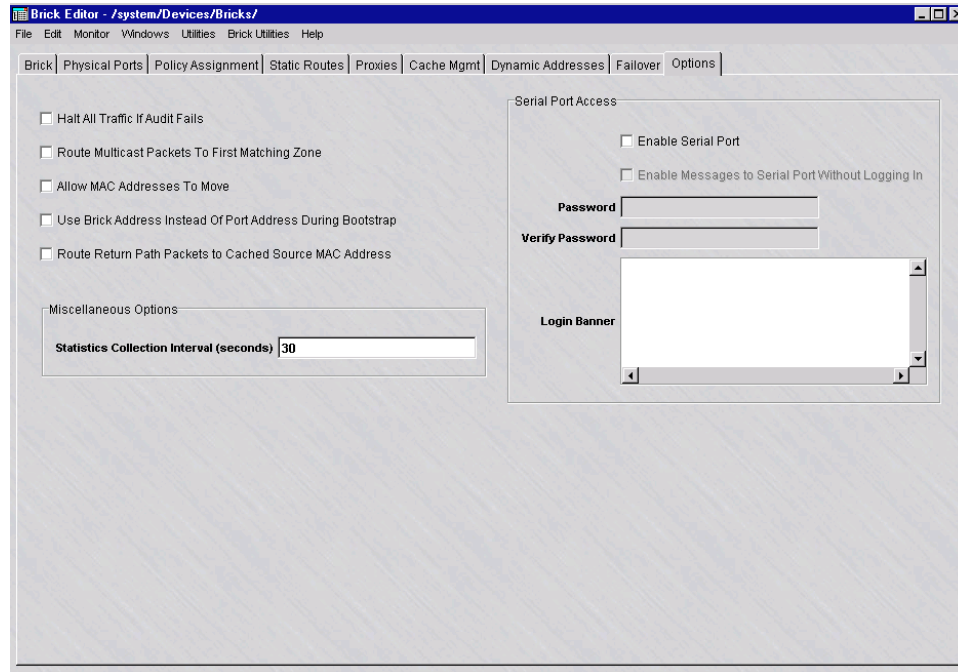
Configuration Options

Once the information requested in the Brick tab has been entered and saved, the Brick can be activated. However, before you activate the Brick, you should decide whether or not to enable certain optional features found in the Brick. These options, which are enabled from the Options tab of the Brick Editor, are the same, regardless of whether you have a Primary LSMS or redundant LSMS pair. In addition, the Options tab allows you to create a password that you can use to log into the serial port of the Brick. The Options tab also allows you to enable Brick failover, if this Brick is part of a failover pair.

The following explains how to use the Options tab:

- 1 Click **Options** to display the Options tab of the Brick Editor. It is shown in [Figure 3-6, “Brick Editor \(Options Tab\)”](#) (p. 3-18).

Figure 3-6 Brick Editor (Options Tab)



- 2 Click the options checkboxes that you want to enable.
The following explains what each checkbox accomplishes.
 1. **Halt All Traffic If Audit Fails**

This checkbox determines how the Brick responds if it loses communication with the LSMS and cannot perform its logging function. The following explains:

- If the checkbox is *unchecked*, the Brick continues to permit traffic through, and will fully enforce all security policies, but will perform no auditing.
- If the checkbox is *checked*, the Brick stops all traffic until communication with the LSMS is re-established.

The LSMS and Brick actively maintain a heartbeat on the audit channel, and can detect failures almost immediately.

By default, this checkbox is unchecked, which means traffic will continue in the event of loss of communication with the LSMS. To stop all traffic, click once to check the checkbox.

2. **NOC Gateway**

This checkbox has to be checked if this Brick is to serve as the NOC Gateway Brick to manage routers securely.

To be the NOC Gateway, this Brick has to be directly connected to the LSMS, and it has to be the endpoint for the management tunnels that the LSMS uses to manage routers securely.

3. **Route Multicast Packets to First Matching Zone**

Normal (unicast) traffic is addressed to a specific host. The Brick determines which zone that host is a member of, and applies that policy. Multicast traffic may be delivered to multiple hosts. If more than one zone has been assigned to a port, hosts in each zone may receive the traffic. If one zone has a rule passing such traffic and another zone has a rule requiring it to be dropped, there is a conflict. This checkbox determines how the Brick resolves such conflicts. The following explains:

- If this checkbox is *unchecked*, the Brick processes multicast sessions in all zones assigned to the port. This means there must be a rule allowing the session in each of these zones for the session to pass through the Brick, even if the session is only intended for one of the zones.
- If this checkbox is checked, the Brick processes multicast sessions in individual zones. This means there must be a zone that protects the multicast address — either explicitly or with the wildcard asterisk. (See "How to Assign a Security Policy to a Port" on page 4-5).

4. **Allow MAC Addresses to Move**

This checkbox has to be checked to allow MAC addresses to float from one Brick port to another. Allowing MAC addresses to float means that when a Brick that has a MAC address associated with a particular port receives an Ethernet frame *from that same MAC address on a different port*, it automatically moves the MAC address assignment and associated firewall sessions to the new port, and marks that MAC entry with the time it was assigned.

5. Use Brick Address Instead of Port Address During Bootstrap

This checkbox allows you to specify that the Brick should use the address specified in the Brick IP Address/mask field to communicate with the LSMS while loading the policy. This is sometimes required when the Brick has only one (public) address for management and for the VBA, but the interfaces on the Brick are all assigned to be private in order to enable more traffic (such as DHCP) to be bridged. After the Brick has loaded its policy from the LSMS, this checkbox has no further effect.

This checkbox is not available if the **Dynamically Learn Address** checkbox is checked and the **Brick/Gateway IP Addressing Method** is a static IP address.

To enable this feature, you have to check this checkbox, and then you have to set the Brick IP address to the VBA of the Brick zone ruleset assigned to the port that is serving as the tunnel endpoint (see "How to Configure a Physical Port" on page 4-1).

6. Route Return Path Packets to Cached Source MAC Address

This checkbox affects how the reverse direction packets are routed in the firewall (in other words, the packets that are routed from the session destination back to the session originator).

This checkbox *only* affects packets that are not bridged. Packets that are bridged are routed according to the destination MAC address and VLAN in the packet. When the checkbox is not checked and the traffic is routed, then the return packets are routed using the ARP table and, possibly, the static route table. When the checkbox is checked and the traffic is routed, then the Brick simply routes the return traffic back to the same interface, VLAN, and MAC address from which the first packet in the session arrived.

However, there is an exception to this handling of return packets. If the Brick has detected that the MAC address should not be used for this purpose (for example, it is a VRRP or HSRP router, then the MAC address is not used for this purpose (regardless of the checkbox setting). The Brick makes this determination by looking at the inner and outer MAC addresses on ARP packets, and, if they are different, marks the outer MAC address as being a VRRP/HSRP MAC address.

Subsequently, the Brick does not return packets back these devices. Unfortunately, experience has shown that this test does not work for all routers.

When this checkbox should be used:

In most cases, the checkbox should be left unchecked. The checkbox should be checked only if traffic is being routed (using static routes) and one of the following applies:

- There is no static route that takes traffic back to the source address
- A session with a particular IP address may arrive at the Brick from more than one router or interface and there is a need to have the return traffic routed the same way

-
- 3 If you intend to use the Brick serial port to access the out-of-band (OOB) command line interface, click the **Enable Serial Port** checkbox, and enter a password twice —

once in the **Remote Password** field and again in the **Verify Password** field. The password can be from 6 - 72 characters (letters and numbers). The password is case sensitive, so capitalization must be consistent.

To access the command line interface, you can connect a terminal or a modem to the Brick serial port. See the *LSMS and Brick Command Line Interfaces Reference Manual* for instructions on how to log into the Brick and use the OOB command line interface. (Serial port parameters are 115,200 baud, no parity, 8 data bits and 1 stop bit.)

If the serial port is enabled, then the console command line will also require the password to log in and access its functions. If the serial port is disabled (and hence no password is given), then the console command will not require the password.

The **Enable Messages to Serial Port Without Logging In** checkbox allows Bricks to send messages to the serial port without having to log in first.

To add a login banner to be displayed when logging into the Brick via the serial port, refer to the procedure [“To Activate a Login Banner on the Brick Serial Port Console”](#) (p. 4-28) in Chapter 4, [“Configuring Lucent VPN Firewall Brick® Device Ports”](#).

-
- 4** In the **Statistics Collection Interval (secs)** field, the default reporting period for Proactive Monitoring (ProMon) statistics is **30** (seconds). You can change or reduce the reporting interval for this Brick to less than 30 seconds. The value entered must be greater than zero (**0**).

For more information about Brick Proactive Monitoring parameters that are logged by the LSMS or Compute Server, refer to the *LSMS Reports, Alarms, and Logs Guide*.

-
- 5** When you are finished, display the File menu and select **Save**.

END OF STEPS



Brick Device Failover

Overview

Brick device failover is based on a simple model: two Brick devices, each connected to the same set of LANs, can share the same identity, including IP address and name. The two Brick devices are administratively treated as one. The first Brick device to boot becomes the active Brick device and behaves like a normal Brick device. The other Brick device is the standby, and waits to take over should the active Brick device fail.

Brick device failover can be set up to be invoked automatically, or it can be performed manually from the LSMS GUI or via the Brick device command line interface (CLI).

How Brick device failover works

In a Brick device failover configuration, both Brick devices issue heartbeats, or keepalive messages, at regular intervals to indicate their operational “health” and the integrity of each of their Ethernet links. (For additional details, refer to the [“Brick device failover protocol and heartbeats”](#) (p. 3-22) section.) A lack of incoming heartbeat messages from the Brick device itself or one of its links can invoke failover to the standby Brick device.

The active Brick device may also yield to the standby if it detects that the standby has better LAN connectivity, or the other Brick is designated to be the *Primary* Brick device in a failover pair. For details about designation of a Primary Brick device, refer to the [“Primary Brick”](#) (p. 3-24) section. Failover Brick device pairs employ state-sharing to maintain sessions despite the failure of one Brick device in the pair.

In addition to heartbeat messages issued by the Brick device, the LSMS provides an option to configure the active Brick device to ping a router or other device on the LAN, at a specified time interval and frequency, to determine if it still has LAN connectivity. In the event that no response is received for the specified number of iterations, the currently active Brick device will initiate a failover if the standby Brick is operational and has link integrity. The standby Brick device must be operational for at least 30 seconds before initiating a failover (to allow state-sharing to occur).

Brick device failover protocol and heartbeats

Both the active and standby Brick devices regularly issue heartbeats. The active Brick device generates ten heartbeats per second for each link.. The standby Brick device is adaptive. During changing conditions, the standby beats once; during stable times, it beats once every 800 milliseconds per link.

Heartbeats serve several functions:

- The heartbeat indicates the presence of an active Brick device. A lack of incoming heartbeats causes the standby Brick device to become active.
- Heartbeats allow Brick devices to determine the relative health of each of their Ethernet links. Without heartbeat messages or some other configured flow, the Brick devices would have to rely on simple link integrity to give a local view of the health of their ports.
- Heartbeats allow Brick devices to share health, status, and priority information with each other. This information is used to make failover and state-sharing decisions.
- Heartbeats also verify that the corresponding ports on each Brick device are connected to the same LANs, essentially for implementing security policies that are based on connection to specific ports.

Heartbeat messages carry authentication and anti-replay information to prevent a local host from shutting down a Brick device by generating forged or previously recorded heartbeats.

Brick device failover states

In terms of operational status and failover response, a Brick device can be in one of the following states:

- *Active*— The operational status is up, has link integrity with its LAN connections, and is available for carrying traffic. Only one Brick device in the failover pair can be in the active state at any time. The currently active Brick device owns the virtual MAC address of the failover pair.
- *Standby*— The Brick is device waiting to take over should the currently active Brick device fail or lose connectivity. This is also the initial operational state on boot-up of the Brick device.
- *X-Wired* — Each of the Ethernet ports of each Brick device must each be assigned to the same LANs. When this is not the case, the standby Brick device goes to the cross-wired state until the situation is resolved. Failover will not work while the two Brick devices are in the cross-wired state.

The `display failover` command, issued from a local or remote Brick device console, can be used to show the failover status of a specific Brick device. Refer to the *LSMS Tools and Troubleshooting Guide* for details on the Brick device CLI and the `display failover` command.

Brick device redundancy for failover

The Brick device failover or redundancy feature has the following characteristics: two Brick devices can be deployed as a failover pair. The Brick devices are identically configured, and share a single IP address.

The Brick device that boots first becomes the *active* Brick device. The other Brick device remains in *standby* state, ready to take over should the active Brick device fail or yield to it. The active Brick device reports to the LSMS on the state of the failover pair. The standby does not have an IP address and therefore cannot report its own state.

When a Brick device that is configured for failover boots, there is a delay while it listens to be sure it does not have a counterpart that is already active. If it finds that the other Brick device in the failover pair to be active, it remains in standby state.

To create a failover pair, the second Brick device is attached to the same LANs as the first one. Each interface on the second Brick device must be attached to the same LAN as the matching interface on the first Brick device.

Both Brick devices must be installed from the same floppy, using the same name, configuration, and authentication credentials. The failover pair is configured as a single Brick device in the Brick Editor, with the **Enable Brick Failover** option checked on the Failover tab. If both Brick devices are not created from the same floppy, they perform as independent devices and form Layer 2 loops.

Primary Brick

One of the Brick devices in a failover pair can be designated as the *Primary* Brick device. The Primary Brick device will be the active Brick device at all times, unless it has experienced some failure or has lost its LAN connectivity. If the Primary Brick device is currently in standby state, and the currently active Brick device detects that the Primary Brick device has been up and running and has LAN connectivity, the active Brick device will initiate failover to the designated Primary Brick device after a provisionable Failback Delay time period has elapsed.

Link health monitoring

Each Brick device continually monitors the health of its network interfaces.

Each Ethernet link can be in one of five states:

- **Down**— No link integrity.
- **Up**— Link integrity but not receiving any frames.
- **Receiving**— Receiving non-heartbeat frames.
- **Unverified** — Receiving heartbeats that do not acknowledge the heartbeats sent on this link.
- **Verified**— Receiving heartbeats that do acknowledge the heartbeats sent.

Note that the heartbeats provide a measure of Brick-to-Brick health checking. The unverified state can come about if the Ethernet link is partially broken and is passing traffic in one direction only.

The default wait period before Brick device failover takes place as a result of a link failure is configured in the **Yield Time** field. If link integrity is restored before the end of the configured wait period, failover is cancelled.

LAN connectivity


Any hub or switch can be used to implement the LANs. Note however, that for optimum failover performance, any switches used should be IEEE 802.1 Spanning Tree-enabled.

Note that if switches are used without enabling spanning tree, after a failover occurs, these switches will continue to send traffic out on the link where the active Brick used to reside until packets are detected arriving from the opposite direction. This could potentially result in a long period of network inactivity

Figure: physical Brick device failover topology

The following figure shows a physically connected redundant Brick device pair that can be configured for failover.

Figure 3-7 Example of Brick Device Failover Physical Topology



No dedicated link between the Brick devices is required. However, a dedicated link is recommended for state-sharing, and can be added, if there is a spare port available. Providing such a link improves redundancy and tolerance of extremely heavy load.

□

To Set Up Brick Device Failover

When to use

Use this task to set up Brick device failover to be invoked automatically.

Before you begin

Before you begin this task, make sure that you have cabled the Bricks to the network in a failover configuration.

Be aware that policies will not be loaded unless this procedure is performed in the proper order. If the active Brick device fails and the LSMS is unreachable, the newly active Brick device may apply an outdated set of policies to network traffic. In a new installation, the failover Brick device may have no policies loaded.

Task

Complete the following steps to set up Brick device failover to be invoked automatically.

- 1 Configure the first Brick device in failover pair by performing Steps 1 to 5 of the task [“Configure a Brick on the LSMS”](#) (p. 3-8).

Make sure that the following fields are completed:

- **Brick Name**
 - **Brick IP Address**
 - **Gateway IP Address**
-

- 2 Click **Options** on the Brick Editor.

Result The Options tab of the Brick Editor is displayed ([Figure 3-6, “Brick Editor \(Options Tab\)”](#) (p. 3-18)).

- 3 Click the **Allow MAC Addresses to Move** checkbox. This field allows the IP address of the Brick failover pair to be bound to a new MAC address in the event of a failover.
-

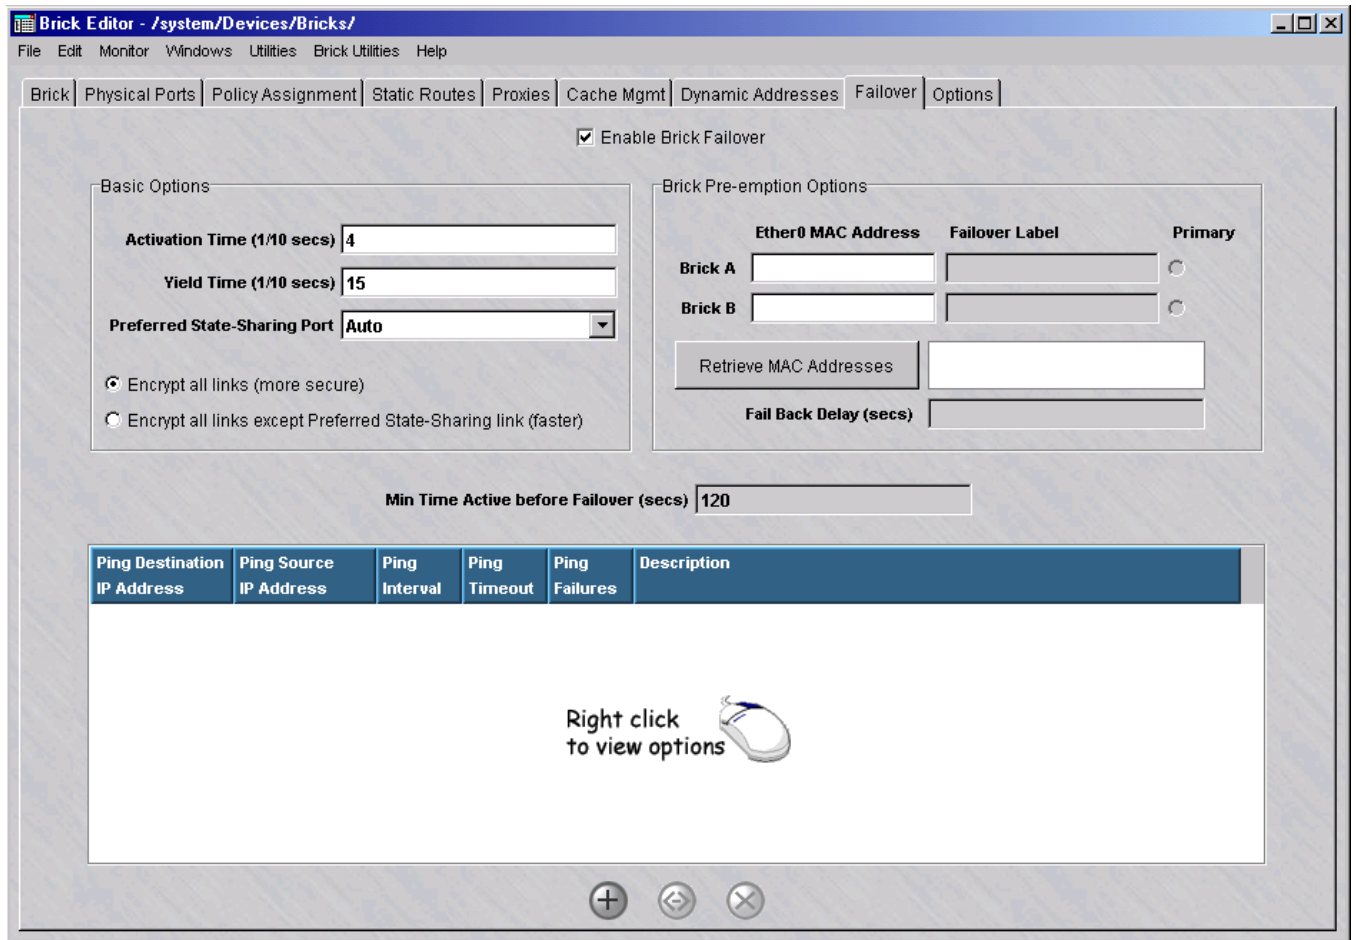
- 4 Click **Failover** on the Brick Editor.

Result The Failover tab of the Brick Editor is displayed.

- 5 Click the **Enable Brick Failover** checkbox to enable the Brick Failover feature.
-

Result The associated fields for the Failover feature are displayed on the Failover tab (

Figure 3-8 Brick Editor (Failover Tab, Brick Failover Enabled)



- 6 If the Brick Failover feature was enabled in [Step 5](#), complete the following fields:
 - **Activation Time (1/10 secs)**—this field specifies the amount of time (in 1/10 seconds) that must transpire after a missed heartbeat by the active Brick before failover to the standby Brick occurs. The default value is **4** (tenths of a second).
 - **Yield Time (1/10 secs)**—this field specifies the amount of time (in 1/10 seconds) that the active Brick must wait before failover to the standby Brick as a result of a link failure. The default value is **15** (tenths of a second, or 1.5 seconds).

- **Preferred State Sharing Port**—click the down arrow next to this field and select the Brick port to be used for state-sharing. It is recommended that the default value **Auto** be used, unless a dedicated crossover connection to one of the Brick ports is being used. If the crossover connection port is not the same on each Brick, you can specify one as the preferred state-sharing port.
- **Encrypt all links (more secure)**—by default, this radio button is selected. For additional speed, click the **Encrypt all links except Preferred State Sharing links (faster)**, which can be used when a particular link between Bricks is the preferred state-sharing link and is completely trusted.

-
- 7 To designate a Primary Brick in the failover pair, enter a MAC address in the related **Ether0 Mac Address** field of both **Brick A** and **Brick B**.

Each MAC address must be specified in hexadecimal octets, separated by a colon (:). (Example: a1:10:04:00:0d:36). A minimum of 2 octets must be specified. Both MAC address fields must be completed and the MAC address of each Brick must be different to designate a Primary Brick.

To obtain the MAC address(es) of the Brick(s), click the **Retrieve MAC Addresses** button. *Note: The Retrieve MAC Addresses button cannot be used during initial provisioning of the failover pair. It will only work after the Bricks have been flopped/booted.* One or two MAC addresses for the Brick pair is displayed in the text area to the right of the button; each address is displayed on a separate line. The first line displays the MAC address of the active Brick, and, if available, the second line displays the MAC address of the standby Brick. Either MAC address can be copied and pasted into the **Ether0 Mac Address** field of **Brick A** or **Brick B**.

Result The **Primary** radio buttons for each Brick and the **Failback Delay** field are activated on the tab.

-
- 8 To designate a Primary Brick in a failover pair, click the **Primary** radio button for **Brick A** or **Brick B** in the Brick Pre-emption Options portion of the tab.

-
- 9 If a Primary Brick was designated in [Step 8](#), in the **Fail Back Delay (secs)** field, enter a time (in seconds) to wait before initiating failover to the designated Primary Brick. The valid range is **20** to **9999**.

Important! This field is required when you designate a Primary Brick.


-
- 10 Enter a label for each Brick to be used for logging failover messages. The default label is the last four hexadecimal digits of the Brick Ether0 MAC address.

-
- 11 In the **Min Time Active Before Failover (secs)** field, enter the amount of time (in seconds) that the active Brick must be up and running and has link integrity before failover can be initiated. The default value is **120**.
-
- 12 Optionally, IP address information can be entered in the IP Tracking table section to test link connectivity to designated links to the currently active Brick. If the designated links are not working, failover to the standby Brick is initiated.

In the IP Tracking table section of the tab, right-click and select **New** from the pop-up menu.

Result The Ping Failover Editor is displayed ([Figure 3-9, “Ping Failover Editor”](#) (p. 3-29)).

Figure 3-9 Ping Failover Editor



-
- 13 In the **Ping Failover Editor**, complete the following fields:
- In the **Ping Destination IP Address** field, enter the IP address of a router or other device to be pinged by the Brick to determine if the associated link is still working.
 - In the **Ping Source IP Address** field, enter the source IP address of the Brick interface from which the ping will originate (either a VBA for the Brick, interface/VLAN address, **PPPoE#1**, **PPPoE#2**, or **DHCP**).
 - In the **Ping Interval (secs)** field, enter the time interval for sending a ping, in seconds. The default value is **10** seconds.
 - In the **Ping Timeout** field, enter the maximum time to wait for a ping response, in seconds. The default value is **1** second.
 - In the **Ping Failures for Failover** field, enter the number of consecutive responses to fail before failover is initiated. The default value is **10**

- If the Brick is configured to show VLAN information (the **Always Show VLAN Information** checkbox is checked on the Brick tab of the Brick Editor), in the **Ping Partition** field, enter the partition from which to initiate the ping. The default value is the partition with local as the VLAN ID.
- In the **Description** field, enter a textual description of this Brick failover configuration. This field can be left blank.

.....

14 Click **OK**.

Result The link entry is displayed in the IP Tracking table section of the tab.

.....

15 Repeat steps 12 through 14 for each link to be pinged to determine if failover should be initiated to the standby Brick.

.....

16 Save and apply the changes to the Brick.

.....

17 Reboot the Brick to put the failover settings into effect. An additional reboot may be necessary.

Important! Steps 18-23 are only required the first time that you set up a Brick failover pair. If any changes are made to the Brick failover configuration after the initial setup, save and apply the Brick changes and reboot the modified Brick to make the changes take effect.

.....

18 Open the Brick Utilities menu and select **Make Package/Floppy**.

Note that this step and the preceding steps are only performed once for the failover pair.

.....

19 Follow the steps in the task [“Make a Brick Boot Floppy or USB Drive on the LSMS Host”](#) (p. 3-35)

.....

20 When the Make Floppy or USB Drive task is completed, insert the boot floppy or boot USB drive in the floppy drive or USB drive and follow the steps in the task [“To Boot the Brick Device”](#) (p. 3-41).

.....

21 Check the LSMS Status Monitor to verify that the Brick is able to connected to the LSMS.

Result When the Brick connects, it synchronizes its clock with the LSMS clock and downloads policies.

- 22** Insert the floppy or USB boot floppy into the floppy or USB boot drive of the second Brick of the failover pair, and power up the second Brick to install the software from the floppy.

Result The second Brick boots up and remains in standby state.

- 23** Reboot the Brick and verify that the second Brick can contact the LSMS.

Result The second Brick synchronizes its clock with the LSMS clock and downloads policies.

- 24** Power up the first Brick and verify that it enters the standby state.

END OF STEPS



To Manually Initiate Failover

When to use

Use this task to manually initiate failover on a Brick in a failover pair.

Related information

A manual failover can also be performed on a Brick by issuing the `failover yield` command through the Brick CLI using the local or remote Brick console. For information about the `failover yield` command and the Brick, refer to the *LSMS Tools and Troubleshooting Guide*.

Before you begin

Before you begin this task, make sure that you have cabled the Bricks to the network in a failover configuration.

Task


Complete the following steps to manually initiate failover of a Brick in a failover pair.

- 1 Click the Bricks folder in the Navigator.

Result The list of currently configured Bricks is displayed in the Contents panel.

- 2 Right-click on the Brick and select **Initiate Failover** from the pop-up menu.

Result The following confirmation dialog box is displayed.



- 3 Click **Yes**.
-

Result The currently active Brick reboots and becomes the standby Brick. The current standby Brick becomes active, based on the **Activation** time parameter set on the Failover tab of the Brick Editor.

Once failover has been manually initiated, the Brick pair remain in this arrangement until a failure event occurs or another manual failover is performed.

END OF STEPS



To Activate a Brick

When to use

Once the Brick device is configured on the LSMS, it is ready to be activated. Essentially, the configuration information for the Brick device must be copied to the Brick device flash memory. When that is done, the Brick device is booted with the information in its flash memory. There are four different options available to accomplish this:

- 1 You can create a floppy disk with the Brick configuration information on the local LSMS host.
- 2 You can create a Brick boot USB drive with the Brick configuration information on the local LSMS host or from a remote host logged into the LSMS. The remote host must be a Windows PC, while you may be logged into either a Windows or Solaris LSMS.

It should be noted that the USB flash drive must be one that is 1) less than 128MB in size, and 2) guaranteed from the manufacturer to be able to function as a bootable device when formatted as FAT. Due to the variety of USB flash drives that are commercially available, and the internal differences between flash drives that seem to be identical from an external view - Lucent Technologies only qualifies that the USB flash drive orderable from Lucent Technologies will work properly. The user assumes all liability when using any USB flash drive other than the one available from Lucent Technologies. Contact your sales representative for further information.

- 3 You can create a floppy disk with the Brick configuration information from a remote host logged into the LSMS. The remote host must be a Windows PC, while you may be logged into either a Windows or Solaris LSMS.

In addition, you can use this remote host procedure to create an encrypted file that you can securely send to another network administrator to create the floppy. This feature is particularly useful if you want to outsource the floppy creation/activation function to staff who are not LSMS or Group Administrators, because these individuals will be able to create the floppy without actually accessing the LSMS.

Important! The make floppy process transfers these files to the floppy disk:

- *tvp.c.zip* (Brick operating system)
- *inferno.ini* (an ASCII configuration file containing the Brick name, IP address, and the IP address of the LSMS)

- *authinfo* (the ASCII certificate file, containing the Brick private and public keys and the LSMS public key)
- *b.com* (the bootstrap loader executable)

In addition, the boot sector on the disk is overwritten with the Brick boot loader code (which is why you cannot simply copy the above files to a disk using the Windows or Solaris copy function).

-
- 4 If you have a connection to the serial port on the Brick (as through a terminal server or a modem), you can activate the Brick without a floppy disk.

END OF STEPS

Make a Brick Boot Floppy or USB Drive on the LSMS Host

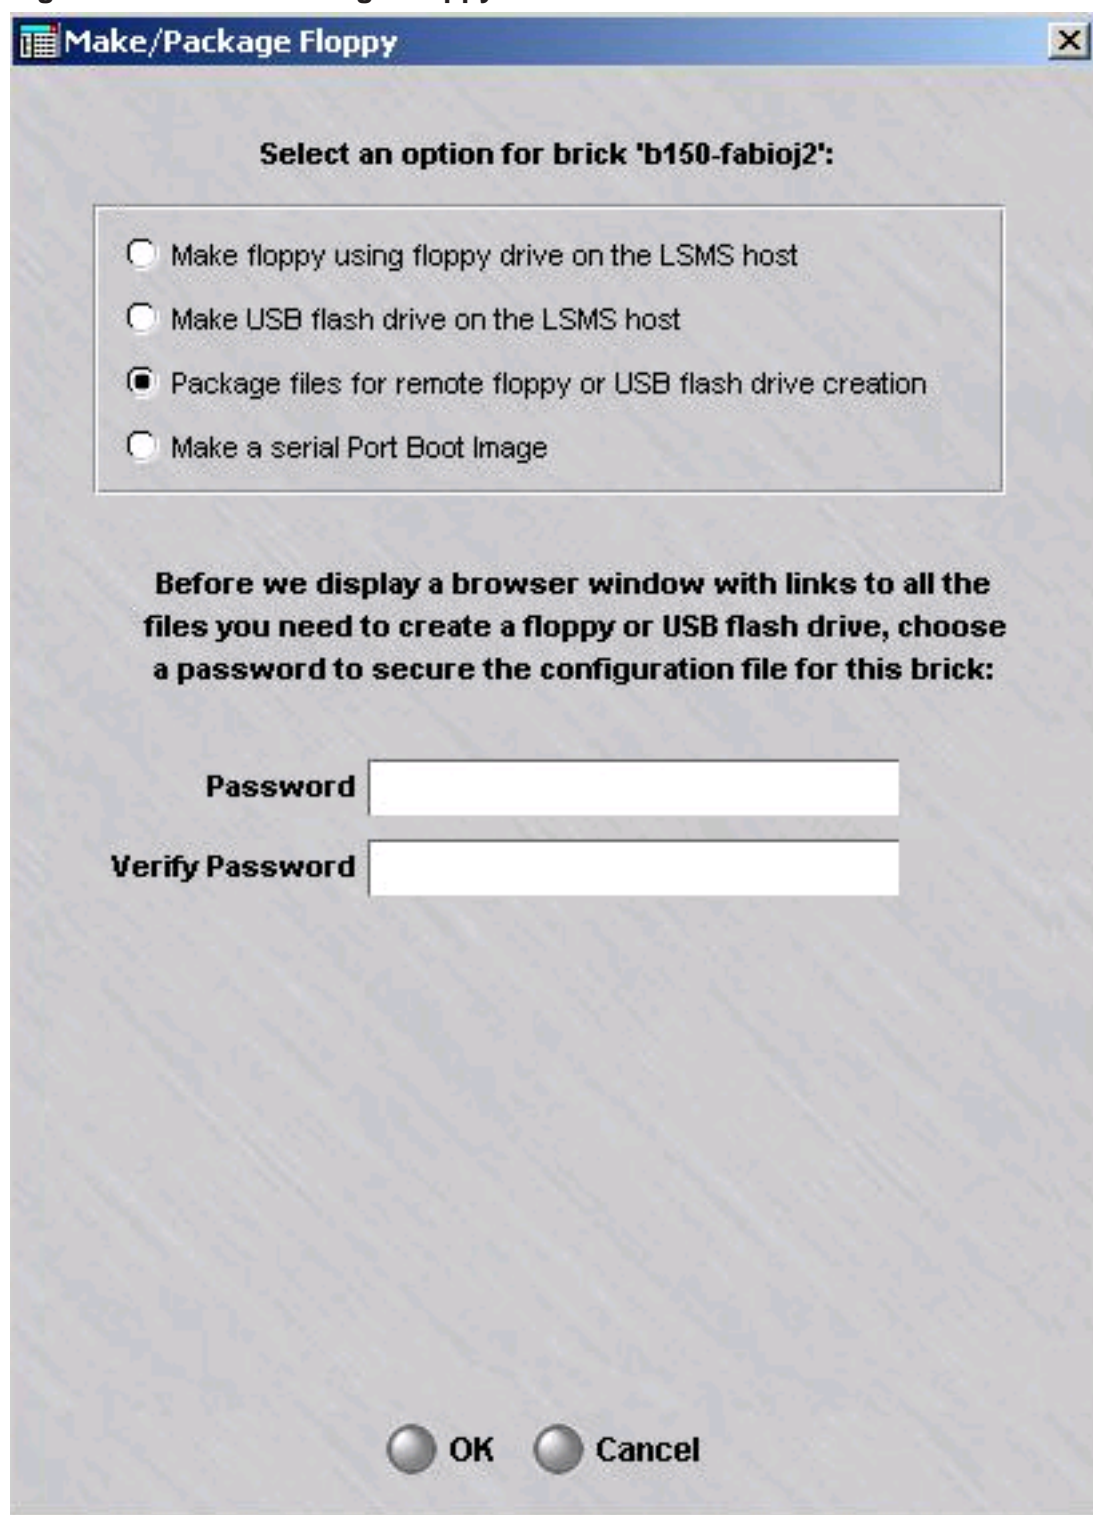
If you are sitting at the LSMS host, follow the steps below to create the floppy disk or boot USB drive:

-
- 1 If the Brick is currently displayed in the Brick Editor, open the Brick Utilities menu and select **Make/Package Floppy**.

If the Navigator window is displayed, open the appropriate group and Devices folders, and click the **Bricks** folder to display all configured Bricks. Right-click the Brick to be activated, and select **Make/Package Floppy** from the pop-up menu.

In either case, the Make/Package Floppy window is displayed (see [Figure 3-10](#), “[Make/Package Floppy Window](#)” (p. 3-36)). By default, the option **Make floppy using floppy drive on the LSMS host** is already selected.

Figure 3-10 Make/Package Floppy Window



-
- 2 If you have selected the option **Make floppy using floppy drive on the LSMS Host** (the default), insert a formatted floppy disk into the disk drive of the LSMS host.

If you have selected the option **Make USB flash drive on the LSMS host**, insert the USB drive into the USB port of the LSMS host machine.

-
- 3 Click **OK**. The LSMS downloads the configuration information to the disk or USB drive. The process takes approximately three minutes.

A pop-up confirmation window is displayed.

-
- 4 When the download is complete, click **OK** in the pop-up window and remove the disk from the LSMS host.

END OF STEPS

Make a Brick Boot Floppy or USB Drive on a Remote Host

If you are logging into the LSMS from a remote host, complete the following steps to create the floppy disk or Brick boot USB drive (Model 50 or Model 150 Brick only):

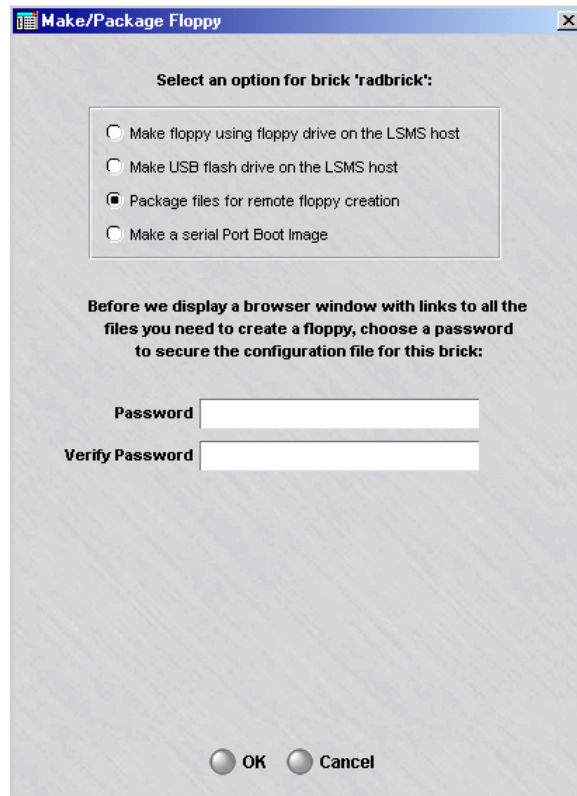
-
- 1 Log into the LSMS from the remote host using the LSMS Remote Navigator.
 - 2 Open the appropriate group and **Devices** folders, and click the **Bricks** folder to display all configured Bricks.
 - 3 Right-click the Brick to be activated, and select **Make/Package Floppy** from the pop-up menu.

Result If you configured the Brick from the remote LSMS, and the Brick is displayed in the Brick Editor, you can open the Brick Utilities menu and select **Make/Package Floppy** instead.

In either case, the Make/Package Floppy window is displayed.

-
- 4 Click the radio button labeled **Package files for remote floppy creation**. The Make/Package Floppy window prompts you to create a password (see [Figure 3-11](#), “[Make/Package Floppy Window \(with Password Fields\)](#)” (p. 3-38)).

Figure 3-11 Make/Package Floppy Window (with Password Fields)



-
- 5 Enter a password in the **Password** field, and then again in the **Verify Password** field. The password must be at least six characters.

Do not use an existing password. Make up a new password specifically for this purpose, and be sure to remember it. This password is used to encrypt and hash the configuration information, and you will need it later to decrypt the information.

If you will be sending the encrypted file to another administrator to make the floppy and boot the Brick, you will have to give that administrator this password.

-
- 6 Click **OK**. A browser window will appear with instructions for making the boot floppy or USB drive. The browser window contains links to all the files you need to make a

boot floppy or USB drive. Figure 3-12, “Browser Window” (p. 3-39) shows the browser window and the position of the links.

Figure 3-12 Browser Window

Brick Boot Files for b150-fabioj2

Download Link	Description
b150-fabioj2.lsp - 1517 bytes	Brick Configuration File
floppypackage.exe - 1083904 bytes	Brick Floppy and USB Flash Drive Creation Software for Windows NT/2000/XP/2003
95floppypackage.exe - 1094144 bytes	Brick Floppy and USB Flash Drive Creation Software for Windows 95/98/Me

How to Make a Brick Floppy or USB Flash Drive

Note: Only the brick 150 currently supports boot USB flash drives.

The following files are required to make a brick floppy or USB flash drive on your local machine:

- **Brick Configuration File** - This encrypted file contains information about a specific brick (its name, IP address, etc).
- **Brick Floppy and USB Flash Drive Creation Software** - This self-extracting ZIP file contains the **makebrickfloppy** executable for Windows and the current version of the brick software.

The **Brick Floppy and USB Flash Drive Creation Software** will be the same for all bricks, so you only need to download this file the first time you make a floppy on a given PC. The next time you need to make a boot device for another brick, you will only need to download the **Brick Configuration File** for that specific brick.

Install Brick Floppy and USB Flash Drive Creation Software:

1. Create a folder on your local machine where you will store the brick boot files. On Windows 95/98 and Windows Me all of the names in the folder path **must** be short names (8 characters or less).
2. Download the **Brick Floppy and USB Flash Drive Creation Software** to your local brick floppy folder:
 - o [floppypackage.exe](#) - 1083904 bytes - for Windows NT/2000/XP/2003
 - o [95floppypackage.exe](#) - 1094144 bytes - for Windows 95/98/Me
3. Double-click on the downloaded file to unzip it. Enter the name of your local brick floppy folder when prompted where to unzip the files.

Make a brick floppy for brick b150-fabioj2:

1. If you haven't already done so, install the Brick Floppy and USB Flash Drive Creation Software (see above).
2. Download the **Brick Configuration File** for brick **b150-fabioj2** to your local brick floppy folder:
 - o [b150-fabioj2.lsp](#) - 1517 bytes
3. Open an MS-DOS window and **cd** to your local brick floppy folder.
4. Run the **makebrickfloppy** command using the following format:
`makebrickfloppy [brick_filename] <password> <floppy_drive>`
 For example:
`makebrickfloppy b150-fabioj2.lsp <your_password> a:`
 Alternatively, simply run the **makebrickfloppy** command by itself, and you will be prompted for all the required information.

Make a brick usb flash drive for brick b150-fabioj2:

1. If you haven't already done so, install the Brick Floppy and USB Flash Drive Creation Software (see above).
2. Download the **Brick Configuration File** for brick **b150-fabioj2** to your local brick floppy folder:
 - o [b150-fabioj2.lsp](#) - 1517 bytes
3. Open an MS-DOS window and **cd** to your local brick floppy folder.
4. Run the **makebrickfloppy** command using the following format:
`makebrickfloppy [brick_filename] <password> <usb_drive>`
 For example:
`makebrickfloppy b150-fabioj2.lsp <your_password> h:`
 Alternatively, simply run the **makebrickfloppy** command by itself, and you will be prompted for all the required information.

-
- 7** If this is the first time you are making a floppy/USB drive on this host, you have to download the Brick floppy/USB drive creation software. This only needs to be done once on a host, so if you have made boot drive files on this host before, you can skip this step and proceed directly to Step 7.

To download this file:

1. Create a folder on the host to store the floppy/USB drive creation file.
2. If the remote host is running Windows 2000 or XP, click **floppypackage.exe** in the browser window. If it is running Windows 95/98/Me, click **95floppypackage.exe** instead. Then, download the file to the directory you just created. The file is a self-extracting Zip file.
3. Double-click the Zip file and extract the files to the folder you created in Step a.

-
- 8** Download the Brick configuration file to the directory containing the Brick floppy/USB drive creation software. This file is already encrypted, so it can be transferred without additional security measures.

Note that:

- If Microsoft Internet Explorer is the default browser, click the **<Brick name>.lsp** file in the browser window and download the file to the appropriate directory.
- If Netscape Navigator is the default browser, right-click the file, select **Save link as** from the pop-up menu, and save the file to the appropriate directory.

Important! *SEND AN ENCRYPTED FILE*

The next two steps explain how to make a floppy/USB drive on the machine you are now using. If your intention is to send the encrypted configuration file to another administrator to create the floppy and boot the Brick, you can stop here.

Instead, email the floppy/USB drive creation executable (*floppypackage.exe* or *95floppypackage.exe*) and the *<Brickname>.lsp* files to the person who will be creating the floppy/USB drive. Make sure that person is also in possession of the password you created in Step 4 — but do not include the password in the email message, as this defeats the purpose of encrypting the file.

-
- 9** Open a Command Prompt window and cd to the directory containing the software you downloaded. Then, double-click the *.exe* file to unzip the files in it.
-
- 10** Insert a formatted disk into the disk drive, or insert the USB drive in the USB port of the LSMS host machine, and enter this command:


```
makeBrickfloppy <Brick filename><password><drive>:
```

where:

- <Brick name>= the name of the Brick configuration file
- <password>= the password you created in Step 4
- <drive>= the disk drive holding the floppy disk (usually *a*) or the USB drive

If you enter makeBrickfloppy without the three parameters mentioned above, you will be prompted to enter them.

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

To Boot the Brick Device

Complete the following steps to copy the files from the floppy disk or USB drive to the Brick flash memory, and then boot the Brick device from the flash memory.

-
- 1** Insert the disk into the disk drive of the Brick device, or insert the USB drive into the USB port of the Brick (for Model 50, Model 150, 700, and 1200 Bricks only). The disk drive of Models 201, 350, 1000 and 1100 is located on the front panel; the disk drive of Models 300 and 500 is located on the back panel; and the disk drive of Models 50, 80, and 150 is an external attachment. Refer to the *User's Guide* for the respective Brick model or refer to the associated Brick datasheet through the following website: <http://www.lucent.com/security>.
.....
 - 2** Power up the Brick by flipping the power switch. The configuration information will be automatically transferred to the Brick flash disk.
You will hear three beeps during the transfer process:
 - A short beep will sound as the process begins,
 - A second short beep will sound approximately 10 seconds later, and a
 - Triple beep will sound within 2.5 minutes, indicating the transfer is complete......
 - 3** When the transfer process is complete, remove the floppy disk from the disk drive, close the front panel, and flip the power switch off and on a second time to re-boot the Brick device from the flash disk. If a USB drive was used to boot the Brick device, remove it from the USB port of the Brick.
Once booted, the Brick attempts to contact its LSMS and download its security policy and advanced configuration.

Important! When booting a Model 700 or 1200 Brick device from a USB Flash Drive that has a USB keyboard connected, a message may display indicating that no keyboard is present. The keyboard can still be used; just press any key on the keyboard and it will function.

-
- 4 If the Status Monitor is not displayed, open it now (see Note above). There should be an entry for the Brick, and within one minute, the Brick status should go from *LOST* to *UP*.

Important! The Certificate Authority (CA) is created on the LSMS when the Brick is initially installed. The certificate for the Brick contains both the public key and private key of the Brick. If an intruder is able to spoof the Brick public address and connect with the LSMS, the Brick policies, including VPN preshared keys and information about addresses/privilege levels might be obtained illegally.

Therefore, if a Brick floppy disk is compromised, the Brick should be deleted from the LSMS and the Brick name should not be reused.

END OF STEPS

Activating the Brick Without a Floppy Disk

In order to activate the Brick without using a floppy disk, you must have a connection to the serial port located on the rear of the Brick. For more information on the exact location of the serial port, please review the appropriate User's Guide for your Brick model. The DB9-DB9 serial cable should be wired for null modem. The Brick port is configured for 115200 baud, no parity bits, 8 data bits and 1 stop bit. Starting in Release 7.0, newly manufactured Bricks come with a special version of Brick software that, when powered up for the first time, allows them to come up to the "bootstrap" state. By using this procedure, an administrator can load a "boot image" (i.e. an encrypted version of the configuration files) to the Brick. Once the image is loaded, the Brick will then automatically reboot and connect to the LSMS via its LAN connection. When connected to the LSMS, the Brick will download the latest version of Brick software to itself, and automatically reboot one more time. At that point, the Brick is in the same state as if you had loaded it with a floppy disk.

Assuming that you have already configured the Brick on the LSMS and that you are connected to the serial port on the Brick, follow these steps:

-
- 1 If the Brick is currently displayed in the Brick Editor, open the Brick Utilities menu and select **Make/Package Floppy**.

If the Navigator window is displayed, open the appropriate group and Devices folders, and click the **Bricks** folder to display all configured Bricks. Right-click the Brick to be activated, and select **Make Package/Floppy** from the pop-up menu.

In either case, the Make/Package Floppy window will appear. Select the last radio button labelled **Make a serial Port Boot Image**. The "boot image" is the encrypted version of the Brick configuration files.

-
- 2 As you can see in [Figure 3-13, "Make/Package Floppy Window" \(p. 3-44\)](#) below, there are several fields to be completed before the boot image is created. You must enter and

verify a password to be used during the Brick bootstrap. Optionally, you may choose to enter a "User Defined Header".

Figure 3-13 Make/Package Floppy Window

Make/Package Floppy

Select an option for brick 'sol_brick3':

- Make floppy using floppy drive on the LSMS host
- Package files for remote floppy creation
- Make a serial Port Boot Image

Password

Verify Password

User Defined Header

Delivery Method Browser FTP

OK Cancel

-
- 3 You can choose the **Browser** or **FTP** option depending on the accessibility of the Brick serial connection to the LSMS platform that you are working on.

If you can connect to the Brick serial port from your LSMS Navigator or LSMS Remote Navigator, you should elect to display the boot image string in your **Browser**. Or, you may choose to **FTP** the boot image string to a more convenient location on another machine.

- 4 If you select **Browser** and click OK, the default browser for your platform will come up and display several fields of information as well as a long encrypted string surrounded by !BEGIN! and !END!, as shown in [Figure 3-14, “Make/Package Floppy Windows \(continued\)”](#) (p. 3-45) below.

Figure 3-14 Make/Package Floppy Windows (continued) (1 of 2)



Figure 3-14 Make/Package Floppy Windows (continued) (2 of 2)

```
! BEGIN!  
gnwa2z+g3f+iFxZzV5xG9BEvZBbc8ScV5XArbZarGBLqDM6jSp00paaIJ2nH/EJXZ1P6NfZJKoLB  
4TWmb4ZqNagPvXQQJF8iHeqKQKnJJ+Q3Xp1EUSdXVsYkpNc5JmCZHgBLdsxtewZrf08EjQqH+p6i  
8F44z4UL71Lsd/kHptrK+yDN/Df7dObCOXrx4WXdPXYzGwzbfGd2G9HSByhUAoeHKtRHq1UAq7Qr  
I9JupNZYI67zvOU2Iro525L/eQe+n+R1Fego8ef/H777pi7WVWkX/sLzcQ8UIFCPmtN1gA3rY+go  
xyJutvtWfali5rR3vXA8IH+OSsTObfyw8ZZB1SrSLpmeLxTsEuiru0+o47Ih9duvm9XCWYDnailk  
GUi29C13WwGT7eBxD7MUhSXnpA1xG7oY521RCCH1h43jnmIQOFgjKvV9rCubDfmY4/xGcQ+zFIU1  
56TS9KzePw70/7KYcLR5bTY+N16dRFEnV1aWz3kMDUTFIDWenWOrIB1PXG/lWieiNIREHB2MJMA7  
3w3+byf/VE1br91Hq5FjKLqiIrnH/eWiHViH8kN12mbzuEk8zQeKMLfokT3PsGx2qXf45Mv8GW1B  
BT2ioMqOJfxVumoKL1W51QsQmKwOZv iw+cSLIjaPwFc3Xp1EUSdXVnZwUqpVshwXxKAx4OVtHZF9  
zY6p5+gSutbp/Mi328Ec6yjkBIRWslRXYU3p3YG+raHAoWwyW+5JJAHWEvWWCzU2tbJXyc7O7A7T  
nVyF61v9yXqpp+UYOf1lpaR9VHPt2oq+BkVWN5vwpJKRkd6NNghC4eFkG7wFZNUIL5JR8mYAtBdO  
rDkeIzJAhCLYHPf1i2vUpNbJLAWYzn6g161zLUpX921+JTfEUxFCz5qKZFrGSNfMj/P+h6ONot64  
E5QJzRb/LC2FSArYDyQn5FqdF6jOOSxIt/RLY21LD4oMfLJE1esFruYgIfEnAXjqng4qDFNrs7S  
6f2aEwtLmTWLH9HjttfQSqe04hu2UZrj7vcrCjwhQlrgej4OheHlsrmU5uwoHE+UCx54h/47WdvU  
SVz8umPlkXRoa1a7mUhbXz14Z/5knISpE+1+oCq/ro3YACgf2rS0Ogl/PCGxFlTqcbOJpkrcQvhh  
VrfXScciC5hxH1L7y7JW7D23oWpu6nymqHbtildRe2Zk6tjANQgNQ+HxbmrRt7sdLsmYj9M4mZp  
9Wfp6L61hsp08R1h7ez7Sq8OcYefaDN4yifPuHpJrd6ddlvhapSvRANhYWv1b8Je2RdU2r4bma9g  
5Vydle41rQMw6SxyRXaIglLZt8zsJz1AfXZn17kY7xwJ/q+QzuBgBUE3xosgPOMQPgkLir14zzPN  
nLBWR8Cy7VjEYpEHJgVDZ1+S0308BFR3w11PBAHIN3C7mi+2aOFMS/xAKftxQTQikoGk+VtlocWX  
Yj0SiSbIBI3RJvqvQQfvE1TfUP0ZU++lvZQLBHS6suVL/UJ1IZnQJ1RY4CJnx0tnZSym56TJjiBs  
bI4OsvcOulmyD4suBm6Uhb rdjGggNcQ+oUutbHlEwcmD+9n38khx imVROtdi1F6aV7n5jtS41KyA  
B2161MDdIGhQCHRgUd+mK5GVq6j2E4WDMhliOMnmAiBBiMK1C+G2R3OW/uod/dVTI1Nr+A5SznDs  
QGtH2Y4KWOD1kEgp5CKYbpHZ1zRg9dzW7iBJEDT9IIOjuQ6oNEhliEwHw10cLtyXS2xiB68CbLb6  
/axleNwHDJQnopyjpbk273h90wRIAOL1GuPhi7krqTuD+pGN8JNBcyis8C5U39EoPhDIKiGpevzM  
inqwPxKhG1YLS38peiFh+LsBOyE/89BxzMNUH8kF/DQPdgXaqvO34fciUqSC42aakzPVM7Aqli9r  
u60wUrReveQcYwkjLwqiKQiVquHZN+Y8V434ucVCFk/13Gf7uwXWc+Ar4sd0qO7rhQ9K4HKvdj9Y  
ERQTCQRGBT2wOudQmDA=  
! END!
```


- 5 If you select **FTP**, additional fields will appear on the lower half of the window, as shown in [Figure 3-15, “Make/Package Floppy Windows \(continued\)”](#) (p. 3-47):

Figure 3-15 Make/Package Floppy Windows (continued)

Make/Package Floppy

Select an option for brick 'sol_brick3':

- Make floppy using floppy drive on the LSMS host
- Package files for remote floppy creation
- Make a serial Port Boot Image

Password

Verify Password

User Defined Header

Delivery Method Browser FTP

FTP User

FTP Password

FTP Host or IP

Destination Directory

You will need to supply valid input for all four fields in order for the ftp transaction to be successful. When you click OK, the LSMS will create the boot image, save it to a "txt" file, and ftp it to the specified location.

-
- 6 Now we are ready to bring up a session (either telnet or hyperterm) to the Brick serial port. Once connected to the Brick, you will see:

```
Brick> - This Brick is in factory-ship state. Bootstrap the Brick
```

-
- 7 Enter bootstrap xxxxxx where "xxxxxxx" is the password that you defined for the boot image. The Brick will respond:

```
OK. Waiting for config file on serial port... (^D to abort)
```

-
- 8 At this point, you need to paste in the long boot image string. You must include !BEGIN! and !END!; if you happen to copy and paste any of the data prior to !BEGIN!, it will be ignored. You may copy and paste the string from either the browser or from the file that you ftp-ed in Step 5.

-
- 9 When you have pasted the boot image string successfully, the Brick will reboot and attempt to connect to the LSMS via its LAN connection. Once it is communicating with the LSMS, the Brick will automatically download the latest version of the Brick software to itself. The Brick reboots itself one more time to ensure that it is running with the latest software.

After the last reboot, the Brick is in the same state as if you had activated it from a floppy disk.

Important! If you want to create a floppy disk that allows the Brick to be restored to the original factory ship mode, there are tools available on the CD to create such a floppy disk.

Please go to the "Tools\Floppyless Boot" directory on the CD and view the readme.solaris.txt file on how to create such a floppy disk on Solaris or the readme.windows.txt file on how to create such a floppy disk on Windows.

.....

END OF STEPS



4 Configuring Lucent VPN Firewall *Brick*[®] Device Ports

Overview

Purpose

This chapter explains how to:

- Configure the Brick physical ports if you want the Brick to perform routing, or if you need to set the ports to a specific speed and mode
- Disable a Brick physical port
- Assign security policies to the ports, where necessary (this includes assigning a special policy to the port connecting the Brick and LSMS)
- Create any necessary static routes so that the Brick can send traffic to LAN segments that are not directly connected to any of its ports
- Configure the Brick intelligent cache management feature so that the Brick automatically frees-up additional memory when its cache usage approaches a preset threshold.

Contents

To Configure a Physical Port	4-2
To Assign a Security Policy to a Port	4-8
Static Routes	4-20
Add a Static Route	4-22
Modify a Static Route	4-25
Activate or Deactivate a Route	4-26
Delete a Static Route	4-27
To Activate a Login Banner on the Brick Serial Port Console	4-28



To Configure a Physical Port

When to use

When the Brick was initially configured, the LSMS automatically assigned the IP address and subnet mask of the Brick to each of its ports (see "How to Configure a Brick on the LSMS" on page 3-7). If you make no changes to the addresses, the Brick will operate in a pure bridge mode.

Routing

If you want the Brick device to perform routing, you can change the IP address\\subnet mask of any of the ports, and enter a different LAN segment. You can also change the Brick port mode from auto-sensing to either half or full duplex.

Disabling a Brick device port

You can also disable a port (interface) on a Brick device. Disabling a Brick device port (interface) through the LSMS GUI is equivalent to disconnecting the wire/fiber attached to the Brick device port. This can be done to ensure that the behavior of the Brick device does not change if someone inadvertently connects that port to a LAN. When a Brick port is disabled, no traffic is allowed to pass in or out of the Brick device through that port.

A Brick port can be disabled by selecting the Physical Ports tab on the Brick Editor, and then selecting the **Disabled** option from the **Mode** field pull-down menu on the Brick Ports Editor.

A Brick port cannot be disabled if the *administrativezone* zone ruleset is assigned to it. The *administrativezone* zone ruleset protects the LSMS that is managing the Brick. A zone ruleset other than the *administrativezone* can be assigned to the disabled port, but if the assigned zone ruleset does not have a VBA, the Brick device is not accessible through that port and a warning message is issued, indicating that the zone rule assignment has no effect.

The **Send/Receive DHCP request on this port** checkbox cannot be checked if a port is disabled. The DHCP interface is typically used by the LSMS to manage the Brick device.

If a Brick is part of a Brick failover pair, at least one Brick port in the pair must be enabled and the **Ignore heartbeat failures on this link** checkbox must be *unchecked*. For additional details about Brick device failover, refer to the "[Brick Device Failover](#)" (p. 3-22) section in [Chapter 3, "Configuring and Activating a Lucent VPN Firewall Brick® Device"](#).

If the disabled port is being used for Brick device state-sharing, the currently active Brick device switches to a different port before the selected port is disabled. If that

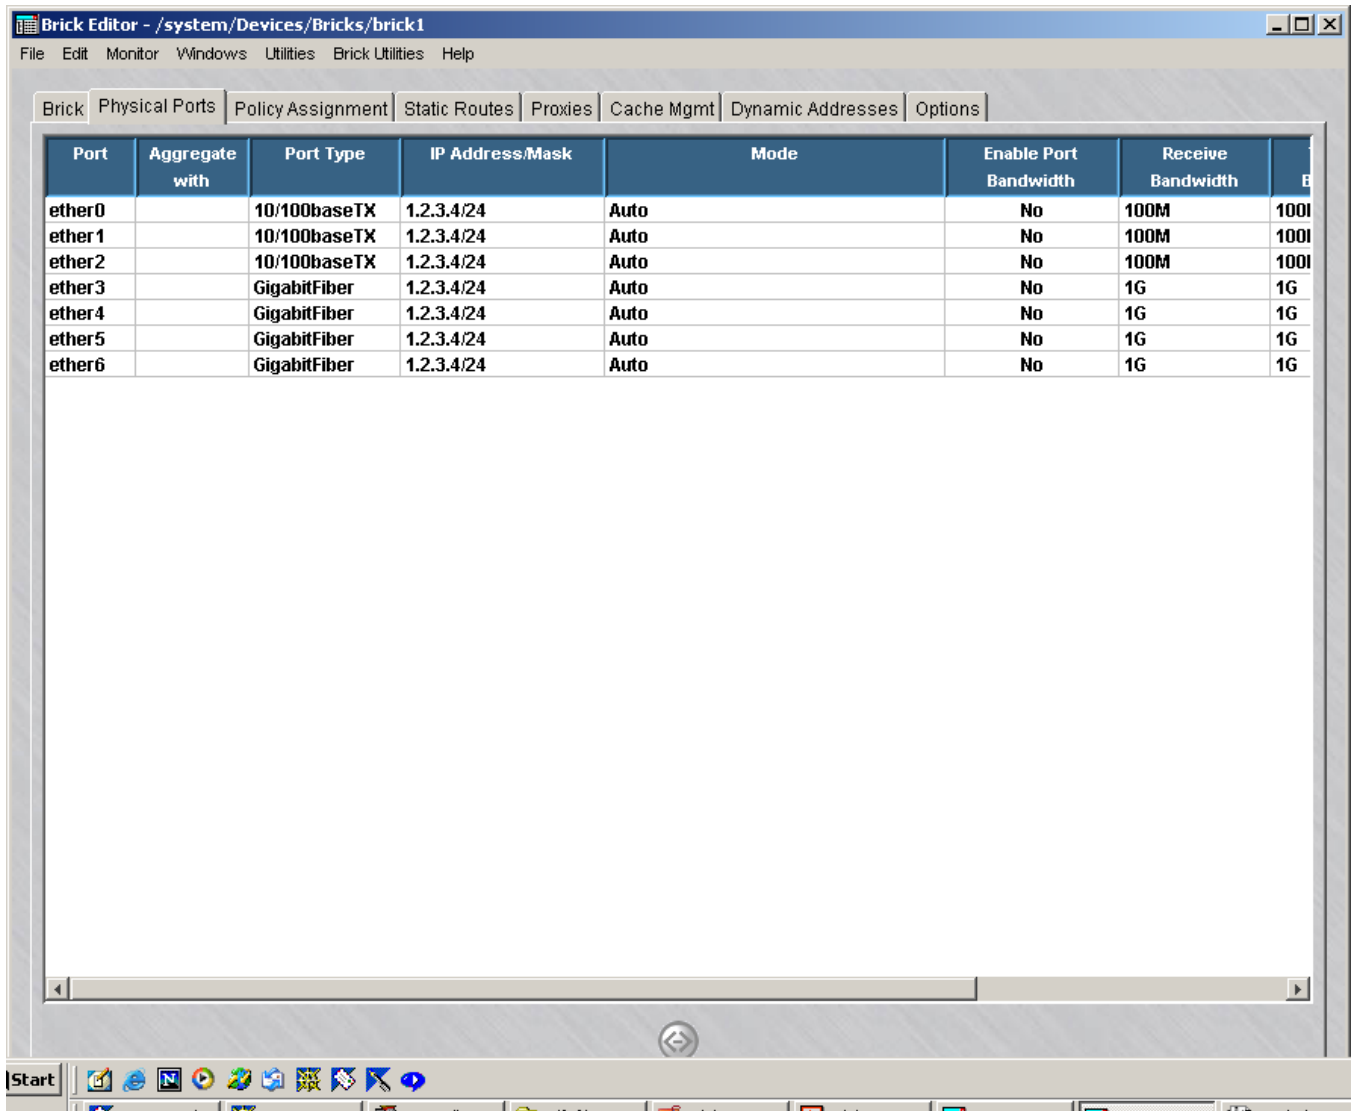
port is later switched from being disabled to enabled, the Brick device will resume using that port once it returns to the verified state.

Task

Complete the following steps to configure a port on a Brick device.

- 1 With the Brick Editor open, click **Physical Ports** to display the Physical Ports tab (see [Figure 4-1, “Brick Editor\(Physical Ports Tab\)”](#) (p. 4-3)) .

Figure 4-1 Brick Editor(Physical Ports Tab)



Note that seven 10/100baseTX Ethernet ports (ethers0 — 6) and two GigabitFiber ports (ethers7 and 8) are shown. This indicates that this is a Model 1000 Brick (7/2 configuration).

If this were a different configuration of the Model 1000, or a different model Brick, the number of ports shown would reflect this.

- 2 Double-click the port you want to configure. The Brick Ports Editor is displayed (see Figure 4-2, “Brick Ports Editor” (p. 4-4)).

Figure 4-2 Brick Ports Editor

Important! VLAN Information

If you clicked the **Always Show VLAN Information** checkbox on the Brick tab of the Brick Editor (see Figure2-5), the Brick Ports Editor will contain additional fields for you to enter the VLAN domain, VLAN default ID, VLAN membership, and receive/transmit formats.

In addition, the Physical Ports tab will not show the IP Address/Mask for each port (as in Figure3-1), but will show the VLAN domain, VLAN default ID, VLAN membership, and receive/transmit formats instead.

Both of these windows are shown with the VLAN information in *Chapter 6. Configuring VLANs on Bricks*. Refer to this chapter if you are configuring one or more ports to handle VLAN traffic.

-
- 3** The Brick supports link aggregation, where two or more physical ports are combined into one logical port so that the logical port handles more bandwidth. To aggregate this port with a logical port, select a logical port with the **Aggregate with** pull-down. When this is done, the port being edited is said to be a *child aggregate*, and the logical port in the **Aggregate with** field is said to be a *parent aggregate*. The ports in the pull-down are restricted to ones having the same **Port Type** as this one, and ones that have not yet been assigned as child aggregates. A child aggregate becomes assimilated with its parent, taking on its parent's name (child aggregates will not appear as choices on the Policy Assignment screen) and all of its parent's port attributes. Port attributes can be changed only in the parent aggregate. A parent aggregate may have more than one child, but a child only one parent.

A blank choice under the **Aggregate with** pull-down allows you to undo link aggregation, at which point both ports have the same values for all attributes. Ports that were previously child aggregates take on the values of its previous parent aggregate for entries in the Policy Assignment screen.

Link aggregation *cannot* be made hierarchical beyond parent and child; parent aggregates may not serve as child aggregates of some other parent. The **Brick Type** may not be changed if any ports are aggregated.

-
- 4** *Optional:* In the **Description** field, enter a textual description of the port.

-
- 5** To have the interface address assigned dynamically, check either **Use DHCP**, **Use PPPoE #1**, or **Use PPPoE #2**, whichever is the appropriate Addressing Method for your environment, otherwise check the Static IP Address/Mask.

-
- 6** To allow the Brick DHCP requests to go out this particular port and replies to come back in, check the **Send/Receive DHCP request on this port** checkbox. By allowing the DHCP request to go out only the port on which the DHCP server is located, you can prevent possible DHCP server spoofing from the other ports. At least one port must have this checkbox checked if a DHCP address is used anywhere on the Brick.

-
- 7 Click the **Enable Port Bandwidth Parameters** checkbox to obtain and display statistics about data packets passing in and out of this Brick port.

This option must be enabled to display Brick port data packet statistics on the Single Brick Zones monitor window, which is accessed by selecting **Monitor > Brick Status > Single Brick Bandwidth Statistics**.

For details about the Status Monitor windows, refer to [Chapter 14, “Using the Status Monitor”](#).

- 8 If you are not assigning the interface address dynamically, then enter the IP address and subnet mask in the **IP Address/Mask** field.
-

- 9 **Transmit Bandwidth** and **Receive Bandwidth** are the “total” bandwidth in each direction. These values limit the total throughput that the Brick will allow out and into the interface.

If you are not using the quality of service features, you can ignore these fields.

- 10 If this is a GigabitFiber port, you can configure the port to handle jumbo frames. To do this, select **Yes** from the drop-down list in the **Enable Jumbo Frames** field. If this is not a Gigafiber port, you will not see this field.
-

- 11 The default for the **Mode** field drop-down menu is **Auto**, which will automatically sense the correct speed for a port; however, you can specify the speed and whether traffic should be evaluated in full duplex or half duplex mode.

To disable a port, select **Disabled** from the **Mode** field drop-down menu.

- 12 **MTU** (Maximum Transmission Unit) is the largest size IP packet that the Brick will transmit on the interface. If left blank, it defaults to 1500 bytes.
-

- 13 You can also decide whether or not to ignore heartbeat failures between redundant Bricks on this link by clicking the **Ignore heartbeat failures on this link** checkbox. By default, this box is not checked. This box should only be checked if a known topology exists which prevents heartbeats from reaching the other Brick.
-

- 14 Click **OK** to dismiss the Brick Ports Editor and return to the Physical Ports tab of the Brick Editor. The changes you just made will appear in the appropriate column.
-

.....
15 Repeat Steps 2 — 13 for each additional port you want to re-configure.
.....

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

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



To Assign a Security Policy to a Port

When to use

Once a Brick device has been configured and activated, and the ports properly configured, you can begin to assign security policies — in the form of a Brick zone rulesets — to the ports on the Brick device. The rules in the ruleset determine which traffic will be permitted through each port, and which will be dropped. No packets will be permitted through the Brick unless they have been examined by at least one ruleset.

The first ruleset you will want to assign to a port is the pre-configured ruleset *administrativezone*. This ruleset is provided with the LSMS application, and its purpose is to protect the LSMS while allowing it to communicate with the Bricks it is managing. The *administrativezone* also prevents accidentally blocking Brick - LSMS intercommunication.

This ruleset should be applied to the port connecting the Brick and LSMS (if the Brick is not directly connected to the LSMS, this is the port connected to the router identified in the **Gateway IP Address** field on the Basic tab of the Brick Editor).

You do not have to assign a ruleset to every port that is in use. For example, frequently administrators do not assign a ruleset to the port connecting the Brick to the Internet. Instead, they rely on the rulesets applied to the ports connecting the internal LANs to the Brick to filter out any unwanted traffic. As long as traffic is processed by at least one zone, it can pass through the Brick. However, if traffic is processed by more than one zone, it must pass all security policies in order to be retransmitted by the Brick.

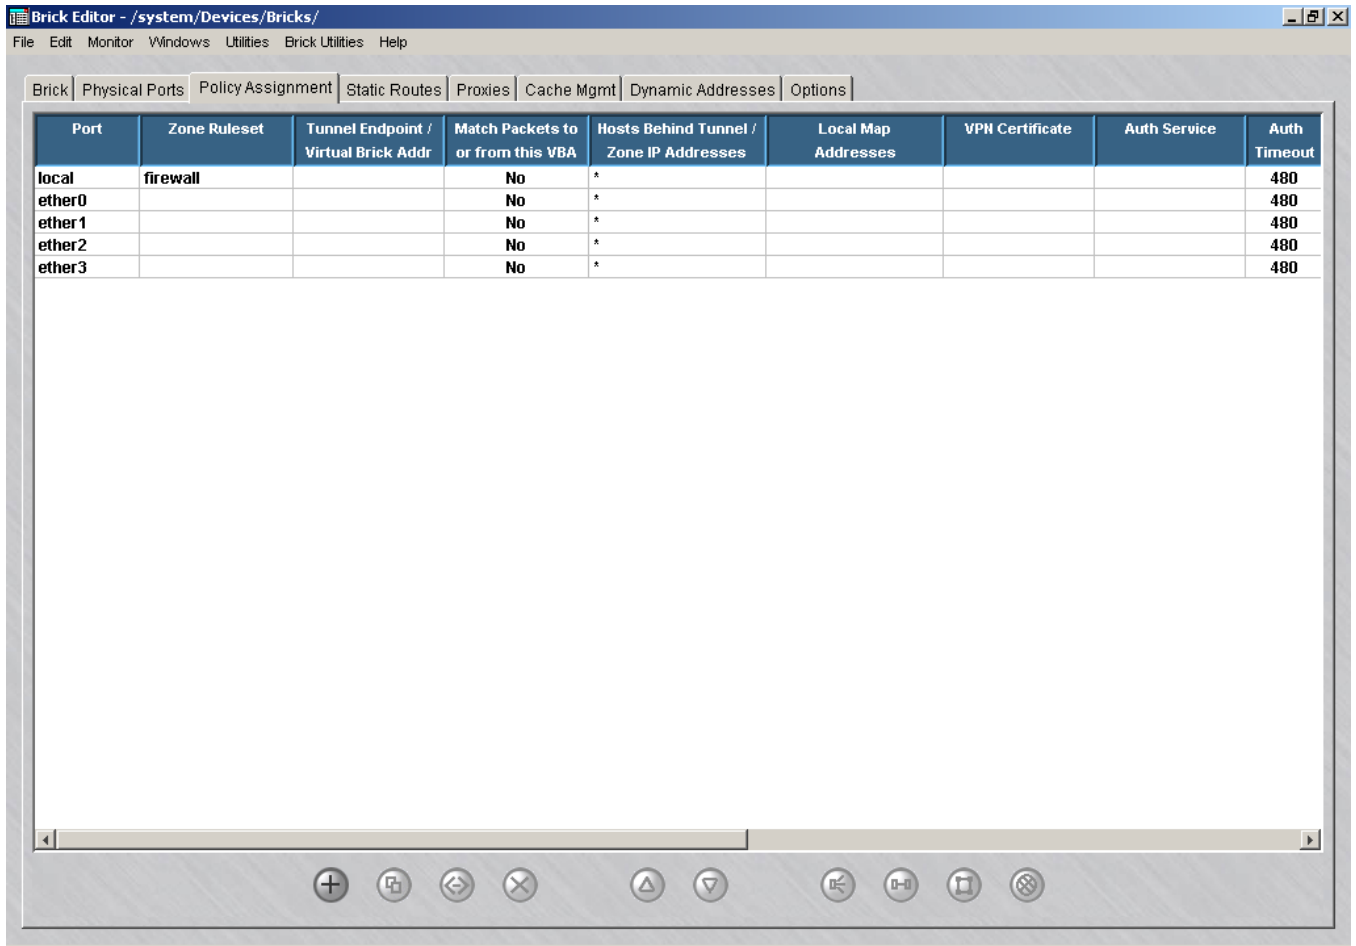
If a port will be serving as a tunnel endpoint, you also have to enter a Virtual Brick Address (VBA). This address can also be used for network address translation purposes. You can also assign IPSec Client tunnel users an address on a local LAN.

Assign a Policy to a Port

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

- 1 With the Brick Editor open, click **Policy Assignment** to display the Policy Assignment tab (see [Figure 4-3, “Brick Editor \(Policy Assignment Tab\)”](#) (p. 4-9)).

Figure 4-3 Brick Editor (Policy Assignment Tab)



-
- 2 Double-click the port you want to assign the policy to. The Brick Policy Assignment Editor is displayed (see Figure 4-4, “Brick Assignment Policy Tab (Basic)” (p. 4-10)).

Figure 4-4 Brick Assignment Policy Tab (Basic)

The screenshot shows the 'Brick Policy Assignment Editor' window with the 'Basic' tab selected. The 'Port' dropdown is set to 'ether1'. The 'Zone Ruleset' dropdown is empty. Under 'TEP/VBA Addressing Method', the 'Static' radio button is selected. Below that, there are two unselected radio buttons for 'Use PPPoE #1' and 'Use PPPoE #2'. The 'Tunnel Endpoint / Virtual Brick Address' field is empty. A checkbox for 'Match Packets to or from this VBA' is unchecked. The 'Hosts Behind Tunnel / Zone IP Addresses' field contains an asterisk (*). The 'Zone VLAN IDs' field also contains an asterisk (*). The 'Local Map Addresses' and 'VPH Certificate' fields are empty. The 'Authentication For External Users' section includes an 'Authentication Service' dropdown, an 'Authentication Timeout (mins)' field set to '480', and an 'Allowed Source IP Addresses' field with an asterisk (*). At the bottom, there are 'OK' and 'Cancel' buttons.

Important! *VLAN Information*

If you clicked the **Always Show VLAN Information** checkbox on the Brick tab of the Brick Editor, the Brick Policy Assignment Editor will contain a Zone VLAN IDs field, in addition to the fields shown above. In addition, the Policy Assignment tab will have an additional column showing the Zone VLAN IDs. See *Chapter 6. Configuring VLANs on Bricks* for a more detailed discussion of VLANs.

-
- 3 In the **Zone Ruleset** field, display the drop-down list and select a Brick zone ruleset, or select **Browse** and then select a ruleset from the Browse window. Here are a few guidelines to assist you:

Important! Ruleset Guidelines:

1. If this Brick is connected to the LSMS via this port, select the *administrativezone* ruleset. This is a pre-configured ruleset provided for this purpose. If you will be managing routers in the clear (not recommended), you can also use this ruleset. (See *Appendix B. Pre-Configured Brick Zone Rulesets* in the *LSMS Policy Guide* for details about all pre-configured Brick zone rulesets).
2. If this port will be used primarily as a firewall, protecting the LAN connected to it from attack via the Internet, select a ruleset you have created. This ruleset should reflect the desired security policy.
3. If this port will be used primarily as a tunnel endpoint, select the *vpnzone* ruleset. This is a pre-configured ruleset that contains all the rules required to quickly establish a tunnel. If you apply it to the ports on both Bricks, you have a functioning tunnel that securely passes all traffic. You can then edit this ruleset to add any additional security rules you require.
4. If you are running the Lucent Proxy Agent and this Brick will be reflecting sessions to one or more proxy hosts via this port, select the *proxyzone* ruleset. This is a pre-configured ruleset provided for this purpose.
5. If this Brick will be the NOC Gateway Brick, and this port will be connected to the LSMS, select the *nocgwzone* ruleset. This is a pre-configured ruleset provided for this purpose. You can use a ruleset of your own, but it must include, as a minimum, the rules found in the *nocgwzone* ruleset.
6. Select the *dhcpzone_on_inside* and *dhcpzone_on_outside* rulesets default zone rulesets. These rulesets are designed for routers using DHCP or PPPoE to get their public facing address. The first zone is used if it is applied on the side of the protected hosts and the second zone ruleset is used if it is applied on the interface connected to the WAN side.

-
- 4** If the **TEP/VBA Address** should be dynamically assigned rather than fixed, check either **Use DHCP**, **Use PPPoE #1**, or **Use PPPoE #2**, whichever is the appropriate Addressing Method for your environment, otherwise check the **Static IP Address/Mask**. Note that only one Zone on a Brick may use a DHCP address, and only one Zone on a Brick may use the same PPPoE address.

If this port will terminate a LAN-LAN or client tunnel, this address is required. This is the address users of the Lucent IPSec Client will enter to enable their tunnels.

If this port will only function as a firewall, this address still may be required if you will be doing network address translation (see *Chapter 5. Network Address Translation* in the *LSMS Policy Guide*).

-
- 5 In the **Host Behind Tunnel/Zone IP Addresses** field, enter the IP addresses of hosts that the ruleset will protect. If multiple rulesets are assigned to the same port, the Brick uses this field to determine which ruleset to apply to each incoming and outgoing packet. If the port will be terminating a LAN-LAN tunnel, the Brick uses this field to determine which hosts can send and receive encrypted traffic through the tunnel.

There are several ways to complete this field:

- Leave the default asterisk (*) in place. This means all hosts connected to the port are protected. The advantage of using the asterisk is that if additional hosts are added to the port, they will automatically be protected by the Brick zone ruleset. If multiple rulesets are assigned to this port, you can leave the asterisk in the *last* ruleset shown in the Policy Assignment tab (in this instance, it means "everything else").
- Display the drop-down list, and select a host group from the list, or select **Browse** and select a host group that is not on the list (because, for example, it is in a subfolder).
- Enter a specific IP address, a range of IP addresses (in the format 1.1.1.1-1.1.1.10), or an address and subnet mask (e.g., 1.1.1.0/24).

-
- 6 In certain situations, you may find it necessary to create a pool of IP addresses that will be used to provide client tunnel users with an address on a local LAN. This feature is especially useful if the client users applications that require opening sessions back to towards the clients (such as X-Windows) or in an environment where routing back to the TEP can only be guaranteed by assigning them from an address pool associated with that TEP.

This is done by entering one or more IP addresses in the **Local Map Addresses** field. This field only becomes active after you enter an IP address in the **Tunnel Endpoint** field. Enter a single IP address, an IP address with subnet mask, or a range of IP addresses. These are the addresses that will be set aside for client tunnel users.

Note that the **Local Map Address** field may also be used when you would like the Brick to respond to certain proxy ARP requests. For more information on this feature, refer to *Chapter 6. Network Address Translation* in the *LSMS Policy Guide*.

For more detailed information on local presence, including instructions on how to set it up, see *Appendix A. Local Presence* in the *LSMS Policy Guide*.

-
- 7 If this port will be used for user authentication (either firewall or tunnel users) and 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 *Chapter 8. User Authentication* in the *LSMS Policy Guide* for a detailed explanation of user authentication.

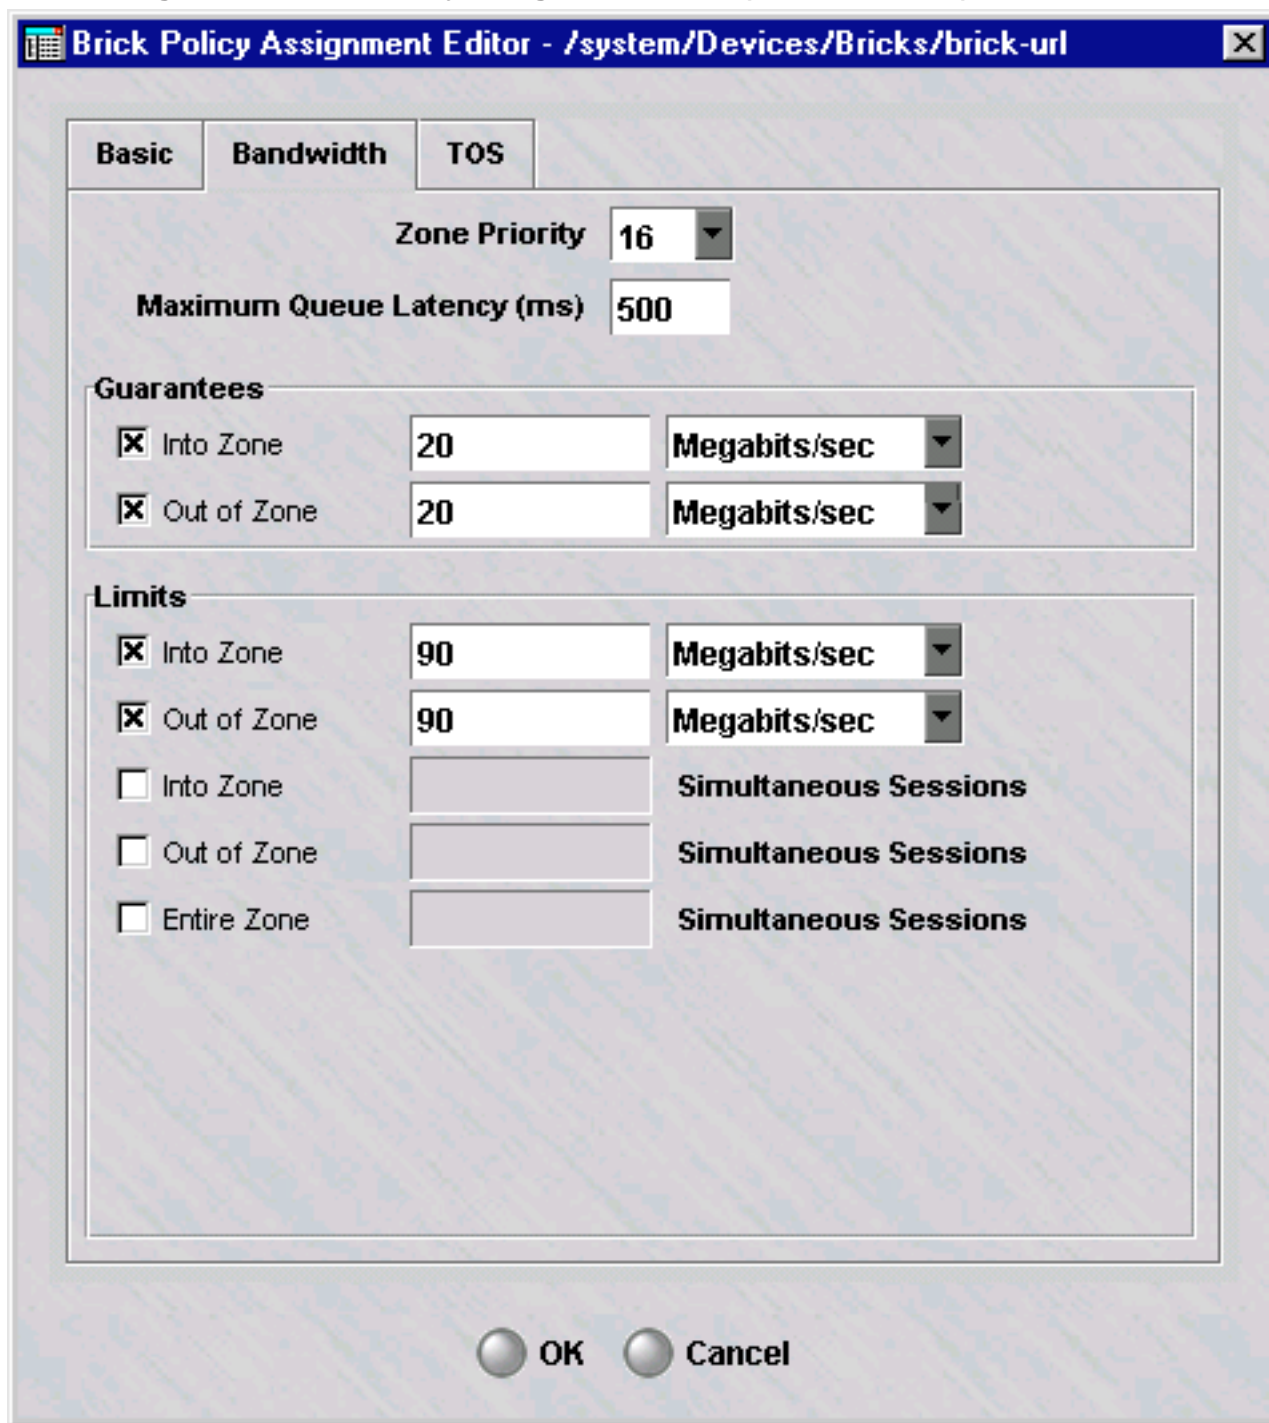
-
- 8** If a digital certificate is being used to authenticate users who will be enabling tunnels to this tunnel endpoint or for LAN-LAN tunnels, click the down arrow to display a drop-down list and select the name of the digital certificate in the **VPN Certificate** field.

If you do not see the certificate in the list, bring up the Certificate Manager, click the Group Assigns tab, and assign the VPN Certificate to the Group of the Zone that is assigned to the Brick interface.

See *Chapter 9. Digital Certificates* in the *LSMS Policy Guide* for an explanation of the procedure to obtain and install a digital certificate.

- 9 Click on **Bandwidth** to display the Bandwidth tab of the Brick Policy Assignment Editor (see [Figure 4-5, “Brick Policy Assignment Editor \(Bandwidth Tab\)”](#) (p. 4-14)). The purpose of this tab is to configure bandwidth parameters for this port.

Figure 4-5 Brick Policy Assignment Editor (Bandwidth Tab)



The parameters on this screen establish the criteria for bandwidth through this zone on this port. They can be set as bit rate guarantees, bit rate limits, and limits on simultaneous sessions. A guarantee sets the minimum acceptable bandwidth when traffic must be reduced here to accommodate higher traffic demands elsewhere. A limit caps the allowable bandwidth or number of simultaneous sessions under high traffic loads. The checkboxes enable or disable a parameter while preserving its configured value.

Note that the total guaranteed bandwidth specified for all zones on a particular physical port must not be greater than the value assigned for 'Transmit Bandwidth' and 'Receive Bandwidth' in the Brick Ports Editor. If there is more than one zone assigned to the same port, the total of the guaranteed bandwidth specified for all of the zones on that port must not be greater than the value assigned for **Transmit Bandwidth** and **Receive Bandwidth** in the Brick Ports Editor.

A limit for the number of simultaneous sessions for the entire zone is provided to allow directionally skewed values (a high value in one direction and a low value in the other), while prohibiting high values in both directions at once.

Zone Priority parameter In order to set the **Zone Priority** 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. **Zone Priority** controls which zone 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).

Maximum Queue Latency parameter This parameter controls the maximum amount of time a packet will be queued in the Brick. 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 0 completely inhibits queuing which harms some applications. Tuning this value is, at best, empirical and guesswork in many applications (especially TCP-based ones). For applications like H.323, this is fairly simple because any latency over 50 msec is considered undesirable.

For additional details, see "How to Add Bandwidth Management to a Rule" in Chapter 2 of the *LSMS Policy Guide*.

- 10 Click **TOS** to display the TOS tab of the Brick Policy Assignment Editor (see [Figure 4-6, "Brick Policy Assignment Editor \(TOS Tab\)"](#) (p. 4-16)). The purpose of this tab is to edit the TOS/DiffServ parameters.

Figure 4-6 Brick Policy Assignment Editor (TOS Tab)



DiffServ is a method (defined in RFCs 2474/2475) that intermediate systems like routers can use to prioritize IP traffic depending on the precedence portion of the ToS field (bits 0, 1, 2) and bits 3, 4, 5 of the ToS field. This tab provides the option for the Brick to change the default priority of an outbound packet as it traverses the public backbone.

These parameters inform the Brick how to set the TOS/DiffServ bits in the second byte of the IP header for packets passing through the port. Several equivalent

representations of the ToS byte 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.

Note that the TOS/Diffserv settings at the Brick level will override the TOS/Diffserv settings in the Brick zone rulesets for a given port. If the settings within the individual rules for TOS/Diffserv are more important, do not enter TOS/Diffserv information for that Brick port.

For additional details, see "How to Add ToS/Alarm Capability to a Rule" in Chapter 2 of the *LSMS Policy Guide*.

-
- 11 Click **OK** to save the changes and dismiss the Brick Policy Assignment Editor.
-
- 12 Display the File menu in the Brick Editor and select one of the **Save** options.

END OF STEPS

Assign Multiple Rulesets to a Port

You can assign more than one Brick zone ruleset to a single Brick port. When a port has more than one Brick zone ruleset assigned to it, the Policy Assignment tab of the Brick Editor has multiple entries for the same port. [Figure 4-7, "Policy Assignment Tab \(Two Rulesets Assigned to Ether1\)"](#) (p. 4-17) shows a Policy Assignment tab with two entries for ether1.

Figure 4-7 Policy Assignment Tab (Two Rulesets Assigned to Ether1)

Multiple entries

Port	Zone Ruleset	Tunnel Endpoint / Virtual Brick Addr	Hosts Behind Tunnel / Zone IP Addresses
local	firewall		*
ether0			*
ether1	sales_zone		100.10.10.1/24
ether1	clerical_zone		*
ether2			*
ether3			*

The hosts in two zones connected to the same port can communicate freely because their traffic does not pass through a Brick port. Both are, however, protected from all other traffic.

If the Zone IP address or range associated with rulesets assigned to the same port is a range of IP addresses, the order of the entries becomes important. This is because the

Brick looks at the ports in order and uses the first match found when directing a packet to a Brick zone ruleset.

Therefore, the entry that is a *subset of the larger entry must appear first*. Otherwise, the Brick will direct packets to the Brick zone ruleset with the larger range first, and packets intended for the Brick zone ruleset with the smaller range will never arrive there.

If one of the Brick zone rulesets is assigned to a port and has its IP address set to the wildcard asterisk (*) — the entry for that zone *must be the last entry for that port*. The asterisk stands for all hosts, and so by definition is the largest range.

To re-order an entry, right-click it and select **Up** or **Down** from the pop-up menu. Repeat until the order is correct.

Assign the Same VBA to Multiple Ports

It is possible to assign the same VBA to more than one port on a given Brick. This allows you to configure your network with redundant paths for VPN traffic.

When multiple ports have the same VBA, all the other policy parameters for those ports (e.g., Brick zone ruleset, hosts behind tunnel, local map address, etc.) must be the same. Therefore, the easiest way to assign the VBA to both ports is to:

1. Assign the VBA to the first port, as described above (see "Assign a Policy to a Port" on page 3-5).
2. Right click the port in the Policy Assignment tab and select **Duplicate** from the pop-up menu (or click the **Duplicate** button).
3. When the Brick Policy Assignment Editor appears, enter the second port in the **Port** field and click **OK**.
4. Repeat Steps 2 and 3 for each additional port you want to assign this VBA to.
5. If necessary, move the new port assignments up or down to ensure that the ports are in correct, ascending order, beginning with *ether0*.
If you attempt to edit the policy associated with one of the ports, you will receive the error message shown in Figure 3-8. The text of the message indicates that all fields must be the same for ports with the same VBA, and it offers you two choices of action: (1) restore the values of this port to the ones used by the other ports, or (2) overwrite the other ports with the new values of this port.

Modify a Policy Assignment

To change an existing policy assignment, follow the steps below:

1. With the Policy Assignment tab displayed (see Figure 3-3), double-click the port you want to edit. The Brick Policy Assignment Editor (see Figure 3-3) will appear, with the policy assignment displayed.
2. Make any changes to the information shown, as necessary.

3. Click **OK** to dismiss the Brick Policy Assignment Editor.
4. Display the File menu and select one of the **Save** options.

Delete a Policy Assignment

You can only delete a policy assignment if more than one policy has been assigned to the same port. In other words, if there are two entries for a given port, one can be deleted. If you have only assigned one policy to a port, and you want to remove the policy, you can only do this by modifying the assignment and removing all the values you entered (zone ruleset, hosts behind tunnel, etc.).

To delete an existing policy assignment, follow the steps below:

1. With the Policy Assignment tab displayed (see Figure 3-3), right-click the port you want to delete and select **Delete** from the pop-up menu. The confirmation window shown below will appear.
2. Click **Yes** to dismiss the confirmation window. The port entry will no longer appear in the Policy Assignment tab of the Brick Editor.

Re-Order the Policy Assignment Entries

If the entries in the Policy Assignment tab get out of ascending order — for example, if the entry for ether3 precedes ether2 — you should re-arrange the entries so that the ports and policy assignments are in the proper ascending order. If multiple zones assigned to the same port get out of order — for example, if a zone with an asterisk for its IP address range precedes a zone with a specific address range — you should rearrange the entries. If the entries are not in their proper order, problems can result.

To re-order existing port entries, follow the steps below:

1. With the Policy Assignment tab displayed (see Figure 3-3), right-click the port entry you want to re-order and select **Up** or **Down** from the pop-up menu. The entry will move one row with each click.
2. Repeat as necessary until the port entries are in the correct order.



Static Routes

Overview

The Brick operates primarily as a bridging device, connecting two pieces of the same LAN segment. However, the Brick has a Static Routing Table, and you can add routes to this table so that the Brick will send traffic to LAN segments that are not directly connected to any of its ports. A static route consists of the network address of the remote LAN and the source address of the next hop segment of the route.

Brick access to the static routing table

A Brick cannot access its Static Routing Table until it is fully booted, is communicating with the LSMS, and has downloaded its policies from the LSMS. If the Brick is not directly connected to the LSMS when it is booting, it must use the value in the **Gateway IP Address** field to route packets to the LSMS.

The Static Routing Table can be loaded from the LSMS or from the Brick flash memory, if the LSMS cannot be contacted.

Default static route

After the Brick has booted and downloaded its policies and static routes from the LSMS, if there is no "default route" entry (0.0.0.0) in the Static Route Table, the Brick uses the **Gateway IP Address** as the default route. If there is a default route entry in the Static Route Table, the Brick uses this static route as the default route for all traffic, including traffic to the LSMS.

Cost-based routing

When multiple routes to the same destination can be defined, the LSMS allows you to assign a cost-weighting value to each of these routes in the Static Route Table. The Brick routes traffic to this destination using the lowest cost available route. Route availability is determined by configuring an IP address to ping at a specified interval to test the route. If a ping to a device on the primary route fails after a specified number of attempts, that route is treated as no longer active and traffic is re-routed by the Brick to the secondary route until that one fails, and so on. This method of defining static routes can be useful for dual PPPoE configurations or for load sharing between static routes.

If multiple routes in a given partition specify the same ping target, the same source address, and the same ping interval, only one ping is sent to the target device.

If a lower cost route is restored, the Brick only switches to this route if the higher cost route becomes unavailable.

When there are two or more routes in the same partition with the same destination subnet and an identical cost value, but different next hop addresses, the Brick transmits data on a round-robin basis across all of these routes, unless one of the routes fails.



Add a Static Route

When to use

Use this procedure to add a route to a Brick Static Routes Table.

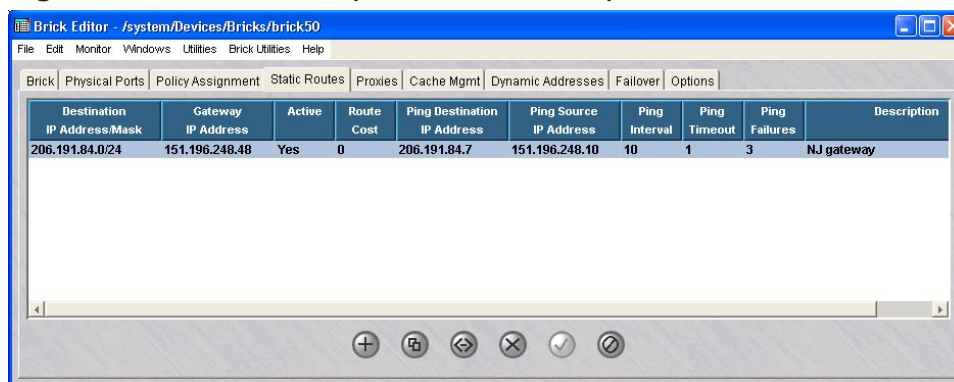
Add a Route

Complete the following steps to add a route to a Brick Static Routes Table.

- 1 If the Brick is currently displayed in the Brick Editor, click **Static Routes** to display the Static Routes tab (Figure 4-8, “Brick Editor (Static Routes Tab)” (p. 4-22)).

If the Navigator window is displayed, open the appropriate Group and Devices folders, and click the **Bricks** folder to display all configured Bricks. Double-click the desired Brick and click **Static Routes** to display the Static Routes tab.

Figure 4-8 Brick Editor (Static Routes Tab)



- 2 To add a route, right-click in the Routes panel and select **New** from the pop-up menu.

Result The Brick Static Route Editor is displayed (Figure 4-9, “Brick Static Route Editor” (p. 4-23)).

Figure 4-9 Brick Static Route Editor

- 3 Click the down arrow next to the **Route Active** field and select **Yes** to make this route active or **No** to define the route but leave it inactive. New routes are active by default.
- 4 In the **Destination IP Address/Mask** field, enter the IP address/subnet of the destination network. You must enter a network address, not a host address.
- 5 If the Gateway should be the Gateway obtained from the DHCP server or one of the two PPPoE sessions, then use the **Gateway IP Address** pull-down to fill in either *DHCP*, *PPPoE1*, or *PPPoE2*, whichever is appropriate. Otherwise, in the **Gateway IP Address** field, enter the IP address of the router that represents the first hop toward the destination. The Gateway address must fall within the IP range of one of the Brick ports.
- 6 In the **Description** field, enter a textual description of the route. This field is optional.

.....

7 In the **Route Cost** field, enter a value (non-negative integer). The value of this field is zero, by default. This field is *required*.

.....

8 To enable route verification, click the **Enable Route Verification** checkbox. This feature is optional and disabled by default.

.....

9 If route verification was enabled in [Step 8](#), complete the following fields:

- In the **Ping Destination IP Address** field, enter the IP address of the router or other device to be pinged by the Brick to determine if this route is still available.
 - In the **Ping Source IP Address** field, enter the source IP address of the Brick interface from which the ping will originate (either a VBA for the Brick, interface/VLAN address, **PPPoE#1**, **PPPoE#2**, or **DHCP**).
 - In the **Ping Interval (secs)** field, enter the time interval for sending a ping, in seconds. The default value is **10** seconds.
 - In the **Ping Timeout** field, enter the maximum time to wait for a ping response, in seconds. The default value is **1** second.
 - In the **Ping Failures for Route Unavailable**, enter the number of consecutive responses to fail before the route is declared to be unavailable. The default value is **3**
-

10 Click **OK** to dismiss the Brick Static Route Editor.

Result The new route is displayed in the Static Routes Table of the Brick Editor.

.....

11 Display the **File** menu in the Brick Editor and select **Save** or **Save and Apply**.

Important! Modified route entries do not become active until a Brick is applied.

.....

END OF STEPS

.....



Modify a Static Route

When to use

Use this procedure to modify an existing static route.


Task

Complete the following steps to modify an existing static route.

-
- 1 With the Static Routes tab of the Brick Editor displayed (see [Figure 4-8, “Brick Editor \(Static Routes Tab\)”](#) (p. 4-22)), double-click the route you want to modify.

Result The Brick Static Route Editor is displayed, with the destination network and gateway displayed.

- 2 Change either field in the window, as necessary.
-

- 3 To duplicate an existing route and modify its parameters, select the route on the Static Routes tab of the Brick Editor and click the Duplicate  button.

Result The Brick Static Route Editor is displayed, with the destination network and gateway displayed. Make any modifications needed.

- 4 Click **OK**. The modified route will appear in the Brick Editor.
-

- 5 Display the **File** menu in the Brick Editor and select **Save** or **Save and Apply**.

Important! Modified route entries do not become active until a Brick is applied.

END OF STEPS



Activate or Deactivate a Route

When to use

Use this procedure to activate or deactivate an existing static route.

Task

Complete the following steps to activate or deactivate an existing static route.

-
- 1 With the Static Routes tab of the Brick Editor displayed (see [Figure 4-8, “Brick Editor \(Static Routes Tab\)”](#) (p. 4-22)), select the route.

 - 2 To activate the route, right-click to display a pop-up menu and select **Activate**, or simply click the Activate button at the bottom of the window.

 - 3 To deactivate the route, right-click to display a pop-up menu and select **Deactivate**, or simply click the Deactivate button at the bottom of the window.

END OF STEPS



Delete a Static Route

When to use

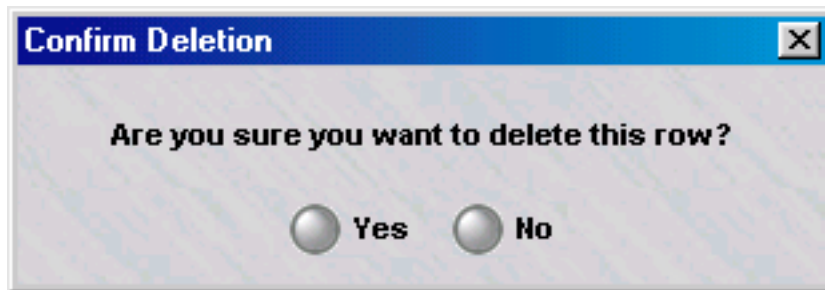
Use this procedure to remove a route from the routing table.

Task

Complete the following steps to remove a route from the routing table.

- 1 With the Static Routes tab of the Brick Editor displayed (see [Figure 4-8, “Brick Editor \(Static Routes Tab\)”](#) (p. 4-22)), right-click the route you want to delete and select **Delete** from the pop-up menu. The confirmation window is displayed (see [Figure 4-10, “Confirm Deletion Window”](#) (p. 4-27)).

Figure 4-10 Confirm Deletion Window



- 2 Click **Yes** to dismiss the confirmation window. The route will no longer appear in the Static Routes tab of the Brick Editor.

END OF STEPS



To Activate a Login Banner on the Brick Serial Port Console

When to use

Use this task to activate a Login Banner on the Brick Serial Port Console.

Before you begin

Before you begin this task, be aware that this procedure can only be performed using LSMS Release 9.0.

Important! If you are upgrading to Release 9.0 from a previous LSMS release, and the Brick Serial Port has already been activated, it will still be active following the upgrade. In this case, simply edit the Brick to add the Login Banner text, save/apply the changes, and reboot the Brick.

The Login Banner text can be modified as desired, but a reboot of the Brick is necessary whenever changes are made.

Task

Complete the following steps to activate a Login Banner on the Brick Serial Port Console.

- 1 Double-click on the desired Brick.

Result The Brick Editor is displayed.

- 2 Select the Options tab.

Result The Options tab is displayed.

- 3 Check the **Enable Serial Port** checkbox.
-

- 4 In the **Password** field, enter a password. Passwords must be a minimum of six alphanumeric characters.
-

- 5 Re-enter the password in the **Verify Password** field.
-

- 6 In the **Login Banner** dialog window of the tab, enter the desired login text.

For example:

THIS SYSTEM IS RESTRICTED SOLELY TO AUTHORIZED USERS FOR LEGITIMATE BUSINESS PURPOSES ONLY. THE ACTUAL OR ATTEMPTED UNAUTHORIZED ACCESS, USE, OR MODIFICATION OF THIS SYSTEM IS STRICTLY PROHIBITED. UNAUTHORIZED USERS ARE SUBJECT TO COMPANY DISCIPLINARY PROCEDURES AND/OR CRIMINAL AND CIVIL PENALTIES UNDER APPLICABLE DOMESTIC AND FOREIGN LAWS.

- 7 From the File menu, choose **Save and Apply**.

Result The Apply Brick window is displayed.

- 8 Click the **OK** button.
-

- 9 **Reboot** the Brick to activate the changes.

To disable the Serial Port Banner, repeat Steps 1 and 2 of this procedure and remove the Login Banner text. Reboot the Brick after saving and applying the changes. It is not necessary to disable the Serial Port to disable the Login Banner.

END OF STEPS



5 Maintaining a Lucent VPN Firewall *Brick*[®] Device Configuration

Overview

Purpose

This chapter explains how to maintain a Brick device configuration. Once a Brick device has undergone its initial configuration, you have to maintain the Brick device on the LSMS to ensure its configuration and software remain up-to-date.

You can modify certain parameters of a Brick device configuration, or delete the Brick device altogether if it is no longer operational. It is also possible to move Brick devices, in case you create additional groups or folders, and want to re-organize the LSMS interface.

You can also reboot the Brick device and refresh its MAC table from the LSMS. Finally, if new software is released, you can download it to the Brick device from the LSMS.

Contents

To View a Brick Snapshot	5-3
To Modify a Brick	5-5
Apply Changes to a Brick	5-6
Delete a Brick	5-9
To Move a Brick Device	5-10
To Reboot a Brick Device	5-11
To Reboot a Brick Device via the LSMS	5-12
Refresh the MAC Table	5-14
ARP and MAC Handling in the Brick	5-16
Static MAC and ARP Assignments	5-18
Initiate a Ping or Traceroute from a Brick	5-20

Download Software to a Standalone Brick	5-22
Download Software to a Failover Brick	5-24
Download Software to Multiple Bricks	5-25
To Configure Intelligent Cache Management	5-28



To View a Brick Snapshot

When to use

Use this procedure to view a snapshot of the current configuration of a selected Brick.

Task

Complete the following steps to view a snapshot of the current configuration of a selected Brick.

-
- 1 Click on the **Bricks** folder in the Folders panel.


Result A list of currently configured Bricks is displayed in the Contents panel.

- 2 Right-click on a Brick and select **View Brick Snapshot**.

Result A snapshot of the selected Brick’s current configuration is displayed.

Figure 5-1, “Brick Snapshot” (p. 5-4) shows an example of a portion of a Brick snapshot.

Figure 5-1 Brick Snapshot



Brick Snapshot for Brick: joesbrick

BRICK - BASIC								
Brick Name	Brick IP	Dynamically Learn Address	Use DHCP Address	Gateway IP	LSMS IP	Description	Show VLAN Info	LSMS Relative Time Offset
joesbrick	135.112.104.28	no	no	135.112.104.1	135.112.104.61		no	0.0

BRICK - TYPE			
Brick Type	Number of 10/100 baseTX Ports	Number of Gig Fiber Ports	Number of Gig Copper Ports
Model 80, 150 or 201	4	0	0

BRICK - LSMS REHOME OPTIONS	
Rehome if Higher Priority LSMS is Available	Rehome Delay (secs)
yes	300

BRICK - HOME LSMS PRIORITY		
Priority	LSMS Name	LSMS IP Address
1	joerich-1 adminServerType=source adminServerAssocLSMS=	135.112.104.61

BRICK - PHYSICAL PORTS												
Port	Aggregate with	Port Type	IP Address/Mask	Use DHCP Address	Send/Receive DHCP Request	Mode	Description	Receive Bandwidth	Transmit Bandwidth	Jumbo Frame	MTU	Ignore Heartbeat Failures
ether0		10/100baseTX	135.112.104.28/24	no	no	auto		100M	100M	no		no
ether1		10/100baseTX	135.112.104.28/24	no	no	auto		100M	100M	no		no
ether2		10/100baseTX	135.112.104.28/24	no	no	auto		100M	100M	no		no
ether3		10/100baseTX	135.112.104.28/24	no	no	auto		100M	100M	no		no

END OF STEPS



To Modify a Brick

When to use

Use this task to modify a Brick device configuration.

You cannot change a Brick device name or IP address. To change either of these, you have to delete the Brick device and re-enter it with the modified information. Then, you have to make a floppy and activate the Brick device again, as you did initially.

To change any other configuration information, follow the steps below. After certain changes (for example, changing the *halt all traffic if audit fails* parameter on the Options tab), you will be prompted to reboot the Brick device for the changes to take effect.

Procedurally:

1. With the Navigator window displayed, open the appropriate Group and Devices folders, and click the Bricks folder to display all configured Bricks.
2. Double-click the Brick to be modified. The Brick Editor is displayed with the configuration of the selected Brick.
3. Make any necessary changes to any of the tabs.
4. Display the File menu and select **Save**.



Apply Changes to a Brick

Task

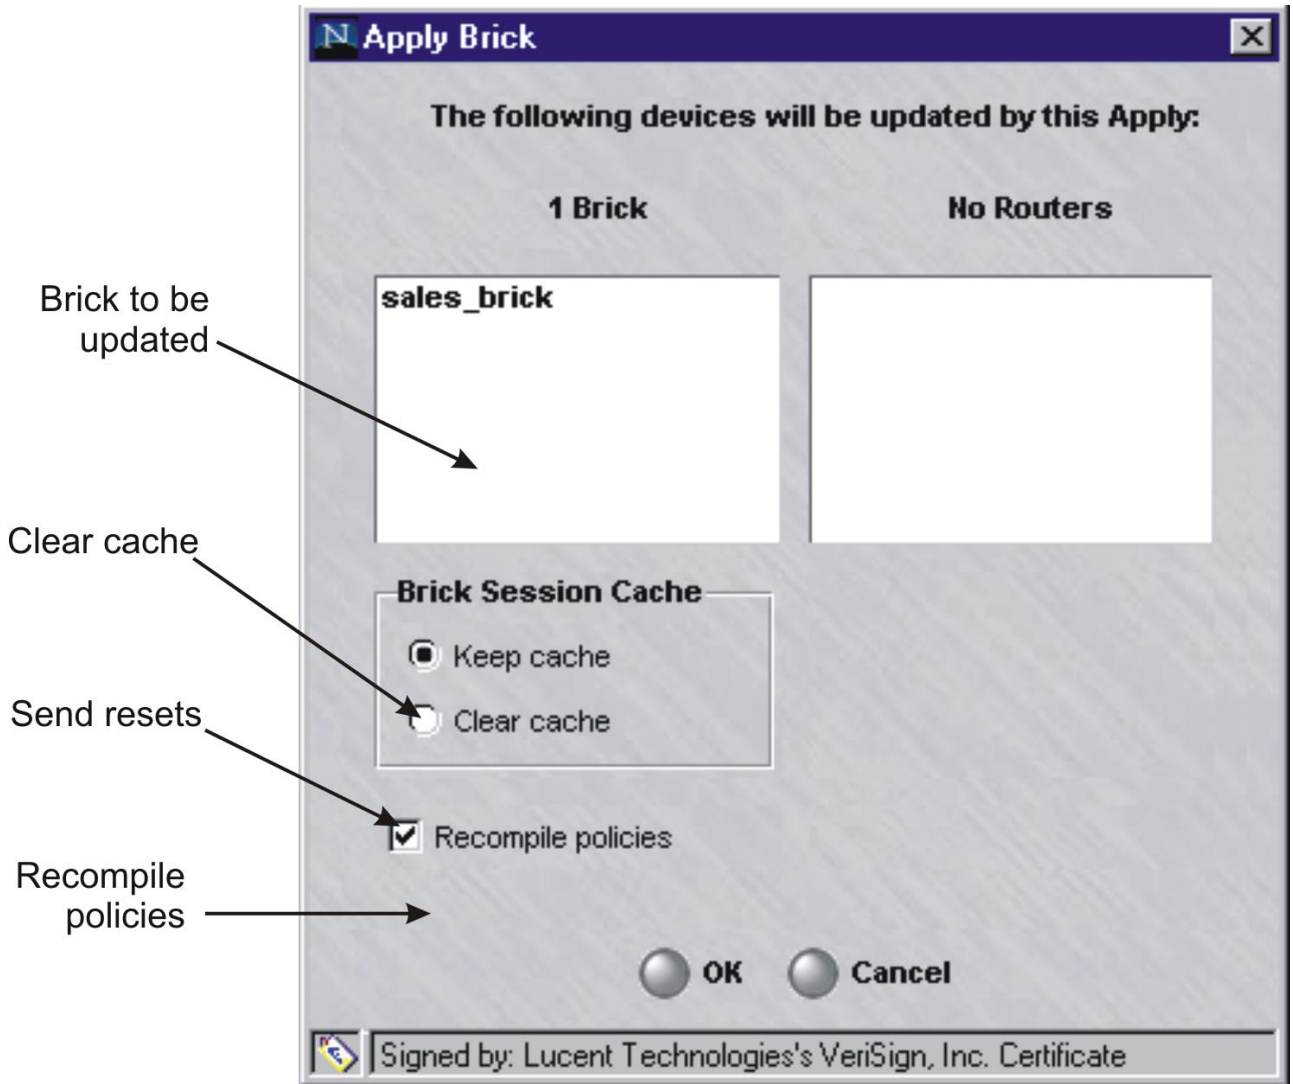
Whenever you make a change to a Brick configuration, you have to apply the changes to the device. To do this, follow the steps below:

- 1 If the Brick is currently displayed in the Brick Editor, display the Utilities menu and select **Brick▶Apply**.

If the Bricks are displayed in the Navigator window, right-click the Brick you want and click **Apply** from the pop-up menu.

The Apply Brick window will appear. It is shown in [Figure 5-2, “Apply Brick Window”](#) (p. 5-7). Note that the Brick you are applying (updating) appears in the left panel on top.

Figure 5-2 Apply Brick Window



- 2 When performing the apply, you have the option of keeping or clearing the Brick session cache. The default is to keep the cache. To clear the cache, click the **Clear Cache** radio button.

The **Keep cache** option will preserve all active sessions while applying the Brick and policy changes to the Brick. Any rule changes will not affect current sessions until the

cache timeout value for that rule is reached. New sessions will be processed using the new ruleset.

The **Clear cache** option will kill all active sessions and immediately implement any policy changes. Some sessions may re-establish automatically. However, strict TCP enforced sessions and other session types may not re-establish without user intervention. In addition, some client tunnels may be disrupted or lost. You will have to contact the client users and instruct them to re-enable their tunnels. You could also disrupt some sessions in progress, including FTP sessions and sessions allowed by rules with dependency masks.

-
- 3 If you select **Clear Cache**, the **Send Resets** checkbox, which was grayed-out, becomes active and is checked by default. If you keep this box checked, the Brick will send TCP resets to all TCP sessions it is clearing. 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, 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.

-
- 4 You have the option of recompiling the policies associated with this Brick before applying the Brick. The default is to recompile the policies. If you do not want to recompile them, uncheck the **Recompile Policies** checkbox.

If you recompile the policies, all changes to the policies will be applied to the Brick. Therefore, if you want to update the Brick configuration, but *not apply any policy changes at this time*, uncheck the **Recompile Policies** checkbox.

-
- 5 When you are ready to begin the apply, click **OK** to dismiss the Apply Brick window. The apply will take place.

Important! If, for some reason, the LSMS is unable to contact the *Brick*® (as during a network outage, for example), the policy is automatically applied when connectivity between the Brick and LSMS is restored.

END OF STEPS



Delete a Brick


Before you begin

If you intend to delete a Brick configuration, you first have to delete any tunnels that terminate on any of the Brick ports. If you do not do this, you will get a message indicating that the LSMS cannot delete this Brick, because something depends on it.

To delete a Brick, follow the steps below:

- 1 With all configured Bricks displayed in the Navigator window, right-click the Brick you want to delete and select **Delete** from the pop-up menu. The confirmation window shown below is displayed.

Figure 5-3 Confirmation Window



- 2 Click **Yes** to dismiss the confirmation window. The Brick will no longer appear in the Navigator window.

END OF STEPS



To Move a Brick Device

When to use

You can move a Brick device from one folder to another in the same group, or to a folder in another group. If you move a Brick device to a folder in a different group, all the policy components associated with the Brick device are copied to the appropriate folders in the new group. This includes Brick zone rulesets, host groups, service groups, application filters, dependency masks, LAN-LAN and client tunnels, and authentication services.

Complete the following steps to move a Brick device.

-
- 1 With the **Bricks** folder displayed in the Navigator window, right-click the Brick device you want to move, and select **Move** from the pop-up menu.

Result A Browse window is displayed.

- 2 Select the folder in the Browse window to which the Brick device is being moved and click **OK** to dismiss the Browse window.

Result The Brick device and its associated policy components are moved to the appropriate folder.

- 3 Use the Navigator window to verify the move.

END OF STEPS



To Reboot a Brick Device

Methods of rebooting a Brick

There are two ways to reboot a Brick device. One way is to power the Brick on and off by toggling the power switch on the Brick device itself. To do this, someone has to be physically present at the Brick location.

The second way to reboot the Brick device is to do it remotely from the LSMS.

Important! Whenever you reboot the Brick device, you should alert any Administrators who are overseeing client tunnels so that they can advise the Lucent IPsec Client users to re-establish their secure connections, which will be terminated by the reboot.



To Reboot a Brick Device via the LSMS

Task

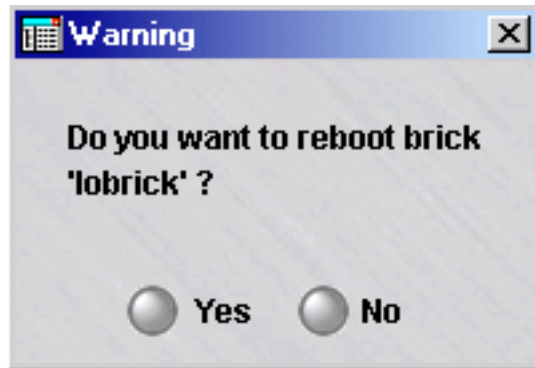
Complete the following steps to reboot a Brick device via the LSMS.

- 1 If the Brick device is currently displayed in the Brick Editor, open the Brick Utilities menu and select **Reboot**.

If the Navigator window is displayed, open the appropriate Group and Devices folders, and click the Bricks folder to display all configured Brick devices. Right-click the Brick device you want to reboot and select **Reboot** from the pop-up menu.

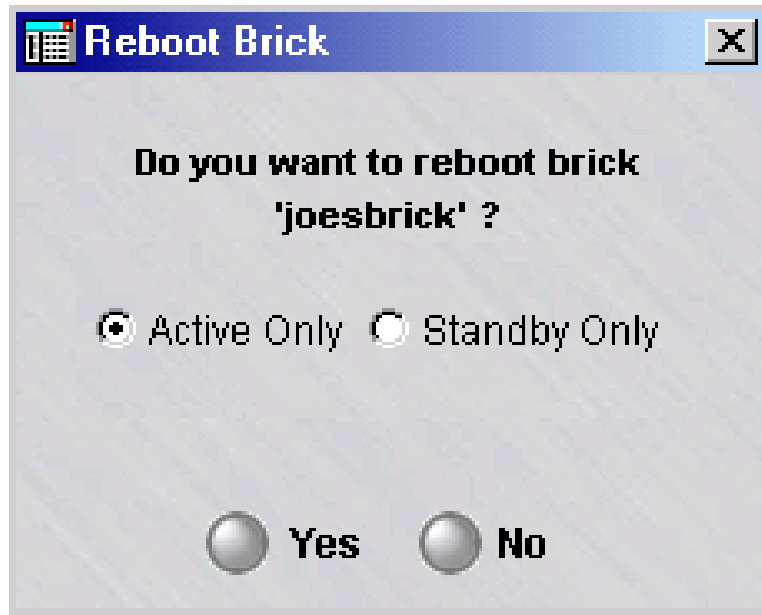
Result If it is a standalone Brick device, a dialog box similar to the following is displayed (Figure 5-4, “Warning Window for Rebooting Standalone Brick Device” (p. 5-12)).

Figure 5-4 Warning Window for Rebooting Standalone Brick Device



If the Brick device is part of a failover pair, a dialog box similar to the following is displayed (Figure 5-5, “Warning Window for Rebooting Brick Device in Failover Pair” (p. 5-13)).

Figure 5-5 Warning Window for Rebooting Brick Device in Failover Pair



-
- 2 If the Brick device is part of a failover pair, and you want to reboot either the Active Brick device or the Standby Brick device, select the one to be rebooted (the Active Brick is the default selection) and click **Yes** to proceed with the reboot.

If it a standalone Brick device, just click **Yes** to proceed with the reboot.

-
- 3 When the reboot is complete, click **OK** to return to the Brick Editor.

Important! The *reboot* function should not be confused with the *make floppy* function. When you initially installed and configured the Brick, you went through the process of making a floppy disk, transferring the information from the floppy disk to the Brick’s flash disk, and then booting the Brick from the flash disk.

The purpose of the *make floppy* function is to transfer encryption and authentication information (i.e., the security certificate) to the Brick so that it can communicate with the LSMS.

The *make floppy* function generally needs to be performed only once, when the Brick is initially installed.

END OF STEPS



Refresh the MAC Table

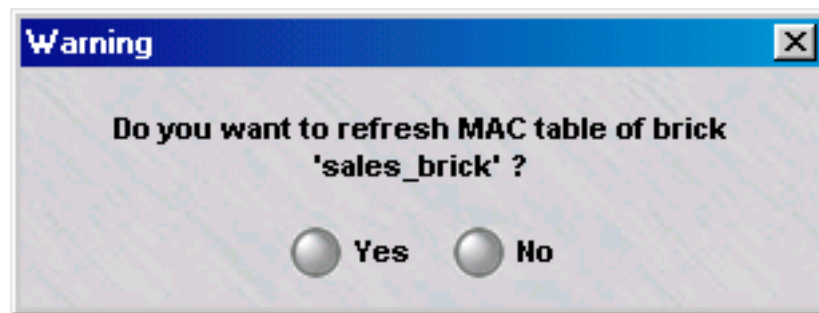
Overview

Media Access Control (MAC) addresses are hardware addresses that are hard-coded in all network interface cards. The Brick contains a table that keeps track of the MAC addresses of the hosts associated with each Brick port. For security reasons, the Brick will not allow you to move a MAC address from one port to another without refreshing the MAC table — unless you have checked the **Allow MAC Addresses to Move** checkbox on the Options tab of the Brick Editor. This prevents spoofing.

To refresh the MAC table:

- 1 If the Brick is currently displayed in the Brick Editor, open the Brick Utilities menu and select **Refresh MAC**.
- 2 If the Navigator window is displayed, open the appropriate Group and Devices folders, and click the Bricks folder to display all configured Bricks. Right-click the Brick you want and select Refresh MAC from the pop-up menu.
- 3 A warning window similar to the one shown in [Figure 5-6, “Warning Window”](#) (p. 5-14) will appear.

Figure 5-6 Warning Window



- 4 Click on **Yes** to confirm the refresh. The table will be refreshed.

-
- 5** When the refresh is complete, click **OK** to return to the Brick Editor.

END OF STEPS



ARP and MAC Handling in the Brick

Overview

MAC addresses are stored when the Brick first receives a packet from a given source MAC address. If that packet is tagged, the MAC address is associated with the VLAN on which the packet is received. The size of the Brick MAC table varies with the Brick model from a few thousand MAC addresses to several hundred MAC entries. Though MAC addresses can be refreshed manually, or allowed to move, the MAC cache is never flushed. Once the MAC table is full, the Brick must be rebooted to clear the table.

The Brick performs discovery of MAC addresses as well as ARP bindings when necessary, for packet forwarding purposes. This occurs in three different ways, using the Address Resolution Protocol (ARP):

1 HOST DISCOVERY

The Brick will transmit a broadcast ARP request in response to a need to forward a packet to a Layer-2 (MAC) address with no information about the location of that MAC address. Since the Brick will cache inbound source MAC addresses, this condition should occur only once for each new host added to the networks directly connected to the Brick.

If the Brick does not know what physical port or VLAN to use to transmit a packet, it drops that packet (under the assumption that upper-layer retransmits will recover it), and instead broadcasts an ARP Request to stimulate the destination host to identify and locate itself. Once a response is received, the MAC address of the recipient will be associated with a particular physical port (and VLAN, if the port is tagged), thereby allowing future packets to that host to be properly forwarded.

This ARP Request will only be sent on the appropriate subnet for that destination IP. Additionally, this ARP request will NOT be sent out on the physical port (and VLAN, if the port is tagged) on which the original packet was received. There is no need to forward the ARP Request out the original port: if both hosts are on that port, then they can communicate directly without the Brick being able to stop them.)

ARP entries used to stimulate MAC discovery *will* be timed out of the ARP cache. However, the MAC address of the host *will* be updated if a packet with that MAC address appears on another physical port, and the "Allow MAC addresses to move" checkbox is checked.

2 LOCAL GATEWAYS AND SPECIAL HOSTS

The Brick transmits a broadcast ARP Request when it needs to send a packet to a Gateway (a router) to perform Layer-3 packet forwarding. This ARP request will

contain the source IP address of the interface or VLAN directly connected to that Gateway, and the source MAC address of the associated physical port.

This ARP request will be sent out over all interfaces on the same subnet as the configured Gateway, or all VLANs containing that subnet. These ARP requests are performed initially when the Brick is booted, and then continuously refreshed afterwards. This action is also performed for any LSMS hosts or LPA hosts on a directly-connected subnet.

ARP Requests of this type will be refreshed periodically using a unicast ARP Request, followed by a broadcast ARP Request later if unanswered. This process continues as long as that address remains statically provisioned in the Brick as a Gateway.

3 LOCAL HOSTS

The Brick transmits a broadcast ARP Request when it needs to send a packet to a locally-connected Host and has performed Layer-3 packet forwarding itself. This ARP request will contain the source IP address of the interface of the VLAN directly connected to that Host.

This ARP Request will be sent out all interfaces on the same subnet as the destination host, or all VLANs containing that subnet.

ARP entries used for local hosts will be timed out of the ARP cache. However, the MAC address of the host will be updated if a packet with that MAC address appears on another physical port. The "Allow MAC addresses to move" checkbox must be checked.



Static MAC and ARP Assignments

Overview

The Brick attempts to learn MAC and ARP binding information automatically based on information available in the surrounding network. However, since the protocols used to do so are inherently insecure, it is possible to attempt to cause the Brick to improperly learn the local network topology. To do so is fairly difficult, since it requires administrator-level access to a directly-connected host, but it is possible. Note that this possibility applies not only to the Brick, but to any network device (such as routers or switches) that rely on MAC addresses or the ARP protocol to help make packet forwarding decisions.

To help mitigate this possibility, the Brick has the ability to create "static" MAC and ARP assignments, which may not be overridden by observed network traffic. If you are concerned that your network may be of such a trust level that hosts directly connected to the Brick may be compromisable, creating static MAC and ARP assignments for directly connected routers (and/or LSMS and LPA hosts) can help.

Task

To create a static ARP or MAC assignment for a given Brick, you need to edit a file on the LSMS host. This requires LSMS operating system host access; therefore, this ability should only be given to the most trusted administrators. Follow these steps:

- 1 In the `inferno.ini` text file, create a new text line, using the following format:

```
staticXX=mac=AA:AA:AA:AA:AA:AA ip=BBB.BBB.BBB.BBB ether=C [vid=D]  
[soft={y | n}][badSrcMac={y|n}]
```

where

XX is an arbitrary number from 0-99 inclusive used to make this entry unique

AA:AA:AA:AA:AA:AA is the MAC address

BBB.BBB.BBB.BBB is the IP address

C is the physical port number (range depends on Brick type)

D is the VLAN ID (1-4094)

A static ARP entry is created by specifying the MAC and IP parameters. A static MAC entry is created by specifying the MAC and port parameters. The two may be combined in a single entry.

The `soft` parameter specifies whether or not the assignment may be updated if new information is discovered. If being used for security purposes, `soft` should be set to `n`, or omitted from the assignment.

The default for vid for static ARP entries is the VLAN ID in the default partition whose subnet includes the IP parameter address. In the case of static MAC entries, the default is the port default VLAN ID.

When the "Route Return Path Packets to Cached Source MAC Address" checkbox on the Brick editor is checked, the Brick will use the source MAC address on the first packet of a session to determine the routing for the return packets of that session. In some cases, such as with some VRRP and HSRP routers, this behavior may cause undesirable routing. The Brick supports the ability to disable this behavior for specific MAC addresses by creating a static MAC address and giving it the "badSrcMac=y" option.

Examples:

1. a static MAC assignment, assigning MAC 01-02-03-04-05-06 to ether3.
`static1=mac=01-02-03-04-05-06 ether=3`
2. a static MAC assignment, assigning MAC 01-02-03-04-05-06 to ether3 and VLAN ID 27, allowed to be overridden
`static5=mac=01-02-03-04-05-06 ether=3 vid=27 soft=y`
3. a static ARP binding MAC 01-02-03-04-05-06 to IP 10.1.1.1
`static7=mac=01-02-03-04-05-06 ip=10.1.1.1`
4. a static combined MAC and ARP binding MAC 01-02-03-04-05-06 to IP 10.1.1.1 on ether4 and VLAN ID 36
`static22=mac=01-02-03-04-05-06 ip=10.1.1.1 ether=4 vid=36`

-
- 2 Change some option in the Brick. Save and Apply the Brick. Change the option back to its original value. Save and Apply the Brick again.

-
- 3 Reboot the Brick.

END OF STEPS



Initiate a Ping or Traceroute from a Brick

When to use

This feature allows an administrator to initiate an outbound ping or traceroute from a Brick to another device. To do this, you must first open a console window on the Brick, either by directly accessing the Brick or by using the LSMS remote console or the LSMS remote navigator console. You do not have to add any rules or make other policy modifications for the ping and traceroute to work.

Ping

To execute a ping, you must include the IP address of the target device in the ping command. The Brick will report each success or failure individually in real-time, with round-trip time (in ms), plus an overall success rate as a count and percentage, including average round-trip time.

The syntax of the ping command is

```
ping [options] target_ip
```

where *target_ip* is the IP address of the target device and the available options are:

- t ttl
- w timeout
- c (for continuous)
- n number of requests
- v vlantag
- i interface# to send ping to
- I interval between pings (in seconds)
- l data size (in bytes)
- s source IP
- o make packet come OUT of the zone
- ! (to bypass rule processing)

The ping command will default to a 64-byte packet sent every second for five seconds.

Traceroute

To execute a traceroute, you must include the IP address of the target device in the traceroute command. The Brick will report each probe round-trip time with increasing TTL until the target is reached, in real-time.

The syntax of the traceroute command is

```
traceroute [options] target_ip
```

where *target_ip* is the IP address of the target device and the available options are:

- t max ttl
- q max queries per hop
- w timeout
- v vlantag
- i interface# to send traceroute to
- U use UDP instead of ICMP
- l data size
- s source IP
- o make packet come OUT of the zone
- ! (to bypass rule processing)

The traceroute command will default to 3 x 64-byte ICMP ping packets sent to every TTL increment, with a one-second timeout.



Download Software to a Standalone Brick

When to use

LSMS upgrades, point releases, and patches will be released periodically via a CD-ROM or from a website. An administrator has to load the new software on the LSMS, using the installation procedures provided with the software.

Task

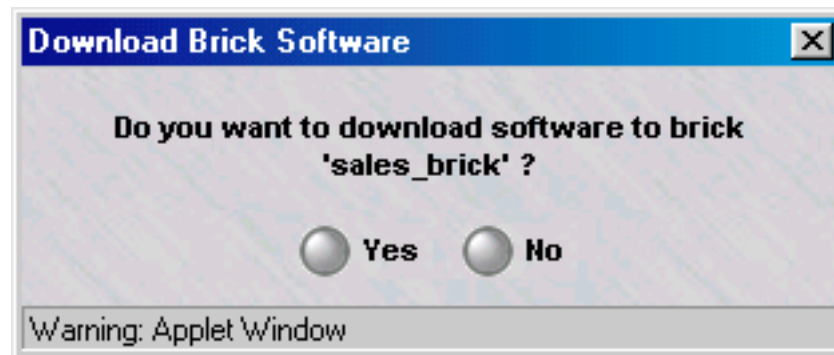
Once this has been done, the administrator has to update the software of each Brick connected to the LSMS. To do this, follow the steps below:

- 1 If the Brick is currently displayed in the Brick Editor, open the Brick Utilities menu and select **Software Download**.

If the Navigator window is displayed, open the appropriate Group and Devices folders, and click the Bricks folder to display all configured Bricks. Right-click the Brick you want and select **Software Download** from the pop-up menu.

A warning window similar to the one shown in [Figure 5-7, “Warning Windows \(Standalone Brick\)”](#) (p. 5-22) will appear.

Figure 5-7 Warning Windows (Standalone Brick)



- 2 Click **Yes** to confirm the download. The software will be downloaded to the Brick. The download is carried out over an encrypted link.

-
- 3** When the download is complete, click **OK** to return to the Brick Editor. A reboot is required to make the new software operational.

END OF STEPS



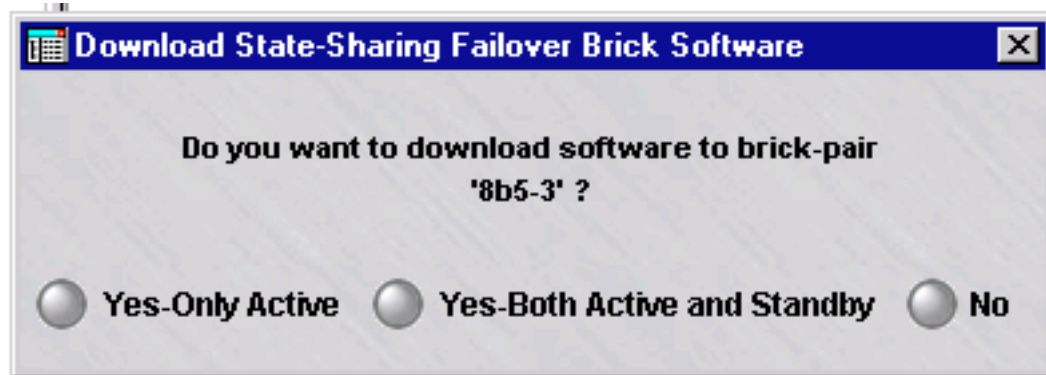
Download Software to a Failover Brick

When to use

If you are downloading the software to a Brick that is part of a failover pair, you have the option of downloading the software to the active Brick only or downloading the software to both the active and standby Bricks.

The procedure is the same as described above for a standalone Brick, except you will see the following message instead of the one shown in [Figure 5-8, “Warning Window \(Failover Brick\)”](#) (p. 5-24).

Figure 5-8 Warning Window (Failover Brick)



Click **Yes-Only Active** to download the software to the active Brick, **Yes-Both Active and Standby** to download the software to both Bricks, or **No** to terminate the download.

□

Download Software to Multiple Bricks

When to use

If you need to download software to more than one Brick, you can do this one Brick at a time, as described above, or you can download the software to all the Bricks in one operation. Since it is necessary to reboot a Brick after the software has been successfully downloaded, you can also include reboot instructions in the operation, so that all the Bricks are automatically rebooted after the download.

When performing the download and reboot, you can specify the Bricks in three ways:

- All the Bricks in a folder
- All the Bricks in a group
- All the Bricks in all groups that you have full device permission on

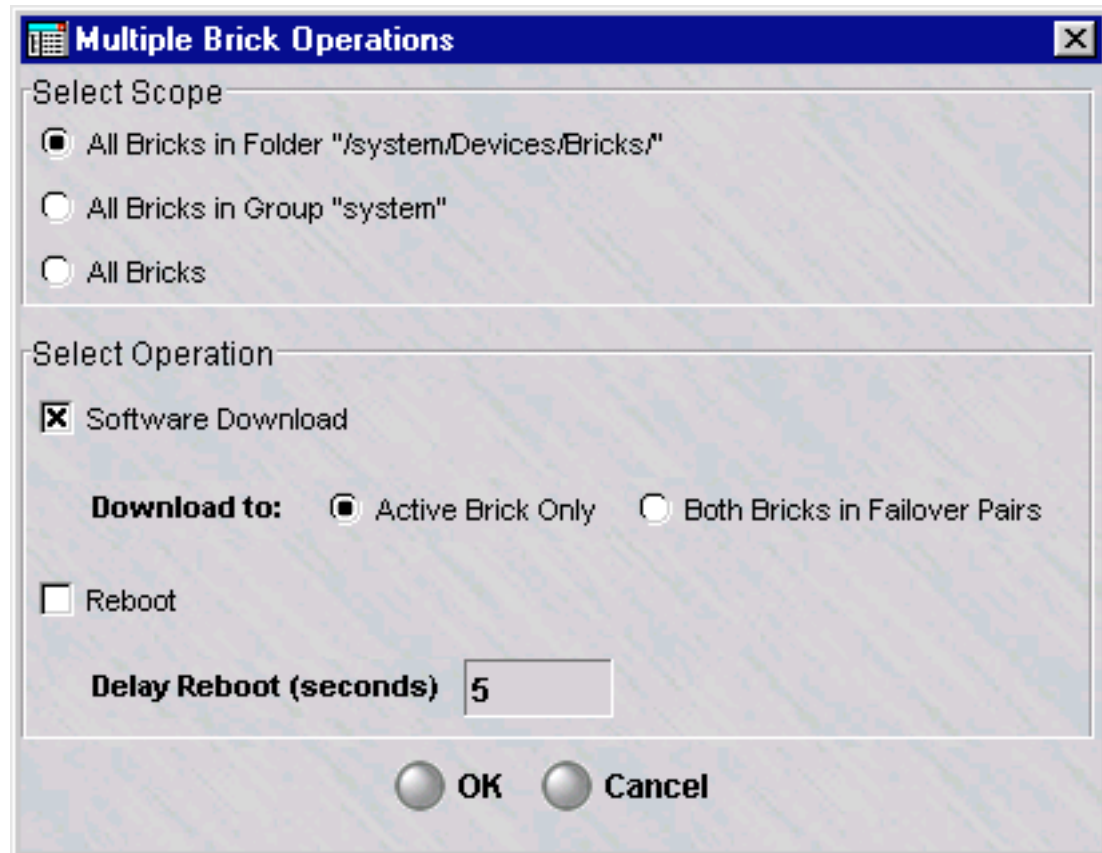
Task

The following explains the procedure:

- 1 With the Navigator window displayed, right-click the Bricks folder and select **Software Download for Multiple Bricks** from the pop-up menu. The Multiple Brick

Operations Window will appear. It is shown in [Figure 5-9, “Multiple Brick Operations Window”](#) (p. 5-26).

Figure 5-9 Multiple Brick Operations Window



Important! You can also select **Software Download then Reboot Multiple Bricks** from the pop-up menu. The only difference is that the screen shown in [Figure 5-9, “Multiple Brick Operations Window”](#) (p. 5-26) will appear with the **Reboot** checkbox already checked.

-
- 2 When the Multiple Brick Operations Window first appears, the **All Bricks in Folder “/system/Devices/Bricks”** checkbox is checked. This means all Bricks in this folder will be rebooted. You can change this to all Bricks in the system group, or all Bricks that you have full device permission on, simply by clicking the appropriate checkbox.
 - 3 By default, the software will only be downloaded to active Bricks. If you have one or more failover pairs, and you want the software downloaded to both Bricks in the pairs, click the **Both Bricks in Failover Pairs** checkbox.

-
- 4** If you want the Bricks to be automatically rebooted after the software has been downloaded, click the **Reboot** checkbox, and then indicate the delay time in the appropriate field (default = 5 seconds).

The delay time determines how long after the download the Bricks will wait before rebooting themselves. This allows the reboot command to apply to all Bricks before any of them actually begin the reboot process. This is important when one of the Bricks is a gateway, and you want to avoid the problem of having this Brick reboot before all the reboot commands reach the Bricks behind the gateway.

For failover Bricks, only the active Brick will be rebooted. If you intend to reboot the Bricks selectively, you may leave the **Reboot** checkbox unchecked.

END OF STEPS



To Configure Intelligent Cache Management

When to use

The LSMS provides a patented intelligent cache management (ICM) feature that allows you to configure the Brick device so that if cache usage approaches a preset threshold, the Brick device automatically purges less important sessions to clear cache memory for new sessions.

The purpose of this feature is to help prevent denial-of-service attacks — in which attackers attempt to flood your network with a sustained stream of high-bandwidth traffic — from flooding the session cache and tying up valuable network resources. (A “session” is not just a TCP connection. The Brick device also treats UDP and ICMP packets as sessions. Some services (such as FTP) require multiple sessions.)

The intelligent cache management feature is one of three features built into the Brick device to protect against denial of service attacks. These features are described in greater length in *Appendix C. Denial of Service Attacks*.

Determine the Threshold Levels

Some considerations when determining the threshold levels are:

- 1 The challenge in determining effective threshold levels for ICM is picking levels that will allow ICM to activate if the network is under attack, while making sure that ICM will not activate under a “normal” traffic load.

The challenge is to determine what a “normal” traffic load is. A reasonable approach is to monitor network traffic for some period of time and attempt to determine how much memory is needed (or how many sessions are created) during that period — particularly at peak traffic times. There are a number of things that can be done to shed light on that question:

- Look in the Proactive Monitor logs to find out what the peak cache memory usage is.
- Attach a protocol analyzer (i.e., a “packet sniffer”) to the network, capture packets, and analyze the captured packets to find out how many ICMP, UDP, and TCP packets there were over a period of time. Use that information (perhaps coupled with reports from the session logs) to determine how the cache class thresholds should be set.

Ideally, the thresholds should be set a little bit above what the expected peak memory usage is. With such settings in place, the Brick should be able to handle normal peak traffic flow without activating ICM; and if an attack causes cache memory usage to exceed the expected levels, ICM will activate to protect the Brick and the network.

END OF STEPS

Set the Threshold Levels

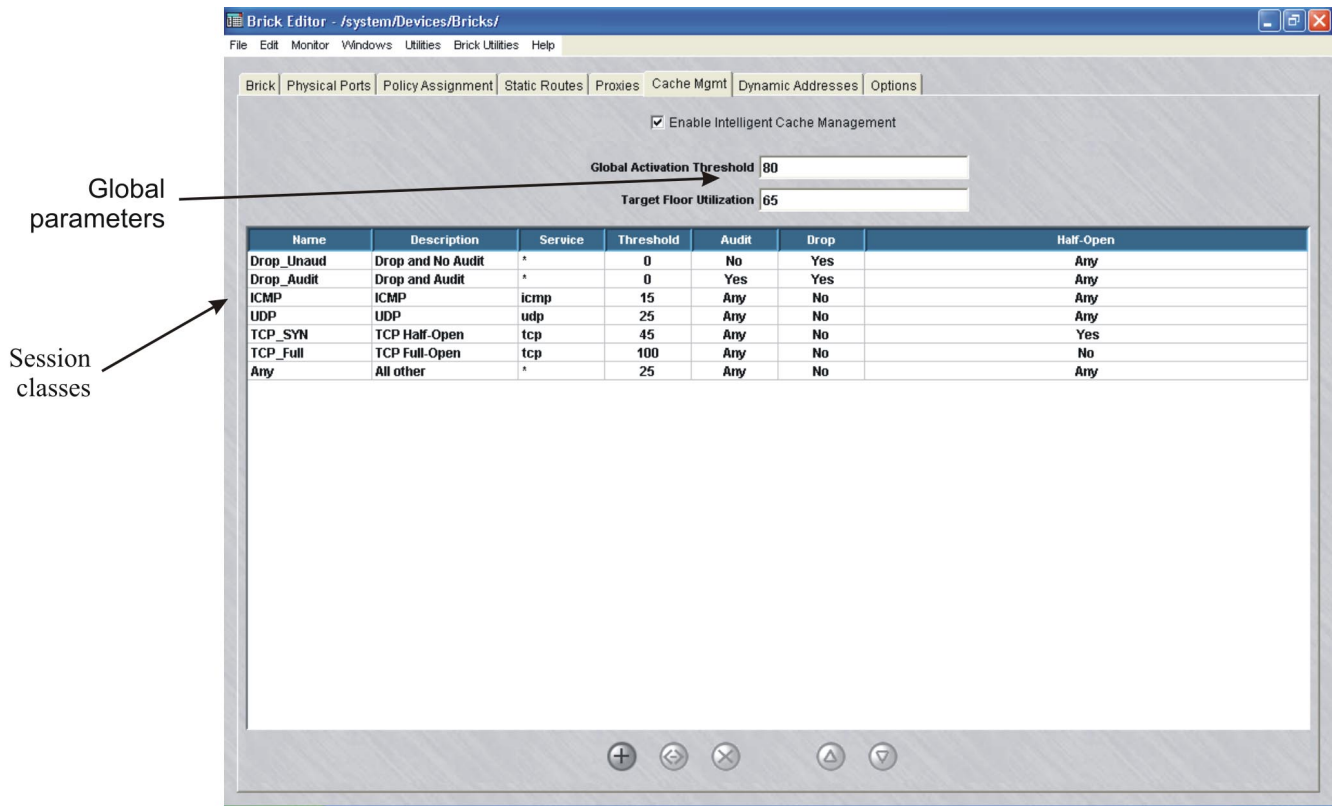
To use the intelligent cache management feature to set global and session thresholds, follow the steps below:

- 1 If the Brick is currently displayed in the Brick Editor, click **Cache Mgmt** to display the Cache Management tab. It is shown in Figure 3-10.

If the Navigator window is displayed, open the appropriate Group and **Devices** folders, and click the **Bricks** folder to display all configured Bricks. Double-click the Brick you want and click **Cache Mgmt** to display the Cache Management tab.
- 2 The **Enable Intelligent Cache Management** checkbox is checked by default. This causes the two global parameters and the session classes to become active. To disable this feature, uncheck this box.
- 3 Change the values in the **Global Activation Threshold** and **Target Floor Utilization** fields, if necessary. The defaults are 80% and 65%, respectively. These figures represent a percent of total cache capacity.

When the global activation threshold is reached, the intelligent cache management feature will begin scanning the session cache to identify sessions that can be cleared.

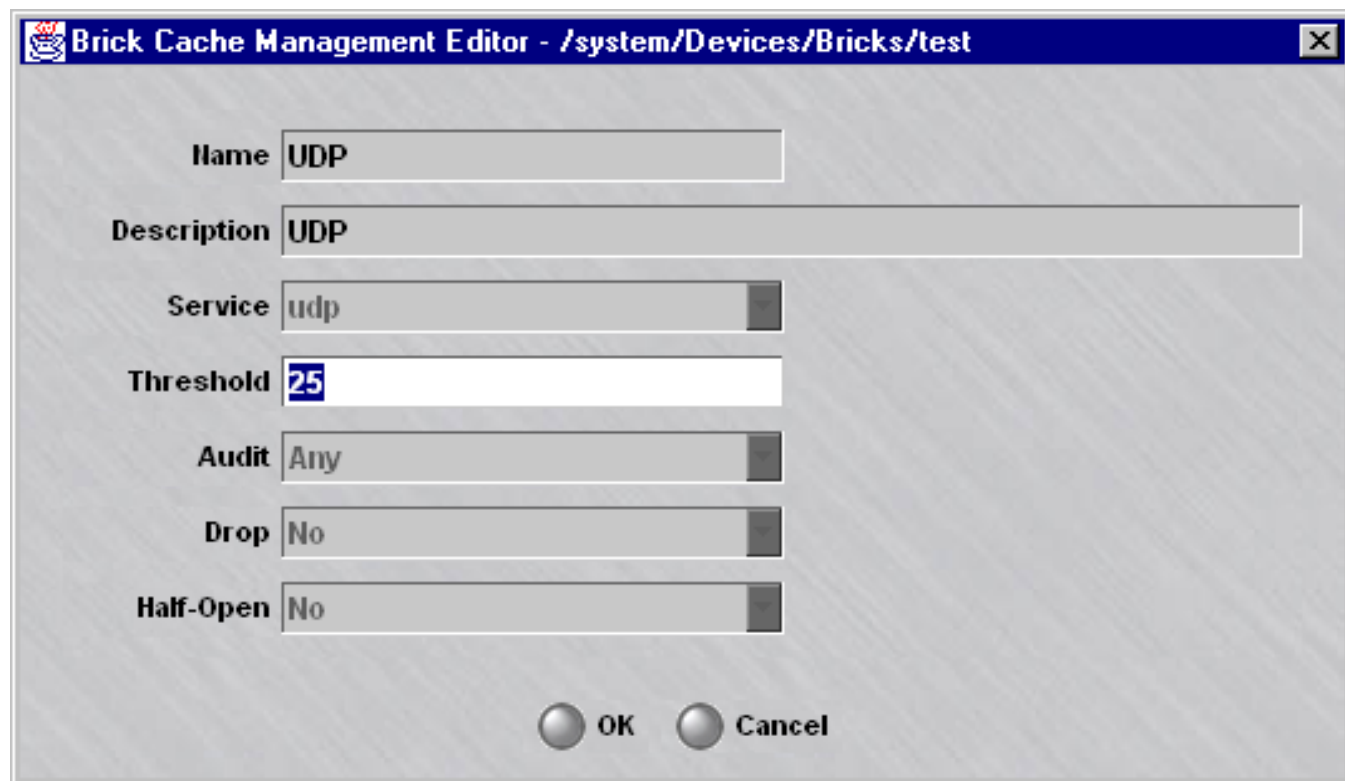
Figure 5-10 Brick Editor (Cache Management Tab)



- 4 Change the threshold for any of the five session classes, as necessary. To change a threshold, double click the session class to display the Brick Cache Management Editor (see Figure 5-11, “Brick Cache Management Editor” (p. 5-31)), and enter a value (percentage of cache capacity) in the **Threshold field**. Then, click **OK** to dismiss the Brick Cache Management Editor.

When the threshold is reached for a given session class, the intelligent cache management feature will begin clearing those types of sessions from cache.

Figure 5-11 Brick Cache Management Editor



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

END OF STEPS

Add an Entry to the Table

To add an entry to the Intelligent Cache Management Table, follow the steps below:

- 1 Right-click in the Cache Management tab of the Brick Editor (Figure 3-13) and select **New** from the pop-up menu. The Brick Cache Management Editor will appear. It is shown in [Figure 5-12, “Brick Cache Management Editor”](#) (p. 5-32).

Figure 5-12 Brick Cache Management Editor

The screenshot shows a dialog box titled "Brick Cache Management Editor - /system/Devices/Bricks/internet_brick". It contains the following fields and controls:

- Name:** A text input field.
- Description:** A text input field.
- Service:** A drop-down menu.
- Threshold:** A text input field.
- Audit:** A drop-down menu with "Any" selected.
- Drop:** A drop-down menu with "Any" selected.
- Half-Open:** A drop-down menu with "Any" selected.
- Buttons:** "OK" and "Cancel" buttons at the bottom.

- 2 In the **Name** and **Description** fields, enter a name and brief description of this entry. The description is optional.
- 3 In the **Service** field, select a service from the drop-down list.
- 4 In the **Threshold** field, enter the threshold.

-
- 5 In the **Audit**, **Drop**, and **Half-Open** fields, select the appropriate value from the drop-down list. The values are **Yes**, **No** and **Any**.
 - 6 Click **OK** to save the new entry and return to the Cache Management tab of the Brick Editor.
 - 7 Open the File menu and select one of the **Save** options to save the new entry.

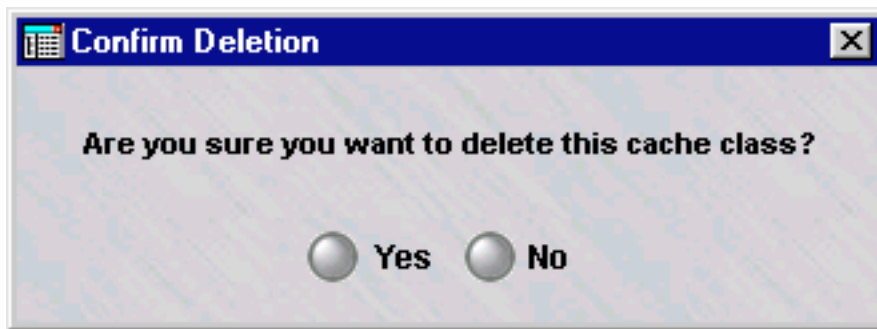
END OF STEPS

Delete an Entry

To delete an entry from the Intelligent Cache Management Table, follow the steps below:

-
- 1 In the Intelligent Cache Management Table, right-click the entry you want to delete and select **Delete** from the pop-up menu. The confirmation window shown in [Figure 5-13, “Confirmation Window”](#) (p. 5-33) will appear.

Figure 5-13 Confirmation Window



-
- 2 Click **Yes**. The entry will be deleted from the table.

END OF STEPS



6 Configuring VLANs on Lucent VPN Firewall *Brick*[®] Devices

Overview

Purpose

This chapter explains how to configure a Brick to recognize, forward and filter VLAN-tagged frames.

Contents

What is a VLAN?	6-2
Why Build VLANs?	6-4
Forwarding Packets and VLAN Boundaries	6-5
Configure and Activate the Brick	6-6
Configure the Brick Physical Ports for VLAN-Tagged Traffic	6-7
Assign a Policy to the Ports	6-13
Associate a Network with a VLAN	6-17
What are VLAN Bridge Groups?	6-20
Enable the Brick to Support VLAN Bridge Groups	6-21
Configuring Bridging Between Specific VLANs	6-22
Save and Apply the VLAN Configuration	6-23



What is a VLAN?


Definition

A VLAN is a collection of hosts on different physical segments of a switched network that communicate with each other as if they were on the same segment. VLANs allow network administrators to define multiple LANs on a single collection of switches.

One useful way to think of VLANs is that the combination of the VLAN and the physical port form a virtual port. From this point of view, a trunk port is simply a collection of many virtual ports

The diagram in [Figure 6-1, “Flat Switched Network \(no VLANs\)”](#) (p. 6-2) shows a typical flat, switched network with no VLANs implemented.


Figure 6-1 Flat Switched Network (no VLANs)



The diagram in [Figure 6-2, “Switched Network \(with two VLANs\)”](#) (p. 6-3) shows the same network after it has been subdivided into two VLANs. The VLANs allow


network administrators to organize network resources by department, function etc., rather than physical connection.

Figure 6-2 Switched Network (with two VLANs)



The diagram in [Figure 6-3, “Switched Network \(with VALNs and Brick on trunk\)”](#) (p. 6-3) shows a Brick situated on a VLAN trunk between two switches. The Brick is able to apply security policies based on VLAN-tagged traffic as it passes between the two switches.

Figure 6-3 Switched Network (with VALNs and Brick on trunk)



□

Why Build VLANs?

The purpose of VLANs

A flat network, as illustrated in [Figure 6-1, “Flat Switched Network \(no VLANs\)”](#) (p. 6-2), is characterized by a large, unsegmented IP space. This architecture differs from a traditional routed IP infrastructure, where networks and subnets define a tree-like hierarchical topology, segregated by routers.

In a VLAN environment, specific ports on each switch are dedicated to specific VLANs, with trunk ports configured to handle traffic between the switches. Traffic on the trunk ports is tagged with a unique VLAN identifier. The switch that initially receives the frame from the host applies the tag to the frame and puts the frame on the trunk, where it is forwarded to the port on another switch that is assigned to that VLAN.

VLANs allow a network administrator to design a logical network so that hosts on physically diverse segments can be given the appearance of sharing a single broadcast segment. Switches on separate floors of a building can be connected by 802.1Q tagged trunks, which allows broadcast traffic to be transmitted on the correct ports of the correct switches, and nowhere else. Contrast this with ordinary switches or hubs where broadcast traffic is repeated on all ports of all devices.

Service providers have started to use VLAN tagging on their internal backbone to enable them to identify packets inbound from a particular customer. This is enabled by routers that convert untagged inbound frames to 802.1Q tagged frames, which traverse the internal ISP backbone tagged by an identifier that can be traced back to a single port on a router. This differs significantly from using the IP address to track down the path of a packet, since IP relies on the sender to correctly fill in its own address. The VLAN tag is supplied by the router, and is based on the inbound port.

When service providers are hosting customers who use the same private addresses, the Brick partition feature can be used to keep the IP address space of each customer distinct. See [Chapter 7, “Configuring Lucent VPN Firewall Brick® Device Partitions”](#) for details.

□

Forwarding Packets and VLAN Boundaries

Overview

When packets enter the Brick, those packets are always associated with a VLAN. If the packets are tagged (using 802.1Q), they are associated with the VLAN ID contained within the VLAN tag, provided that the VLAN ID is configured for the physical port. Otherwise, the packets are associated with the default VLAN assigned to the physical port at which they arrived. Default VLAN assignment occurs even if the Brick window is set not to display VLAN information (the LSMS creates VLANs automatically for the user).

The Brick forwards packets according to Layer-2 information (MAC addresses) where possible, and using Layer-3 information (IP addresses and routes) otherwise. broadcast packets (such as ARP requests) are restricted to the VLAN or VLAN bridge group on which they entered the Brick.

Normally, packets can only traverse or change VLANs via Layer-3 forwarding. That is, they must be routed using IP information associated with the Brick, either in the Static Routes table, or by using local VLAN or Interface IP addresses. While this is reasonable to set up, it requires allowing the Brick to participate in routing a packet at Layer-3, which may necessitate a large static routing table.

It may sometimes be advantageous to allow the packet to enter the Brick on a given VLAN, and be moved to a different VLAN before forwarding, without static IP routing. VLAN bridge groups allow this to occur. This process further enhances the Brick functioning as a transparent bridge, since it allows the Brick to facilitate passing packets at Layer-2, even when moving packets across VLANs. (See [“What are VLAN Bridge Groups?”](#) (p. 6-20) on [“What are VLAN Bridge Groups?”](#) (p. 6-20) for a description of VLAN bridge groups).

□

Configure and Activate the Brick

When to use

To set up a Brick to process frames bearing VLAN tags, the first thing you have to do is configure and activate the Brick, just as you would any Brick.

If this has not already been done, refer to [Chapter 3, “Configuring and Activating a Lucent VPN Firewall Brick® Device”](#) for complete instructions on how to configure and activate a Brick, and ensure it is communicating with the LSMS.



Configure the Brick Physical Ports for VLAN-Tagged Traffic

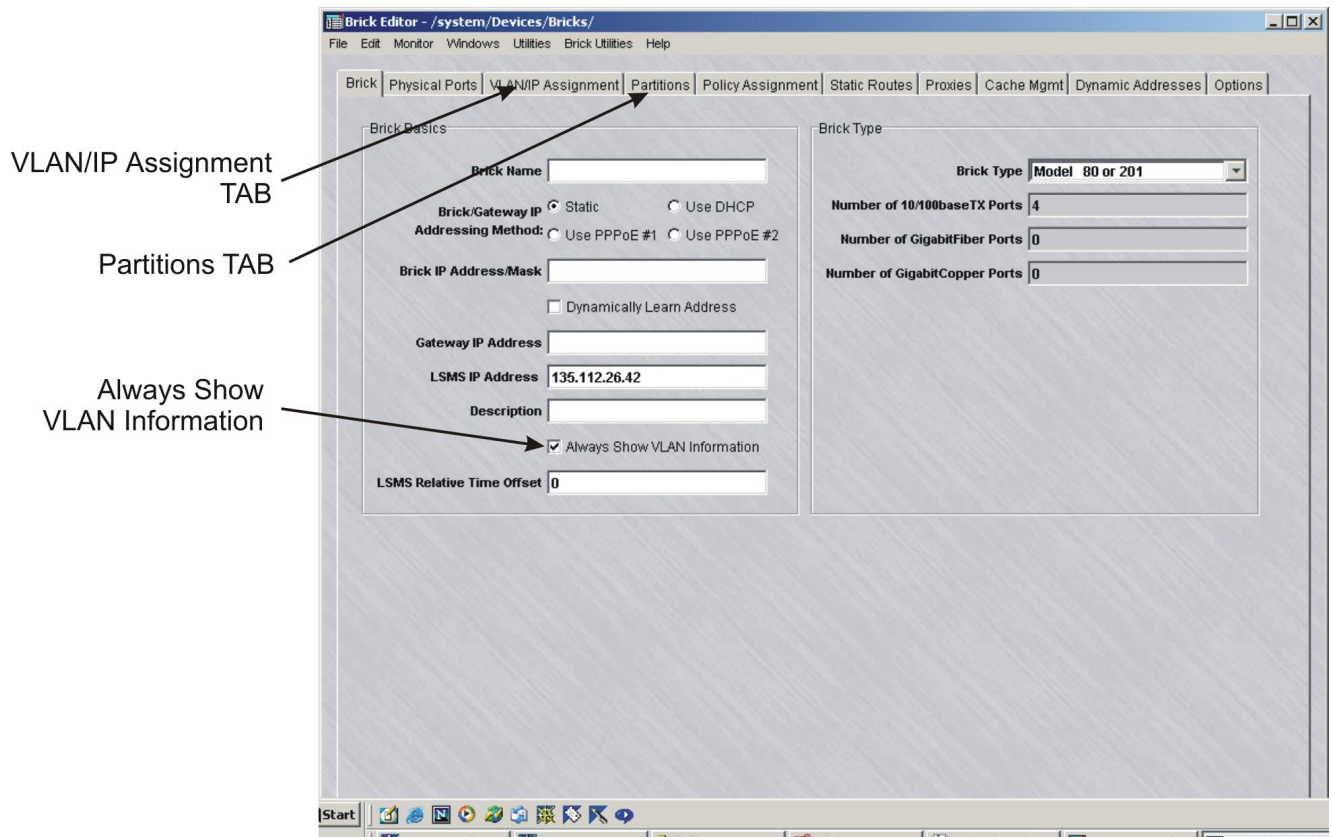
Task

To configure the ports on the Brick that will be accepting VLAN-tagged traffic, follow the steps below:

- 1 With the Navigator window displayed, open the appropriate group, Devices and Bricks folders, and double-click the Brick you want. The Brick Editor (Brick tab) will appear, with the configuration of the Brick you selected displayed.
- 2 Click the **Always Show VLAN Information** checkbox. Additional columns will now be added to the tables in the Physical Ports, Policy Assignment and Static Routes tabs, and two new tabs entitled **VLAN/IP Assignment** and **Partitions** will appear.

Figure 6-4, “Brick Editor (VLAN View)” (p. 6-7) shows the Brick tab of the Brick Editor after the checkbox has been selected.

Figure 6-4 Brick Editor (VLAN View)



Important! By default, the **Always Show VLAN Information** checkbox is unchecked. This is so those administrators who do not use VLANs will not see the VLAN fields and the VLAN/IP Assignment tab.

If you save the Brick configuration with the checkbox checked, VLAN information will be permanently displayed for that Brick. The checkbox will be grayed-out (as in [Figure 6-4, “Brick Editor \(VLAN View\)” \(p. 6-7\)](#)), and you will not be able to uncheck it.

As long as no **save** operation is performed while this box is checked, you can toggle back and forth between a VLAN and non-VLAN view. Therefore, if you do not want to permanently view VLAN information, close without saving.

The fact that you are viewing VLAN information does not mean you are committed to using VLAN's, only that you are committed to viewing the VLAN fields and VLAN/IP Assignment tab.

Once this box has been checked, you must delete and re-create the Brick to not display VLAN information.

-
- 3 Click **Physical Ports** to display the Physical Ports tab, and double-click a port that will be receiving VLAN traffic. The Brick Ports Editor will appear. It is shown in [Figure 6-5, “Brick Ports Editor”](#) (p. 6-9).

Figure 6-5 Brick Ports Editor

Brick Ports Editor - /system/Devices/Bricks/

Port: ether0

Aggregate with: [dropdown]

Port Type: 10/100baseTX

VLAN Domain: [text box]

Default VLAN ID: 1

VLAN Membership: [text box]

Receive Format: Untagged

Transmit Format: Untagged

Send/Receive DHCP request on this port

Enable Port Bandwidth Parameters

Transmit Bandwidth: 100 Megabits/sec

Receive Bandwidth: 100 Megabits/sec

Mode: Auto

MTU: [text box]

Ignore heartbeat failures on this link

OK Cancel

4 Enter the following information in the Brick Ports Editor.

Field	Description
VLAN Domain	<p>The purpose of the domain identifier is to distinguish the VLANs assigned to this port. It allows the Brick to treat identical VLANs from different trunks differently.</p> <p>Leave this field blank unless the Brick is positioned between two switches with different VLAN numbering schemes.</p> <p>If this is the case, enter a lower-case letter from a to n to identify the VLAN domain associated with this port.</p>
Default VLAN ID	<p>The default VLAN ID is the ID number that will be assigned to all untagged frames entering this port.</p> <p>The default is 1. You can change the default to any number from 0 — 4094.</p>
VLAN Membership	<p>Identifies all the incoming VLAN tags that will be permitted to access this port. The default VLAN ID for this port is automatically included in the VLAN membership, and should not be manually entered.</p> <p>You can leave this field blank, or you can enter any of the following:</p> <ul style="list-style-type: none"> • A single VLAN ID • Multiple VLAN IDs separated by commas (1,2,3) • A range of VLAN IDs using a dash, with individual ranges separated by commas (1-3, 7-10) • An asterisk (*), indicating all VLAN IDs are acceptable <p>If both the receive and transmit formats (see below) are untagged, you must leave this field blank.</p>
Receive Format	<p>The frame format to be allowed into this port. The choices are Untagged (default), 802.1Q, or Any.</p>

Field	Description
Transmit Format	<p>The frame format the Brick will send out of this port. The choices are Untagged (default), 802.1Q, 802.1Q Except Default, or Preserve.</p> <p>Only certain combinations of receive and transmit formats are supported:</p> <ul style="list-style-type: none"> • When the receive format is Untagged, the transmit format must be Untagged. • When the receive format is 802.1Q or Any, the transmit format can be any of the above <i>except</i> Untagged. <p>The normal setting for untagged links is Untagged, Untagged. The normal setting for tagged links is Any, Preserve.</p> <p>802.1Q Except Default specifies that frames on the port default VLAN will be sent untagged, while frames on any of the other VLAN members for the port will be sent with 802.1Q tags.</p> <p>Preserve means that when a frame is bridged (forwarded based on its MAC address), it will be sent as it was received (with or without a tag). If the frame is forwarded based on its IP address (i.e., is "routed"), this option behaves the same as 802.1Q Except Default.</p>
Send/Receive DHCP request on this port	<p>Check this box to allow the Brick's DHCP requests to go out this particular port and replies to come back in. By allowing the DHCP request to go out only the port on which the DHCP server is located, you can prevent possible DHCP server spoofing from the other ports. At least one port must have this checkbox checked if a DHCP address is used anywhere on the Brick.</p>
Transmit Bandwidth Receive Bandwidth	<p>Transmit Bandwidth and Receive Bandwidth are the "total" bandwidth in each direction.</p> <p>The value entered here restricts the maximum throughput that the Brick will transmit/accept on the interface. If the value is equal to or higher than the physical capacity of the link, then it serves only to bound the guarantees on the zones assigned to this physical port..</p>
Mode	<p>By default, a port will auto-sense the correct speed. However, you can specify the speed of the port and whether traffic should be configured in full duplex or half duplex mode on that port. Gigabit Fiber optic links are fixed at 1000 Gbps Full duplex, but autonegotiation may be disabled and flow control may be enabled or disabled with this option.</p>
MTU	<p>Maximum Transmission Unit is the largest size IP packet that the Brick will transmit on the interface. If left blank, it defaults to 1500 bytes.</p>

Field	Description
Ignore heartbeat failures on this link	Checking this box results in ignoring heartbeat failures between redundant Bricks on this link. It should be checked only if a known topology exists which prevents heartbeats from reaching the other Brick.

-
- 5 When you have finished entering the information above, click **OK** to dismiss the Brick Ports Editor and return to the Physical Ports tab of the Brick Editor. The information you entered will appear in the appropriate columns.

-
- 6 Repeat [Step 3](#) — [Step 5](#) for each additional port you need to configure.

END OF STEPS

.....



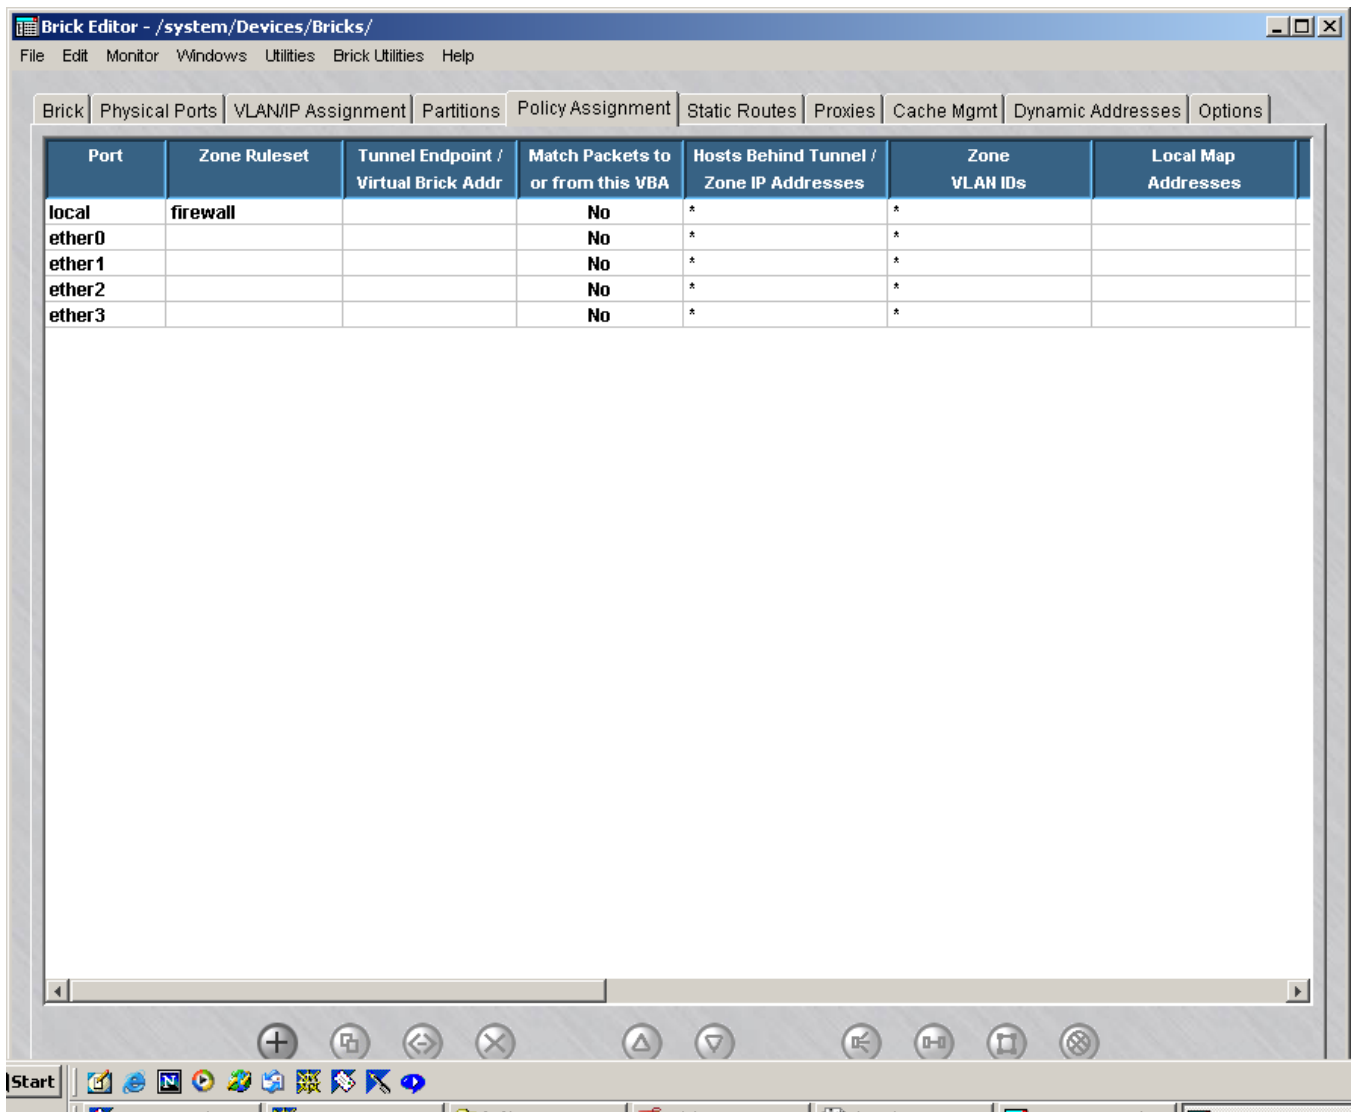
Assign a Policy to the Ports

Task

Once the physical ports have been configured to recognize VLAN-tagged frames, you need to assign a security policy to the VLANs on each port. To do this, follow the steps below:

- 1 With the Brick Editor displayed, click **Policy Assignment** to display the Policy Assignment tab. It is shown in [Figure 6-6, “Brick Editor \(Policy Assignment Tab\)”](#) (p. 6-13).

Figure 6-6 Brick Editor (Policy Assignment Tab)

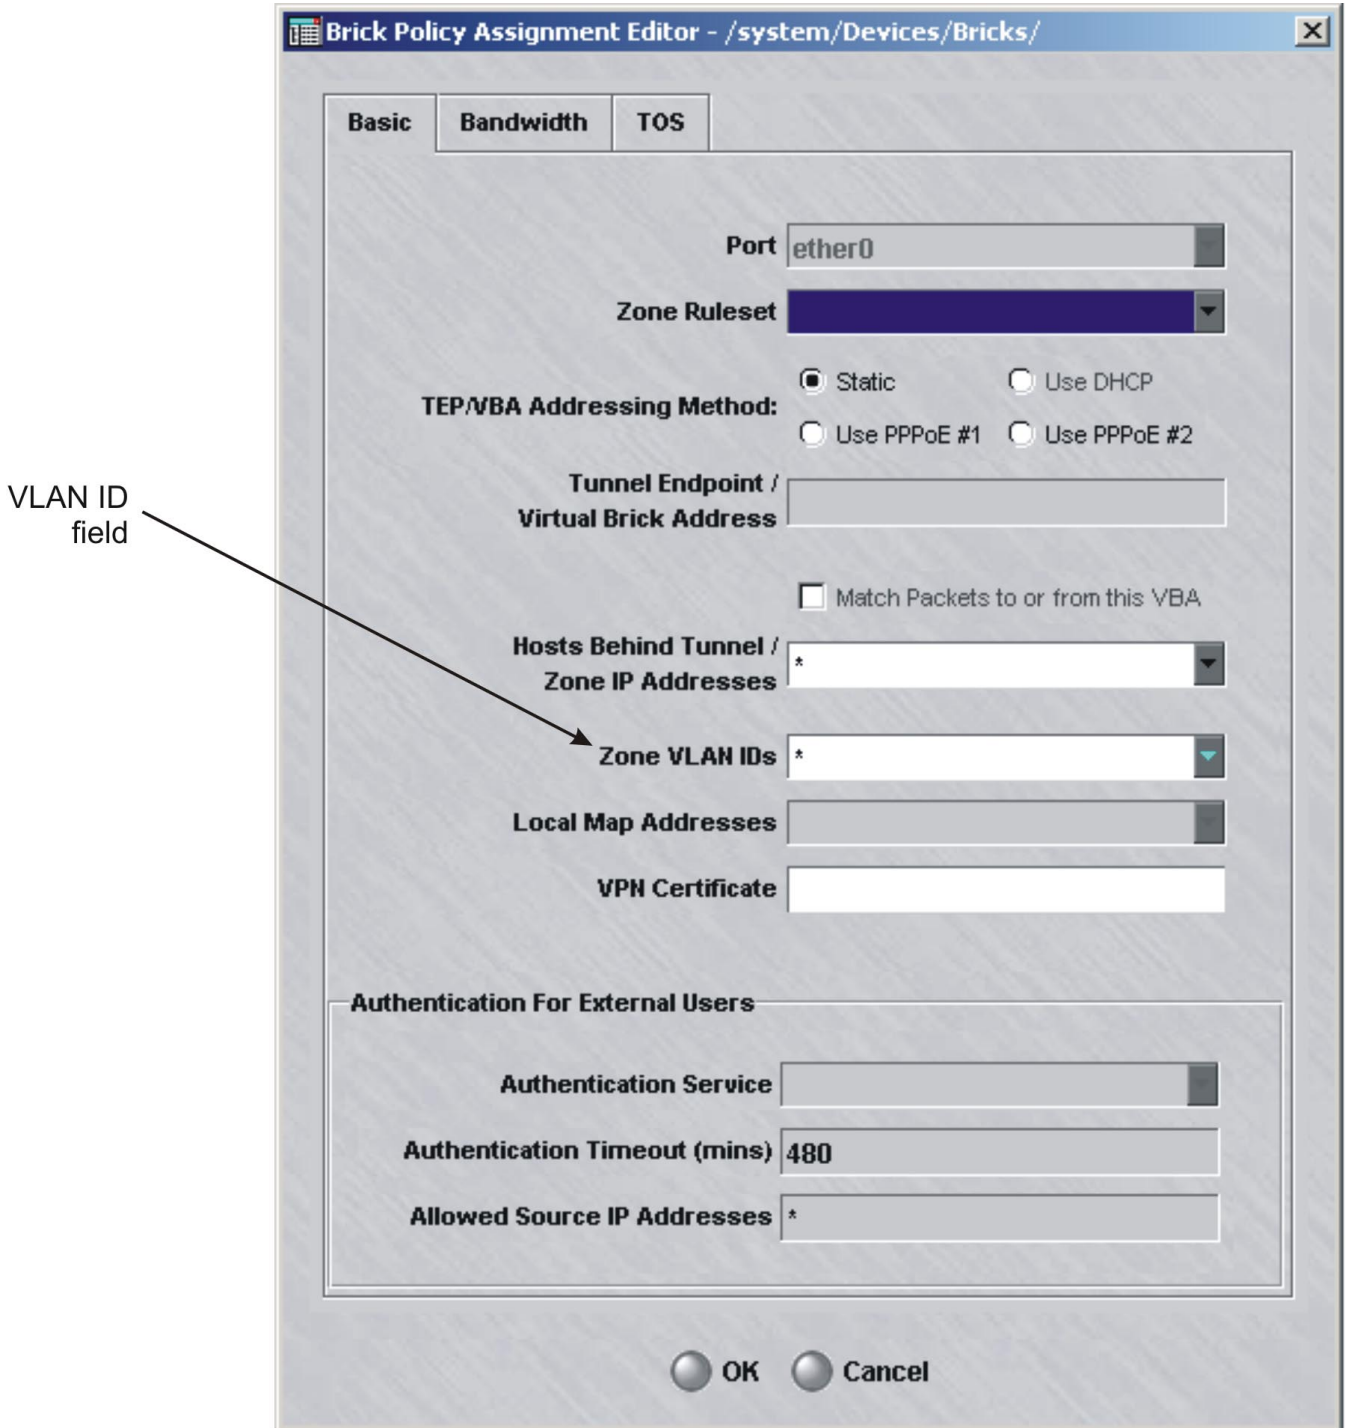


-
- 2 Double-click a port that will be receiving VLAN traffic. The Brick Policy Assignment Editor will appear.

This is the window you use to assign a security policy (i.e. a Brick zone ruleset) to a Brick port (see [“To Configure a Physical Port” \(p. 4-2\)](#) in Chapter 3, [“Configuring and Activating a Lucent VPN Firewall Brick® Device”](#)).

When the **Show VLAN View** checkbox is checked, this window has a VLAN ID field. The Editor and field are shown in [Figure 6-7, “Brick Policy Assignment Editor” \(p. 6-15\)](#).

Figure 6-7 Brick Policy Assignment Editor



-
- 3 Fill in all the fields as you ordinarily would when assigning a ruleset to a port, as explained in [Chapter 4, “Configuring Lucent VPN Firewall Brick® Device Ports”](#). In the **Zone VLAN ID** field, enter all the VLAN IDs on this port to which you want to assign this Brick zone ruleset. You can 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.

The IDs you enter must be IDs defined for this port. IDs defined for this port include the default VLAN ID and all IDs in the VLAN membership. You can choose the entire membership, or you can enter a subset of the membership. For example, if the membership is 1—99, you could enter 1—10 in this field.

You can also select **Port Default** and the “asterisk” from the drop-down list. Port default is the default VLAN ID for this physical port as defined on the Physical Ports tab of the Brick Editor (see [Chapter 4, “Configuring Lucent VPN Firewall Brick® Device Ports”](#), [Figure 5-2, “Apply Brick Window”](#) (p. 5-7)). The asterisk means the policy applies regardless of the VLAN ID.

Note that the drop-down arrow is blue, indicating that more than one entry can be selected for this field.

Important! Only one VLAN ID can be associated with a given Virtual Brick Address (VBA). Therefore, if there is a VBA defined for this port, make sure no more than one VLAN ID is selected in this field.

-
- 4 When you have finished entering the information above, click **OK** to dismiss the Brick Policy Assignment Editor and return to the VLAN Policy Assignment tab of the Brick Editor. The information you entered will appear in the appropriate columns.

END OF STEPS



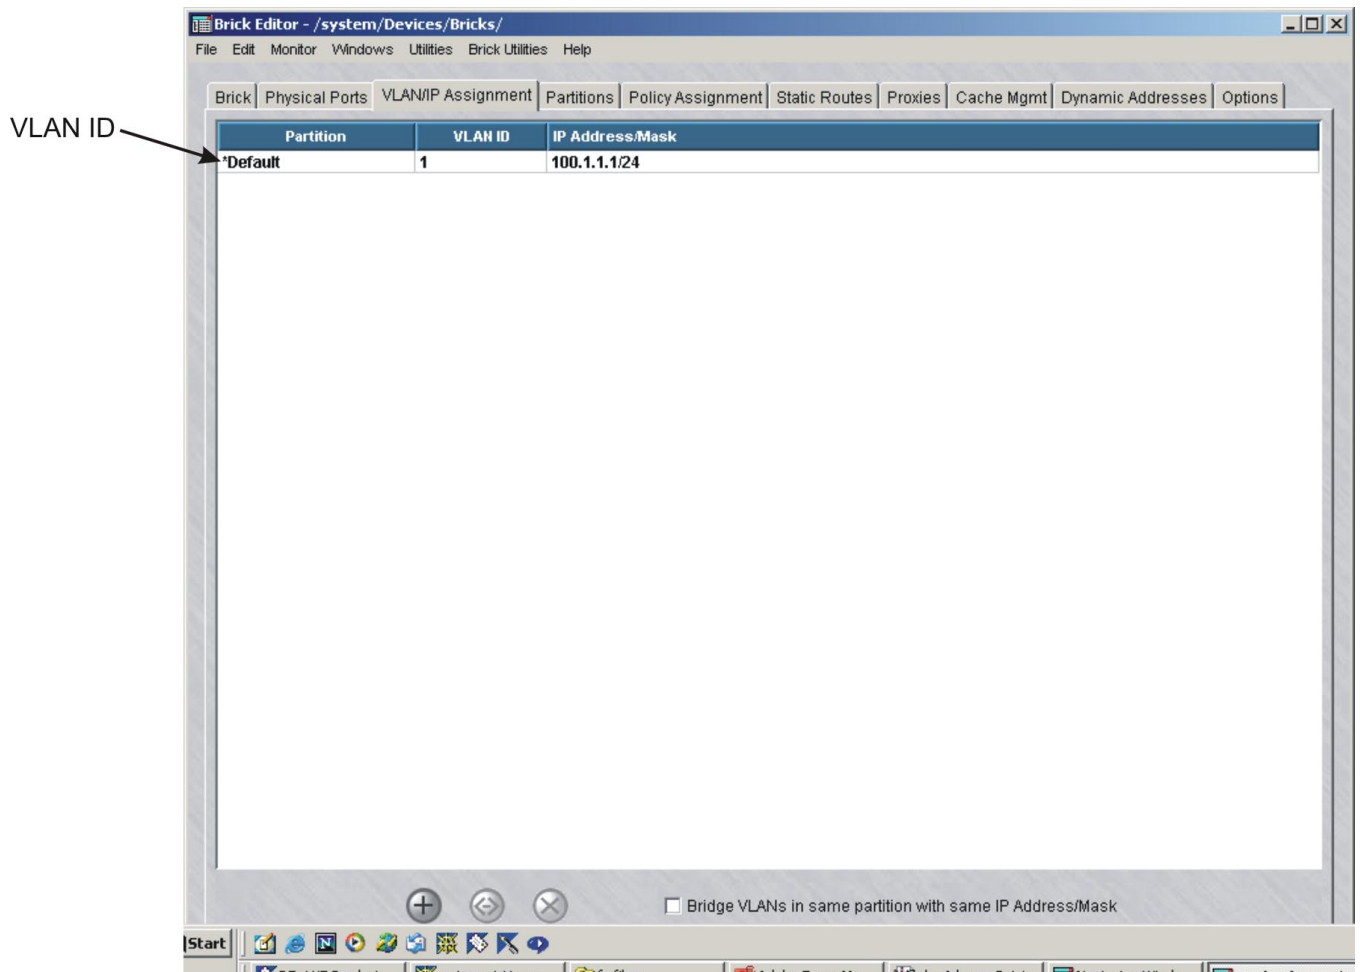
Associate a Network with a VLAN

Task

To assign a network (IP address/mask) to a VLAN (identified by its VLAN ID), follow the steps below:

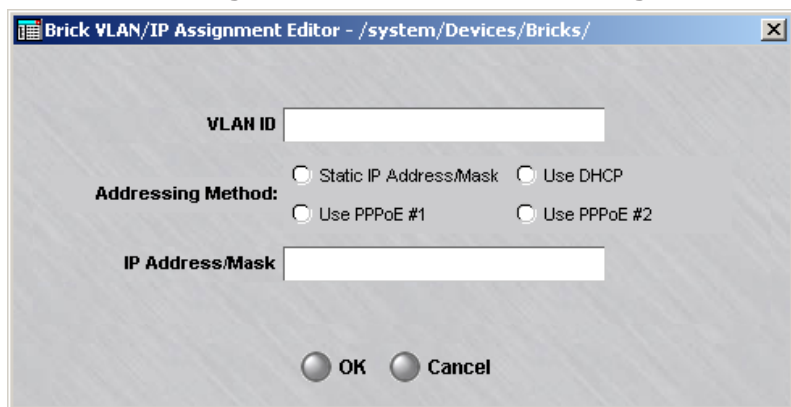
- 1 With the Brick Editor displayed, click **VLAN/IP Assignment** to display the VLAN/IP Assignment tab. It is shown in [Figure 6-8, “Brick Editor \(VLAN/IP Assignment Tab\)”](#) (p. 6-17). As this figure shows, the VLAN ID is automatically displayed and assigned the IP address of the Brick (including the appropriate subnet mask).

Figure 6-8 Brick Editor (VLAN/IP Assignment Tab)



-
- 2 To add a new VLAN ID and IP address/mask, right-click in the display area and select **New**. The Brick/VLAN IP Assignment Editor will appear. It is shown in [Figure 6-9](#), “Brick/VLAN IP Assignment Editor” (p. 6-18).

Figure 6-9 Brick/VLAN IP Assignment Editor



-
- 3 Enter the following information in the Brick/VLAN IP Assignment Editor. Multiple subnets can be assigned to the same VLAN by adding multiple entries with the same VLAN ID.

Field	Description
VLAN ID	Enter the domain and VLAN ID. Use the letter of the domain followed by a dot and the number of the VLAN ID (0 — 4094). The domain is optional.

Field	Description
IP Address/Mask	If the Address is statically assigned, then enter the IP address and mask to be associated with this VLAN ID. The mask is required.
Addressing Method	Check either the Use DHCP, Use PPPoE #1, or Use PPPoE #2 box (whichever is appropriate) if the Brick address on this VLAN is to be acquired dynamically rather than statically assigned.

Important! VLAN/IP assignment entries may now contain network addresses, in addition to host entries. This is useful if you want the Brick to be aware of the networks to which it is attached, but without using an IP address on the Brick itself. Network addresses are those that have a zero in the host portion of the address; for example: 10.1.1.0/24 is a network address, whereas 10.1.1.15/24 is a host address. Also, 192.168.15.128/28 is a network address, whereas 192.168.15.130/28 is a host address (since the host piece of the address with a 28-bit mask is the last 4 bits)

-
- 4 When you have finished entering the information above, click **OK** to dismiss the Brick/VLAN IP Assignment Editor and return to the VLAN IP Assignment tab of the Brick Editor. The information you entered will appear in the appropriate columns.

END OF STEPS



What are VLAN Bridge Groups?

Definition

A VLAN Bridge Group is a set of VLAN IDs that can mutually bridge packets. That is, Layer-2 forwarding can be used among all members of a VLAN Bridge Group to cause packets to forward from one VLAN ID to another, without requiring Layer-3 routing.

Packets are routed using their destination MAC address. The Brick looks in its MAC cache to determine if the destination MAC address is known, and to determine to which virtual port that MAC address is bound. If that virtual port happens to be a different VLAN ID than the one on which the packet already exists, the Brick will modify the VLAN ID information on that packet such that it will be associated with the correct destination VLAN ID.

All members of a VLAN Bridge Group must have one or more addresses and IP subnets in common, since the Bridge Group essentially functions as a single VLAN. Layer-2 Broadcast packets and multicast packets are also sent to all members of a VLAN Bridge Group.

VLAN Bridge Groups are mostly used when separating security domains using VLANs on external devices connected to the Brick via 802.1Q tagged VLAN trunks. Note this may involve as few as one switch, with the Brick acting as a "one-legged firewall", scaling up to as many switches as necessary. (It is necessary to configure such attached switches to prevent packets from crossing VLAN boundaries; consult your switch vendor to determine the correct configuration for your switch, if applicable.)

Once enabled, VLAN Bridge Groups are created implicitly in the VLAN/IP Assignment tab, by simply assigning the same IP address and mask to multiple VLAN IDs. There are NO explicit VLAN Bridge Group objects that appear in the LSMS Navigator.

VLAN Bridge groups never include a VLAN whose address is assigned via DHCP or PPPoE.



Enable the Brick to Support VLAN Bridge Groups

Task

To enable a Brick to support VLAN bridge groups as well as configuring bridge groups themselves, follow the steps below:

- 1 With the Navigator window displayed, open the appropriate group, Devices and Bricks folders, and double-click the Brick you want. The Brick Editor (Brick tab) will appear, with the configuration of the Brick you selected displayed.
- 2 Click **VLAN/IP Assignment to** display the VLAN/IP Assignment tab. Click the **Bridge VLANs in same partition with same IP/Mask** checkbox. This step needs only be performed once, regardless of the number of VLAN bridge groups set up on a given Brick. Once the checkbox is checked, you can skip this step.

END OF STEPS



Configuring Bridging Between Specific VLANs

Task

A VLAN bridge group is configured implicitly in the VLAN/IP Assignment screen, once the feature is enabled.

- 1 In the VLAN/IP Assignment tab, either choose an existing VLAN entry or create a new one by following the standard procedure.

- 2 To add a new VLAN bridge group, simply create a new VLAN/IP assignment with the same network or networks as specified in the chosen VLAN/IP assignment entry, and a different VLAN ID. IP address/mask pairs must match precisely. An error condition will result if an identical IP address is entered with a different subnet mask. If some, but not all, of the networks are identical, only packets destined to the identical networks will be part of the Bridge Group. Also note that all VLAN IDs in a VLAN Bridge Group must be in the same Brick Partition.

- 3 Repeat Step 2 for as many VLANs as you wish to be in the VLAN Bridge Group.

- 4 You may also create as many different VLAN Bridge Groups. Each VLAN Bridge Group is identified by the shared IP networks that span it.

END OF STEPS



Save and Apply the VLAN Configuration

Task

Once the ports have been configured for VLAN traffic, you have to save and apply the configuration. Open the File menu and select **Save and Apply**.



7 Configuring Lucent VPN Firewall *Brick*[®] Device Partitions

Overview

Purpose

Brick partitions allow true virtual firewalls to be implemented on the LVF Brick. Each virtual firewall has its own routing information, its own set of IP addresses, its own firewall policies, and so forth. In fact, each partition maintains its own session table, to ensure complete uniqueness and isolation for the virtual firewall.

For example, in a service provider environment, it is frequently necessary to ensure that packets *cannot* cross VLANs unless they are explicitly allowed. This circumstance may arise when a set of conflicting or "overlapping IP addresses" are used by two disjointed networks connected to two different physical or virtual interfaces. In this case, Brick Partitions ensure that sessions are distinct, regardless of IP addressing, as well as ensuring that packets CANNOT cross from one VLAN to another, unless an explicit route is created.

Brick partitions are created using the VLAN capabilities inherent to the LVF Brick. Each partition is created by choosing a set of VLANs that belong to it, and giving the Partition a name. (It is NOT necessary to use 802.1Q VLAN tagging to support Brick Partitions, since the Brick always relates packets to VLANs internally.)

The Brick partition feature requires that the Brick be displaying VLAN information. This is accomplished by checking the **Always Show VLAN Information** checkbox (see "[Configure a Brick on the LSMS](#)" (p. 3-8)" in [Chapter 3, "Configuring and Activating a Lucent VPN Firewall *Brick*[®] Device"](#)).

Contents

What are Brick Partitions?	7-3
Configure Brick Partitions	7-4
Use Static Routes with Partitions	7-6

Allow Partitions to Intercommunicate with Static Routes	7-7
Save and Apply the Brick Configuration	7-10
Interpreting IP Addresses When Brick Partitions Are Configured	7-11



What are Brick Partitions?

Definition

Brick partitions are used to create true virtual firewalls, with no potential of confusion or ambiguity in rulesets or session cache entries. Brick partitions are also designed to allow traffic to only traverse VLANs where explicitly allowed. Without Brick partitions, packets in a given VLAN are free to route to other VLANs where policy and configuration allows.

For example, it is possible to conceive of a network where a service-provider edge connects to two different customers. Both customers use the 10.0.0.0/8 network internally, though they have legal, registered addresses externally.

Brick partitions allows the Brick to treat each entity as completely separate at Layer-2, regardless of the fact that their IP address ranges overlap. With Brick partitions, traffic from one customer's host 10.1.2.3 is treated as completely distinct from traffic on the other customer's host 10.1.2.3, even if they are accessing the same destination server, on the same TCP ports.

Once a zone is assigned to a VLAN on one or more ports on a Brick, it is effectively assigned to that VLAN's partition as well. The same zone may not be assigned to more than one partition on a Brick. There is a restriction on Virtual Brick Addresses (VBAs). VBAs must be unique on the Brick. Therefore the same VBA cannot be used for different zones even if the zones are in different partitions.

Brick partitions must be used in conjunction with static routes to allow partitions to intercommunicate. Additionally, network address translation must be configured if the partitions contain overlapping IP addresses, as in the above example.

Brick partition objects are created per-Brick, and not shared across Bricks.




Configure Brick Partitions

When to use

Brick partitions are configured as named objects. Each set of Brick partitions is created for a given *Brick*®. For each Brick, you must configure the desired set of Brick partitions, along with appropriate static routing information. To do this, follow the steps below:

- 1 With the Navigator window displayed, open the appropriate group, Devices and Bricks folders, and double-click the Brick you want. The Brick Editor (Brick tab) will appear, with the configuration of the Brick you selected displayed.
- 2 Click the **Partitions** tab to display a list of Brick partitions currently configured for that Brick. [Figure 7-1, “Brick Editor \(Partition Tab\)”](#) (p. 7-4) shows this display.

Figure 7-1 Brick Editor (Partition Tab)

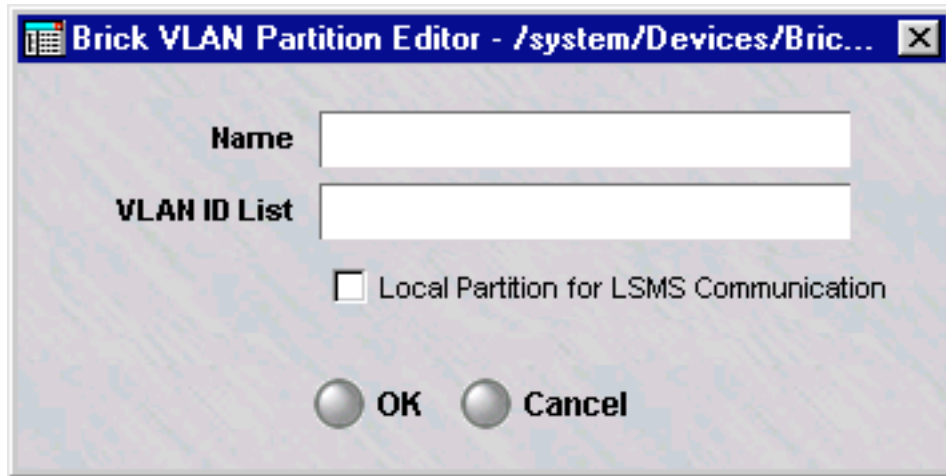


Note that in [Figure 7-1, “Brick Editor \(Partition Tab\)”](#) (p. 7-4), the only entry is “*Default.” If no partitions are configured, a single entry called “*Default” with a “*” assignment will appear in the Partitions tab. This indicates that all VLANs are in the *Default* partition; with this configuration, the Brick is “unpartitioned.”

The entry for the default partition is not editable. It is there as a reminder that any VLAN ID not mentioned explicitly in the Partitions table belongs to the default partition.

- 3 To add a Brick partition, right-click in the display area and select **New** from the pop-up menu. The Brick VLAN Partition Editor will appear. It is shown in [Figure 7-2, “Brick VLAN Partition Editor”](#) (p. 7-5).

Figure 7-2 Brick VLAN Partition Editor



- 4 In the **Name** field, enter a name for the partition. The name must be unique for each Brick.
- 5 In the **VLAN ID List** field, enter the VLAN IDs that will be included in the partition. A given VLAN ID may be assigned to only one partition. If a given VLAN ID is not explicitly assigned to a given Brick partition, it is implicitly a member of the *Default* partition.
- 6 Check the **Local Partition for LSMS Communication** checkbox if this partition is to be used for LSMS — Brick communication. Exactly one partition must be used for this purpose. If this checkbox is not checked on any partition, the *Default* partition is assumed.
- 7 Click the **OK** button to return to the Partitions tab of the Brick Editor ([Figure 7-1, “Brick Editor \(Partition Tab\)”](#) (p. 7-4)).

END OF STEPS



Use Static Routes with Partitions

Overview

Each static route has a partition indicator. This is the partition to which that route applies. In other words, any packet in a given partition can *only* use the static routes for *that* partition. When static routes are created, it is important to pay attention to the partition in which the route is created, since that route will only apply to packets within that partition.

Default routes are also applied per-partition. It is frequently useful to have a 0.0.0.0/0 route for each partition, to specify a default gateway so that partition can get to the outside if all else fails.

To create a static route, follow the procedure outlined in the section entitled “[Add a Static Route](#)” (p. 4-22) in Chapter 4, “[Configuring Lucent VPN Firewall Brick® Device Ports](#)”.



Allow Partitions to Intercommunicate with Static Routes

Overview

By default, traffic on VLANs in one partition cannot traverse to VLANs in any another partition. To allow traffic from one partition to another, a static route must be created.

Create Static Routes with a Partition as Gateway

Each static route you create can now have a partition as its gateway, as well as a VBA. Using a partition as the next hop causes the Brick to change the current partition of the packet and to continue searching for a static route using the entries for the new partition. Such a route may be helpful in directing a packet from another partition to an initial zone within this partition.

All communication between partitions should be bi-directional: If partition "A" has a static route with partition "B" as its next hop, then there must be one or more static routes with "A" as their next hop. If there is no return route, then packets from partition B are only able to return to partition A if the **Route Return Path Packets to Cached Source MAC Address** option is checked. Even in this case, return traffic may be halted if the cache for Partition B is cleared or if a MAC is moved. A warning message is displayed if you omit the return route.

Using a VBA as the next hop causes the packet to be processed by the zone associated with that VBA, then forwarded.

Figure 7-3 Brick Static Route Editor

Brick Static Route Editor - /system/Devices/Bricks/radbrick

Basic Routing

Route Active **Yes**

Partition ***Default**

Destination IP Address/Mask

Next Hop Gateway IP Address Partition

Gateway IP Address

Description

Route Cost **0**

Route Verification

Enable Route Verification

Ping Destination IP Address

Ping Source IP Address

Ping Interval (secs) **10**

Ping Timeout (secs) **1**

Ping Failures for Route Unavailable **3**

OK Cancel

To create a static route, follow the procedure described in the “[Add a Static Route](#)” (p. 4-22) section in [Chapter 3, “Configuring and Activating a Lucent VPN Firewall Brick® Device”](#).

Examples

In the simplest case of inter-partition communication, all the hosts that need to communicate have unique public addresses. In this case all that is needed is to add static routes in each partition pointing to the hosts in the other partition.

As a more complicated case, consider a service provider who hosts two customers, one in Partition1, the other in Partition2. Each customer uses addresses in the 10.0.0.0/8 subnet, but they wish to set up a private application between them. One way to do this is pick two other subnets, one for each partition — for example, 192.168.1.0/24 for Partition1 and 192.168.2.0/24 for Partition2.

The first address in each subnet will be arbitrarily reserved for a VBA. Zone1 in Partition1 will be assigned VBA 192.168.1.1 and Zone2 in Parttion2 will be assigned VBA 192.168.2.1. To allow clients in Partition1 to originate traffic to servers in Partition2, you must do the following:

1. Add outbound rules in Zone1 with destination addresses in the 192.168.2.0/24 subnet. These rules should source NAT to the VBA for Zone1.
2. In the static route table, add these two static routes:
Partition: Partition1 Destination: 192.168.2.0/24 Next hop: Partition2
Partition: Partition2 Destination: 192.168.2.0/24 Next hop: VBA 192.168.2.1
3. For return traffic, add these static routes:
Partition: Partition2 Destination: 192.168.1.0/24 Next hop: Partition1
Partition: Partition1 Destination: 192.168.1.0/24 Next hop: VBA 192.168.1.1
4. Add inbound rules in Zone2 that perform destination network address translation (NAT) to map the relevant 192.168.2.0/24 addresses to the appropriate addresses in the 10.0.0.0/8 space of Partition2.

At this point host 10.10.10.10 can have a dialogue with host 10.10.10.10. The 10.10.10 client in Zone1 thinks it is connecting to, say 192.168.2.134, while the 10.10.10 server in Zone2 thinks it is being contacted by 192.168.1.1.

If there are hosts in Partition2 that must originate traffic to hosts in Partition1, then perform the corresponding steps 1 and 4 for this direction. The static route table does not have to change.

□

Save and Apply the Brick Configuration

Task

Once Brick partitions and static routes are created, you have to save and apply the configuration. Open the File menu and select **Save and Apply**.



Interpreting IP Addresses When Brick Partitions Are Configured

Overview

Whenever you look at an IP address in a Brick configuration, you must bear in mind what partition is applicable. If you see an IP address that you recognize, you must be careful to ask yourself if it refers to the endpoint you are thinking of, or to a totally different endpoint in another partition.

For example, in the VLAN/IP assignment table, the same Brick address can be assigned to two different VLAN IDs. If both VLAN IDs belong to the same partition, then these entries form a bridge group (and the "Bridge VLANs in same partition with same IP/Mask" checkbox must be checked). If they are in different partitions, however, there is no relationship between the addresses.



8 Creating Groups and Administrators

Overview

Purpose

This chapter discusses the concept of a group, describes the **System** group (a special group provided with the LSMS), and explains how to create new groups.

This chapter also describes the two types of administrators, LSMS Administrators and Group Administrators, and explains how to create new administrator accounts. It describes the different privileges that can be given to Group Administrators.

Contents

What is a Group?	8-2
To Create a Group	8-5
To Maintain Groups	8-7
What are LSMS and Group Administrators	8-9
To Create Administrator Accounts	8-10
To Assign Groups and Privileges	8-17
To Maintain Administrator Accounts	8-21
To Use the LSMS Messenger	8-25



What is a Group?

Definition

A group is a collection of objects that are managed as a whole. In general, no object can exist in more than one group. However, some objects can be made globally visible to all the other groups. In addition, a LAN-LAN tunnel may have endpoints in two different groups.

If you are a Managed Service Provider, groups typically represent all devices, policies, tunnels, and users of a specific customer. (See *Chapter 2. Getting Started* for a more detailed discussion of group characteristics.)

Organization

The objects in a group are organized into folders and subfolders. The folders and subfolders that comprise each group are:

- Devices
 - Subfolders
 - Bricks
- Policies
 - Subfolders
 - Brick zone rulesets
 - Host groups
 - Service groups
 - Application filters
 - Dependency masks
- VPNs
 - Subfolders
 - LAN-LAN tunnels
 - Client tunnel endpoints
 - VPN Defaults
- User Authentication
 - Subfolders
 - Users
 - User groups
 - Authentication services

New Groups

LSMS Administrators have the choice of using the **System** group that is provided with the LSMS application, or creating additional groups.

The *System* group is a special group. It is the home of the NOC gateway Brick, and it is automatically populated with five pre-configured Brick zone rulesets, two router main rulesets and two router tunnel rulesets. In addition, three host groups are also included. (See *Appendix B. Pre-Configured Brick Zone Rulesets* in the *LSMS Policy Guide* for a detailed description of these rulesets).

Each new group you create will have the same folder/subfolder structure as the **System** group, but it will only contain four Brick zone rulesets, one router tunnel ruleset, and two host groups. The *nocgwzone* Brick zone ruleset, the *mgmt-tunnel* router tunnel ruleset, and the *Bricks* host group are not included, because they are only needed to set up the NOC gateway, which is always in the *System* group.

The names of the rulesets will be different in each new group you create. The LSMS will automatically add the @ symbol and the group's name after the name of the ruleset. So, for example, if you created a group called XYZ, you would see these rulesets: *administrativezone@xyz* and *main-encrypt@xyz*.

Administration

Groups can be administered by both LSMS Administrators and Group Administrators. LSMS Administrators have full privileges over all groups, which means they can access all folders in all groups and make any additions, modifications, or deletions they deem necessary.

Group Administrators, on the other hand, can only access the specific groups to which they are assigned. In addition, Group Administrators can be given three levels of privilege over the folders in their groups: None, View and Full.

This means you can create multiple Group Administrators for a group, each with different privileges, to administer different aspects of the group's operations. For example, one Group Administrator could have Full privileges over devices, but only View privileges over policy, while a second Group Administrator could have View privileges over devices and Full privileges over policy.

All valid LSMS and Group Administrators must have an administrator account in the LSMS. When creating new groups, you can create the Group Administrator account

first, and then create the group, or you can create the group first, and then create the administrator account. The following explains the difference:

- *Create Group First*

If you create a new group before any Group Administrator accounts have been created, you simply enter a name and optional description in the Group Editor, and the group is created. The complete procedure is given in "How to Create a Group" below.

You must then create a Group Administrator account, and assign the administrator to the group you just created. If you edit the group after creating and assigning the new administrator, you will see the administrator's account listed in the Group Editor.

- *Create Administrator First*

If you create a Group Administrator account before creating the new group, you will be asked to enter the administrator in the Group Editor when creating the group. You will also be able to indicate the administrator's privileges over the group.

If you then edit the administrator's account, you will see the new group listed in the Administrator Editor. If you later delete the administrator's account, it will automatically be removed from the Group Editor.



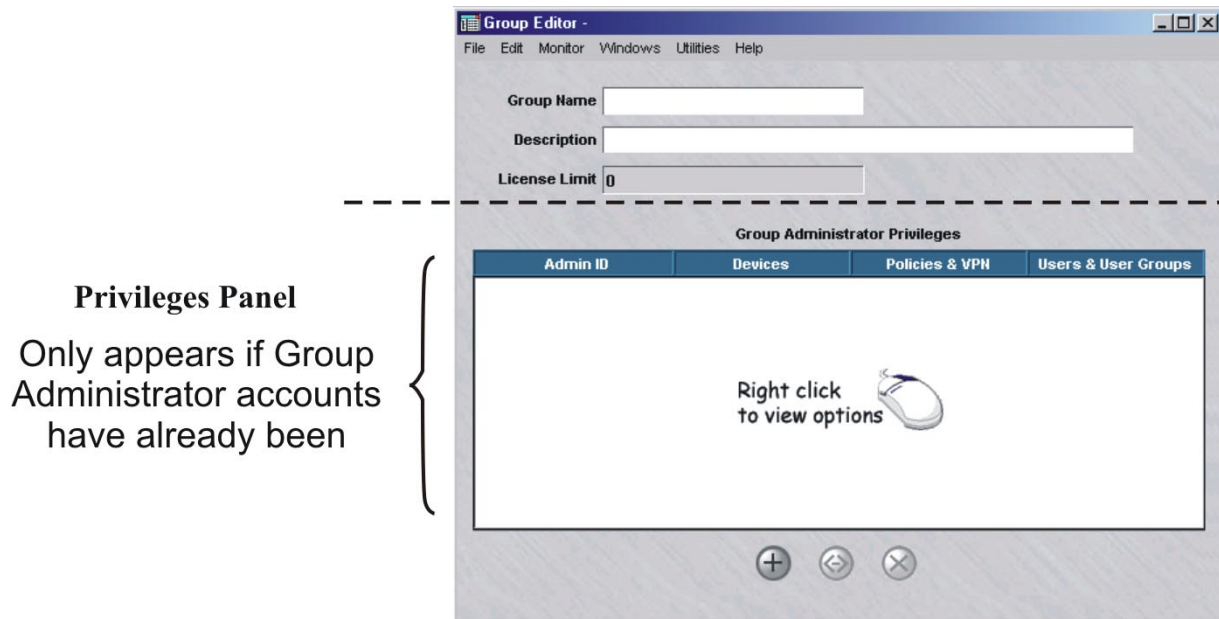
To Create a Group

Task

To create a new group, follow the steps below:

- 1 With the Navigator window displayed, right-click any group folder and select **New Group** from the pop-up menu. The Group Editor will appear. It is shown in Figure 7-1.

Figure 8-1 Group Editor



- 2 In the **Group Name** field, enter a unique name. The name can contain 1 to 40 characters.
- 3 In the **Description** field, you can enter an optional description of the group. The description can contain up to 80 characters.
- 4 The License Limit field is a read-only field that shows the number of licenses allocated to the Group. The value of this field can be changed by right-clicking on a Group folder and selecting **Allocate Licenses**.

-
- 5** At this point, you can display the File menu and select one of the **Save** options to save the new group.

If Group Administrator accounts have already been created and the Privileges panel is displayed, you can assign administrators to this group before performing the save operation. Right-click in the **Group Administrator Privileges** box and select **New** from the pop-up menu.

END OF STEPS



To Maintain Groups

When to use

LSMS administrators can edit and delete existing groups. You cannot change a group's name, but you can change its description, the administrators, and their privileges. For example:

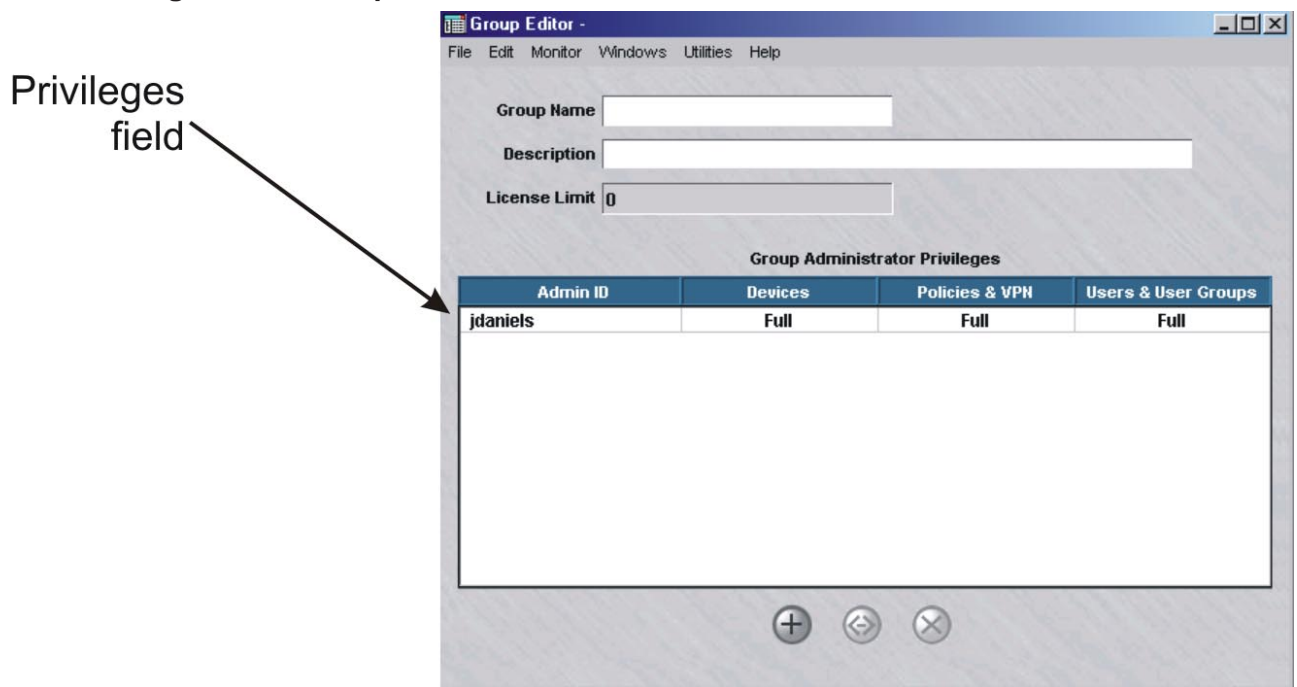
- You can add a new Group Administrator or delete an existing one.
- You can change the privileges of an existing administrator.

Edit a Group

To edit an existing group, follow the steps below:

- 1 With the Navigator window displayed, right-click the group folder and select **Edit** from the pop-up menu. The Group Editor window is displayed, as shown in Figure 8-2.

Figure 8-2 GroupEditor Window



- 2 The **Group Name** field is greyed-out and cannot be changed, but you can change the privileges.

-
- 3** You can assign a new Group Administrator, modify the privileges of an existing administrator, or delete an administrator.

To do this:

- To assign a new administrator, right-click in the Privileges panel and select **New**. Then, enter the administrator ID and privileges, and click **OK**.
- To modify an administrator's privileges, right-click the administrator in the Privileges panel and select **Edit**. Change the administrator's privileges and click **OK**.
- To delete an administrator, right-click the administrator in the Privileges panel and select **Delete**. Click **Yes** in the pop-up window to confirm the deletion. The administrator will no longer appear in the Privileges panel. However, the administrator's account is not deleted, only the administrator's ability to manage this group. The account can still be seen by clicking **Administrators** on the Navigator window.

END OF STEPS

Delete a Group

An LSMS administrator can only delete a group if the group is empty. This means you must remove all devices in the Devices folders, all administrator-created rulesets in the Policies folders, and all LAN - LAN tunnels before deleting the group. Related objects such as host groups, service groups, user authentication components, and the pre-configured rulesets will be automatically removed.

Once the group is ready, follow the steps below to delete it:

-
- 1** With the Navigator window displayed, right-click the group and select **Delete** from the pop-up menu.
-
- 2** Click **Yes** in the pop-up box to confirm the deletion. The group will be removed from the Navigator window.

END OF STEPS



What are LSMS and Group Administrators

Overview

The LSMS supports two types of administrators — LSMS Administrators and Group Administrators. There can be multiple LSMS Administrators and Group Administrators. Every administrator must have a valid administrator account in the LSMS.

Only LSMS Administrators can create other LSMS Administrators and Group Administrators, and assign privileges to Group Administrators.

Administrator logins may be authenticated in one of three ways:

- Local Password
- RADIUS Server
- RSA SecurID Server

LSMS Administrators

An initial LSMS administrator account is created during installation of the LSMS application. This account can then be used to log onto the LSMS and begin setting things up.

A LSMS administrator has full privileges over *all* folders in all groups. This set of privileges entitles them to create, edit, delete, and apply all devices, rulesets and tunnels, as well as set up and maintain user authentication.

Group Administrators

Group Administrators are created by LSMS administrators and assigned to one or more groups. The functions that Group Administrators can perform in their groups are determined by the privileges they have been assigned by the LSMS administrator.

Group Administrator functions are broken down into three areas:

- Devices (Bricks)
- Policies (rulesets, tunnels, authentication services)
- Users and user groups

For each area, Group Administrators can be assigned one of three privilege categories:

- Full (Create, Edit, Delete, and Apply).
- View (Display but not Create, Edit, or Delete)
- None (the folders will not be visible)

In addition, Group Administrators can view their assigned privileges and change their name, password, e-mail, phone number, and pager information.



To Create Administrator Accounts

Task

To create a new LSMS or Group Administrator, you have to create an account in the LSMS. If the account is for a Group Administrator, you also have to assign groups and privileges to the account for it to take effect.

To create an administrator account, follow the steps below:

-
- 1 With the Navigator window displayed, right-click the **Administrators** folder and select **New Administrator** from the pop-up menu. The Administrator Editor will appear with the Administrator tab highlighted, as shown in [Figure 8-3, “Administration Editor \(Administrator Tab\)”](#) (p. 8-11).

- 2 By default, the **Group Administrator** radio button is clicked in the **Role** box. This means this account is for a Group Administrator. If you intend this administrator to be an LSMS Administrator, click the LSMS **Administrator** radio button.

Figure 8-3 Administration Editor (Administrator Tab)

Administrator Editor -

File Edit Monitor Windows Utilities Help

Administrator | Authentication | Group Privileges

Administrator Information

Enable Administrator

Role LSMS Administrator Group Administrator

Admin ID

Full Name

Description

Email

Telephone

Pager Info

- 3 Enter the information requested in the fields. **Admin ID** and **Full Name** are required fields. The **Enable Administrator** field is checked by default. An LSMS Administrator can modify this checkbox for any user but himself / herself, while a Group Administrator cannot modify this checkbox.

The table below describes each field under the Administrator tab:

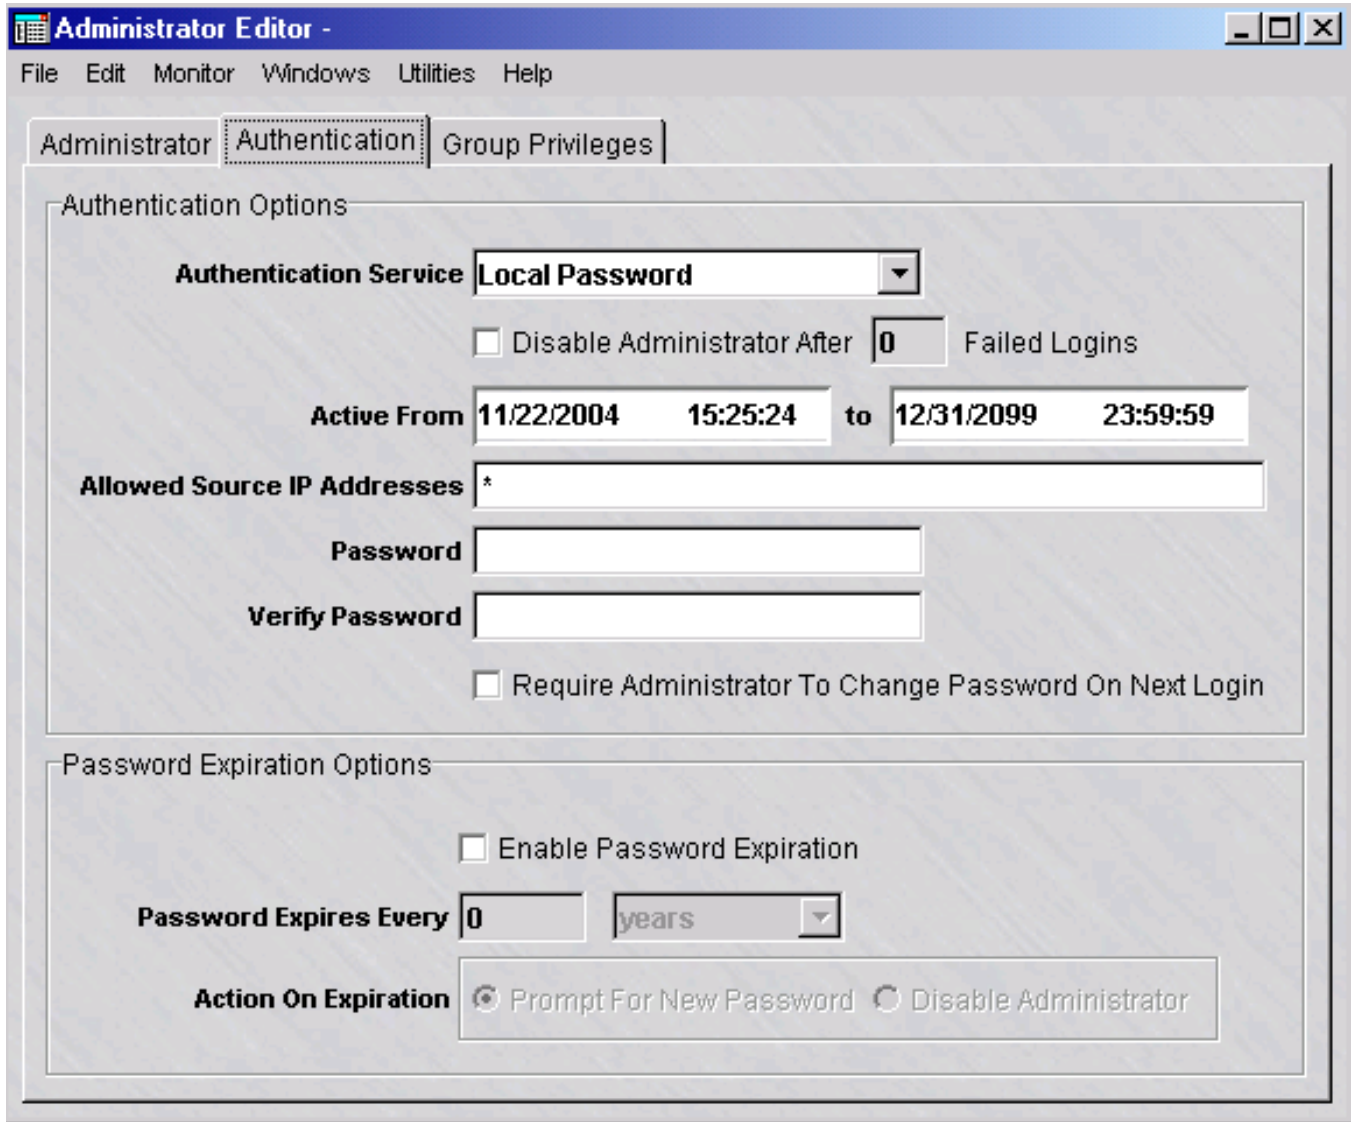
Field	Description
Admin ID	The administrator login. Can contain 1 to 16 characters (letters and numbers) and must be lowercase. This is a required field.
Full Name	The administrator's name. Can contain 1 to 80 upper and lower case characters (letters and numbers). This is a required field, but it does not appear on-screen when you display administrators in the Navigator window.
Description	A description of the administrator. Can contain 1 to 80 upper and lower case characters (letters and numbers). It is displayed in the Navigator window, so you can enter the administrator's name, role (LSMS/Group), or any other information that might be helpful to see in the Navigator window.
E-mail Address	The administrator's e-mail address. The contents of this field is used in the e-mail alarm action. See the <i>Reports, Alarms, and Logs Guide</i> .
Telephone Number	The administrator's office telephone number.
Pager Info	The PIN of the administrator's paging service (e.g., SkyTel, PageNet, MetroCall). The contents of this field is used in the Direct Page alarm action. See the <i>Reports, Alarms, and Logs Guide</i> .

4 Next, click on the Authentication tab.

- **Authentication Service** - Default choice is "Local Password". If you have already configured a RADIUS or SecurID server to provide authentication for your firewall or VPN users *in the "system" group*, it will be also listed here as an option. You may initially create a user with a local password and update it later to use a different authentication method.

The choice of Authentication Service affects the options available under the Authentication tab. If you chose RADIUS or RSA SecurID, proceed to Step 5. If you selected Local Password, you will see the following screen:

Figure 8-4 Administrator Editor (Authentication Tab)



The table below describes each field listed for Local Password:

- **Disable Administrator After** - You have the option to disable an administrator after "n" number of failed logins.
- **Active From** - By default, a user is active from the time of creation until 12/31/2099. You may redefine that period as desired. This only applies to Group Administrators.

- **Allowed Source IP Addresses** - By default, this field is set to "*" (allow any source IP). You may narrow the scope of allowed addresses as desired. If you are using the LSMS Remote Navigator, you can only connect to the LSMS from addresses listed in this field.
- **Password** - The password required to validate the login. The password is case-sensitive. The minimum password length is six characters (or the minimum password length setting for the **Local Password** Authentication Service) and the maximum length is 42 characters long. This is a required field if you are using a local password.

When a new password is set for an administrator that is authenticated using Local Password authentication, or an existing local password is changed, if the strong password (SOX compliance) option is enabled (the default) via the Configuration Assistant, stricter password requirements would apply. In this case, the password:

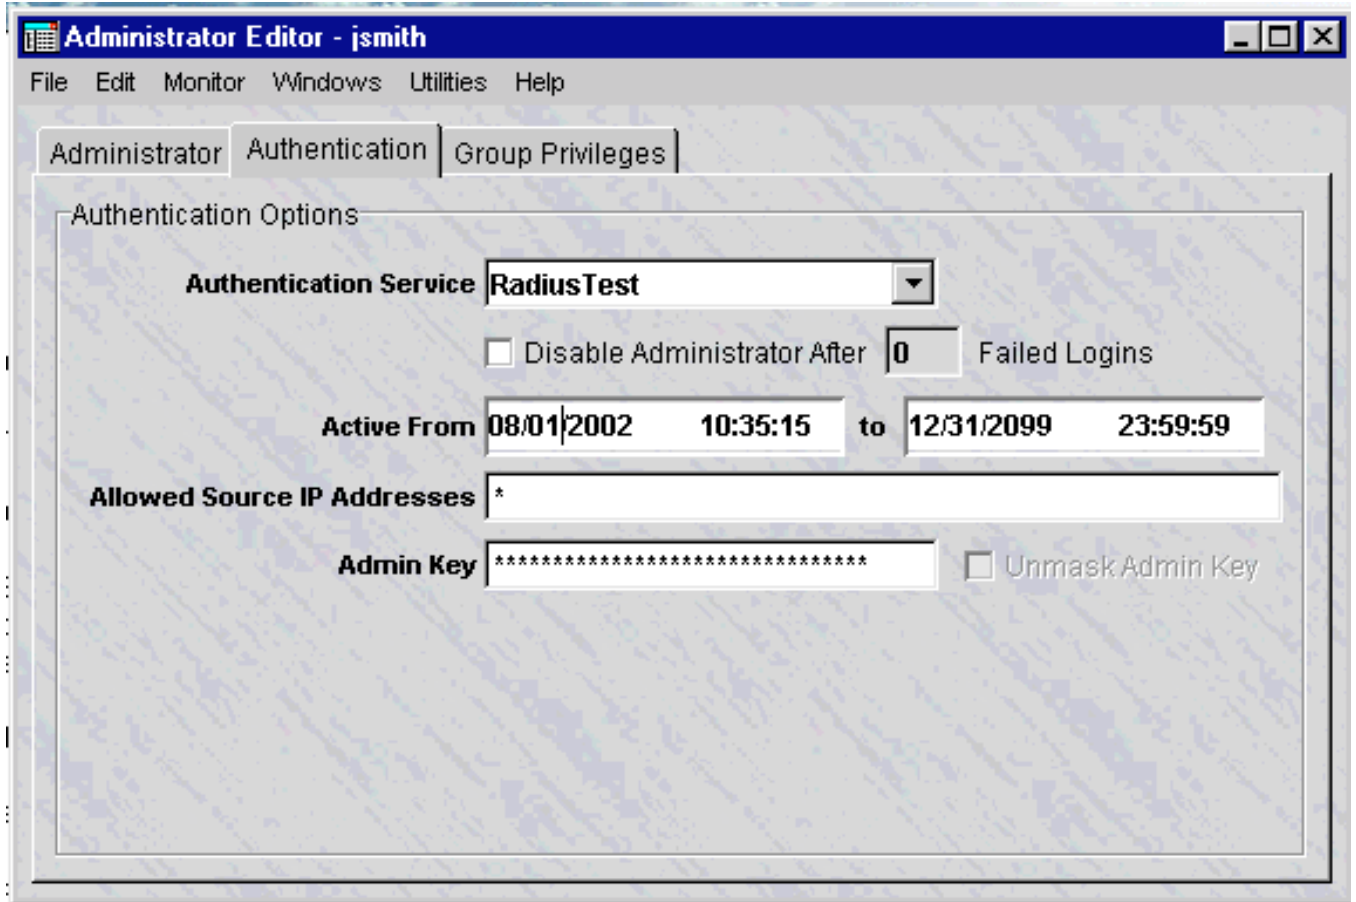
- Must be a minimum of eight characters, or the **Minimum Password Length** set for the **Local Password** Authentication Service, whichever is greater
- Must contain at least one alpha character and one non-alpha character (0-9, special characters, no restrictions)
- Cannot contain three or more repeated 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 details about enabling or disabling the strong password (SOX compliance) option, refer to the description of the User Authentication Parameters settings of the Configuration Assistant in [Chapter 11, "Using the Configuration Assistant"](#).

- **Verify Password** - The password, entered exactly as above. If using capital letters, ensure that you are capitalizing consistently. Also a required field for local password admins.
- **Require Administrator to Change Password on Next Login** - Check off if desired.
- **Password Expiration Options** - If "Enable Password Expiration" is checked, you may define the length of time before a password expires, and also choose the **Action on Expiration** - whether to prompt the user for a new password or to disable the administrator.

- 5 Under the Authentication tab, if you selected an Authentication Service of "RADIUS" or SecurID", you will see the following screen:

Figure 8-5 Administration Editor (Authentication Service)



The list below describes each field shown:

- **Disable Administrator After** — You have the option to disable an administrator after "n" number of failed logins.
- **Active From** — By default, a user is active from the time of creation until 12/31/2099. You may redefine that period as desired. This only applies to Group Administrators.
- **Allowed Source IP Addresses** — By default, this field is set to "*" (allow any source IP). You may narrow the scope of allowed addresses as desired. If you are using the LSMS Remote Navigator, you can only connect to the LSMS from addresses listed in this field.

- **Admin Key** — During the login process, this key must be provided, either automatically via a RADIUS attribute, or entered manually by the user, to authenticate the administrator. See "Administrator Authentication with RADIUS and SecurID" on page 8-14.
 - **Unmask Admin Key** — While the Admin Key is initially created, if this box is checked, the admin can see the key as it is typed. If left unchecked, each keystroke is masked by a "*". However, administrators cannot see the key of other administrators once it is created.
-

- 6** If you are creating an LSMS Administrator account, display the File menu and select one of the **Save** options.

If this is a Group Administrator account, you have to assign this administrator to at least one group and specify the administrator's privileges before you save the account. The procedure is described in the next section.

END OF STEPS



To Assign Groups and Privileges

When to use

Group Administrators have to be assigned to the groups they will manage, and they have to be given privileges.

Privileges

Privileges determine the extent of a Group Administrator's control over the group and its components. Group administration is broken down into three functional areas:

- Devices (Bricks)
- Policies & VPN (rulesets, tunnels, authentication services)
- Users and user groups

For each functional area, a Group Administrator can be assigned one of three privilege categories:

- Full
- View
- None

The table below shows the result of assigning each privilege to the three functional areas.

Area	Privilege	Result
Devices	None	Does not see Devices folder
	View	Sees all configured Bricks. Can make no additions, deletions, or changes.
	Full	Can configure Bricks. Can edit and delete configured devices. Can apply devices.
Policies & VPN	None	Does not see Policies folder, VPN folder, and Authentication Services folder
	View	Sees all configured rulesets, host groups, and service groups. Sees all configured tunnels and authentication services. Can make no additions, deletions, or changes.
	Full	Can create, edit, and delete Brick rulesets, host groups, and service groups. Can configure LAN-LAN tunnels and client tunnel endpoints. Can apply policy, LAN-LAN tunnels, and client tunnel endpoints.

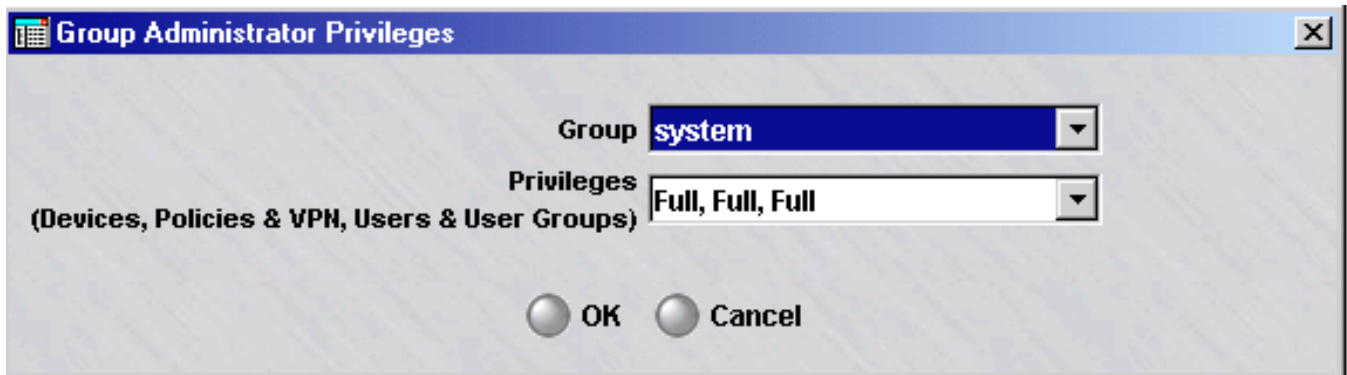
Area	Privilege	Result
Users & User Groups	None	Does not see Users and User Groups folders.
	View	Sees all user accounts and user groups that have been created. Can make no additions, deletions, or changes.
	Full	Can create, edit, and delete user accounts and user groups.

Assign Groups and Privileges

To assign groups and privileges to a Group Administrator account, follow the steps below:

- 1 Proceed to the Group Privileges tab in the Administrator Editor. Right click in the white area and select **New** from the pop-up menu. The Group Administrator Privileges window will appear. It is shown in Figure 7-6.

Figure 8-6 Group Administrator Privileges Window



- 2 In the **Group** field, select a group from the drop-down list.
- 3 In the **Privileges** field, enter the privileges this administrator will have over the group you selected. You must select one of the privilege combinations on the drop-down list.

The table below explains each combination:

Privilege Combination	Description
Full, Full, Full	Can manage all aspects of the group's operations.

Privilege Combination	Description
Full, Full, View	Can manage devices and policies, including all tunnels and user authentication services. Can only see user accounts and user groups.
Full, View, View	Can manage devices, but can only see policies, user accounts, and groups. Can only apply devices.
View, Full, Full	Can view all configured devices. Can create policies, user accounts, and user groups. Can apply policy, LAN-LAN tunnels, and client tunnel endpoints.
View, Full, View	Can view all configured devices, and all user accounts and user groups. Can manage policies, including LAN-LAN tunnels, client tunnel endpoints, and authentication services. Can apply policy, LAN-LAN tunnels, and client tunnel endpoints.
View, View, Full	Can view all configured devices and all policies. Can manage user accounts and user groups. Cannot apply.
View, View, View	Can view all aspects of a group's operations. Can make no changes. Cannot apply.
None, Full, Full	Can manage policies, user accounts, and user groups. Does not see Devices folder in Navigator window. Can apply policy, LAN-LAN tunnels, and client tunnel endpoints.
None, Full, View	Can manage policies. Can view user accounts and user groups. Does not see Devices folder in Navigator window. Can apply policy, LAN-LAN tunnels, and client tunnel endpoints.
None, View, Full	Can manage user accounts and user groups. Can view policies. Does not see Devices folder in Navigator window. Cannot apply.
None, View, None	Can view the Policies and VPN folders. Cannot view the Devices or User Auth folders in Navigator window. Cannot apply.
None, None, Full	Can manage user accounts and user groups. Does not see Devices or Policies folders. Cannot apply.

.....

4 When you are finished, click **OK** to dismiss the Group Administrator Privileges window and return you to the Administrator Editor. The new administrator's privileges will appear in the Administrator Editor.

.....

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

.....

END OF STEPS

.....

Administrator Authentication with RADIUS and SecurID

If LSMS or Group Administrators using Local Password authentication enter the correct password as stored in the LSMS database, they are able to login successfully.

If an administrator is authenticating via an external RADIUS server, there are several scenarios to consider:

- If the RADIUS authentication service in the LSMS is configured with the Admin Key attribute, the users in RADIUS must be configured to return the attribute, otherwise the login will fail. The value of the Admin Key that is returned from RADIUS must match the value that is entered for that administrator in the Administrator Editor, in order for the login to succeed.
- If the authentication service in the LSMS is NOT configured for Admin Key, then the administrator will be prompted to enter the Admin Key during the login process.

If an administrator is authenticating through a SecurID server, this is similar to RADIUS authentication where the LSMS parameter "Admin Key" has NOT been configured. After the SecurID authentication is complete, the user will be prompted to enter their Admin Key. If the SecurID authentication was successful and the Admin Key is correct, the user is authenticated and the LSMS or Group Administrator login is successful.

Administrators and LSMS Installations and Upgrades

Although LSMS and Group Administrators have the option to authenticate through external RADIUS or SecurID servers, it is imperative to ensure that at least one LSMS Administrator could log into the LSMS Navigator *if all external servers are unavailable*.

The initial LSMS administrator created during the installation of the LSMS application will be assigned the Local Password authentication service. To ensure that there is always an administrator account available to install software upgrades, the initial LSMS administrator account created during installation cannot be deleted although the password for this account can be set up to expire after a set period of time. This account is assigned the Local Password authentication service, and the authentication service cannot be changed.

When upgrading the LSMS application, an LSMS administrator ID and password must be provided. The login ID and password of any administrator that is assigned to the Local Password authentication service may be used to install upgrade software. Even if a Local Password administrator is disabled via the **Disable Administrator after x Failed Logins** checkbox on the Administrator Editor, and they can no longer log in to the GUI to administer the LSMS, that login ID and password may still be used to upgrade the LSMS application.



To Maintain Administrator Accounts

When to use

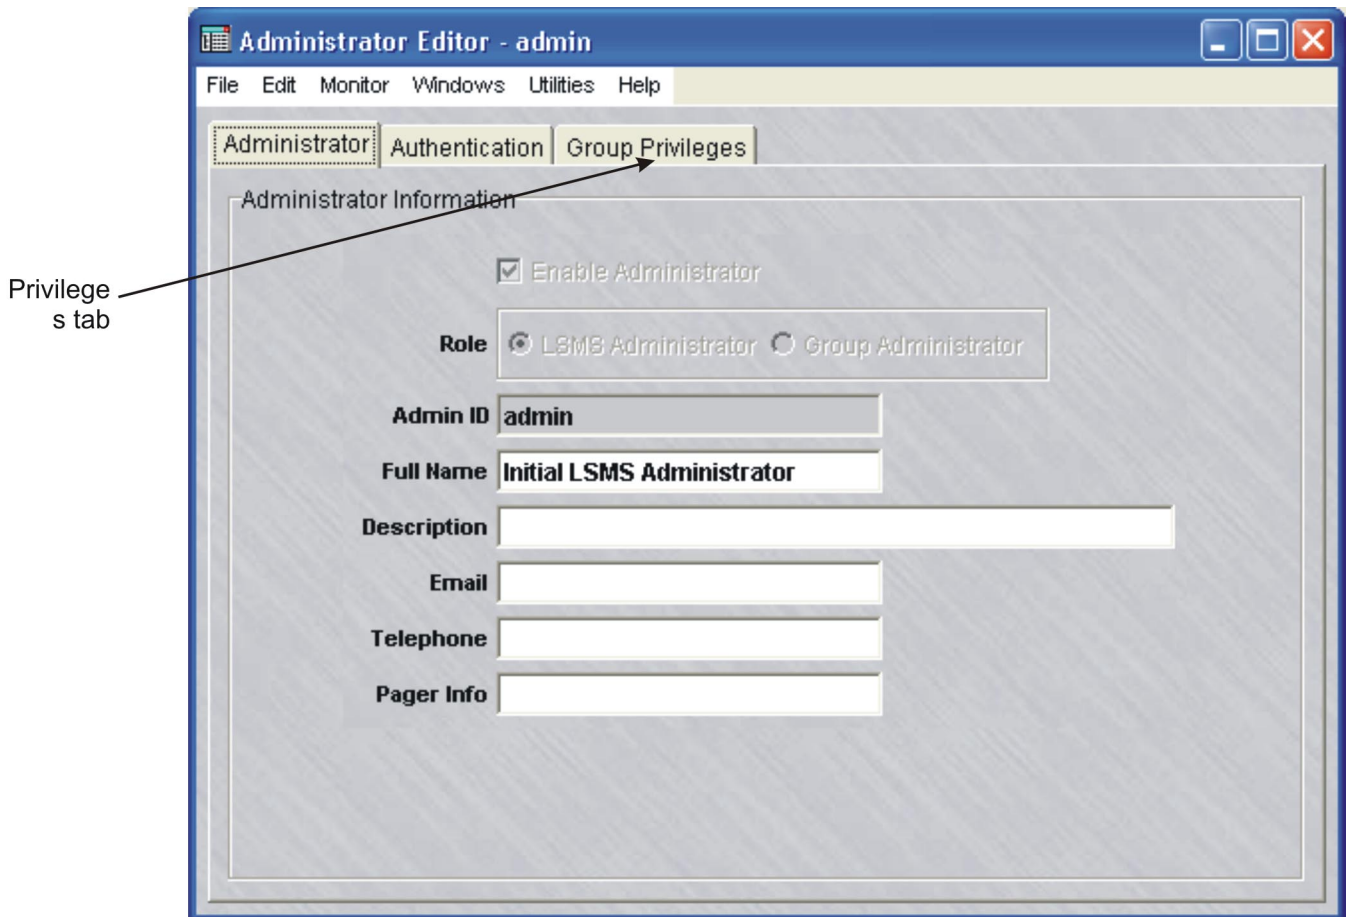
Once an administrator has been created, an LSMS administrator can edit the account when information or privileges change, or delete it when it is no longer necessary.

Edit an Administrator Account

To edit an administrator account, follow the steps below:

- 1 With the Navigator window displayed, click the **Administrators** folder to display all administrator accounts in the Contents panel.
- 2 Highlight and right-click or double-click the account you want to edit. It will appear in the Administrator Editor, as shown in Figure 7-7.

Figure 8-7 Administrator Editor (Edit Mode)



-
- 3 Make any changes to the information shown. You cannot change the Admin ID. The password can be changed under the Authentication tab, but it never displays.
-
- 4 If this is a Group Administrator account, you can add groups to the account, and delete groups. You can also change the administrator's privileges for an existing group. The following explains how:

To do this:

- To add a new group, proceed to the Group Privileges tab, right-click in the Privileges panel and select **New** from the pop-up menu. Enter the group and privileges, and click **OK**.
 - To delete a group, right-click the group and select **Delete** from the pop-up menu. Click **Yes** to confirm the deletion.
 - To change the administrator's privileges, double-click the entry in the Privileges panel, enter the new privileges and click **OK**.
-

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

Important! *When Changes Take Effect*

If a Group Administrator is currently logged in, and the privileges are changed by an LSMS administrator, the changes will not take effect until the Group Administrator logs out and logs back in again.

The current session of the Group Administrator is not terminated when changes are made.

END OF STEPS

View the Status of an Administrator Account

- 1 To view the status of an Administrator account, follow the steps below, with the Navigator window displayed, click the **Administrators** folder to display all

administrator accounts in the Contents panel (see [Figure 8-8, “Contents Panel \(Administrators folder\)”](#) (p. 8-23)).

Figure 8-8 Contents Panel (Administrators folder)

Folder: /Administrators/ - Total 4 item(s)				
Name	Status	Admin	Last Modified	Description
admin	enabled	setup	2004-11-08 15:47:07	
elobelo	enabled	admin	2004-12-07 14:00:41	Isms administrat
jdaniels	enabled	admin	2005-06-03 10:42:24	group administrat
tbanks	disabled	admin	2005-06-03 10:42:34	group administrat

The list below describes each field shown:

- **Name** — The Admin ID of the administrator.
- **Status** — A field that shows whether the administrator account is currently **enabled** or **disabled**. The status of an administrator can be changed only by an LSMS administrator by checking/unchecking the **Enable Administrator** checkbox on the Main tab of the Administrator Editor window. For details about the **Enable Administrator** checkbox, refer to the section [“To Create Administrator Accounts”](#) (p. 8-10).
- **Last Modified** — A field that shows the date and time that any changes were made to the administrator account.
- **Description** — A text field that can be used to provide any additional descriptive information about the administrator account.

END OF STEPS

Delete an Administrator Account

To delete an existing administrator account, follow the steps below:

- 1 With the Navigator window displayed, click the **Administrators** folder to display all administrator accounts in the Contents panel.
- 2 Right-click the administrator and select **Delete** from the pop-up window.
- 3 Click **Yes** in the pop-up window to confirm the deletion.

Important! The initial LSMS administrator account created during installation cannot be deleted. Only accounts created afterwards can be deleted.

Group Administrators cannot delete accounts, but they can edit information in their own account.

END OF STEPS



To Use the LSMS Messenger

Overview

The LSMS Messenger is a feature that allows LSMS and Group administrators to send short messages to other administrators who are already logged in. Using this feature, an administrator can send a message to one or more specific administrators or to all active administrators.

If you have a redundant or multi-site LSMS configuration, or a Compute Server configuration, you can send messages to administrators logged into the primary LSMS, secondary LSMS, or LSCSs.


Send messages using the LSMS Messenger

Complete the following steps to send a message to one or more administrators using the LSMS Messenger:

- 1 From any LSMS window, select **Utilities > LSMS Messenger**.

Result The LSMS Messenger is displayed ([Figure 8-9, “LSMS Messenger” \(p. 8-25\)](#)).

Figure 8-9 LSMS Messenger



- 2 Type the message in the **Message** field at the top of the window.
To enter a multi-line message, type `\t` wherever you want a line break.

Example: the system will be going down in 5 minutes\tplease terminate all activities\tthe system will be back up at 22:00

3 In the left-hand panel, indicate the recipient of the message.

There are three ways to do this:

- Click **Send to all Active Admins** to send the message to all active administrators (both LSMS and Group Administrators).
 - Click **Send to Active LSMS Admins** to send the message to all active LSMS Administrators (not Group Administrators)
 - Select one or more administrators, which are displayed under the checkboxes. You can select multiple administrators by holding down the **[Ctrl]** key and clicking each one, or by holding down the **[Shift]** key and clicking two administrators (all administrators between the two will also be selected).
-


4 Click the **Send** button.

Result The message is sent to the selected administrator(s).

A record of the message sent is displayed in the main panel of the LSMS messenger. The record includes the time the message was sent, the administrator who sent the message, the intended recipients of the message, and the text of the message.

Figure 8-10, “Records of Sent Message” (p. 8-27) shows a typical display of a message that can be sent.

Figure 8-10 Records of Sent Message



To clear the display, click the **Clear All** button.

END OF STEPS


Receive messages using the LSMS Messenger

Complete the following steps to receive a message sent using the LSMS Messenger:

- 1 Log onto the LSMS using one of the procedures described in “Log On and Off the LSMS Server or Compute Server” (p. 1-2) in Chapter 1, “Getting Started”.

Result If a message has been sent to you using the LSMS Messenger, the system displays a yellow envelope icon the drops repeatedly from the “clouds” at the top right of the Navigator ([Figure 8-11, “Messenger Envelope”](#) (p. 8-28) shows an example).

Figure 8-11 Messenger Envelope



This is the same area where a yellow bell icon is displayed when an LSMS or Brick alarm condition has been detected.


- 2 Click on the yellow envelope icon or select **Utilities > LSMS Messenger**.

Result The LSMS Messenger is displayed (see [Figure 8-10, “Records of Sent Message”](#) (p. 8-27)).

- 3 Click on the record of the message displayed in the **Instant Message** portion of the window.

Result The message sent is displayed in the bottom panel of the window [Figure 8-12, “LSMS Messenger \(Message Sent\)”](#) (p. 8-28) shows an example).

Figure 8-12 LSMS Messenger (Message Sent)



To clear the display, click the **Clear All** button.

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



9 Compute Servers

Overview

Purpose

This chapter discusses the concept of Compute Servers as an alternative means of collecting log information from Lucent VPN *Brick*® devices that are managed by LSMS. It also describes how to configure Compute Servers.

Contents

What is a Compute Server?	9-2
How to Configure a Compute Server	9-6



What is a Compute Server?

Overview

To maximize the scalability of the LSMS/Brick security solution, LSMS provides the option of adding a separate set of servers called Compute Servers (LSCSs), which are associated with an LSMS server or redundant pair of LSMSs and act as collection points for Brick log traffic. Using an LSCS to collect Brick log data frees up computing resources on the LSMS itself and extends the number of Bricks and total log traffic that can be handled. Each Brick managed by the LSMS can be homed to one of the associated LSCSs or the managing LSMS for logging purposes.

An LSCS provides most of the same functionality as an LSMS, but does not have its own database. The database is centrally located on the Primary LSMS. The Brick log data is collected and stored in files on the associated LSCS.

Quantity of Compute Servers and Bricks Supported by an LSMS

One LSMS server can support up to five LSCSs. Each LSCS can collect log data from up to 1,000 Bricks. A redundant pair of LSMSs, in a Primary LSMS/Secondary LSMS arrangement, can support up to 10 LSCSs and manage up to 10,000 Bricks.

Redundancy

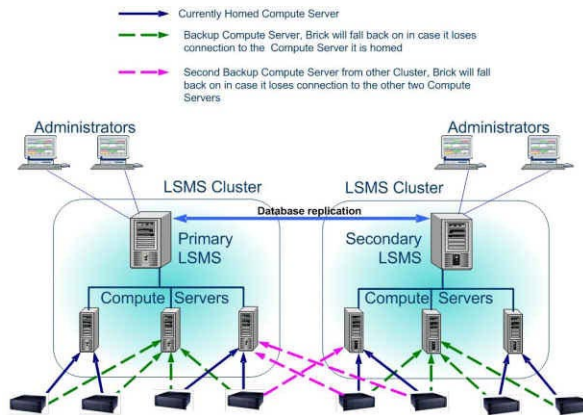
If a group of LSCSs has been added and configured, a Brick can be homed to one of the LSCSs to serve as the logging server, with the remaining LSCSs configured as backup log collection points in the event that the Brick connection to the first LSCS fails. The failover priority of LSCSs can be defined for a Brick, using the Brick Editor window (refer to [Chapter 3, “Configuring and Activating a Lucent VPN Firewall Brick® Device”](#) for details on how to configure a Brick).

LSCSs can be configured to work with a single Primary LSMS, but a redundant LSMS pair is recommended for a large network of Bricks to evenly distribute the collection of log data and to maintain connectivity.

An LSCS can be associated with either the Primary or Secondary LSMS in a redundant pair. The LSMS can still be used to collect log data and perform configuration functions on the associated Brick regardless of which LSMS is currently active.

Figure 9-1, “LSMS Cluster Arrangement of Computer Servers and Redundant LSMS Pair” (p. 9-3) depicts a typical cluster arrangement of Compute Servers (LSCSs) employed in collecting Brick log data managed by redundant LSMS servers.

Figure 9-1 LSMS Cluster Arrangement of Computer Servers and Redundant LSMS Pair



Compute Servers can be geographically dispersed and still communicate securely with the LSMS. However, to minimize the complexity of the configuration and reduce the amount of log data to be sent over expensive WAN links, it is recommended that the Compute Servers be placed locally with the associated LSMS. A single Compute Server or group of Compute Servers should be protected by a Brick to ensure that only authorized traffic reaches the LSMS cluster.

Accessing a Compute Server

An LSCS need to be configured on the LSMS by an LSMS Administrator before it can be brought up. For instructions on how to configure an LSCS, refer to the procedure “How to Configure a Compute Server” (p. 9-6).

Logging Into a Computer Sever

LSMS Administrators can log into an LSCS from the LSMS console or remotely using the Remote Navigator. Only LSMS Administrators have access permissions to add, modify, or delete LSCSs. Group Administrators do not have access permissions to LSCSs and can only view these servers through the Brick Editor window.

An LSMS Administrator can reconfigure an LSCS while logged into the server directly from the LSMS console or through the Remote Navigator.

LSMS Tools on the Compute Servers

Most of the functions that are available on the managing LSMS can be performed while logged into an LSCS, including creation and update of Bricks and Policies on the Bricks that are associated with the LSCS.

Alarms and Logs

Alarms that are generated on a specific LSMS or Compute Server can be viewed by all Administrators logged into all LSMSs/LSCSs in the network while logged into an LSMS. Distributed reports can also be run to obtain a consolidated picture of all LSMS/LSCS activities.

An LSCS does not support the SNMP application. SNMP-based information retrieval must still be viewed from the managing LSMS.

Only the local LSCS log data can be viewed while logged into the LSCS.

To view the raw logs on an LSCS, an LSMS Administrator must log into that server and view the Brick data rows in a flat file or use the local Log Viewer (LSMS function).

LSMS Services

The DataBase (DB), SNMP, and VGC services do not run on an LSCS. The LSCS communicates over the network with its associated LSMS for DB access. Since the VGC service does not run on an LSCS, Compute Server environments only support IKE on the Brick for VPN service.

Status Monitor

An Administrator can monitor the activity and status of an LSCS using the Status Monitor window.

On the Status Monitor window, the status of each LSCS is shown on a separate, indented line below its associated LSMS and provides the following details:

- Name
- IP Address
- Status (Up, Down, Lost)
- LSMS association
- Number of Bricks Connected

For complete information about the LSMS Status Monitor window, refer to

Operating System

LSCSs supported by the LSMS run on the *Windows*[®] XP operating system. An LSCS can communicate with an LSMS running on the Solaris operating system, since it is a separate server that resides in front of the LSMS and is protected by a Brick.



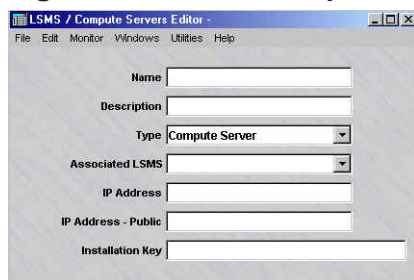
How to Configure a Compute Server

Task

To configure a new Compute Server, follow the steps below:

- 1 With the Navigator window displayed, right-click the **LSMS and LSCSs** folder and select **New LSMS and Compute Servers** from the pop-up menu. The LSMS / Compute Servers Editor window is displayed ([Figure 9-2, “LSMS/Computer Servers Editor Window”](#) (p. 9-6) shows a sample window).

Figure 9-2 LSMS/Computer Servers Editor Window

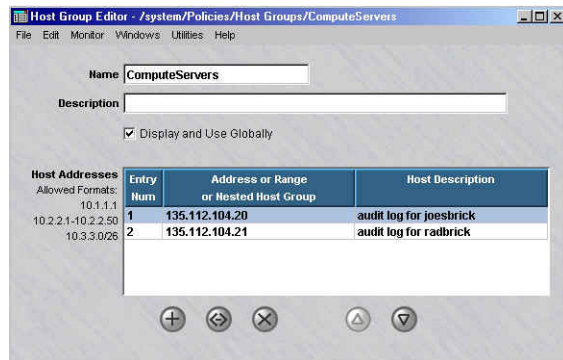


- 2 Enter values in the following fields:
 - **Name** - The name of the Compute Server, 1-45 characters.
 - **Description** - A textual description of the Compute Server (Bricks supported, types of logging data collected, and so forth).
 - **Type** - Click the down-arrow next to this field and select **Compute Server** from the drop-down list.
 - **Associated LSMS** - Click the down-arrow next to this field and select the Primary or Secondary LSMS associated with this Compute Server.
 - **IP Address** - The real IP address of the Compute Server. On private networks, this is the real IP address of the server that can be mapped to a virtual address using NAT.
 - **IP Address - Public** - This is the Virtual Brick Address (VBA) of the Brick protecting the Compute Server. This field is optional.
 - **Installation Key** - The license key used to install the Compute Server and associate it with an LSMS.

- 3 From the File menu, select **Save and Close**.

The new Compute Server is configured. A entry for the new Compute Server appears in the ComputeServers Host Group (Figure 9-3, “Host Group Editor window (ComputeServers Host Group)” (p. 9-7) shows a sample Host Group Editor window).

Figure 9-3 Host Group Editor window (ComputeServers Host Group)



END OF STEPS



10 Remote Administration

Overview

Purpose

This chapter will explain all the steps needed and options available to administer the LSMS from a remote PC or workstation, including:

- Installation and supported platforms
- LSMS policy changes needed to permit remote administrator access
- Login procedure
- Tasks that may be performed remotely

Contents

What is the Remote Navigator?	10-2
Install the Remote Navigator on Windows	10-3
Install the Remote Navigator on Solaris	10-7
Permitting Remote Administration on the LSMS	10-10
Create the Host Group	10-11
Create the Security Rules	10-12
Log in from a Remote Host	10-15
What Can a Remote Administrator Do?	10-17



What is the Remote Navigator?

Overview

The LSMS Remote Navigator application allows you to run the LSMS application from a remote host instead of the LSMS console. It is provided on the LSMS CD-ROM, and may also be retrieved remotely from the LSMS via a browser.

An Administrator can log into the LSMS remotely using a host running either *Windows*[™] or *Solaris*[®] operating systems. Interoperability across platforms is supported, so Windows remote clients can access a Solaris LSMS, and Solaris remote clients can connect to a Windows LSMS. The following gives the software requirements for both platforms.

Windows

The remote host must be running the following software:

- Windows 2000 Professional Windows 2000 Server, Windows XP Professional, or Windows Server 2003
- Microsoft Internet Explorer 5.0 or greater
- Adobe Acrobat Reader 4.0 or above, to read the on-line documentation.

Solaris 2.8, 2.9, 2.10

The remote host must be running the following software:

- Solaris 2.8, 2.9, or 2.10
- Netscape Communicator 4.7 or greater
- Adobe Acrobat Reader 4.0 or above, to read the on-line documentation.



Install the Remote Navigator on Windows

Overview

There are two ways to install the LSMS Remote Navigator— directly from the LSMS CD-ROM, or by downloading the software from the LSMS.

Important! LSMS Version 8.0 Remote Navigator is not compatible with LSMS 7.0.

If you need to access both releases of the LSMS, you must install the Version 8.0 Remote Navigator in a separate directory on your PC. Then, use the Version 7.0 Remote Navigator to access LSMS 7.0 and the Version 8.0 Remote Navigator to access LSMS 8.0.

From the CD-ROM

To install the LSMS Remote Navigator from the CD-ROM, follow the steps below:

- 1 With the LSMS CD in the CD-ROM drive, open the Windows Explorer and locate this directory on the CD-ROM:

\Nt\RemoteNavigator

- 2 Double-click the file

lsmsremnav-8.0.xxx.exe

where *xxx* is the version number of the software. This is the installation program. The installation process will begin.

Important! If you are upgrading from an earlier release of the Remote Navigator, a window will pop up and indicate that an earlier release is loaded on the host you are using. Click **Yes** to overwrite the older release.

- 3 The first window to appear is the Welcome window. Read the text in the Welcome window, and when you are finished, click **Next** to continue with the installation.
-

- 4 The Choose Destination Location window will appear. This window allows you to specify where the LSMS Remote Navigator software will be installed. The default is:

c:\LSMSRemNav8.0

We recommend you accept the default. Click **Next** to do this, or click the **Browse** button and enter a new destination location before clicking **Next**.

-
- 5 The Select Program Folder window will appear. This window allows you to select the folder in which the LSMS Remote Navigator program selections appear on the Windows Start menu, Programs. The default is:

Lucent Security Management Server

We recommend you accept the default. Click **Next** to do this, or select another folder from the Existing Folders panel before clicking **Next**.

- 6 Installation of the files will now begin. When the installation is finished, the Setup Complete window will appear. Click **Finish** to complete the installation. It is not necessary to reboot your machine. The LSMS Remote Navigator may be run immediately.

END OF STEPS

From the LSMS

To download the LSMS Remote Navigator from the LSMS, follow the steps below:

- 1 Create a temporary directory on the hard drive of your computer to hold the downloaded file.
-
- 2 Open a browser and enter the URL of the LSMS. The URL consists of the LSMS IP address, its port number, and the directory in which the LSMS application is stored. Depending on whether the LSMS web server is HTTP or HTTPS, the URL will look like one of the following:

http://<ip_address>:<port>/LSMS

— or —

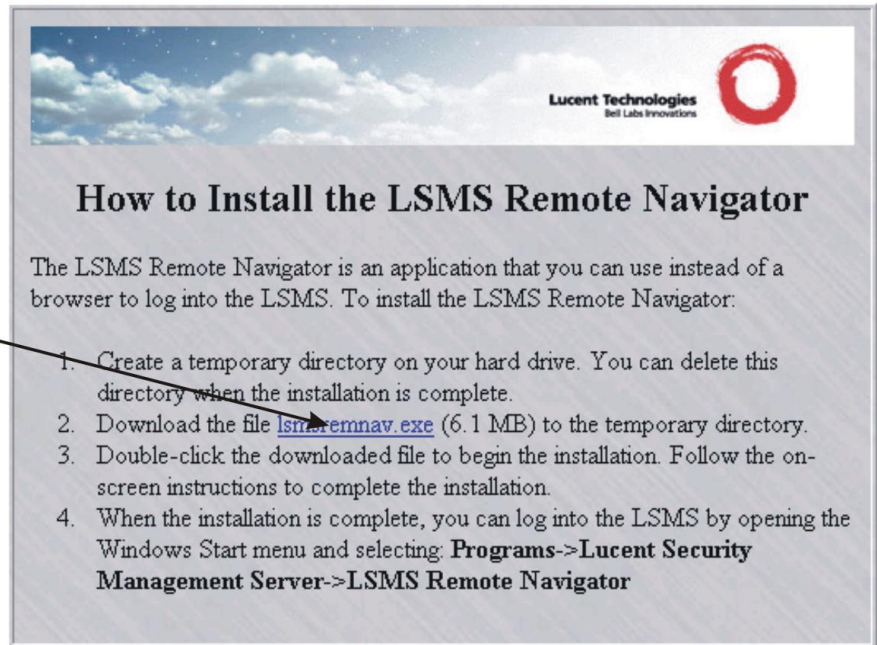
https://<ip_address>:<port>/LSMS

Important! Before the browser is displayed with the link to download the Remote Navigator software, you may be required to enter an "Authentication Key." For additional information about configuring this key, see the LSMS web server parameter in [Chapter 11, "Using the Configuration Assistant"](#).

-
- 3 The How to Install the Remote Navigator window will appear (see Figure 8-1). Click the **lsmsremnav.exe** link. You may want to leave the How to Install the Remote Navigator window up as a reference.

Figure 10-1 How to Install the Remove Navigator Window

lsmsremnav.exe
Link



-
- 4 The File Download window will appear. The **Save this program to disk** option is the default and should already be selected. Click **OK** to save the program.
-
- 5 A Save As dialog box will appear. Navigate to the temporary directory you created, and click the **Save** button. A progress indicator dialog box will appear.
-
- 6 When the dialog box indicates that the download is complete, click **Open** to begin the installation process.

-
- 7 When the installation is finished, the Setup Complete window will appear. Click **Finish** to complete the installation. It is not necessary to reboot your machine. The LSMS Remote Navigator may be run immediately.

END OF STEPS



Install the Remote Navigator on Solaris

Methods of installation

There are two ways to install the LSMS Remote Navigator — directly from the LSMS CD-ROM, or by downloading the software from the LSMS.

Important! LSMS Version 8.0 Remote Navigator is not compatible with LSMS 7.0.

If you need to access both releases of the LSMS, you must install the Version 8.0 Remote Navigator in a separate directory on your PC. Then, use the Version 7.0 Remote Navigator to access LSMS 7.0 and the Version 8.0 Remote Navigator to access LSMS 8.0.

From the CD-ROM

To install the LSMS Remote Navigator from the CD-ROM, follow the steps below:

- 1 Create a temporary directory on the hard drive of the your computer to hold the installation file.

- 2 With the CD-ROM in the CD-ROM drive, locate this directory on the CD-ROM:
/Solaris/RemoteNavigator

- 3 Copy the file
LSMSRemNav.tar
to the temporary directory.

- 4 Make the temporary directory the present working directory and issue the command:
`tar xvf LSMSRemNav.tar`

- 5 When you have successfully untarred the files, change directories to the /lmf directory created by the tar extraction process and enter the following to install the program:
./install

END OF STEPS

From the LSMS

To download the LSMS Remote Navigator from the LSMS, follow the steps below:

- 1 Create a temporary directory on the hard drive of your computer to hold the downloaded file.

- 2 Open a browser and enter the URL of the LSMS. The URL consists of the LSMS IP address, its port number, and the directory in which the LSMS application is stored. Depending on whether the LSMS web server is HTTP or HTTPS, the URL will look like one of the following:
http://<ip_address>:<port>/LSMS
— or —
https://<ip_address>:<port>/LSMS

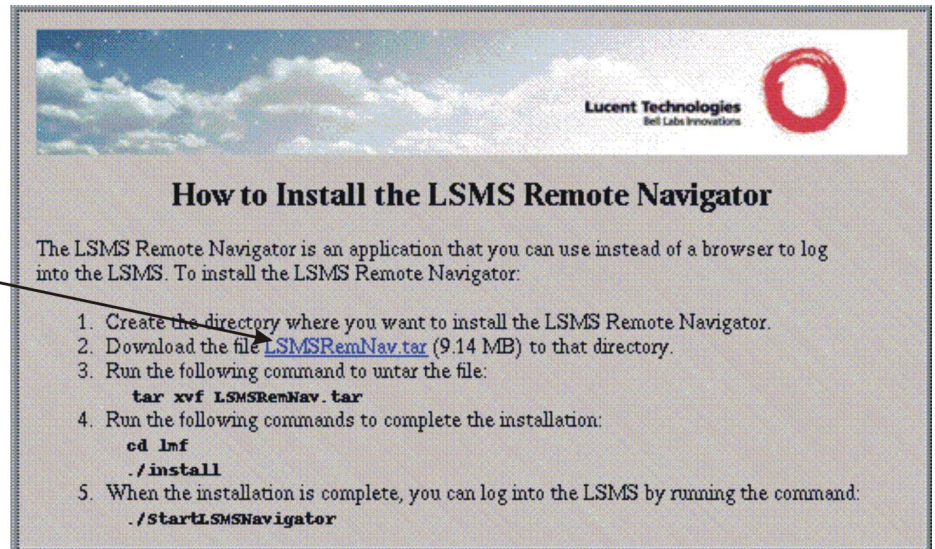
- 3 The LSMS login window will appear. Enter your Admin ID and password in the appropriate fields, click the checkbox labeled **Download the LSMS Remote Navigator**, then click the **OK** button.

- 4 The How to Install the Remote Navigator window will appear (see Figure 8-2). Click the **LSMSRemNav.tar** link and download the .tar file that contains the LSMS Remote

Navigator software to the temporary directory. You may want to leave the How to Install the Remote Navigator window up as a reference.

Figure 10-2 How to Install the LSMS Remote Navigator Window (solaris)

LSMSRemNav.tar
Link



-
- 5 Make the temporary directory the present working directory and issue the command:
- ```
tar xvf LSMSRemNav.tar
```
- 
- 6 When you have successfully untarred the files, change directories to the /lmf directory created by the tar extraction process and enter
- ```
./install
```
- to install the program.
- END OF STEPS
-



Permitting Remote Administration on the LSMS

Purpose

This section explains how to enable Administrators to log into the Lucent Security Management Server (LSMS) remotely when the LSMS is protected by a Lucent VPN FirewallBrick[®] device.

To set up remote administration, you have to modify the security policy of the Administrative Zone or the NOC Gateway Zone. These zones were created automatically during the installation of the LSMS software. For details, refer to *Appendix B. Pre-Configured Brick Zone Rulesets* in the *LSMS Policy Guide*.

One of these policies has to be changed to allow Administrators to log into the LSMS from remote hosts outside one or more of these Bricks. This involves creating two new rules and a host group, and adding them to the security policy.

This section provides information to perform the following:

1. Create a host group containing the IP addresses of all the hosts that will be permitted to access the LSMS remotely.
2. Create two rules allowing sessions to and from these remote hosts to pass through the Brick.

□

Create the Host Group

When to use

The first thing you should do to set up remote administration is to create a host group that contains the IP addresses of the remote hosts that each Administrator will be using to log into the LSMS.

If you do not create this host group, you will have to create separate security rules for each Administrator. However, with this host group, you will need only two rules to cover all Administrators.

If a new Administrator is introduced, or an existing Administrator leaves, you simply add the new IP address to, or delete the old IP address, from the host group. The security rules are not affected.

Create a New Host Group

To create a new host group, follow the steps below:

- 1 Open the Policies folder.
.....
- 2 Right-click the Host Groups folder and select **New Host Group**.
.....
- 3 The Host Group Editor appears.
.....
- 4 Enter a name that uniquely identifies this host group for example, *remote_admins*, in the **Name** field. The name can contain up to 45 alphanumeric characters.
.....
- 5 Enter a brief description of the host group in the **Description** field. The description is optional. It can contain up to 80 alphanumeric characters.
.....
- 6 Enter the IP addresses of the Administrators. Enter each IP address on a single line, or enter a range of addresses.
.....
- 7 Display the File menu and choose one of the **Save** options.

END OF STEPS



Create the Security Rules

When to use

Two rules must be added to the security policy of the Administrative Zone, or NOC Gateway Zone if the Brick is configured as a NOC Gateway, to permit the hosts in the host group you just created to access the LSMS from outside a Brick.

The two rules are needed to allow two-way communication between the remote hosts and the LSMS. One rule allows the Administrator's remote session *into* the Administrative Zone or NOC Gateway Zone, and the other rule allows a session initiated by the LSMS *out of* the Administrative Zone or NOC Gateway Zone.

Important! If the computer the Administrator is using to log in remotely is on a LAN in a zone that is assigned to another port on the Brick, these same rules have to be added to the security policy of that zone — *with the directions reversed*.

Create the First Rule

The first rule to be considered is the rule that allows sessions initiated by the remote Administrators to pass through a Brick into the Administrative Zone (or NOC Gateway Zone). To create this rule, follow the steps below:

- 1 Open the Policies folder.
- 2 Open the Brick Zone Rulesets folder and double-click **administrativezoneornocgw-zone**.
- 3 In the Brick Zone Ruleset Editor, right-click any rule and select **New**.
The Brick Zone Rule Editor appears.
- 4 In the **Direction** field, select **IN TO ZONE** from the drop-down list. This rule will now apply to sessions initiated outside the Administrative Zone.
- 5 In the **Source** field, click **Host** and select the host group from the drop-down list that contains the IP addresses of the Administrators who will be logging in remotely.

This ensures that only sessions initiated by those IP addresses are permitted to pass through the Brick.

-
- 6 In the **Destination** field, click **Host** and enter the IP address of the LSMS. You could also enter an asterisk, since the LSMS should be the only host in the Administrative Zone.
-

- 7 In the **Service or Group** field, select ****BROWSE****. In the Browse: Select a Service Group window, select:

secure_remote_admin_to_SMS

from the drop-down list. This is a Service Group (tcp/7000/* and tcp/443/*) provided with the LSMS specifically for this purpose.

Important! The entry 'tcp/443/*' implies that the LSMS web server has been configured for HTTPS. An entry of 'tcp/80/*' can be added to this service group if the LSMS web server is set up for HTTP. However, using HTTPS is preferred because it is more secure.

- 8 In the **Action** field, select **Pass** from the drop-down list.
-

- 9 In the **Description** field, enter an optional description, if necessary.
-

- 10 Click the **Advanced** tab and increase the value for "Session timeout".

It is recommended that the default value of 300 seconds be increased because the Remote Navigator session will be removed from the Brick after five minutes of idle time. The administrator login will still be active, but when the user attempts to reuse the GUI after, say, 30 minutes, the Brick no longer has an entry for this session. The resumed interaction looks to the Brick like a new session, but one that is starting without the normal TCP connection "handshake. The new TCP validations (configured under the Advanced tab) will block the session.

With the session timeout set to a more convenient threshold, the Remote Navigator session can be preserved for a longer period with the protection afforded by the Brick's strict TCP enforcement. For example, if you set the timeout to 3600 seconds, you will have one hour of idle time.

- 11 Click the **OK** button to temporarily store the rule on the LSMS.

END OF STEPS

Create the Second Rule

The purpose of the second rule is to permit a session initiated by the LSMS out of the Administrative Zone or NOC Gateway Zone. This rule is the same as the first rule, except:

- **Direction** is *out of* the zone
- **Source** is the IP address of the LSMS
- **Destination** is the host group containing the IP addresses of the remote Administrators
- **Service or Group** is *secure_remote_admin_from_SMS*.

Click the **Advanced** tab, and ensure that all TCP validations are unchecked. Otherwise if the Remote Navigator is logged in and idle for awhile, its session cache will time out. The next time the administrator attempts to use the Remote Navigator, it will need to (transparently) open a new session, and the TCP validation may block it.

Create this rule as you created the first rule, and then click the **OK** button to store it temporarily on the LSMS. From the File menu in the Brick Zone Ruleset Editor, select Save and Apply to apply them to the Brick.

□

Log in from a Remote Host

When to use

Important! It is possible to install the LSMS Remote Navigator on the LSMS host so that you are running both the LSMS Navigator and the LSMS Remote Navigator on the same machine. The only reason to do this is if you intend to use the LSMS host to log into another LSMS remotely.

Task

To log in from a remote host using the LSMS Remote Navigator, follow the steps below:


- 1 If the remote host is running Windows, click the **Start** menu and select:
Programs ➤ **Lucent Security Management Server** ➤ **LSMS Remote Navigator**

If the remote host is running Solaris, go to the installation root directory (*/opt/isms/lmfif* if you used the defaults during installation) and enter:

./StartLSMSNavigator

from the command line. In either case, the login window shown in Figure 9-3 will appear.

Figure 10-3 LSMS Remote Navigator Login Window



-
- 2 Enter your **Admin ID** and **Password**. The Admin ID and password are the ones that were created during LSMS installation, or those given to you by another administrator.

If you want to access the Status Monitor without logging into the rest of the LSMS, you can check **Status Monitor Only Login**. This is useful if:

- You have a special monitoring room with large screens, and you want to display the graphs to monitor the health of the system, or
 - If you want to provide someone with the ability to monitor the system, but you do not want this person to be able to view or change the system's configuration.
-

- 3 Enter the URL of the LSMS or LSCS. The URL is either:

http://<IP_address>:<port_number>/LSMS

— or —

https://<IP_address>:<port_number>/LSMS

where *<IP_address>* is the IP address of the LSMS or LSCS and *<port_number>* is the port the web server is listening on. Ports 80 and 443 are the standard ports for HTTP and HTTPS, respectively. The port your web server is using was assigned during installation; if another port was entered, use it instead.

Each URL you enter will be placed in the drop-down list in the LSMS or LSCS URL field, so that each time you enter this URL after the initial entry, you can simply select it from the drop-down list instead of typing it in. You can store multiple URLs in this list in the event that you need to log into more than one LSMS remotely.

- 4 Click **Connect**. Initially the remote host and the LSMS will exchange keys to set up a 3DES tunnel. Once the tunnel is in place, the LSMS will authenticate the administrator ID and password. If the ID and password are indeed valid, another 3 DES tunnel is enabled to maintain maximum security throughout the session.

When you have successfully logged in, the Navigator window is displayed.

Important! LSMS Version 8.0 Remote Navigator is not compatible with LSMS 7.0.

If you need to access both releases of the LSMS, you must install the Version 8.0 Remote Navigator in a separate directory on your PC. Then, use the Version 7.0 Remote Navigator to access LSMS 7.0 and the Version 8.0 Remote Navigator to access LSMS 8.0.

END OF STEPS



What Can a Remote Administrator Do?

Remote administration capabilities

Once an administrator has successfully logged into the LSMS Remote Navigator, everything that an LSMS administrator (LA) or group administrator (GA) can do from a local LSMS - provision new Bricks, apply policy changes, create VPN tunnels, and so forth, can be done from the Remote Navigator.

The remote administrator can benefit from a number of options through the **Utilities** menu bar on the Navigator. The options are:

1. Configuration Assistant - (LSMS Administrators only) - This tool allows an administrator to easily configure a number of system wide parameters. For more information, see *Chapter 10. Using the Configuration Assistant*.
2. New Feature Setup - (LSMS Administrators only) - The New Feature Setup utility allows an administrator to install a new license key for optional features or to increase the management capacity of the LSMS. For more information, refer to [Appendix E, "New Feature Setup"](#).
3. Restart Services - (LSMS Administrators only) - This tool stops and starts LSMS Services on the LSMS that you are remotely logged into.
4. LSMS Service Status - (LSMS Administrators only) - This utility presents a graphical summary of all LSMS services. This may be helpful for system monitoring or for troubleshooting purposes. For more information, see the *LSMS Tools and Troubleshooting Guide*.
5. LSMS Log Viewer - (All Administrators) - The remote administrator may monitor a variety of LSMS activities in "real time". This utility is particularly helpful for troubleshooting. For more information, see the *Reports, Alarms and Logs Guide*.
6. LSMS Messenger- (All Administrators) - All administrators currently logged in may now exchange short messages with this tool. It can be used for a two way exchange or a broadcast to all administrators. For more information, see *Chapter 8. Creating Groups and Administrators*.

Often it is important to monitor the status of individual Bricks. In this release, LSMS administrators can access the Brick console directly from the Navigator. Once the "remote Brick console" is displayed, the admin can issue any of the commands that can be executed from a local Brick console. For more information, see *Chapter 5 Introduction to the Brick Command Line Interface* in the *Command Line Interfaces User Guide*.

When a Brick is first created on the LSMS, its initial configuration must be loaded on the device. In this release, the user has the three choices to transfer the information to the Brick:

- Create a floppy with the Brick configuration on the local LSMS and load the floppy on the device.
- Create a floppy with the Brick configuration on any remote PC or workstation and load the floppy on the device.
- If a serial connection on the Brick is available from the network (as via a terminal server), the Brick configuration may be loaded onto the device without a floppy.

Of course, once the floppy has been loaded on the Brick, the LSMS and Brick can communicate directly for subsequent configuration updates, new software downloads, policy changes, etc. For more information on configuring and loading Bricks, see Chapter 3 *Configuring Bricks* in the *LSMS Administration Guide*.

7. Edit LSMS Parameters (LSMS Administrators only) - this utility allows an administrator to enable or disable the Concurrency Control feature and configure additional options related to this feature. For more information, refer to [“Concurrency Control”](#) (p. 1-22) in Chapter 1, [“Getting Started”](#).
8. View LSMS Parameters (LSMS Administrators only) - this utility allows an administrator to view the current Concurrency Control feature settings. For more information, refer to [“Concurrency Control”](#) (p. 1-22) in Chapter 1, [“Getting Started”](#).

□

11 Using the Configuration Assistant

Overview

Purpose

This chapter explains how to use the Configuration Assistant to set a number of parameters that affect the system's operation and performance.

The Configuration Assistant may be accessed either directly at the local LSMS host or remotely while logged in via the LSMS Remote Navigator. In order to run the Configuration Assistant from the Remote Navigator, you must be logged in as an LSMS Administrator.

Contents

What is the Configuration Assistant?	11-3
Alarms	11-9
Detailed Policy Audit	11-11
Direct Paging	11-13
GUI and Status Monitor Parameters	11-18
Log Files	11-20
Log Transfer	11-24
Login Banner	11-27
LSMS Web Server	11-29
Reports	11-32
SNMP Agent	11-34
Software Download	11-37
TL1 Alarms	11-43
Tunable Parameters	11-45

User Authentication	11-48
Strong Passwords	11-51
VPN Debugging	11-53



What is the Configuration Assistant?

Definition

Configuration Assistant is the name given to the software that is used to set a variety of parameters that affect the way the overall system operates.

The installation software for the LSMS gives you the opportunity to display the Configuration Assistant immediately after the LSMS software has been installed. If you choose to do this, you can set any of the parameters at this time.

An alternative is to operate the system for awhile using the default parameters, and then decide whether or not to change any of the parameters. The Configuration Assistant can be activated at any time to make changes.

Starting the Configuration Assistant

The Configuration Assistant can be started from either the local LSMS host or while logged into the LSMS Navigator or LSMS Remote Navigator. The following explains how to do this on both the Windows and Solaris platforms:

Windows

- 1 Click **Start** on the Windows taskbar and select:

Programs  **Lucent Security Management Server**  **Utilities** 
Configuration Assistant

to display the Configuration Assistant window (see Figure 10-1)

-OR -

- 2 Log into either the LSMS Navigator or Remote Navigator as an LSMS administrator. From the Utilities Menu Bar, click on System Utilities, then select the Configuration Assistant.

END OF STEPS

Solaris procedure

Solaris

- 1 Open a window and enter
`cd <install_directory>`

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

This makes the installation root directory the present working directory.

2 Enter

`./configurationAssistant`

to display the Configuration Assistant window (see Figure 10-1).

- *OR* -

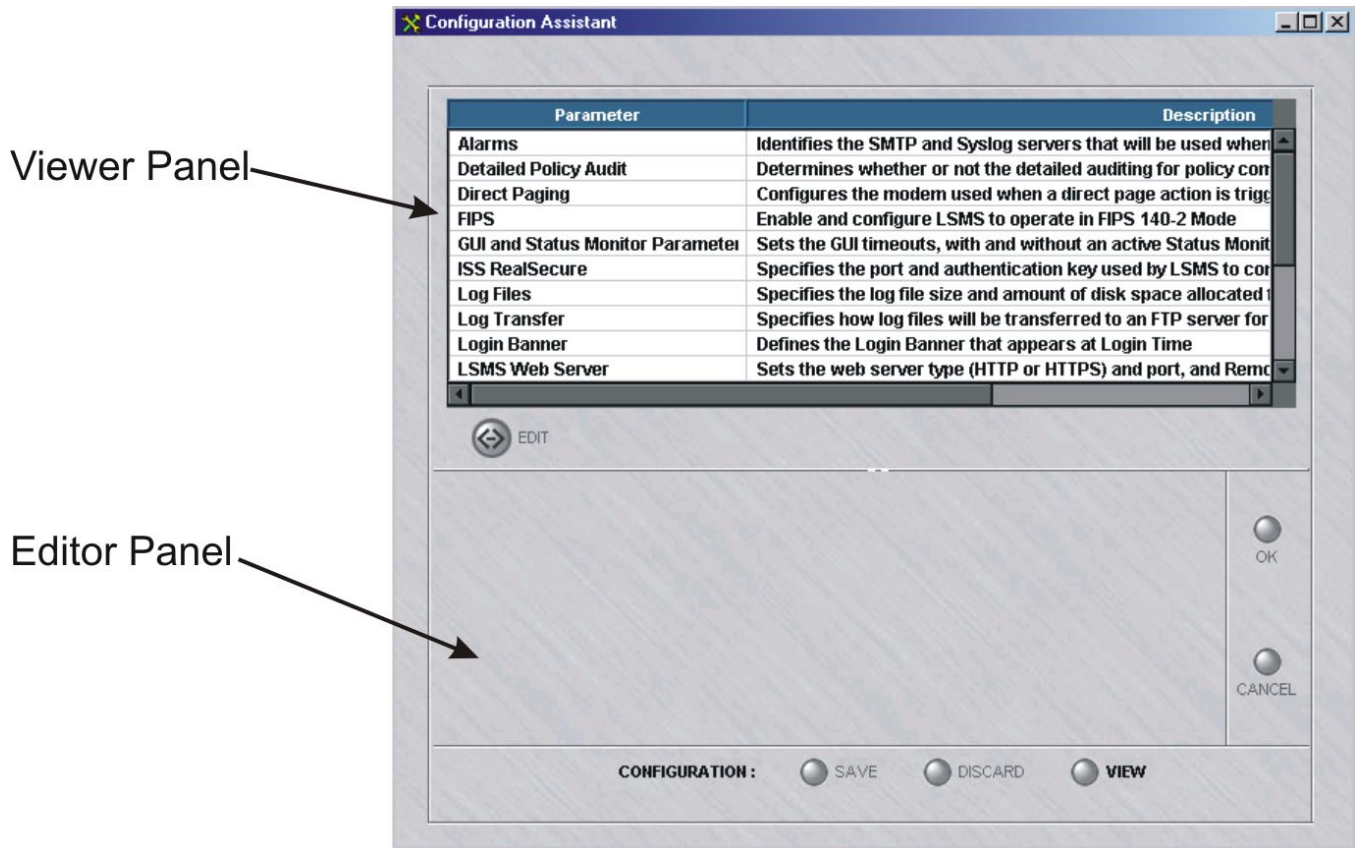
3 Log into either the LSMS Navigator or Remote Navigator as an LSMS administrator. From the Utilities Menu Bar, click on System Utilities, then select the Configuration Assistant.

END OF STEPS

Configuration Assistant Window

- 1 [Figure 11-1, “Configuration Assistant Window” \(p. 11-5\)](#) shows the Configuration Assistant window.

Figure 11-1 Configuration Assistant Window



The Viewer Panel of the Configuration Assistant window contains 18 groups of parameters. The following is a brief explanation of each group:

- **Alarms**
Identifies the SMTP and Syslog servers that will be used when an alarm triggers an e-mail or Syslog message.
- **Detailed Policy Auditing**
Allows the LSMS to store archive copies of changes to Lucent VPN Firewall *Brick*® devices, rulesets, and other policy components.
- **Direct Paging**
Configures the modem used when a direct page action is triggered by an alarm.
- **FIPS**
Allows you to enable and configure the LSMS to operate in FiPS 140-2 mode.

- **GUI and Status Monitor Parameters**

Sets the GUI timeouts, with and without an active Status Monitor. Permits real time Brick status to be displayed on the Status Monitor for both LSMSs in a redundant pair.
- **Log Files**

Sets the log file size and the amount of disk space allocated for log files. Can also define an alternate directory to store log files, as well as specify the maximum number of users that can simultaneously use the Log Viewer.
- **Log Transfer**

Specifies how log files will be transferred to a designated FTP server for creating reports.
- **Login Banner**

Specifies an optional text message that is displayed to administrators after a successful login to the LSMS. The message can be a legal notice, security policy notification, disclaimer, etc.
- **LSMS Web Server**

Sets the web server type (HTTP or HTTPS), and the port. It also sets an optional authentication key for downloading Remote Navigator software. If this option is selected, administrators will have to enter a key to download Remote Navigator software.
- **Reports**

Sets the disk space used by reports, the maximum record size of a report, and the number of lines on a report in both portrait and landscape formats.
- **SNMP Agent**

Allows you to change the default port for the LSMS SNMP Agent. This port was initially set during installation of the LSMS. You may also define an alternate System SNMP Agent Listening Port for other non-LSMS/Brick SNMP events. You can also change SNMP Read Community parameter.
- **Software Download**

Determines how the IPSec client software will be deployed to users across the public Internet.
- **TL1 Alarms**

Allows you to enable and configure TL1 alarms.
- **Tunable Parameters**

Allows you to adjust certain *maxHeap* parameters in the configuration file (*config.ini*) so that the Brick better handles environments with large numbers of Bricks, client tunnels, and audit records.

- **User Authentication**
Sets the authentication server type (HTTP or HTTPS), and the port.
- **VPN Debugging**
Determines whether or not VPN connections will be audited and recorded in log files.

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

Set a Parameter

Using the Configuration Assistant window, you can set the individual parameters that are part of each parameter group. After setting a parameter, you will be instructed to stop and restart services, if necessary to activate your changes.

To set a parameter, follow the steps below:

- 1 Click the parameter group in the Viewer Panel.
The parameter group is highlighted, indicating it has been selected. The parameter's values are displayed in the Editor Panel below.
.....
- 2 Click the **EDIT** button to cause the fields and buttons in the Editor Panel to become active.
.....
- 3 Make any necessary additions or changes to the fields shown.
.....
- 4 Click the **OK** button to temporarily store the changes on the LSMS.
.....
- 5 Click the **SAVE** button to save the changes permanently on the LSMS.
A pop-up window will appear and indicate the date and time the configuration file was updated.
.....
- 6 Click **OK** to dismiss the pop-up window that confirms saving the selection.

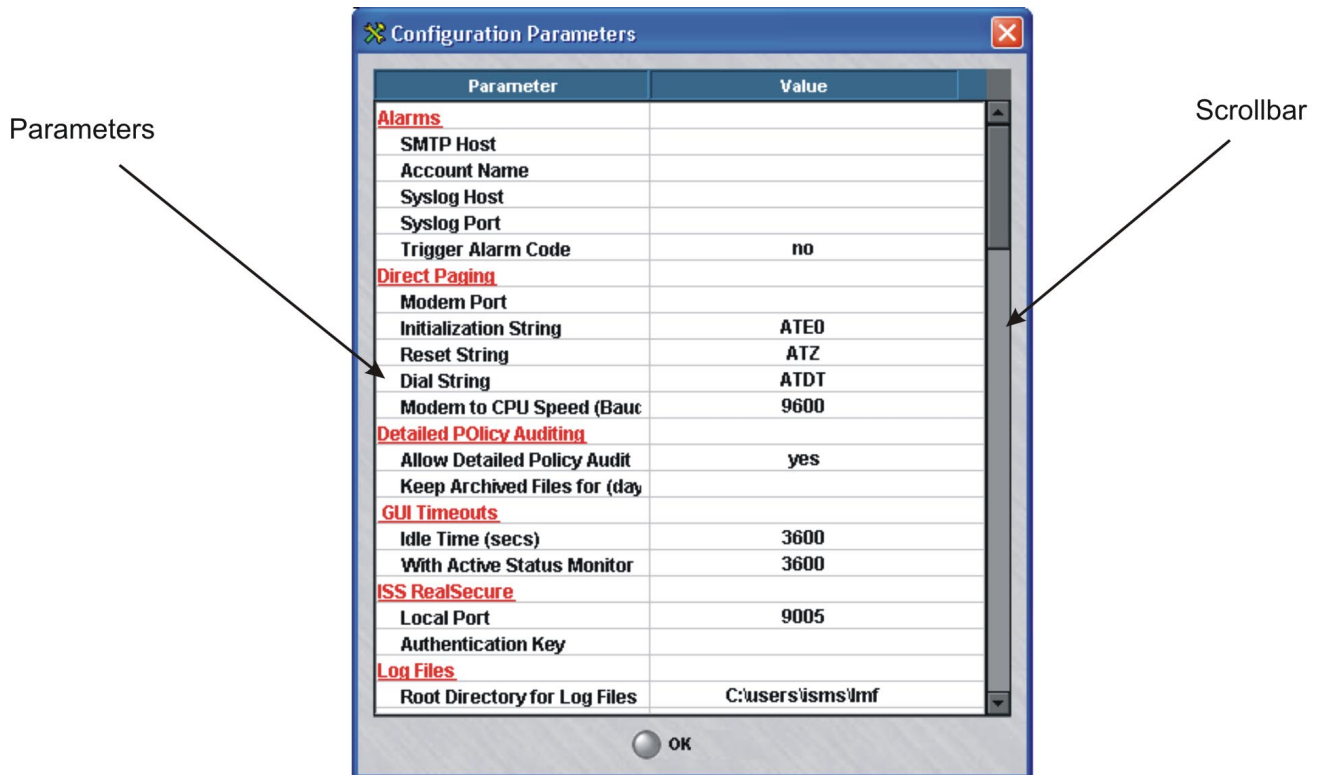
.....
E N D O F S T E P S
.....

View the Parameters

To view the values of all the parameters in one window, click the **VIEW** button to display the Configuration Parameters window. An example of this window is shown in [Figure 11-2, “Configuration Parameters Window”](#) (p. 11-8).

Note that not all parameters can be displayed at once, and a scrollbar has been provided to display the parameters that are not shown.

Figure 11-2 Configuration Parameters Window



The remainder of this chapter explains in greater detail how to set any of the parameters shown.



Alarms

Overview

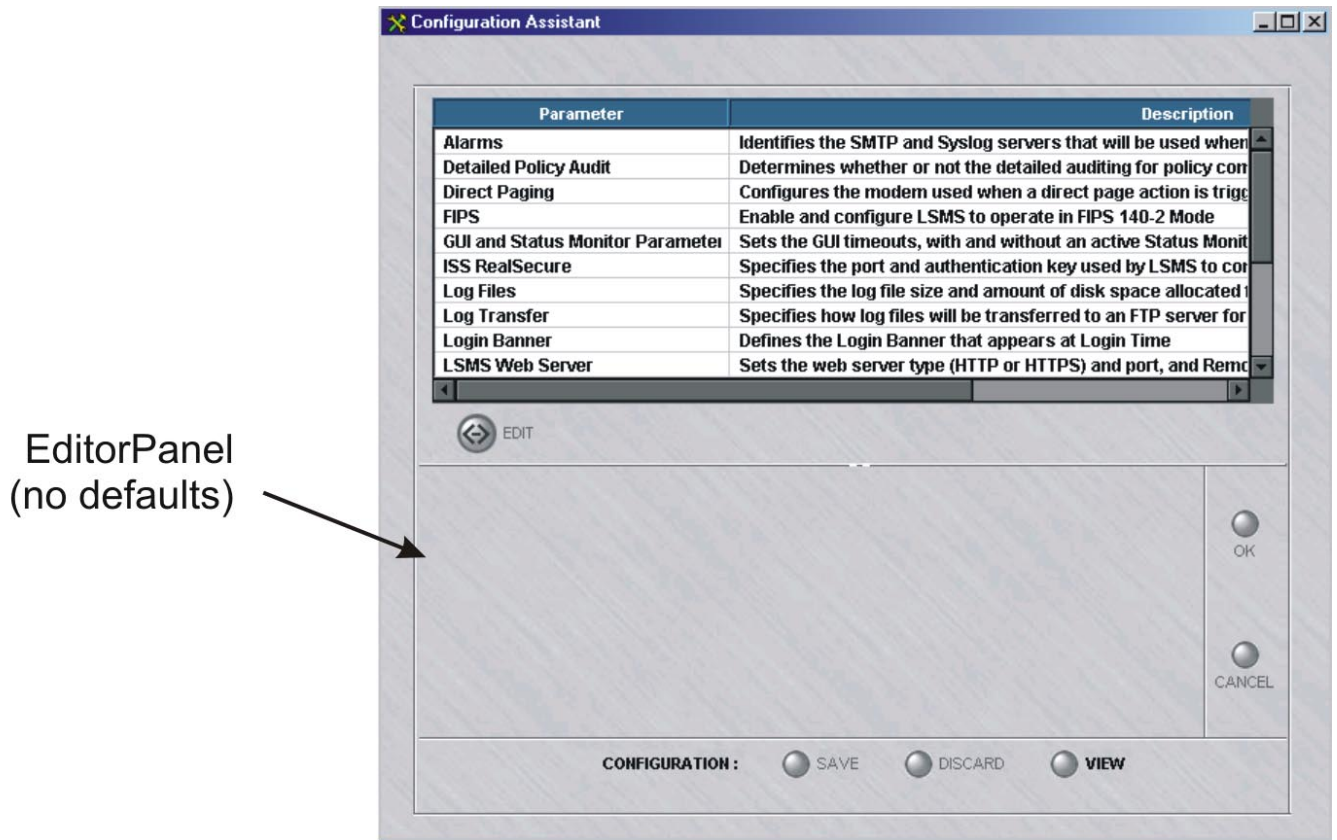
The Alarms parameters allow you to indicate the SMTP and Syslog servers that will be used when an alarm triggers an e-mail or Syslog message.

If you modify any of these parameters, you will have to stop and then restart all the LSMS services.

Default Values

There are no default values provided for the Alarms parameters. As [Figure 11-3, “Alarms Parameter Group”](#) (p. 11-9) shows, you have to enter the appropriate information for the first time.

Figure 11-3 Alarms Parameter Group



SMTP Host

The **SMTP Host** field identifies the Simple Mail Transport Protocol server(s) that will process e-mail requests generated by alarms that have been configured to trigger e-mail messages.

You can enter one SMTP host in this field, or you can enter multiple hosts, separating each host with a comma.

If DNS is accessible from the LSMS, you can enter the machine name(s) of the SMTP host(s). If not, you have to enter the IP address(es).

Account Name

The **Account Name** field identifies the sender of the e-mail message that is triggered by an alarm. This e-mail address is used to allow the LSMS to send e-mail to the specified server.

You can enter one account, or you can enter multiple accounts, separating each account with a comma.

Syslog Host

The **Syslog Host** field identifies the Syslog server(s) that will process the Syslog messages generated by alarms that have been configured to trigger Syslog messages.

You can enter one Syslog host in this field, or you can enter multiple hosts, separating each host with a comma.

If DNS is accessible from the LSMS, you can enter the machine name(s) of the Syslog host(s). If not, you have to enter the IP address(es).

Syslog Port

The **Syslog Port** field identifies the port on which the Syslog server will be listening to receive Syslog messages.

Typically, the Syslog server listens on port 514.

Trigger Alarm Code

Click this checkbox to turn the alarm code feature on. When it is on, you can include an alarm code in a rule, so that when the rule is invoked by an inbound or outbound session, an alarm is triggered (see *Chapter 2. Brick Zone Rulesets* in the *LSMS Policy Guide*).

Turning the alarm code feature on could cause a degradation in performance to process all of these alarms, depending on how many rules have alarm codes configured and how often traffic passing through the Brick matches those rules.

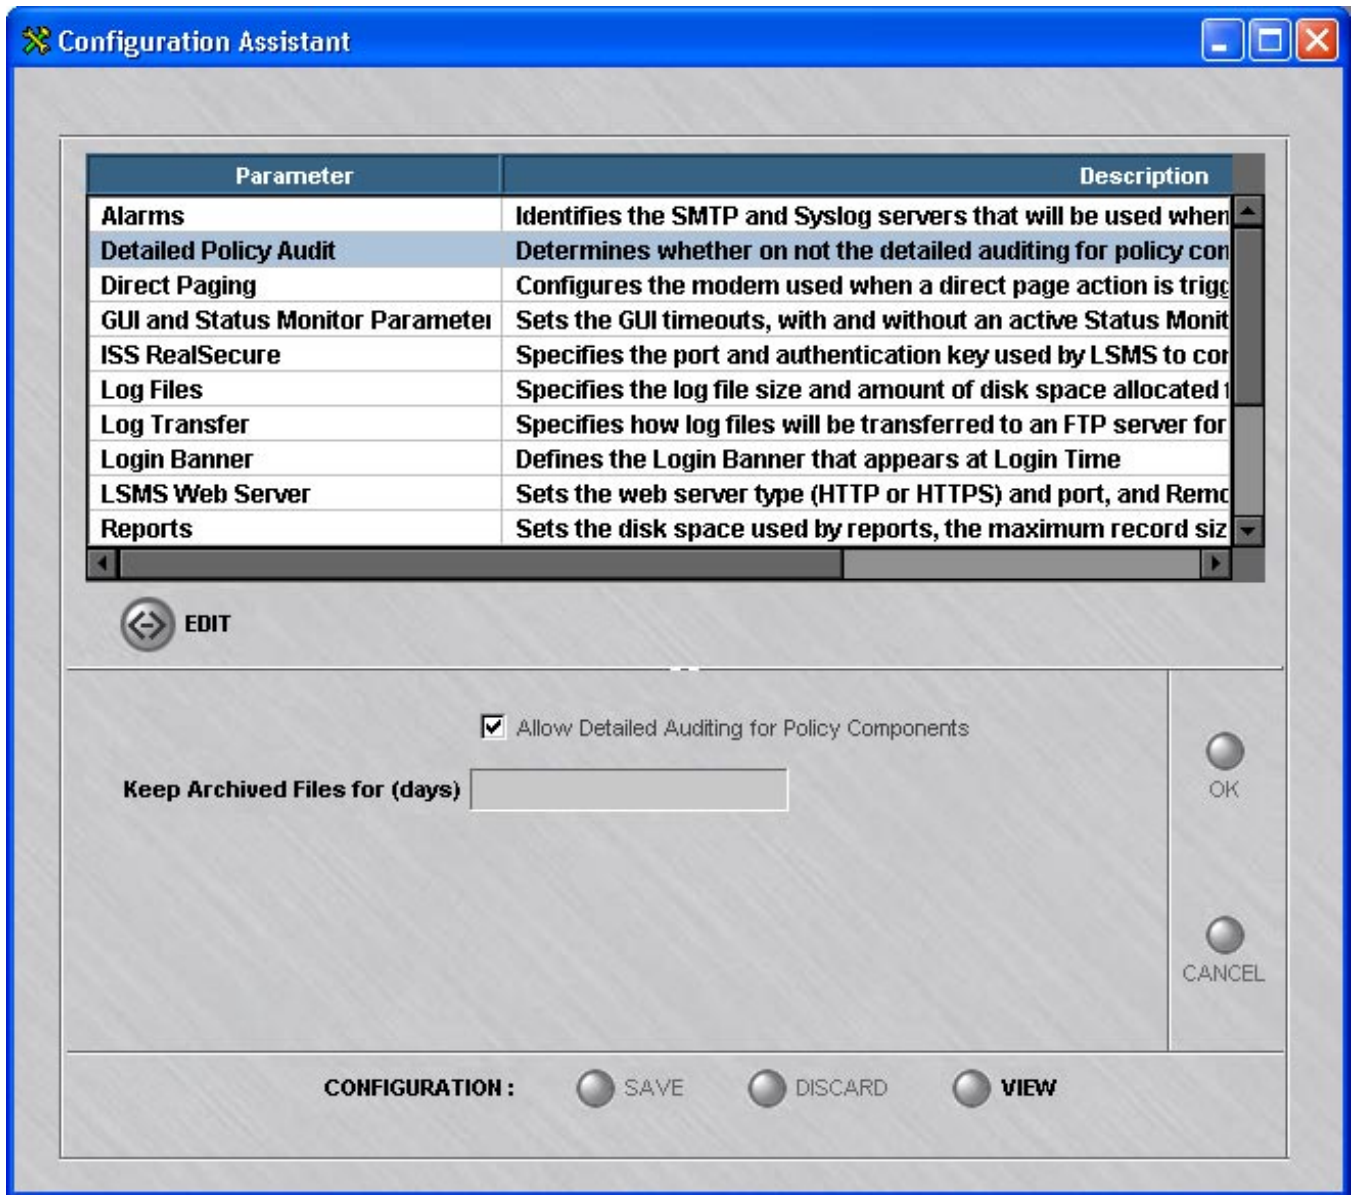


Detailed Policy Audit

Overview

The Detailed Policy Audit parameters allow you to determine whether you wish to preserve archive copies in the LSMS installation directory of all changes to Bricks, Brick Zone Rulesets, Host Groups, Service Groups, Application Filters and Dependency Masks. If necessary, the user may use the LSMS Command Line Interface to restore an archived version of a Brick or policy component.

Figure 11-4 Detailed Policy Audit Parameter Group



Allow Detailed Auditing Checkbox

To activate this feature, click this checkbox.

Keep Archived Files For (Days)

You may specify how long you wish to preserve your archive files. If no value is noted, the archive files are kept indefinitely.



Direct Paging

Direct paging parameters

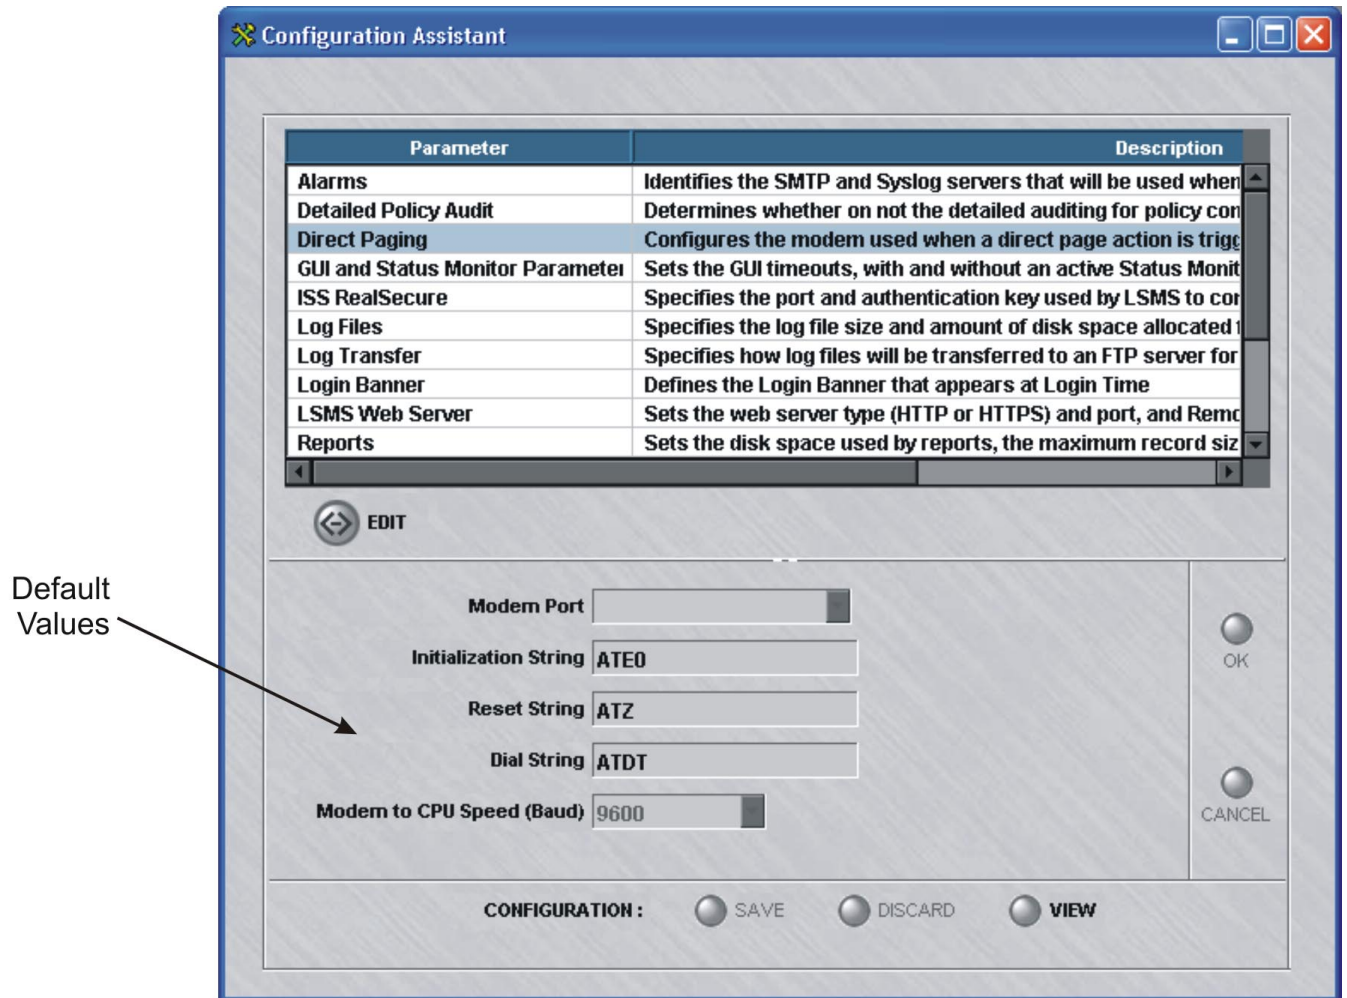
The Direct Paging parameters allow you to configure the modem that will be used when an alarm is set up to page an administrator.

The page is sent via a PSTN/modem-based connection that is made to a wireless pager service provider, such as SkyTel.

Default Values

Figure 10-5 shows the default values for the Direct Paging parameters.

Figure 11-5 Direct paging parameter Group



Modem Port

The **Modem Port** field specifies the physical port on the LSMS to which the modem is connected.

- For Windows, a COM port is used.
- For Solaris, the modem is a serial device in the UNIX file system. The default is */dev/cua/a*.

Initialization String

The **Initialization String** field contains a Hayes-modem string that is used to initialize the modem every time a direct page is transmitted.

Refer to the Hayes Command Set documentation of your modem manufacturer for the exact syntax of this string.

This string is not required and can be left blank or the default (ATE0) can be accepted.

Reset String

The **Reset String** field contains a string that is sent to the modem if the modem is unresponsive.

For example, if the modem does not respond to the LSMS's commands, the reset string will be sent.

This string is not required and can be left blank or the default (ATZ) can be accepted.

Dial String

The **Dial String** field contains a string that is sent to the modem to cause the modem to dial. It defaults to Hayes touch-tone dial command.

Examples include:

- ATDT *nnn,
Where nnn is an access code. Used for touch-tone dialing to dial-out.
- ATDP *9,
Used for pulse-signaling. The *9 is included to dial-out.

Modem Speed

The **Modem to CPU Speed (baud)** field specifies the speed at which the CPU of the LSMS talks to the modem.

This string is not required and can be left blank, or the default, 9600 bits per second, can be entered.

Important! You may need to set the **Modem to CPU Speed (baud)** field to match the speed of the paging service provider. This may be especially necessary if the value is very low (300 bps) or very high (56K).

FIPS

All communications between an LSMS and a Brick are encrypted. Many US government agencies require that the LSMS and devices communicate with each other in compliance with several Federal Information Processing Standards (FIPS):

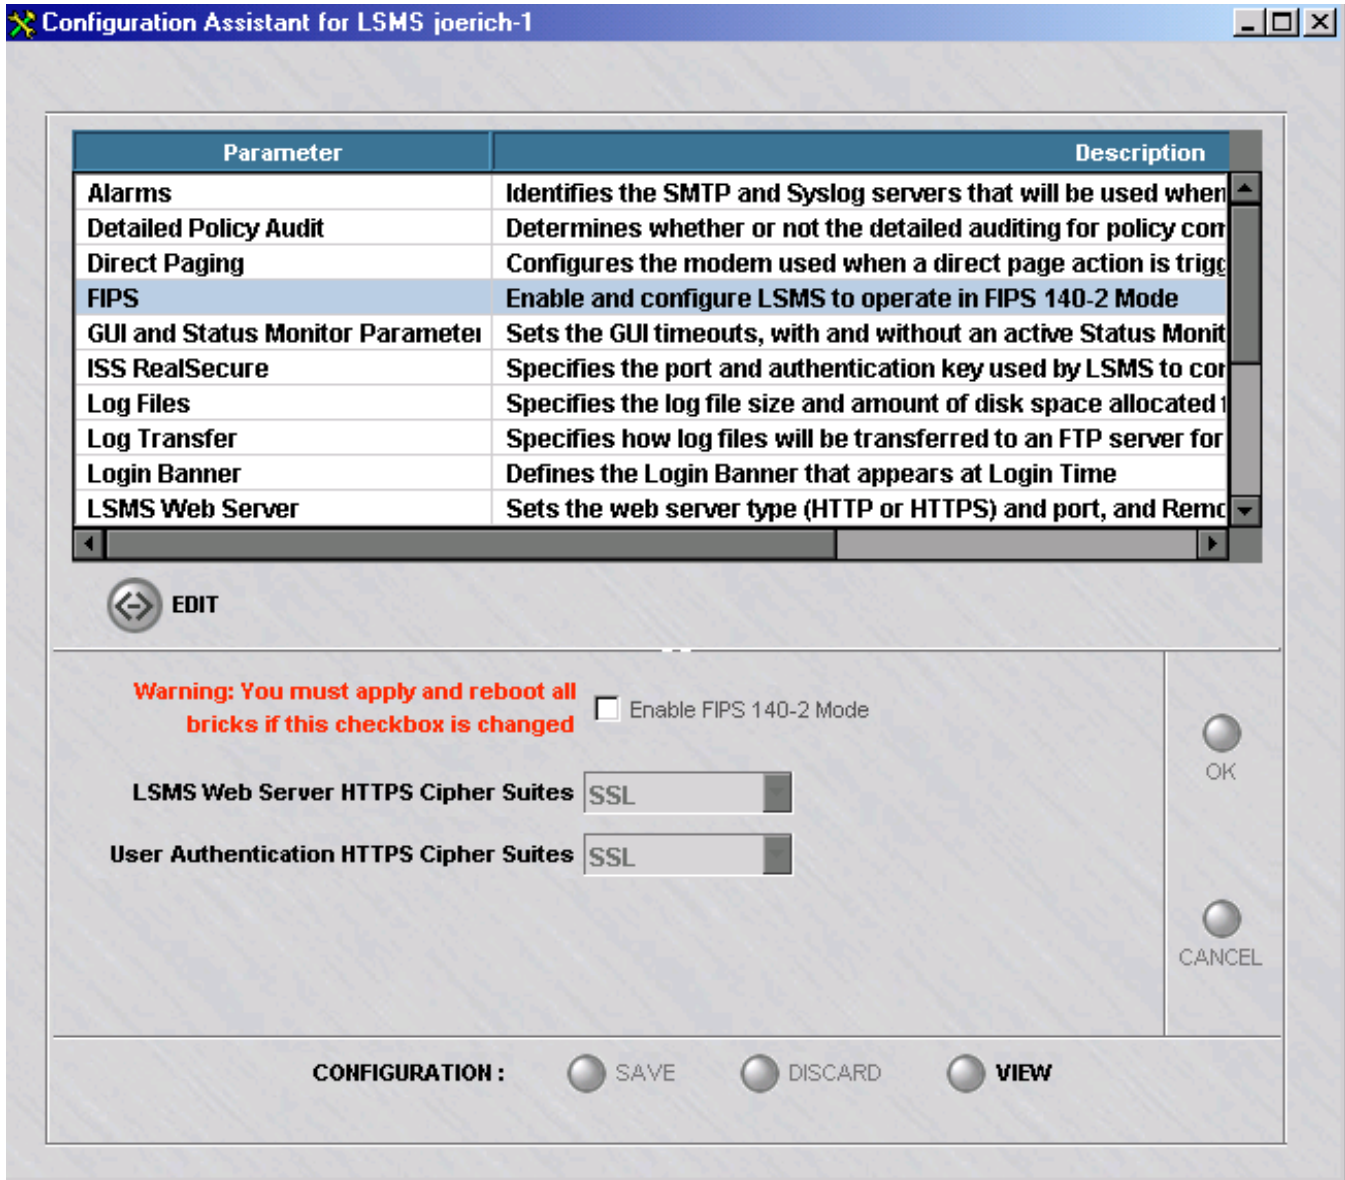
- FIPS Pub 180-1 (Secure Hash Algorithm, SHA-1)
- FIPS Pub 140-1 (Security Requirements for Cryptographic Modules)
- FIPS Pub 186-2 (Digital Signal Standard)

If you are a US government department, it is recommended that you click the **Enable FIPS 140-2 Mode** checkbox. When FIPS is enabled, it is enabled for the LSMS and all managed Bricks. If the value of this checkbox is changed, all Bricks must be applied and rebooted.

Default Values

Figure 11-6, “FIPS Parameters Group” (p. 11-16) shows the default values for the FIPS parameters.

Figure 11-6 FIPS Parameters Group



LSMS Web Server HTTPS Cipher Suites

If FIPS is enabled, the cipher suites for the LSMS Web Server HTTPS is set to the FIPS-compliant value **TLS** (Transport Layer Security). **SSL** (Secure Sockets Layer) is not FIPS-compliant, but is provided as a choice for customers that prefer to use this

option. When FIPS is enabled, HMAC MD5 is not available as a choice for ISAKMP and IPsec Proposal Authentication Type for Client tunnel Endpoints and LAN-LAN Tunnels.

User Authentication HTTPS Cipher Suites

If FIPS is enabled, the cipher suites for the User Authentication HTTPS is set to the FIPS-compliant value **TLS** (Transport Layer Security). **SSL** (Secure Sockets Layer) is not FIPS-compliant, but is provided as a choice for customers that prefer to use this option. When FIPS is enabled, HMAC MD5 is not available as a choice for ISAKMP and IPsec Proposal Authentication Type for Client tunnel Endpoints and LAN-LAN Tunnels.



GUI and Status Monitor Parameters

Overview

The GUI Timeout parameters allow you to set the LSMS's graphical user interface (GUI) timeout feature. Note that when the GUI times out, the LSMS session is still active, but in a "locked" state.

There is also a checkbox that will allow an option for the GUI to remain open if the status monitor is opened.

If you modify any of these parameters, you will have to stop and then restart all the LSMS services.

Default Values

Figure 11-7, "General Parameters" (p. 11-18) shows the default values for the GUI Timeout parameters.

Figure 11-7 General Parameters

The screenshot shows the 'Configuration Assistant for LSMS joerich-1' window. It features a table with the following parameters and descriptions:

Parameter	Description
Alarms	Identifies the SMTP and Syslog servers that will be used when an alarm triggers
Detailed Policy Audit	Determines whether or not the detailed auditing for policy components is to be enabled
Direct Paging	Configures the modem used when a direct page action is triggered by an alarm
FIPS	Enable and configure LSMS to operate in FIPS 140-2 Mode
GUI and Status Monitor Parameter	Sets the GUI timeouts, with and without an active Status Monitor
ISS RealSecure	Specifies the port and authentication key used by LSMS to communicate with the ISS RealSecure device
Log Files	Specifies the log file size and amount of disk space allocated for log files
Log Transfer	Specifies how log files will be transferred to an FTP server for creating reports
Login Banner	Defines the Login Banner that appears at Login Time
LSMS Web Server	Sets the web server type (HTTP or HTTPS) and port, and Remote Navigator default

Below the table is an 'EDIT' button. The main configuration area includes:

- Directory for Log Files: C:\users\isms\lwf
- Simultaneous LoggerViewers (Per log): 10
- A table for log file settings:

	LogFile Rollover Interval (mins)	Max LogFile Size (Mb)	Max Allocation Size (Mb)	Halt Traffic if Log Full
Session		10	1000	<input type="checkbox"/>
Admin Events		1	100	<input type="checkbox"/>
Proactive Monitor		10	200	<input type="checkbox"/>
User Auth		1	100	<input type="checkbox"/>
VPN		1	100	<input type="checkbox"/>

At the bottom, there are 'CONFIGURATION :', 'SAVE', 'DISCARD', and 'VIEW' buttons. An arrow from the text 'Default Values' points to the 'Admin Events' row in the table above.

GUI Timeouts

The GUI timeout feature determines how long an administrator can remain logged into the LSMS without performing any GUI-related activities in the Navigator window. If GUI-related activities are not detected within the specified time period, the administrator's session will be "locked". To unlock the session, simply re-enter the admin's password.

Idle Time

The **Idle Time** field determines the length of time that must elapse without GUI activity before an administrator's session is locked.

The default is 3600 seconds (one hour).

With Active Status Monitor

The purpose of the **With Active Status Monitor** field is to allow you to set a second GUI timeout that only applies when the Status Monitor is active.

This allows an administrator to remain logged in for a longer period of time when performing only monitoring duties, such as when keeping an eye on the Status Monitor but not initiating any GUI-related activities.

For this reason, the value of this field should always be equal to or greater than the value in the **Idle Time** field. The default is 3600 seconds.



Log Files

Log files parameters

The Log Files parameters allow you to determine four sets of parameters:

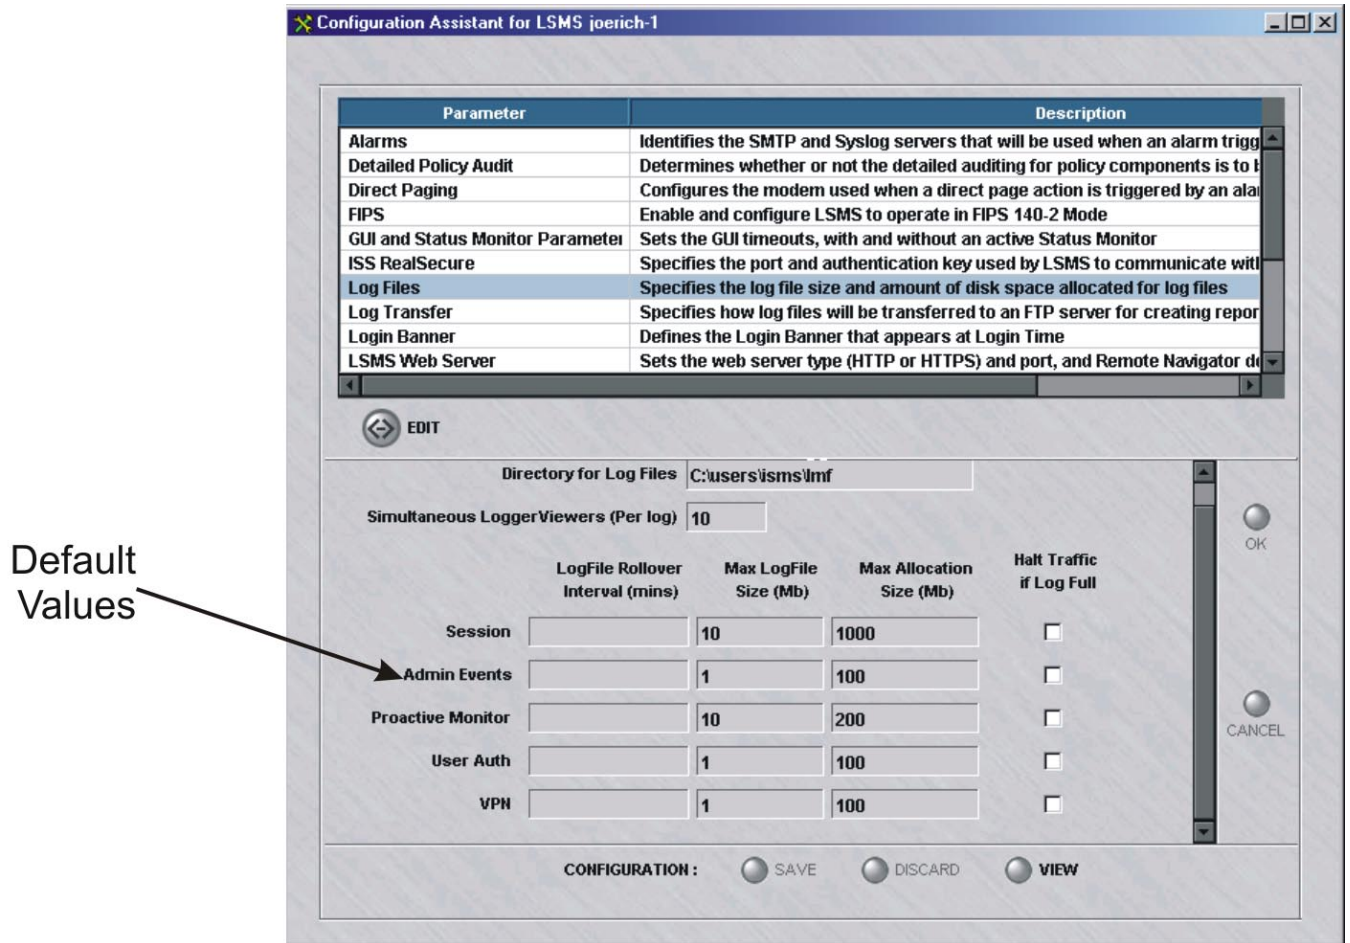
- The directory used to store LSMS log files
- The maximum number of users that can simultaneously access the LSMS Log Viewer
- The log file rollover interval, based on a specified time interval or filesize constraint
- The maximum size of the various log files and the amount of disk space to be allocated for log files. These parameters also allow you to indicate whether or not you want all traffic through the Brick halted when the disk space allocated for log files is exhausted.

If you modify any of these parameters, you will have to stop and then restart all the LSMS services.

Default Values

Figure 11-8, “Log Files Parameters” (p. 11-21) shows the default values for the Log Files parameters.

Figure 11-8 Log Files Parameters



Directory for Log Files

If you wish to store all of the LSMS log files on the LSMS but not in the default LSMS installation directory (for Windows, \users\isms\lmf and for Solaris /opt/isms/lmf), specify the new directory path here.

If you would like to store the LSMS log files on a remote machine, please review the *Log Transfer* parameter later in this chapter.

Simultaneous Logger Viewers (Per log)

Administrators may review real time system activity through the LSMS Log Viewer, whether they are logged into the LSMS locally or remotely. In order to preserve bandwidth, you may elect to restrict the number of simultaneous connections to any of the four logs.

By default, this value is set to 10 users.

Logfile Rollover Interval (min)

If an interval is specified (in minutes), log files will rollover to a new file name after the elapsed time interval or the maximum log file size is met, whichever comes first.

Maximum Log File Size

The LSMS maintains five logs, and for each of these logs, the **Max LogFile Size** fields determine how large each log file will be permitted to grow before a new file is begun.

The old log files are saved until the maximum disk space allocated for that log is reached (see below). At that time, the LSMS begins to delete the files to free up space, with the oldest files deleted first.

The following shows each log and its default maximum file size:

Log File	Maximum File Size
Session Log	10 Mb
Admin Events Log	1 Mb
Proactive Monitor Log	10 Mb
VPN Log	1 Mb
User Auth Log	1 Mb

Maximum Disk Allocation

The **Max Allocation** field determines the amount of disk space that will be allocated to each log.

The following shows each log and its default disk allocation:

Log	Maximum Disk Allocation
Session Log	1000 Mb
Admin Events Log	100 Mb
Proactive Monitor Log	200 Mb

Log	Maximum Disk Allocation
VPN Log	100 Mb
User Auth Log	100 Mb

If you have upgraded the LSMS software from an earlier version, the default allocations will be based on the allocations you were using in the earlier version.

Halt All Traffic

A checkbox labeled **Halt Traffic if Log Full** appears to the right of each log.

By default, the box is not checked. This means that when the disk space allocated to this log is exhausted, the LSMS will begin to delete old log files to create space for new ones, and the Brick will continue to pass traffic.

However, if you do not want the old log files deleted, you can click the box to check it. Then, when the disk allocation is exhausted, no files are deleted, and the Brick halts all traffic.

Important! If you click the checkbox to halt all traffic, you should also make sure the checkbox labeled **Halt All Traffic if Audit Fails** on the Brick Editor is also checked.

The Brick Editor is the window you use to configure a Brick (see *Chapter 3. Configuring and Activating a Brick*).



Log Transfer

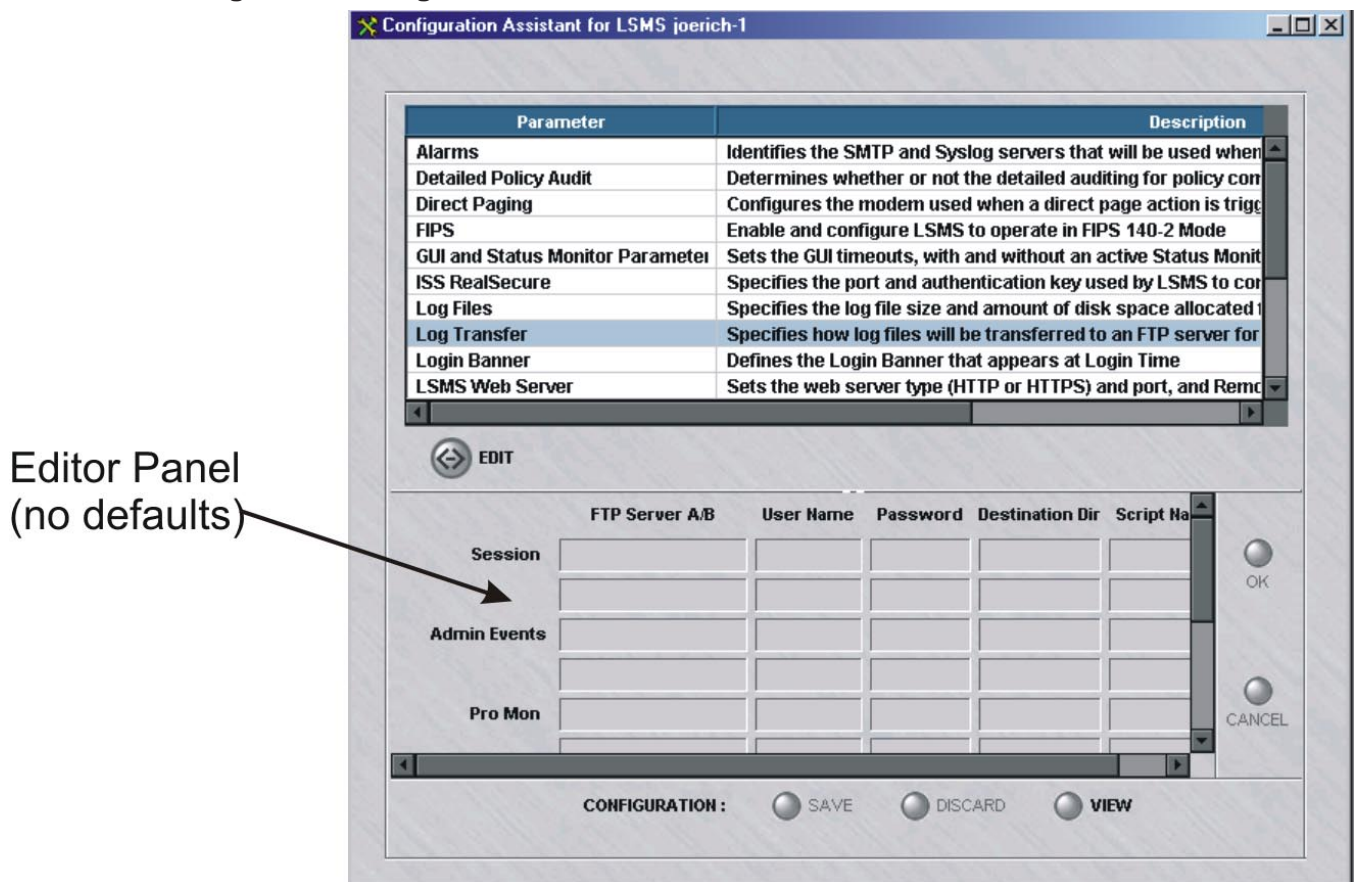
Log transfer parameters

The Log Transfer parameters allow you to specify where and how the various log files will be transferred for use in creating reports.

Default Values

There are no default values provided for the Log Transfer parameters. As [Figure 11-9](#), “Log Transfer Parameters” (p. 11-24) shows, you have to enter the appropriate information for the first time.

Figure 11-9 Log Transfer Parameters



FTP Server

For each of the five logs, the **FTP Server A/B** field indicates the host(s) that will receive the transferred log files. The host(s) must be an FTP server. Up to two hosts per log can be entered.

If DNS is accessible from the LSMS, you can enter the machine name of the host. If not, you have to enter the IP address.

User Name

The user name must be a valid user account on the FTP server.

Password

The password must be the valid password for the user account entered in the **User Name** field.

Important! Even if you log on as anonymous, you must still enter a password.

Destination Directory

The destination directory is the directory on the FTP server in which the log files will be placed.

Enter either an absolute path name or a path name that is relative to the user's home directory.

Script Name

Scripts can be created to pre-process the log files before they are transferred to the FTP server. You could, for example, create a script to compress the files to make the transfer more efficient.

If you will be making use of a script, enter the path and name of the file that contains the script for each log in the **Script Name** field.

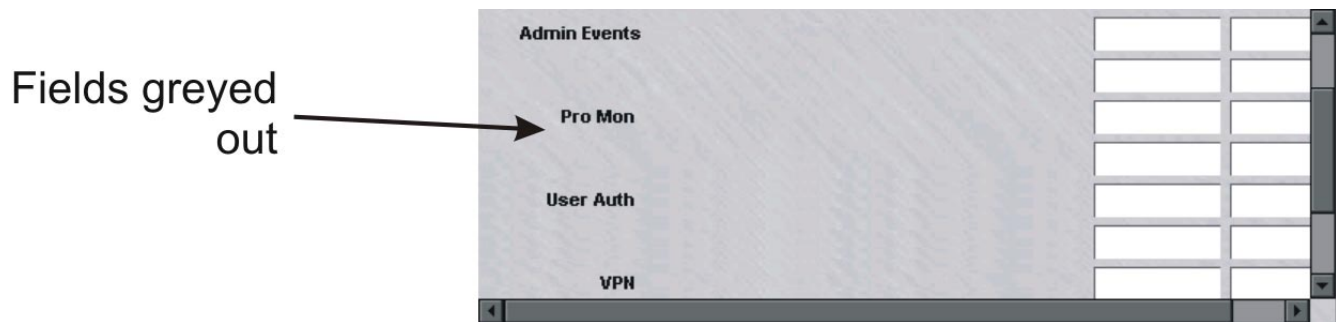
Same FTP Server

If you want to send all the logs to the same FTP server and account, the Configuration Assistant provides an easy way to do this.

Click the checkbox labeled **Send all logs to same FTP server with same user/password**.

This will cause the **FTP Server A/B**, **User Name**, and **Password** fields for the last four logs to become greyed-out, as shown in Figure 10-10.

Figure 11-10 Same FTP Server



Enter an FTP server, user name, and password for the Session log. The other three logs will automatically be sent to the same FTP server, using the same user name and password.

Note that you must enter a different destination directory and script for each of the four logs.

Delete Files

If you want to delete the log files after they are transferred, click the checkbox labeled **Delete files after transfer**.

Login Banner

Overview

When this feature is enabled, a login banner is displayed immediately after an administrator successfully logs into the LSMS Navigator. Refer to Figure 10-12 for the default notice (computer misuse act notice). The feature is enabled by placing a check mark in the check box labeld **Enable Login Banner** on the bottom of the screen. The notice can be edited by clicking on **Edit** and editing the text in the edit box (See Figure 11-11, “Login Banner Parameters” (p. 11-27)).

Default Values

Figure 11-11 Login Banner Parameters

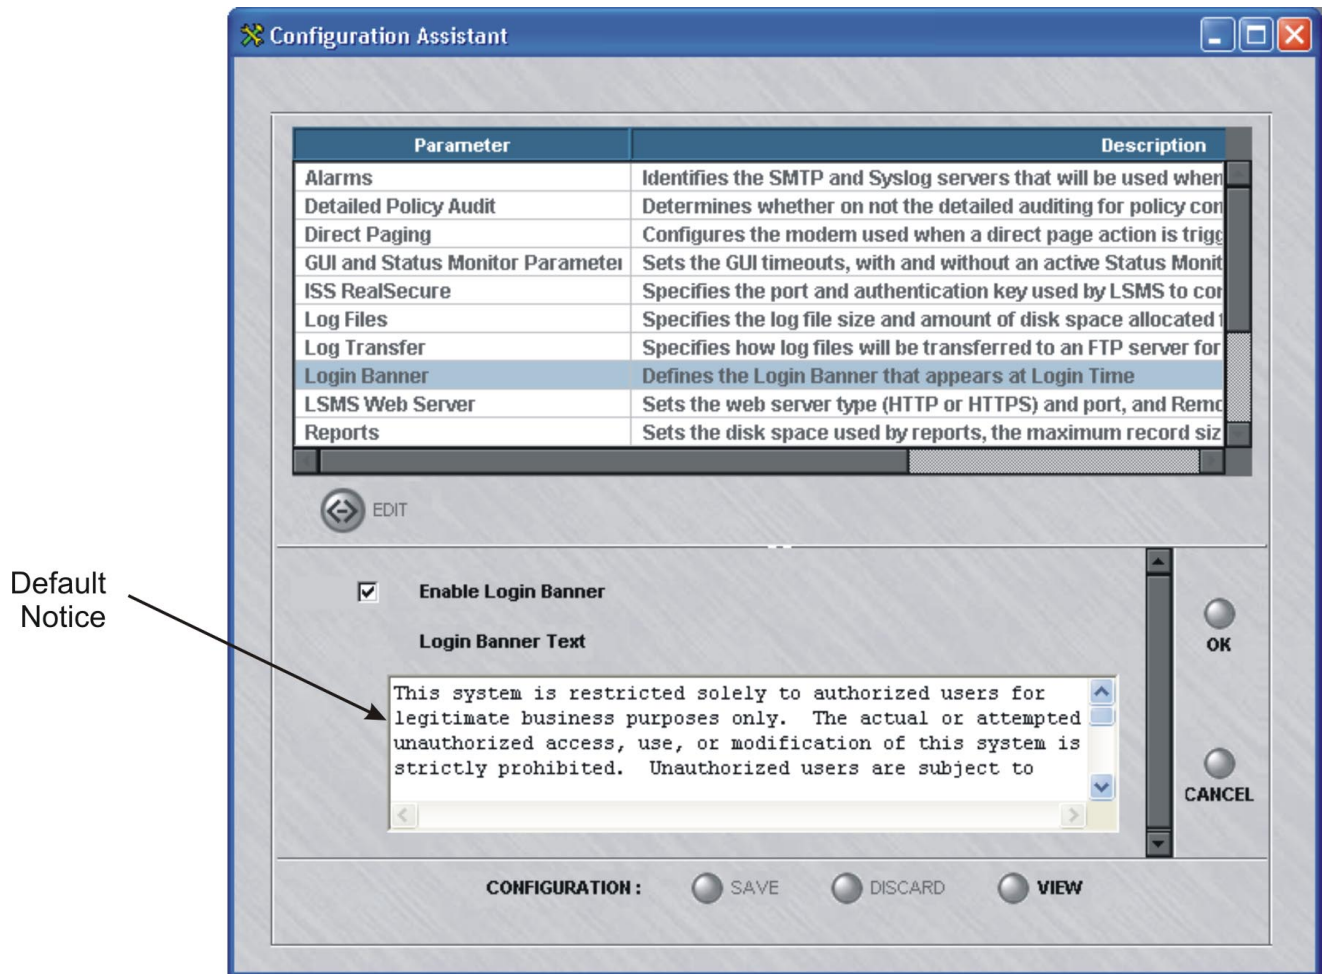
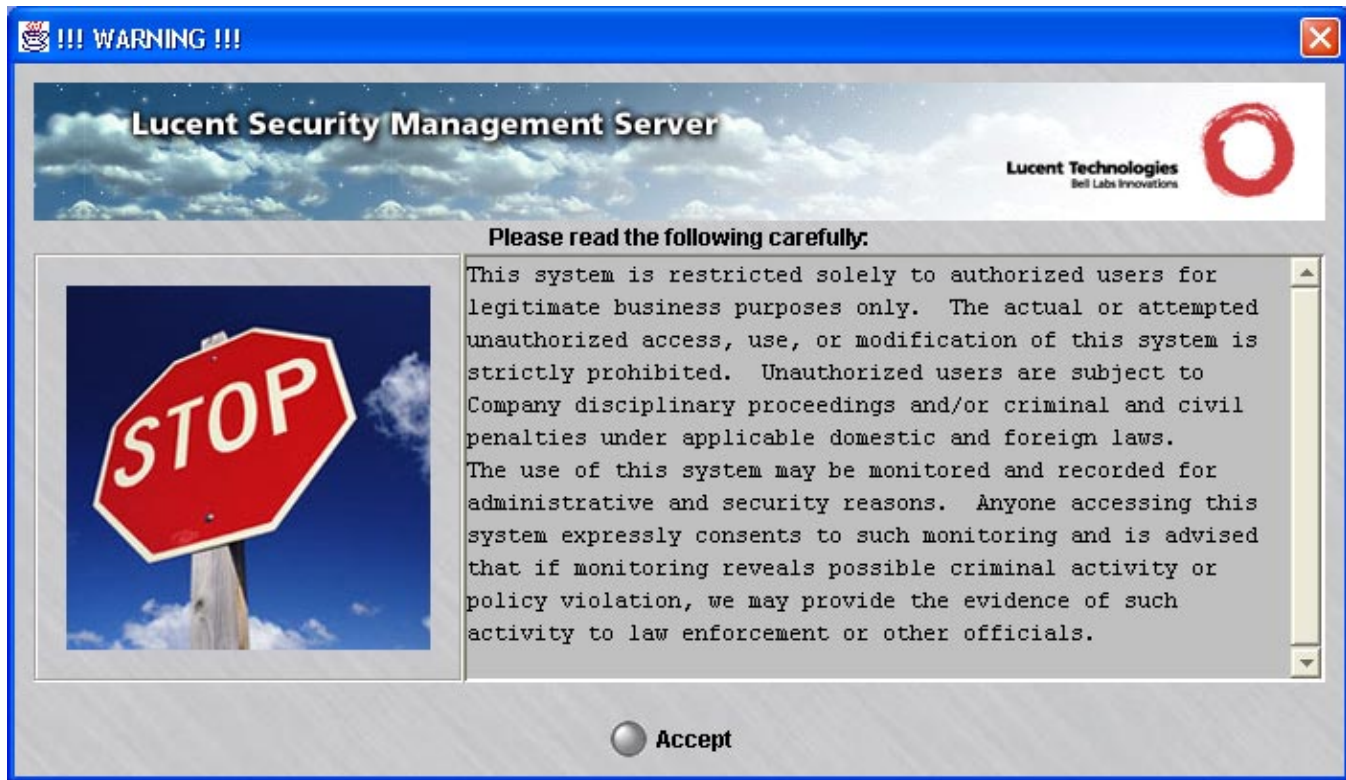


Figure 11-12 Login Banner Parameters



□

LSMS Web Server

Web server parameters

This LSMS Web Server parameters allow you to indicate the type of Web server (HTTP or HTTPS) running in the LSMS and the port on which it is listening.

If you have obtained a digital certificate from Verisign or another certificate authority, the Web server should be HTTPS; if you have not, the Web server should be HTTP.

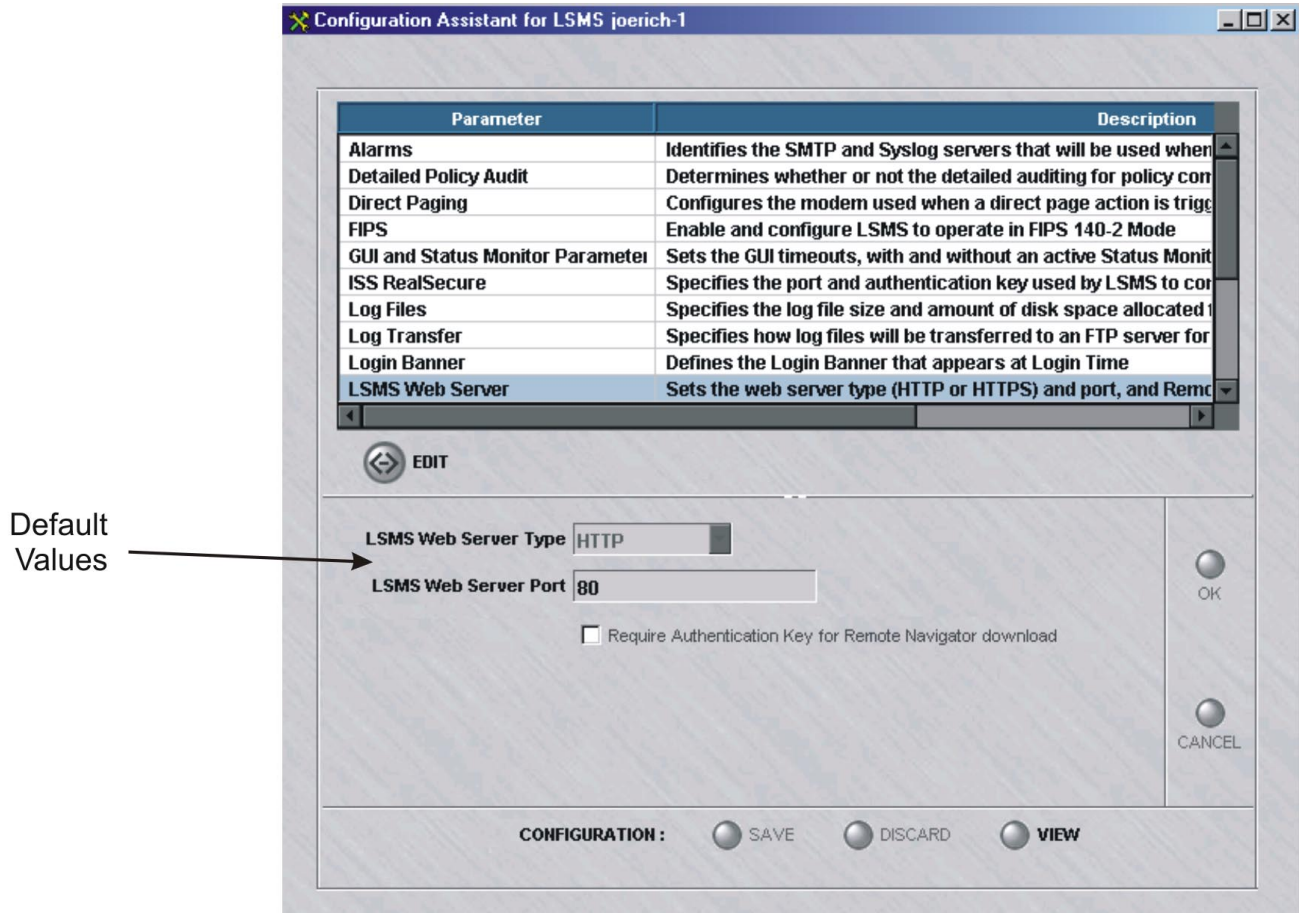
If you modify any of these parameters, you will have to stop and then restart all the LSMS services.

Default Values

The default values for these parameters are determined during the installation of the LSMS software.

Figure 11-13, “LSMS Web Server Parameters” (p. 11-30) shows the default values. The default is a HTTP Web server listening on port 80, the standard HTTP port.

Figure 11-13 LSMS Web Server Parameters



LSMS Web Server Type

The type of web server will either be HTTP or HTTPS, depending on the type that was entered during the installation of the LSMS software.

If you will be accessing the LSMS remotely, you need to consider the security of the LSMS web server. The Remote Navigator provides secure, encrypted access to the LSMS.

The Remote Navigator does not use the LSMS web server during login, or to perform LSMS administration. However, the web server is used to view reports, online help, and download Brick floppy packages and new Remote Navigator software. Of these, the only one that may contain sensitive information is reports.

If you need to view reports remotely over the public internet, and you will not be using the Lucent IPSec Client to establish a secure encrypted tunnel, it is strongly recommend that you set up the LSMS web server for HTTPS.

LSMS Web Server Port

The port number will be the port that was entered during the installation of the LSMS software.

Port 443 is standard for HTTPS and port 80 is standard for HTTP. However, non-standard ports can be entered.



Reports

Reports parameters

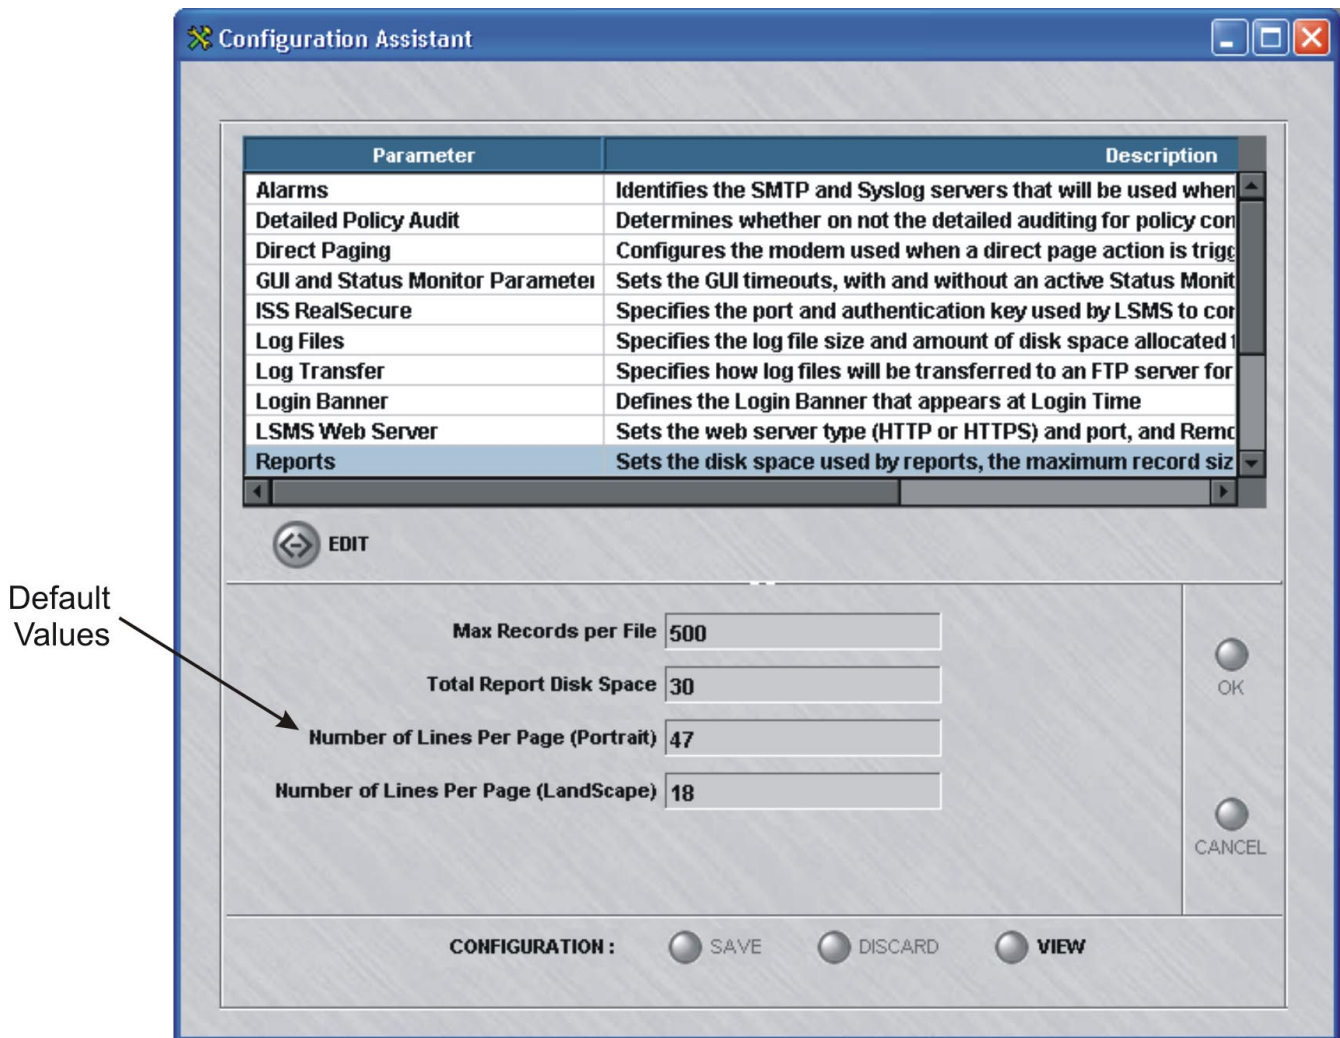
The Reports parameters allow you to specify the maximum number of records in a report file and the total amount of disk space to be allocated for reports.

They also allow you to set the number of lines per page in both portrait and landscape formats.

Default Values

Figure 11-14, “Reports Parameters” (p. 11-32) shows the default values for the Reports parameters.

Figure 11-14 Reports Parameters



Maximum Records Per File

The **Max Records per File** field specifies the maximum number of records permitted in any HTML report file.

The default is 500. The minimum value is 50.

Total Report Disk Space

The **Total Report Disk Space** field specifies the amount of disk space (in megabytes) to be allocated for report output.

The default is 30 megabytes.

Lines Per Page

The **Number of Lines per Page (Portrait)** and **Number of Lines per Page (Landscape)** fields specify the number of lines per printed page in each format.

The default for portrait is 47 and for landscape 18.



SNMP Agent

SNMP agent parameter

The SNMP Agent parameter allows you to change the port on which the SNMP agent listens. This port was initially set during installation of the LSMS.

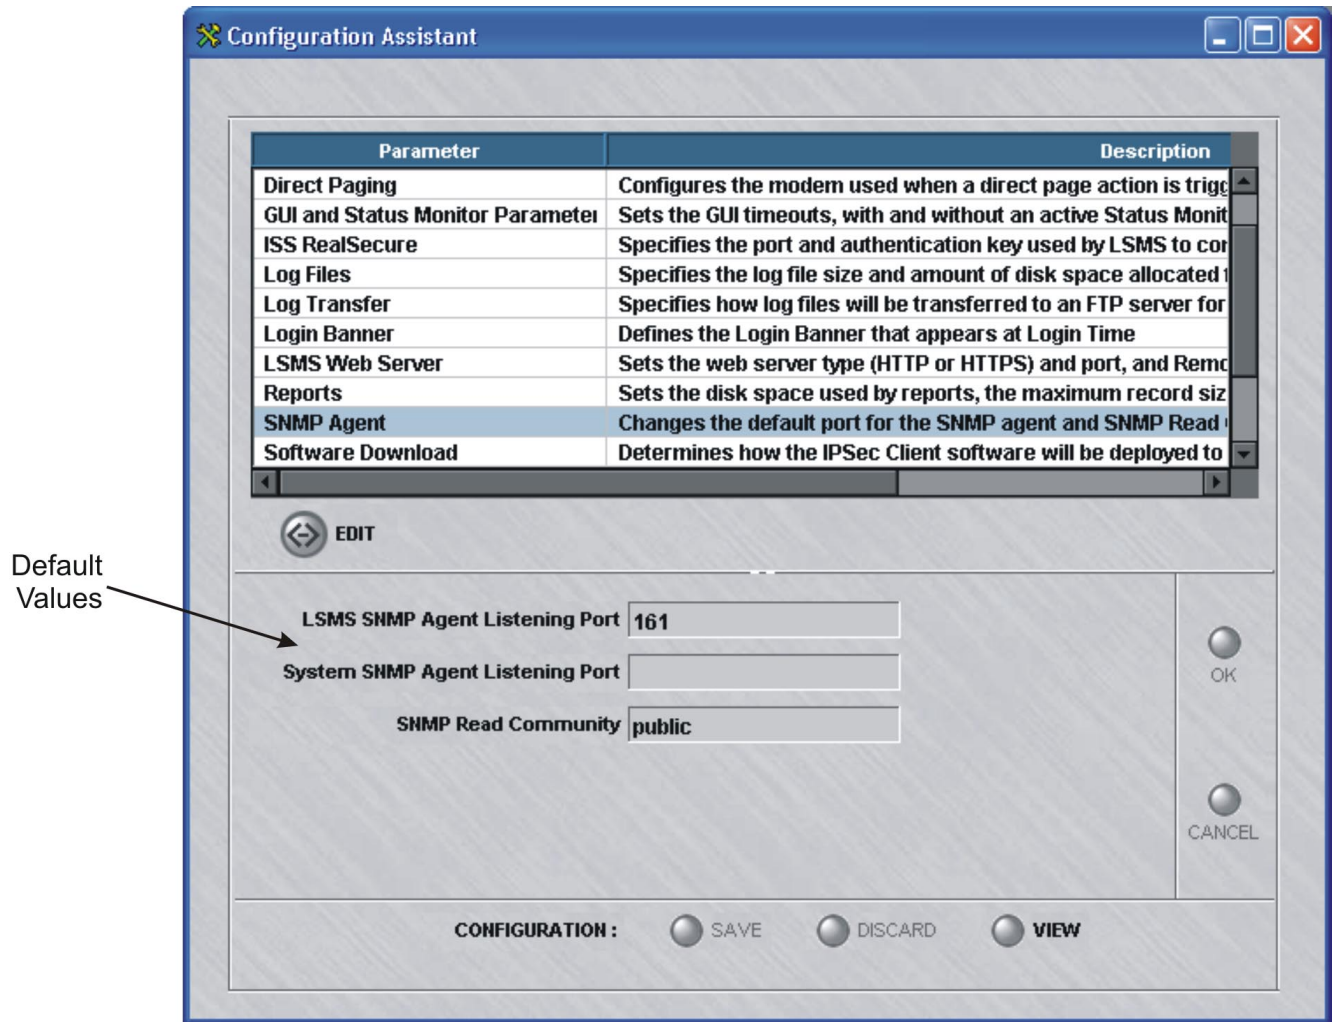
The SNMP agent runs continuously on the LSMS and responds to queries initiated by a Network Management Station (NMS) administrator. The LSMS returns data about the health and state of the Bricks and the LSMS software itself to the NMS after being polled, so that the NMS is able to monitor trends in the network (see *Appendix A. SNMP* in the *Reports, Alarms, and Logs Guide*).

If you modify this parameter, you will have to stop and then restart all the LSMS services.

Default Values

Figure 11-15, “SNMP Agent Parameters” (p. 11-35) shows the default values for the SNMP Agent parameters.

Figure 11-15 SNMP Agent Parameters



LSMS SNMP Agent Listening Port

The port the LSMS SNMP agent is listening on is set to 161 if you selected the default value during LSMS installation. If you selected a different port during installation, that port will appear as the default.

You can change the port in the **LSMS SNMP Agent Listening Port** field. The port that you choose must be different from the value specified for the System SNMP Agent Listening Port.

System SNMP Agent Listening Port

If you have a third party SNMP package running on your LSMS, all SNMP GET requests for objects that are not defined in the LSMS MIB will be forwarded to the port that you designate here.

You can change the port in the **System SNMP Agent Listening Port** field. The port that you choose must be different from the value specified for the LSMS SNMP Agent Listening Port.

SNMP Read Community

By default, the **SNMP Read Community** field is set to **public**. This value is used to control the SNMP network management station's access to the LSMS.



Software Download

IPSec client CD-ROM

The IPSec client CD-ROM provides software that can be directly installed on a client user's laptop or distributed to many IPSec client users across the public Internet.

When distributing the software across the Internet, the software on the CD-ROM must be copied to either the LSMS, an FTP server, or a Web server. If you suspect that simultaneous downloads from the LSMS may degrade performance, you may want to consider distributing the software from an FTP server or Web server instead.

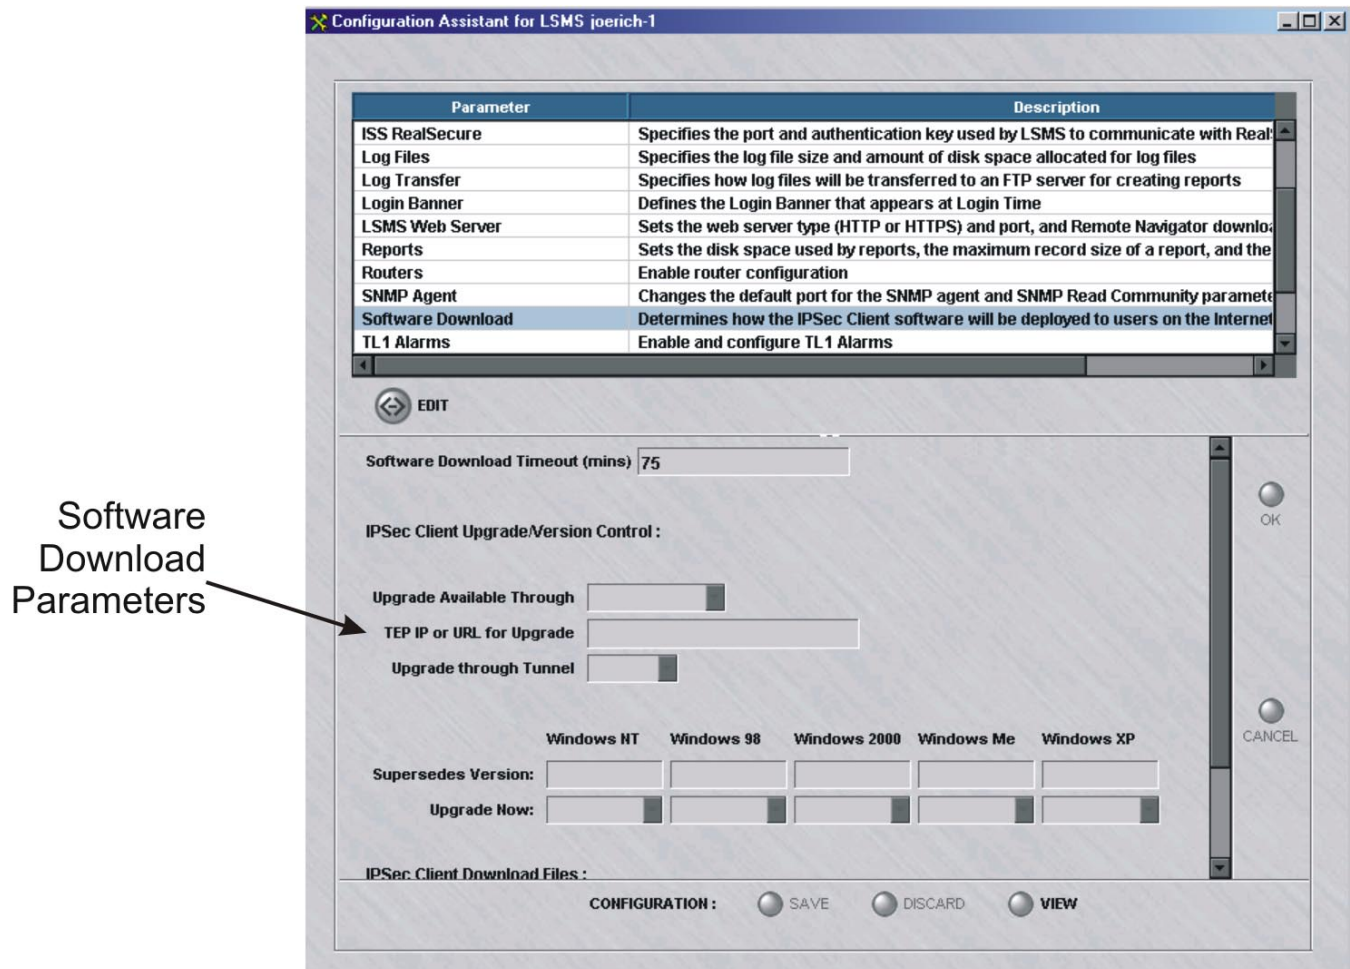
If you modify any of these parameters, you may have to stop and then restart all the LSMS services.

Default Values

There is one default value provided for the Log Transfer parameters (software download timeout). As [Figure 11-16, "Software Download Parameters" \(p. 11-38\)](#)

shows, for the rest of the parameters you have to enter the appropriate information for the first time.

Figure 11-16 Software Download Parameters



LSMS Software Download

When distributing the software from the LSMS that is "behind a Brick", rules need to be applied to the interfaces of the Bricks through which the software must pass.

Examples of the required rules are provided in the *vpnzone* Brick zone ruleset. Rules #305 and #306 can be used as a template when the software is distributed through a tunnel. Rule #310 can be used as a template when the software is distributed outside of a tunnel.

FTP/Web Server Download

To provide a level of security when distributing software from publicly accessible FTP server or Web servers, you should force the IPSec client users to be authenticated.

When distributing the software from an FTP server or Web server that is "behind a Brick", a rule needs to be applied to the interfaces of the Bricks between the FTP server or Web server and the client. For example, the rule could be defined as follows:

- Source = All_Users
- Destination = <FTP/Web Server>
- Service = FTP
- Action = Pass

Before You Begin

To set up the IPsec client software so it can be distributed from the LSMS, an FTP server, or a Web server, do the following:

- 1 Insert the IPsec client CD-ROM into the disk drive on the LSMS.
- 2 Copy the *clientSWversion.info* file into the <LSMS Root>/fac/docroot directory on the LSMS.
- 3 If distributing the software from the LSMS, copy the *6.0.xxx.exe* file (where *xxx* is the build number) and the *releasenotes.txt* file into the <LSMS Root>/fac/docroot directory on the LSMS.

For example, if you accepted the default directory during installation, on a Windows LSMS, copy the above files to

c:\users\isms\lmf\fac\docroot.
- 4 If distributing the software from an FTP server or Web server, insert the IPsec client CD-ROM into the disk drive on the FTP server or Web server and copy the *6.0.xxx.exe* file (where *xxx* is the build number) and the *releasenotes.txt* file into a directory.

END OF STEPS

Software Download Timeout (minutes)

The **Software Download Timeout** field sets the amount of time that is allocated for downloading the software to the IPsec client host. The timer starts after the IPsec client user has successfully been authenticated and downloading the software has been initiated.

If the period of time expires and the software download is not complete, the download will terminate. The IPSec client user will then have to be re-authenticated.

The default timeout is 75 minutes. This should provide an ample amount of time for downloading the software (~7 MB) using a 14.4 kbps modem connection. However, if necessary, you can increase the timeout.

Important! This timeout value overrides the authentication timeout value (see *Chapter 9. User Authentication* in the *LSMS Policy Guide* for details).

IPSec Client Upgrade/Version Control

Specify values for the fields, as the following table explains:

Field	Selection	Description
Upgrade Available Through	Specific TEP	Applicable for both methods of download. Implies the software will be downloaded from a specific tunnel endpoint (TEP) associated with a Brick interface. Enter the IP address of the TEP in the TEP IP or URL for Upgrade field. This address was entered when the ruleset containing the rules establishing the tunnel is assigned to a Brick interface (see <i>Chapter 4. Configuring Brick Ports</i>).
	Any TEP	Applicable for both methods of download. Implies the software will be downloaded from the currently active tunnel endpoint being used by the client. In this case, entering the IP address of the TEP in the TEP IP or URL for Upgrade field is not required.
	URL	Applies only when the software will be downloaded from an FTP server or Web server. This requires entering the full pathname (including the file name) in the TEP IP or URL for Upgrade field as described below.

Field	Selection	Description
TEP IP or URL for Upgrade	If Specific TEP was selected...	Enter the IP address of the TEP.
	If Any TEP was selected...	This field is not editable and "any" is displayed.
	If URL was selected...	Then enter the full pathname of the Web server or FTP server (including the file name). For example: <i>ftp://105.39.90.100/Upgrade/6.0.410.exe</i> <i>http://www.your_company.com/ipsec_client/upgrade.html</i>
Upgrade Through Tunnel	Yes	If the software can be upgraded through a tunnel, the IPsec client user is informed of the availability of new software at the end of tunnel establishment.
	No	If the software cannot be upgraded through a tunnel, the IPsec client user is informed of the availability of new software at the end of disabling the tunnel.

Supersedes Version

For each platform, the version of software that can be superseded by the new software on the LSMS, an FTP server, or Web Server is displayed.

This indicates that Windows NT and Windows 2000 clients who have Version 6.0.39 or earlier would benefit from this upgrade.

The versions of software that are displayed are pulled from the *clientSWversion.info* file that was copied to the LSMS as explained in "Before You Begin".

For each platform, decide if you want client users to be notified that the new software is available. Choose either:

- **Yes**
After a user has successfully been authenticated, they will immediately receive notification when new software is available.
- **No**
After a user has successfully been authenticated, they will NOT receive notification when new software is available.

IPSEC Client Download Files

In these fields, enter the name and size of the software to be downloaded and optionally, the release notes that may accompany it.

These fields only require values if the software will be distributed from the LSMS.

- **Client Software Name/URL**

Note that:

- In the **File Name/URL** field, enter the name of the software file to be downloaded from the LSMS. For example, `6.0.410.exe`.

This file must reside in the `<LSMS Root>/fac/docroot` directory, for example, `c:\users\isms\lms\fac\docroot`, and must match the name that was copied from the CD-ROM as explained in "Before You Begin".

- In the **File Size(bytes)** field, enter the size of the file.
For example, 800000 indicates that the file size is approximately 8 MB.

- **Release Notes File Name/URL**

Note that:

- In the **File Name/URL** field, optionally enter the name of the Release Notes file that will accompany the software. For example, `releasenotes.txt`.

This file must reside in the `<LSMS Root>/fac/docroot` directory, for example, `c:\users\isms\lms\fac\docroot`, and must match the name that was copied from the CD-ROM as explained in "Before You Begin".

Even though this is optional, it is recommended to provide Release Notes so that users are aware of the new features and bug fixes that are included in the upgrade software.

- In the **File Size(bytes)** field, enter the size of the file.
For example, 1000 indicates that the file size is approximately 1 K.



TL1 Alarms

Overview

The TL1 Alarms parameters allow you to enable and configure the Transaction Language 1 alarm reporting interface.

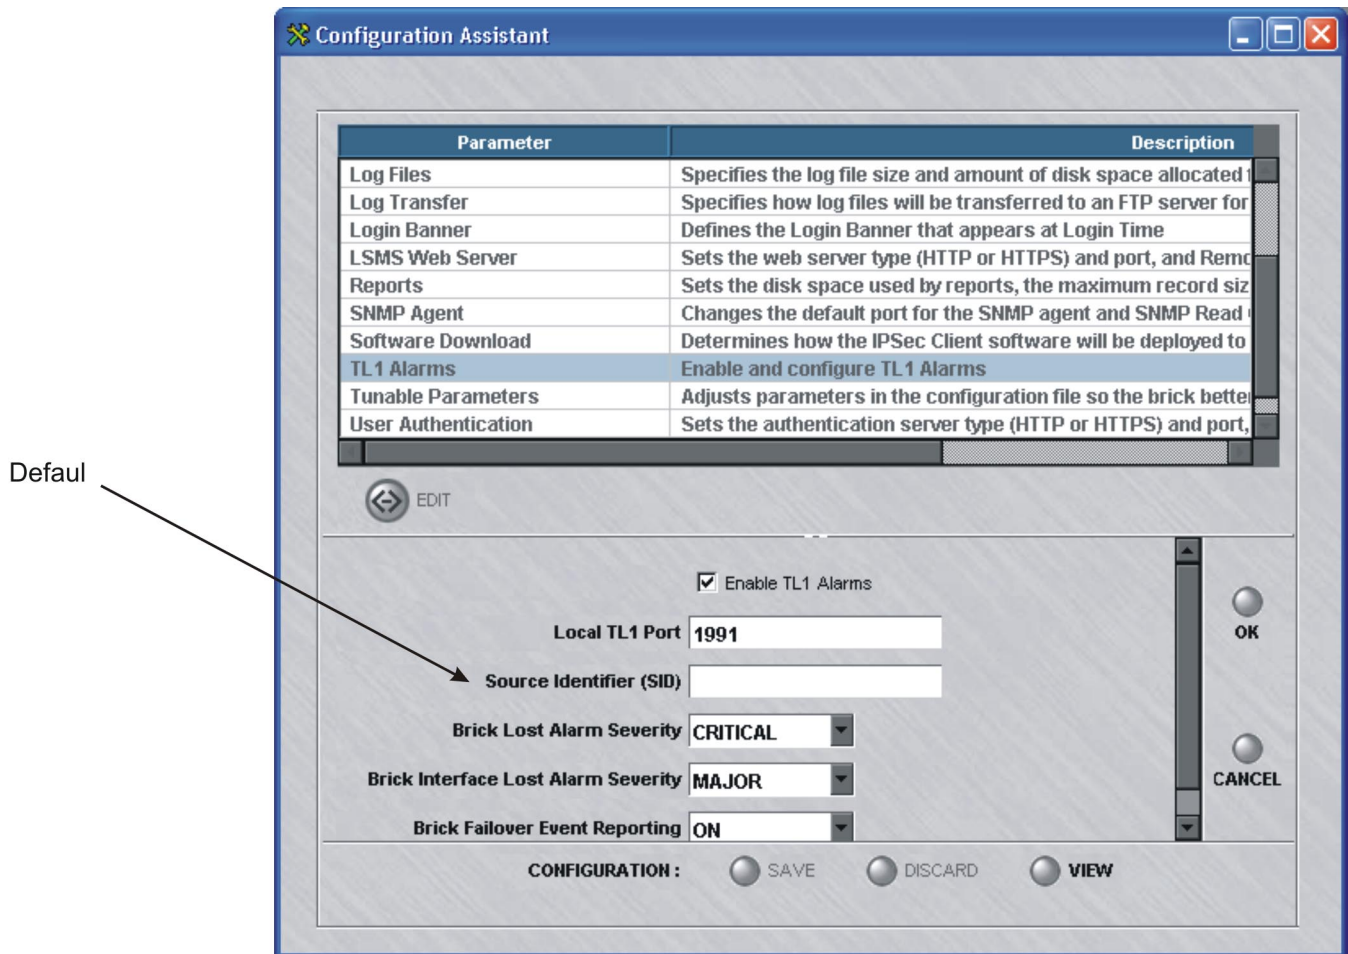
If you modify any of these parameters, you will have to stop and then restart all the LSMS services.

To activate this feature, click on the **Enable TL1 Alarms** checkbox on the Configuration Assistant screen (see Figure 10-17 below).

Default Values

TL1 Alarms are disabled as the default. [Figure 11-17, “TL1 Alarm Parameters” \(p. 11-43\)](#) shows the default values for the *TL1Alarm* parameters when it is enabled.

Figure 11-17 TL1 Alarm Parameters



Local TL1 Port

The **Local TL1 Port** field identifies the port on which the TL1 server will be listening to receive TL1 Alarm session requests.

By default, the TL1 Server listens on port 1991. If you wish to change the default value, you should run a networking command, like *netstat -a*, on the LSMS computer to ensure the port you are designating is not already in use.

Source Identifier (SID)

The **Source Identifier** field is used to specify the name of the LSMS in TL1 command responses and autonomous messages.

It is recommended that for traditional telecommunications applications, the **Source Identifier** should include the Common Language Location Identifier for the LSMS.

Brick Lost Alarm Severity

The **Brick Lost Alarm Severity** field specifies the TL1 alarm severity for Brick Lost alarms.

A choice of **CRITICAL**, **MAJOR**, **MINOR**, and **NONE** is given. If **NONE** is chosen, no Brick Lost alarms will be sent. The default value is **CRITICAL**.

Brick Interface Lost Alarm Severity

The **Brick Interface Lost Alarm Severity** field specifies the TL1 alarm severity for Brick Interface Lost alarms.

A choice of **CRITICAL**, **MAJOR**, **MINOR**, and **NONE** is given. If **NONE** is chosen, no Brick Interface Lost alarms will be sent. The default value is **MAJOR**.

Brick Failover Event Reporting

The **Brick Failover Event Reporting** field specifies whether Brick Failover events will be reported via the TL1 Alarms interface.

A choice of **ON** and **OFF** is given. The default value is **ON**.



Tunable Parameters

Overview

The Tunable parameters allow you to adjust certain *maxHeap* parameters in the configuration file (*config.ini*) so that the LSMS better handles environments in which there are:

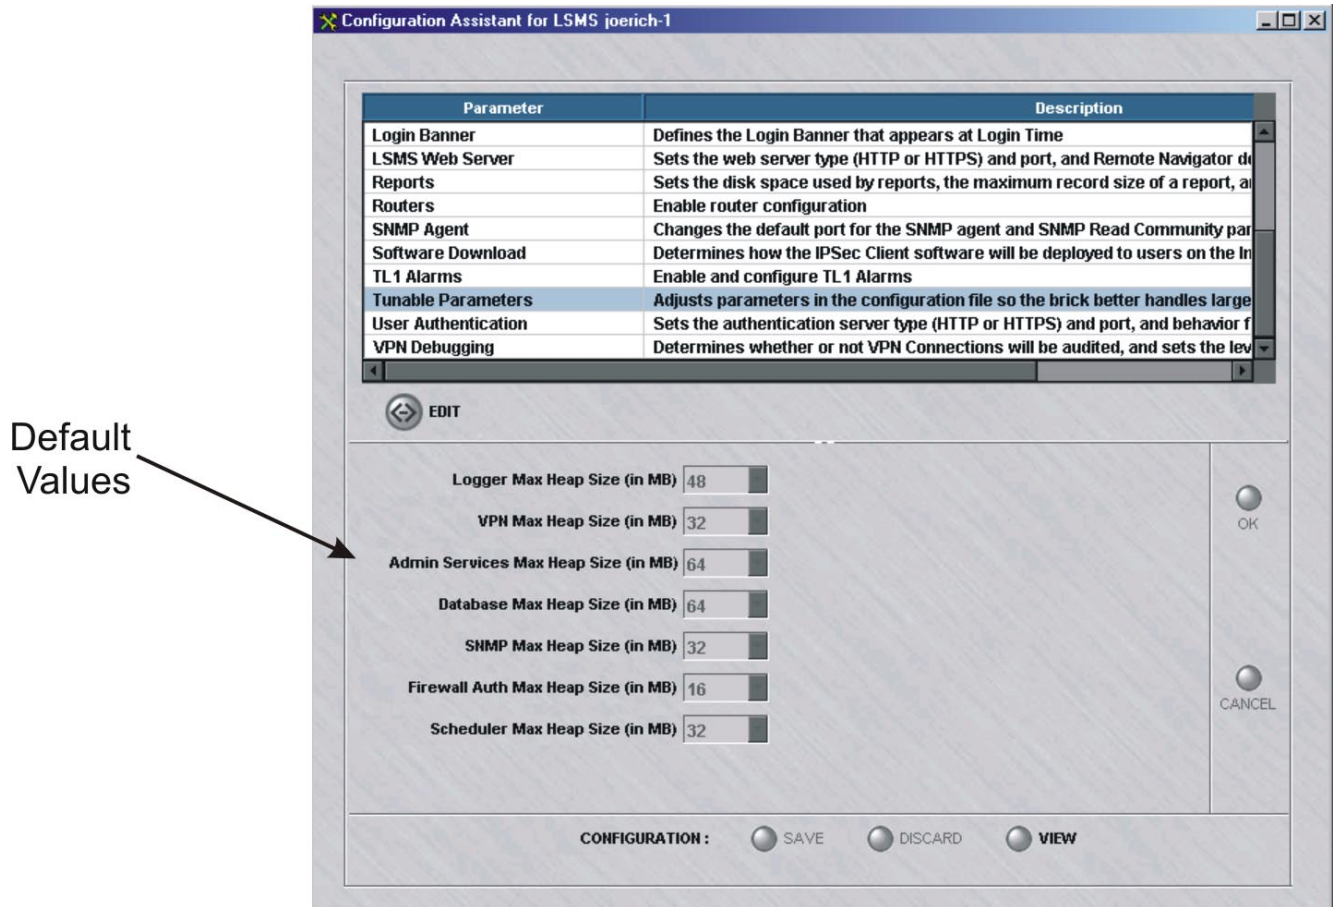
- A large number of Bricks managed by the LSMS,
- A large number of Lucent IPsec client users enabling client tunnels to Bricks managed by the LSMS, and/or
- A large number of audit records being generated per Brick/per second.

A single LSMS can manage up to 1000 Bricks, handle up to 8000 client tunnels, and process up to 15 audit records per second/per Brick (a Sunfire V210 with at least 512 MB of RAM or an equivalent platform is required to manage such a load). If your environment approaches these high-end figures, you need to re-set certain *maxHeap* parameters. The following explains.

Default Values

Figure 11-18, “Tunable Parameters” (p. 11-46) shows the default values for the *maxHeap* parameters.

Figure 11-18 Tunable Parameters



Admin Services Max Heap Size

The **Alarms Max Heap Size** parameter is set by default to 32 Mb. If this LSMS is managing more than 400 Bricks, set this parameter to 48 Mb.

Logger Max Heap Size

The **Logger Max Heap Size** parameter is set by default to 48 Mb. If this LSMS will be managing between 200-400 Bricks, set this parameter to 64 Mb. If it will be managing between 400-600 Bricks, set the parameter to 80 Mb. If it will be managing more than 600 Bricks, set the parameter to 96 Mb.

VPN Max Heap Size

The **VPN Max Heap Size** parameter is set by default to 32 Mb. If this LSMS will manage from 3 - 5000 client tunnels, set this parameter to 48 Mb. If it will manage more than 5000 tunnels, set the parameter to 64 Mb.

Other Max Heap Size Parameters

The defaults of the other max heap size parameters should generally not be changed.



User Authentication

User authentication

User Authentication is used to authenticate remote users whose IP addresses are unknown, allowing them access. The User Authentication parameters allow you to define the protocol (HTTP or HTTPS) and port that will be used by the authentication server.

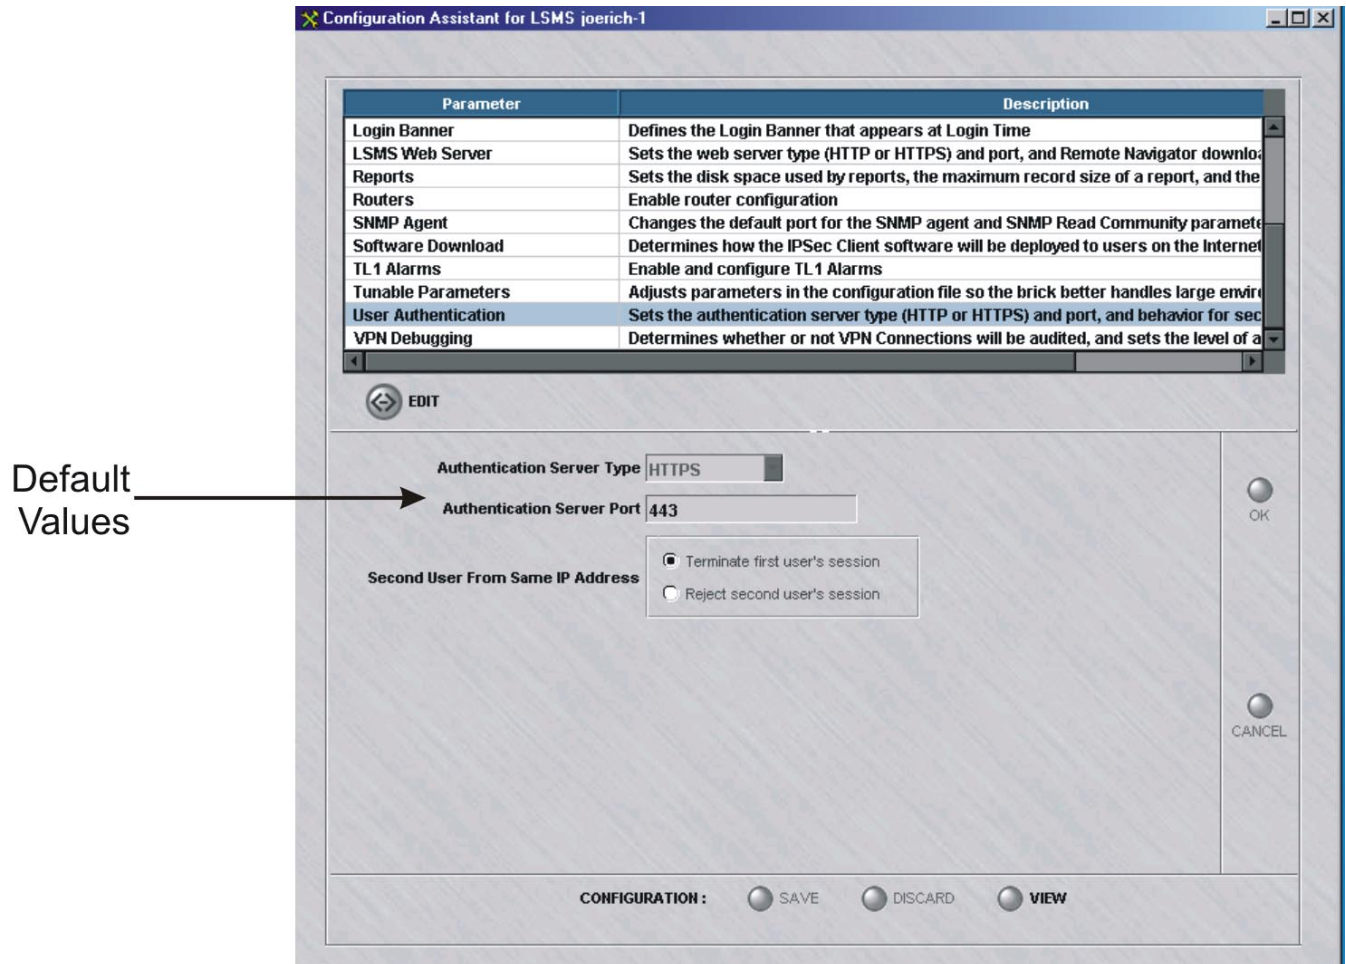
The authentication server is the LSMS web server that is used for user authentication. The authentication server is only used for firewall Authentication, not VPN authentication

If you modify any of these parameters, you will have to stop and then restart all the LSMS services.

Default Values

Figure 11-19, “User Authentication Parameters” (p. 11-49) shows the default values for the User Authentication parameters.

Figure 11-19 User Authentication Parameters



Authentication Server Type

The **Authentication Server Type** field determines whether the connection to the authentication server will be HTTPS (secure) or HTTP (non-secure).

Authentication Server Port

The **Authentication Server Port** field identifies the port to which the end users connect for application user authentication. (See *Chapter 9. User Authentication* in the *LSMS Policy Guide* for additional details.)

If the authentication server is HTTPS (as highly recommended), the port is usually 443.

Second User From Same IP Address

This set of options defines the action to be taken if a second user tries to authenticate from the same IP address. This could happen, for example, if there is a shared workstation, and user1 is authenticated and then leaves the workstation without ending the authenticated session by logging off. If a second user comes to the workstation and attempts to log in, this option defines what should be done.

There are two options for the action to be taken:

- **Terminate first user's session.** This is the default option. If the second user is successfully authenticated, the first user's session is terminated. The IP address of the workstation is removed from any User Groups to which user1 belongs and is added to any User Groups to which user2 belongs. User Groups are used in Brick Zone Rulesets to allow or deny access to authenticated users for certain resources protected by the Brick device. Also, any active sessions in the Brick cache for user1 are terminated.
- **Reject second user's session.** The second user's login will be rejected, even if they enter the correct password. For security reasons, they are not told why their authentication fails. A message will be logged to the User Auth Log, which is viewable by the LSMS Administrator, which indicates that the authentication failed because another user already has an active session from that IP address. If this option is chosen, user1's session remains active until user1 logs out or until the authentication times out, whichever comes first.



Strong Passwords

Strong passwords

The Strong Password option, when enabled, enforces stricter requirements when creating new passwords or modifying existing passwords in the LSMS to comply with Sarbanes-Oxley (SOX) password requirements.

The Strong Passwords feature is enabled, by default.

When a new password is set for a user or administrator that is authenticated using Local Password authentication, or an existing local password is changed, if the strong password (SOX compliance) option is enabled (the default) via the Configuration Assistant, stricter password requirements would apply for password 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

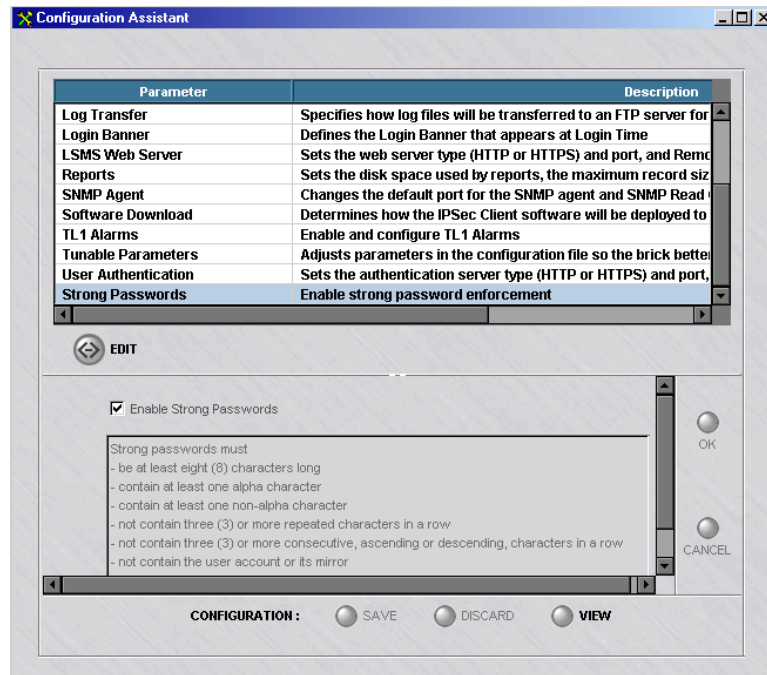
The strong password (SOX) requirements, when enabled, apply to new or changed passwords for:

- A clean installation password of the master user
- Local passwords
- User passwords
- Administrator passwords
- User login passwords for the Brick device console
- Admin key (additional key) which is needed when SecurID or RADIUS authentication methods are used

Strong passwords panel

Figure 11-20, “Enable Strong Passwords Panel” (p. 11-52) shows an example of the Strong Passwords panel with the Enable Strong Passwords checkbox checked, by default.

Figure 11-20 Enable Strong Passwords Panel



Enable strong passwords

The **Enable Strong Passwords** checkbox is checked, by default. This enables the Strong Passwords feature. To disable this feature, click the checkbox to remove the check.

Strong passwords restrictions

The read-only portion of the panel is a scrollable list of password restrictions that are enforced if the Strong Passwords feature is enabled.



VPN Debugging

Overview

The VPN Debugging parameter allows you to determine whether or not VPN connections will be audited and recorded in the VPN log file.

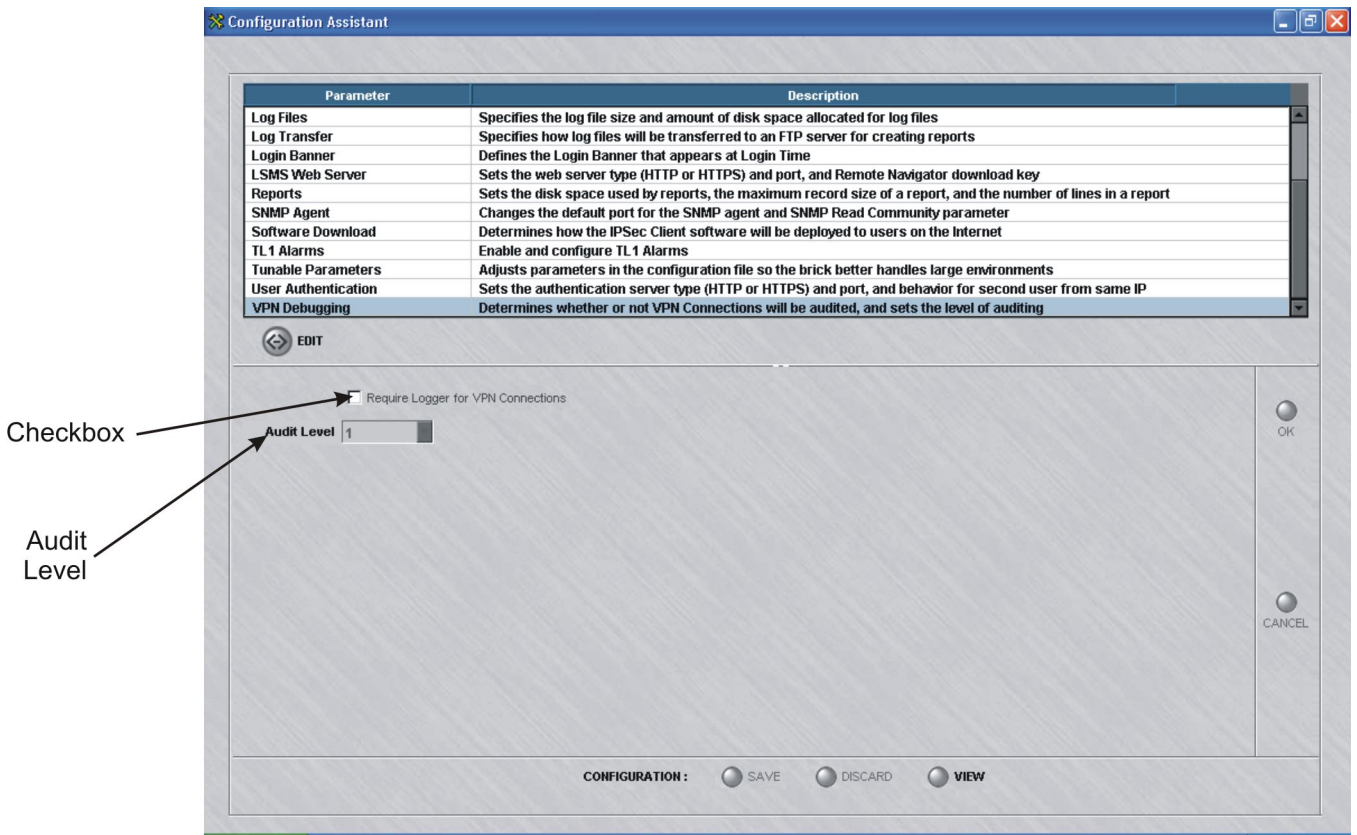
You also can decide on the verbosity of auditing to be performed. The higher the number (0-3), the greater level of detail that will be recorded.

If you modify any of these parameters, you will have to stop and then restart all the LSMS services.

Default Values

The VPN Debugging parameter consists of a single checkbox and pulldown menu, as shown in [Figure 11-21, “VPN Debugging Parameters”](#) (p. 11-53).

Figure 11-21 VPN Debugging Parameters



Require Logger for VPN Connections Checkbox

The checkbox labeled **Require Logger for VPN Connections** determines whether or not VPN connections will be logged. By default, the box is not checked, which means logging will not take place.

If you want all the VPN connections to be recorded in the VPN log files, then click the checkbox.

If the box is checked and the logger is down, then all VPN connections will be denied. Each time a VPN connection is attempted, it will check again for the presence of the logger. Until the logger comes back up, no VPN connections will be allowed.

Audit Level

If the checkbox is checked, so that VPN connections are logged, you can determine the level of auditing to take place.

From the drop-down list, choose a number between 0 and 3. The higher the number, the more verbose the messages are when written to the VPN log file.



12 Backing Up and Restoring Data

Overview

Purpose

This chapter explains how to back up and restore data on both a standalone LSMS and a redundant pair of LSMS (primary and secondary).

The data on an LSMS that has to be backed up consists of a database and a number of configuration files. The database contains the configuration data that is managed by the LSMS, such as policies, devices, and VPN tunnels. The configuration files contain less essential data, such as the *config.ini* file and the digital certificates used for authentication. It is not necessary to back up log trace files.

Contents

Automatic Backup	12-2
Manual Backup	12-3
Scheduled Backups	12-6
Restore Procedure	12-7
Restore Scenarios on Redundant LSMS	12-10
Other Restore Scenarios	12-11



Automatic Backup

Overview

By default, all standalone and primary LSMS automatically back up their databases and configuration files each night at 2:00 am. The backed-up database is stored in the sub-directory "db/backups" under the directory in which the LSMS application is installed. In addition, subdirectories labeled "1 - 7" are created. One backup is recorded in each directory nightly. After a week's time, directory 1 is overwritten. In addition, after the nightly backup, a copy of the backup is written to the secondary LSMS under `installdir/db/primary_backups`. It is to be used in the event that there is a failure on the primary LSMS and its nightly db backup is not available.

No automatic backup is performed on a secondary LSMS because it is not necessary to back up the database on a secondary LSMS. Whenever the secondary LSMS is upgraded or restored, the database is copied directly from the primary LSMS. It is necessary, however, to back up the configuration files on the secondary LSMS, and this has to be done manually.

□

Manual Backup

Overview

The automatic backup ensures that the database on each standalone and primary LSMS is backed up at least once a day. However, you can also perform manual backups so that your database is backed up more frequently.

This is important, because each administrative change (add, modify or delete) alters all or portions of the configuration data in the database. The ability to restore the database becomes critical if you make a change, and then discovers the change has to be reversed — especially if it creates a security violation.

Do not back up your files to a directory in *installdir* tree (i.e., to any subdirectory under the installation directory). This is for safety reasons so that your backup is not inadvertently lost if you need to uninstall and reinstall the LSMS application. You might also consider moving the backup to a tape or another machine.

For the backup to be complete, do not back up the database when the LSMS is being used to add, modify, or delete configuration data. Otherwise, the backup will not have integrity. Also, the system is collecting real-time session statistics and should not be taken off-line to perform backups.

If you need to obtain an estimate of the disk space that is required for the backup, look at the */db/LSMS* directory under the installation root directory. This will be the approximate amount of disk space required for the backup.

Even though other configuration data from other directories are included in the backup, they are minor compared to this directory that holds the database.

Manual Backup Procedure

Even though the LSMS services can be running while backing up, the backup process should not be done during a period of system change, such as adding or deleting Lucent VPN Firewall *Brick*[®] devices or users. Backing up should be performed when there is a light processing load on the LSMS, for example, at midnight.

To backup the database, follow the steps below:

-
- 1 Open a command prompt window.

2 Change directories to the installation directory as follows:

- On Windows, this is `c:\users\isms\lmf` if you selected the default directory during installation.
- On Solaris, this is `/opt/isms/lmf` if you selected the default. You also need root privileges on the Solaris hosts.

3 At the command prompt, enter

```
local/bin/backup <backup_directory>
```

where:

`<backup_directory>` is the mountable destination directory used to record the database. This directory will be created for you if it does not already exist. It can be a mapped directory or an external drive such as a Zip drive. It can include fully-qualified path names or be relative to the installation root directory.

Important! When entering the command on a Windows platform, you have to include the letter indicating the disk drive, as in the example below:

```
c:\backups\2000-05-22
```

On a Solaris platform, you do not have to include the disk drive letter. The following is an example:

```
../backups/2000-05-22
```

4 During the course of the backup, the following messages will appear on the screen:

- *Standalone/Primary Backup*

```
# ./local/bin/backup /export/home/tmp/backup_0726
Backing up database...
Backing up non-database files...
Backup is complete and successful #
```

The message is for:

- *Secondary Backup*

```
# ./local/bin/backup /export/home/tmp/sec_backup_0726  
Backing up non-database files only...  
Backup is complete and successful #
```

END OF STEPS



Scheduled Backups

LSMS task scheduler

The LSMS Task Scheduler allows you to run a database backup command on a scheduled basis through a separate GUI function. A database backup can be scheduled to be run once, a specific number of times, or periodically at a set time (hourly, daily, weekly, monthly, or yearly).

For details about how to use the Task Scheduler, refer to [Chapter 9, “Compute Servers”](#).



Restore Procedure


When to use

While restoring the database, the LSMS services must be stopped. The restore process should not compete with other processes for CPU and memory resources. Restoring files should be performed when there's a light processing load on the LSMS, for example, at midnight.

Important! Since the entire database is saved when it is backed up, the entire database is also restored with the restore utility. Individual files or a portion of the database cannot be selectively restored with this utility.

Task

To restore the database that was previously backed up, do the following:

- 1 Stop the LSMS services as follows:
 - On Windows, click the **Start** menu and select **Programs ► Lucent Security Management Server**  **Stop Services**
 - On Solaris, from the installation root directory enter `./stopServices`.

- 2 Open a command prompt window.

- 3 Change directories to the installation root directory as follows:
 - On Windows, this is `c:\users\isms\lmsf` if you selected the default directory during installation.
 - On Solaris, this is `/opt/isms/lmsf` if you selected the default directory. You also need root privileges on Solaris hosts.

- 4 At the command prompt, enter
`local/bin/restore <backup_directory>`
where:
`<backup_directory>` is the source directory where the database was backed up. It can be a mapped directory or an external drive such as a Zip drive. It can include fully-qualified path names or be relative to the installation root directory. Examples include:
`D:\backups\2000-05-22`

../bkup/2000-05-22

5 During the course of the restore, the following messages will appear on the screen:

- Standalone/Primary Restore

```
# ./local/bin/restore /export/home/tmp/bu_test1
Be sure the LSMS Services are not running.
If they are running, STOP THEM NOW.
When they are stopped, hit 'Enter'

Restoring database...
Restoring non-database files...
Removing old publication...This may take a few
minutes...

Database Schema is up-to-date...no changes made...

Creating updated publication...This may take a few minutes...

Done with successful restore
- Please re-start LSMS services
- Then run restore on the secondary LSMS #
```

The message that appears for Secondary Restore is as follows:

```
# ./local/bin/restore /export/home/tmp/bu_test1
Be sure the LSMS Services are not running.
If they are running, STOP THEM NOW.
When they are stopped, hit 'Enter'

Restoring non-database files only...
Secondary LSMS Setup started...

Enter Primary LSMS IP Address [10.20.10.155] :
Connecting.....done

Performing Handshake with Primary.....done

Getting Version Info from Primary LSMS.....done

Validating Installation Key on Primary LSMS.....done

Validating Installation Key on Secondary LSMS.....done

Downloading database information from Primary
LSMS.....done

Creating the Database for Secondary LSMS...This may
take a few minutes.....done

Closing connection with Primary.....done

Done with successful restore
- Please re-start LSMS services
```

6 Restart the LSMS services as follows:

- On Windows, Click **Start** and select **Programs > Lucent Security Management Server > Start Services**.
- On Solaris, from the installation root directory enter `./startServices`.

END OF STEPS



Restore Scenarios on Redundant LSMS

Overview

There are several possible scenarios that you have to be aware of when performing a restore procedure on a redundant pair of LSMS:

- *Scenario #1:* Restore has been performed on a primary LSMS.
When the services on the primary LSMS are restarted, the primary and secondary LSMS are not "connected" because their database certificates are not in sync. These are the certificates required to encrypt the transfer of the database between the two machines to ensure their security.
When a restore is done on the secondary LSMS, its configuration files are restored and a dbsetup is done. In the dbsetup, the database is copied from the primary. Any changes made on the secondary LSMS since it lost connection with the primary LSMS have been lost.
- *Scenario #2:* The primary database has not changed, but a restore is needed on the secondary.
Before doing a restore on the secondary LSMS, you must issue the following command from the installDir on the primary LSMS:
`local/bin/allowSecondarySetup`
This resets the database certificate on the primary LSMS to the default and allows the secondary LSMS to copy the database from the primary LSMS. The restore command on the secondary LSMS will then restore the configuration files and do the dbsetup (copy database from the primary LSMS).
Please note that this is the only way to restore the database on the secondary LSMS. *Do not* attempt to restore the backup from the primary LSMS directly onto the secondary LSMS.
- *Scenario #3:* The user does not need to restore the configuration files on the secondary LSMS, but needs to copy the database from the primary LSMS.
This situation could arise if the primary and secondary LSMS have not been connected for more than a week. As in Scenario #2, you have to issue the `allowSecondarySetup` command on the primary LSMS. On the secondary LSMS, you have to stop services and then type
`local/bin/dbsetup`
to copy the database from the primary LSMS. After the dbsetup is done, you have to restart the services.

□

Other Restore Scenarios

Overview

The following additional scenarios are also possible:

- *Scenario #4:* The user has recently upgraded their machine with an LSMS patch. After the patch upgrade, the user has discovered a need to restore the database, but only an older version of the database backup is available. The restore utility will automatically upgrade an older version of a database to the current release, if needed.
- *Scenario #5:* The user has decided to move an existing LSMS onto a different machine. Prior to restoring the database onto the new machine, the LSMS application must already be installed. The LSMS installation paths must be the same on both machines. After the restore has been completed, other steps may be necessary before the LSMS services can be restarted on the new machine. If the IP address on the new LSMS is different than the old LSMS, you must run the "changeIP" utility and update the Bricks. For more information, check *Appendix D* in the *LSMS Administration Guide*. If the machine name on the new LSMS is different than the old LSMS, you must run the "changeName" utility. For more information, please review *Chapter 6. Database Utilities* in the *LSMS Administration Guide*. If the new LSMS is a "primary" LSMS, you must run a *restore* or a *dbsetup* on the secondary LSMS to synchronize the two environments. See the steps mentioned earlier in Scenario #1.
- *Scenario #6:* The secondary LSMS needs to be reinstalled, and it is running a patch version of the LSMS software. The first step in performing a "clean" install on a secondary LSMS is to uninstall the existing version of the LSMS and to delete the "lmf" directory (by default, \users\isms\lmf on Windows or /export/home/isms/lmf on Solaris). The user needs to install the "gold" version of the LSMS on the secondary first. The difficulty arises when the administrator attempts to copy the database from the primary. Since the primary is running the LSMS patch version, by definition its database version will be out of sync with the secondary. Therefore the database cannot be copied from the primary to the secondary. To overcome this obstacle, the administrator must execute the following steps:
 - Install the gold version of the LSMS on the secondary. During the database portion of the installation, define this LSMS as a "Standalone" LSMS, NOT a redundant LSMS.
 - Install the same patch version of the LSMS on the secondary as is running on the primary.

- On the primary, from a DOS or terminal window, "cd" to the LSMS installation directory. Type "local/bin/allowSecondarySetup". This temporarily resets the database encryption certificate used for communication between the primary and secondary databases.
- On the secondary, stop all LSMS services. From a DOS or terminal window, "cd" to the LSMS installation directory. Type "local/bin/dbsetup". This repeats the database portion of the LSMS installation. You can now redefine this LSMS as a redundant, secondary LSMS. The database from the primary will be copied to the secondary during this procedure.
- At the completion of the dbsetup, restart the LSMS services on the secondary.



13 Task Scheduler

Overview

Purpose

This chapter discusses the LSMS Task Scheduler.

Contents

What is the Task Scheduler?	13-2
Schedule Editor	13-3



What is the Task Scheduler?

Definition

The LSMS Task Scheduler allows you to run commands (perform tasks), such as database backups or log transfers, at scheduled times via a GUI window. A command (task) can be scheduled to run once, a specific number of times, or periodically at a set time (hourly, daily, weekly, monthly, or yearly).

The LSMS Task Scheduler is installed with just one task scheduled (backing up the LSMS database), but you can schedule other commands (tasks) such as transferring log files.



Schedule Editor

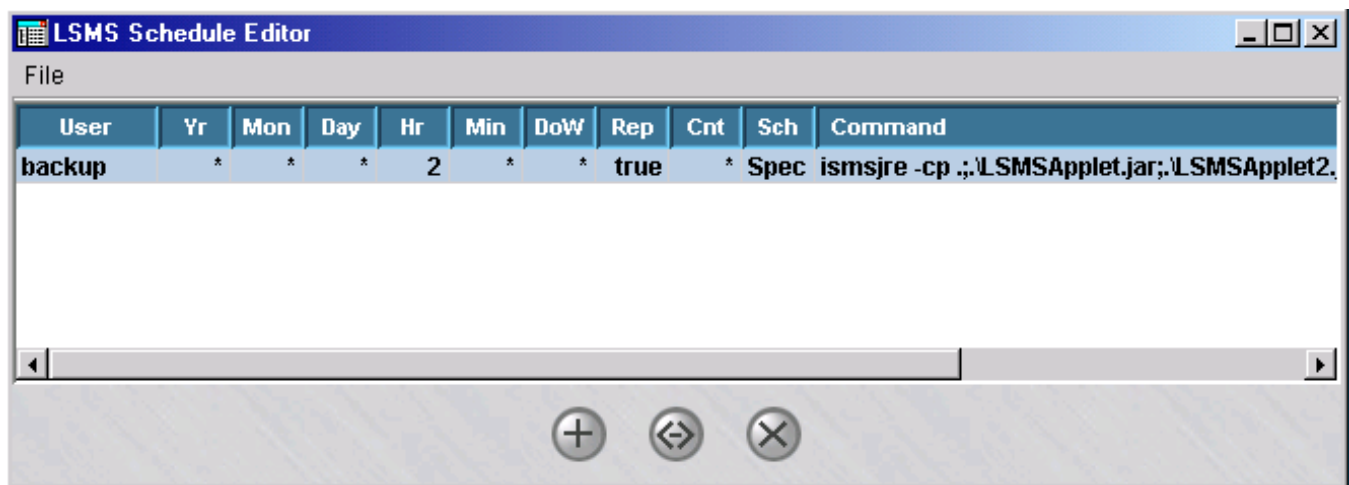
Overview

The Schedule Editor, which is the editing window accessible through this tool, allows you to modify the actions of the Task Scheduler without disrupting the system by stopping and restarting LSMS services. It also provides a less error-prone method of modifying the scheduler configuration files than editing them manually using a text editor.

LSMS Schedule Editor Window

Figure 13-1, “LSMS Schedule Editor Window (Initial View)” (p. 13-3) shows a sample of the LSMS Schedule Editor window..

Figure 13-1 LSMS Schedule Editor Window (Initial View)



Each tabular row in the window displays a task and the details about when it is scheduled to be performed, as follows:

- **User** - identifies the user account who scheduled the command.
- The **Yr** (year), **Mon** (month), **Day**, **Hr** (hour), **Min** (minute) and **DoW** (day of the week) fields provide the scheduling particulars of the task.
- The **Rep** (repeat) and **Cnt** (count) fields specify how many times the command is to be executed. If the Rep entry is true, then the command is run repeatedly the number of times shown in the **Cnt** field. A non-numeric or entry of **0** in the **Cnt** field indicates that the command will be run indefinitely.
- The **Sch** field indicates whether the time fields indicate a specific date/time (**Spec**) for the command to be run or a time interval (**Freq**) between execution of the command, if it has been scheduled to be performed repeatedly.

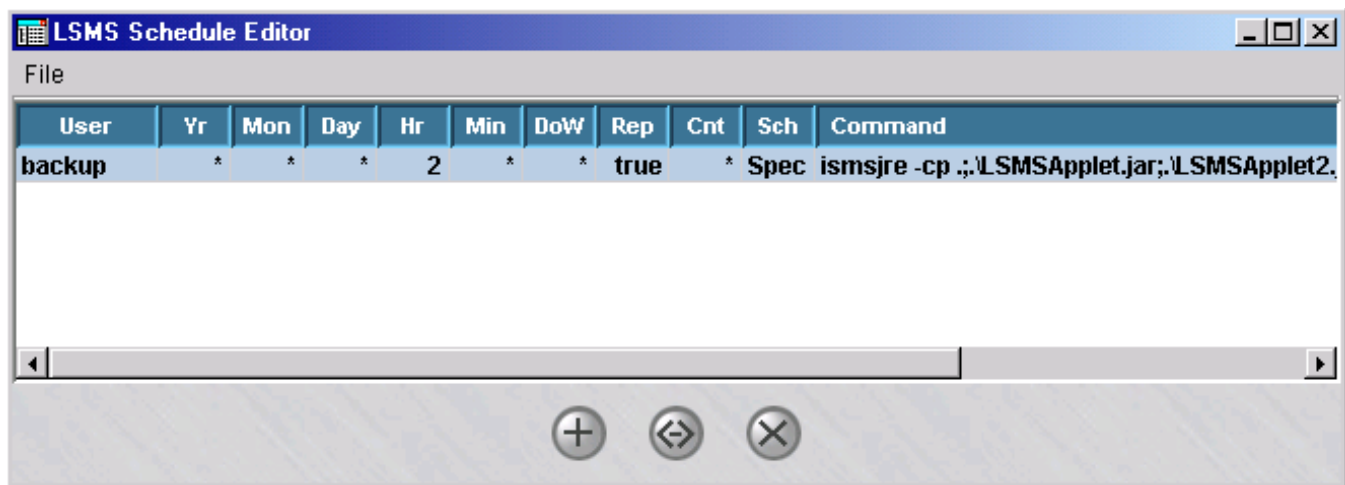
Scheduling a Command

- 1 To schedule a command (task), follow the steps below: If the remote host is running Windows, click the **Start** menu and select:

Programs > Lucent Security Management Server > Utilities Lucent Security Management Server Utilities > LSMS Schedule Editor

The LSMS Schedule Editor window is displayed with the backup task pre-selected in the task list, which is pre-set when the LSMS software is installed (Figure 13-2, “LSMS Schedule Editor Window (Initial View)” (p. 13-4) shows a sample window).

Figure 13-2 LSMS Schedule Editor Window (Initial View)



- 2 To schedule a backup, or any other command to be run, click the Edit (↔) button, or just double-click on the backup task in the task list.

The Edit Command Schedule window is displayed (Figure 13-3, “Edit Command Schedule Window” (p. 13-5) shows a sample window).

Figure 13-3 Edit Command Schedule Window

- 3 To schedule the backup (or another command) to run, edit the fields as follows:
 - **Command** - this field is pre-set with the command to run an LSMS database backup. If you want to schedule a backup, leave the contents of this field as is. To schedule a different command, enter the complete command line of the task to be scheduled to run.
 - **User** - Enter the user login who is scheduling the command/task.

- **Run command once** - click this radio button to run the command only once at the scheduled time interval.
 - **Run command n times** - click this radio button and specify the number of times to run the command at the scheduled time interval.
 - **Run command repeatedly** - click this radio button to run the command repeatedly at the scheduled time interval.
If the option is selected to run the command repeatedly, specify the waiting time (number of years, months, days, hours, minutes) between each command run.
 - **Run command at specific time** - If the command is being run only once, specify the exact time to run the command (year in yyyy format; month in mm format; day, 0=Sunday, 1=Monday, and so forth; hour, in hh format, minutes, in mm format)
-

- 4 After making your schedule settings, click the **OK** button.

The LSMS Schedule Editor window is displayed, showing the newly scheduled command on the tasklist. The command/task will be scheduled to run at the selected frequency or time interval.

END OF STEPS



14 Using the Status Monitor

Overview

Purpose

This chapter explains how to use the LSMS Status Monitor. The Status Monitor provides a mechanism for monitoring the status of all Lucent VPN Firewall *Brick*[®] devices, VPN tunnels, and LSMSs at varying levels of summary and detail. In addition, it shows all LSMS and Group Administrators currently logged into the LSMS you are logged into, and it displays all console alarm messages.

Administrators can use the Status Monitor to:

- Monitor the current health of all LSMS-related network components to ensure that they are operating smoothly
- Track and analyze long range traffic patterns through all installed Bricks and tunnels
- Compare different time periods' worth of network activity and analyze the differences.

Contents

How to Access the Status Monitor	14-2
How to Interpret the Status Monitor	14-3
Status Overview Window	14-6
Administrators Window	14-10
Brick Status Windows	14-14
Console Alarms Window	14-27



How to Access the Status Monitor

Methods of access

You can access the Status Monitor using either the LSMS Navigator or the LSMS Remote Navigator.

Display the Status Monitor

There are two ways to display the Status Monitor:

- *Status Monitor Only*

This method allows you to log into the Status Monitor without also logging into the LSMS. To display the Status Monitor without opening the LSMS, click the **Status Monitor Only Login** checkbox on the Login window when logging into the LSMS. (The Status Overview window is displayed instead of the LSMS Navigator window.)

This method is *not* recommended for LSMS or Group Administrators. It is intended primarily for use in network operations centers, or to enable an individual to view network status without having the ability to view or change any of the configuration parameters.

- *Status Monitor and LSMS*

To display the Status Monitor from the LSMS, log into the LSMS without clicking the **Status Monitor Only Login** checkbox. Then, open the **Monitor** menu on any LSMS window and select the Status Monitor component that you want.

Status Monitor Components

The Monitor menu has four options on it, each of which corresponds to one of the components of the Status Monitor. The Status Monitor consists of these components:

- Status Overview window
- Administrators window
- LSMS and LSCS (Compute Server) window
- Brick Status window
- Console Alarms window.

An administrator can keep more than one Status Monitor window open at the same time. For example, you could keep the Status Overview window open to provide a high-level view of network operations, and at the same time keep open windows displaying individual Brick status, as well as a window listing the Bricks that are lost.



How to Interpret the Status Monitor

Overview

The Status Monitor provides a variety of data in both tabular and graphical form.

Status Monitor Data

The bulk of the data displayed in all Status Monitor windows except the Console Alarms window is gathered from the Proactive Monitoring Log. This log contains information about Brick events, logger events, VPN Gateway Controller (VGC) events, and Firewall Authentication Controller (FAC) events. (For a more detailed explanation of the Proactive Monitoring Log, see *Chapter 2. Audit Logs* in the *LSMS Reports, Alarms and Logs Guide*.)

By default, the LSMS automatically refreshes the Status Monitor windows every 30 seconds to ensure the data displayed is current. However, you can change the refresh interval according to your requirements. This is done from the window's toolbar (see "[Toolbar](#)" (p. 14-4)" on "[Toolbar](#)" (p. 14-4) for an explanation of how to change the refresh interval.)

If you cannot wait until the next system refresh, you can perform a manual refresh at any time. You can also turn off the system refresh feature, so that the Status Monitor only refreshes itself when you perform a manual refresh. This is also done from the toolbar.

Important! Every 30 seconds, a new batch of data is taken from the Proactive Monitoring Log and held by the LSMS. The last 30 seconds worth of data that was collected is the data that is sent to the Status Monitor when a refresh (either system or manual) is performed.

Hence, when a refresh is performed, it is not the last 30 seconds worth of "real-time" data that is displayed, but the last 30 second interval that was cached by the LSMS.

Brick States

The Status Monitor indicates the current status of a Brick by giving the current state of the Brick. The state of a Brick depends on the condition of the Brick itself, as well as the condition of the standby Brick, if the Brick has been configured as part of a failover pair.


If the Brick is *not* part of a redundant pair, it can be in one of two states : *Up* (the Brick is healthy) or *Lost* (the Brick is down). If the Brick *is* part of a redundant pair, it can be in one of four states. The table below describes these states.

State	Active	Standby	Explanation
Up-Up	Up	Up	Both Bricks are up, equivalently healthy, and communicating with each other
Up-Lost	Up	Lost	The active Brick is up and the standby has either not transitioned into ready mode yet or is down.
Up-Unhealthy	Up	Unhealthy	Both Bricks are up, but the standby Brick is reporting suboptimal health compared to the active Brick.
Up-X-wired	UP	X-wired	Both Bricks are up, but the standby is cross-wired

Toolbar

Every Status Monitor window except the Console Alarms window has a toolbar across the top, directly below the menubar. [Figure 14-1, “Status Window Toolbar” \(p. 14-4\)](#) below shows the toolbar with each component labeled.

Figure 14-1 Status Window Toolbar



The following table explains how to use the components of the toolbar:

Component	What it does
Refresh button	Manually refreshes the data in the Status Monitor window.

Component	What it does
Refresh interval	<p>Determines the amount of time between system refreshes. The alternatives are:</p> <ul style="list-style-type: none"> • Manual only • Every 30 second (default) • Every 1 minute • Every 2 minutes • Every 5 minutes • Every 10 minutes • Every 15 minutes • Every 30 minutes
Time elapsed since last refresh	<p>Displays the amount of time in minutes and seconds since the last refresh (either system or manual). The counter returns to 00:00 after each refresh.</p>
Monitored button	<p>Displays the list of monitored Brick.</p> <p>The Bricks in the Monitored Bricks List are those that you chose to include. You can use this list to monitor any Bricks that you want to keep an eye on. See “Brick Lists” (p. 14-14) on “Brick Lists” (p. 14-14) for additional details.</p>
Monitored list checkbox	<p>This checkbox only appears on Brick Status windows.</p> <p>To add a Brick to the Monitored Bricks List, select the Brick in any Brick Status window and click this checkbox. To remove a Brick from the list, select the Brick and uncheck the checkbox.</p> <p>If you are not sure if a particular Brick has been added to the list, select the Brick in any Brick Status window and see if this checkbox is checked.</p>
Current time	<p>Displays the current time, according to the LSMS clock. This time will most likely differ from your own PC time.</p>
Console Alarms button	<p>Opens the Console Alarms window.</p> <p>The bell becomes yellow when a new alarm has been recorded.</p>



Status Overview Window

Overview

The Status Overview window displays summary information for all the Bricks you have permission to view.

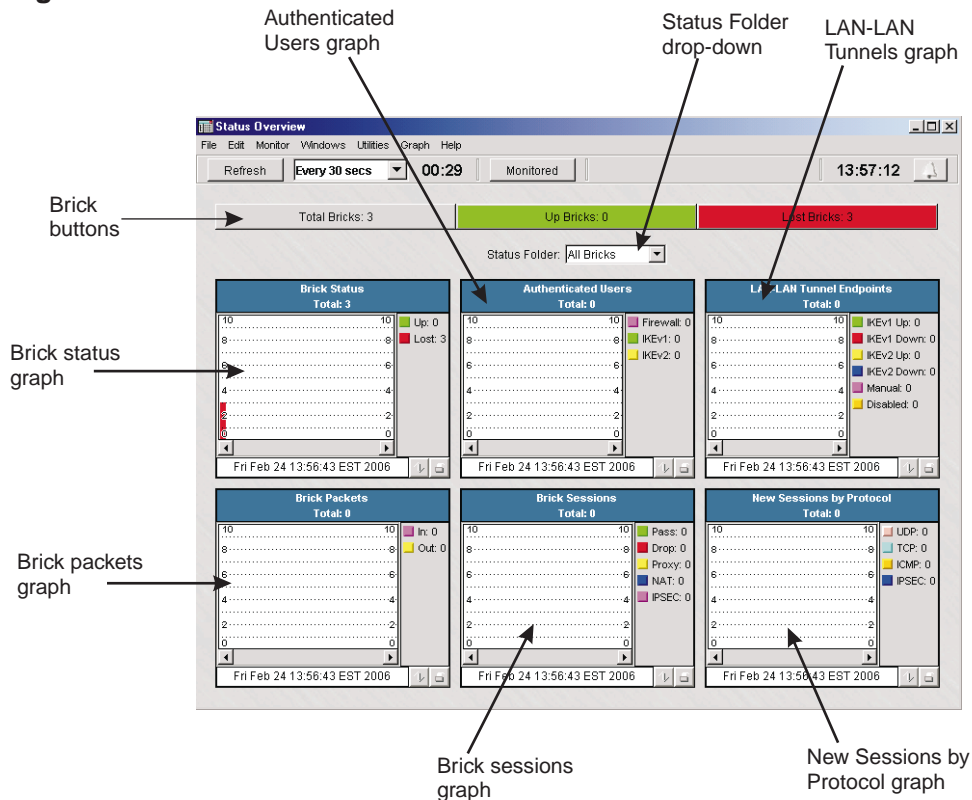
If you are an LSMS Administrator, you will see all Bricks in all groups. If you are a Group Administrator, you will only see Bricks in groups over which you have permission, and you must have at least *Device/View* permission (see [Chapter 9, “Compute Servers”](#) for an explanation of administrative permissions).

Window Components

To display the Status Overview window, open the Monitor menu and select **Status Overview**. (If you clicked the **Status Monitor Only Login** checkbox when logging in, this is the first window that appears.)

[Figure 14-2, “Status Overview Window”](#) (p. 14-6) shows a typical Status Overview window, with its major components labeled.

Figure 14-2 Status Overview Window



Brick® Buttons

Three Brick buttons appear at the top of the Status Overview window:

- Total Bricks (gray button)
- Up Bricks (green button)
- Lost Bricks (red button)

The number on each button indicates the number of Bricks currently configured in all groups in that category (such as total, up or lost). When you click one of the three buttons, a Brick Status window (also known as a "Brick List") appears with the appropriate Bricks listed in the window (see the "[Brick Status Windows](#)" (p. 14-14)" section for a more detailed description of Brick Status windows and Brick Lists).

Status Folder Drop-Down

The **Status Folder** drop-down list lets you select the Bricks you want displayed in the Status Overview window. The default is **All Bricks**.

You can change the default by selecting **Browse** from the drop-down list, and then selecting another Bricks folder. Once you select a folder this way, it will be permanently added to the drop-down list, so that in the future you will be able to access this folder directly from the list, without having to browse. In addition, all the data in the graphs will be re-drawn.

Brick Graphs

The Status Overview window contains six graphs that provide different views of network activity. The title of each graph can be found in a blue bar across the top, along with the total for that graph. The total is the total for the last datapoint displayed in the graph. A legend appears below each graph. Using a button below the graph, you can hide the legend if you wish.

The graphs are stacked bar graphs in that the values for each point at a given time are stacked on top of one another. The graph region contains the graphic representation of all of the data points in the history of this graph.



If the data point is the first data point in the history, its width will be all the remaining displayable space to the left of the graph. As a new data point is received, it is added at the right of the graph, pushing the oldest point (on the left of the graph) out of the viewable area.

You can modify the width of the bars and the spacing between the bars by right-clicking on a graph and selecting an option from the pop-up menu. The width of the bars has no meaning other than cosmetic.

The table below describes what each graph shows. The Brick Status graph shows all Bricks managed by your LSMS, regardless of whether the LSMS is a Primary LSMS or part of a redundant LSMS pair.

Graph	What it shows
Brick Status	Shows the number of up and lost Bricks
Authenticated Users	Shows the number of firewall and Client VPN users currently authenticated
LAN-LAN Tunnels	Shows the number of LAN-LAN tunnels currently up, down, manual and up (no heartbeat).
Brick Packets	Shows the number of packets passing through the Bricks
Brick Sessions	Shows all sessions that were passed, dropped, proxied, NATed and IPsec (tunneled)
New Sessions by Protocol	Shows all new sessions through the Bricks

Each graph has two buttons to the right of the date:

- The **Legend** button  on the left acts as a toggle. Click it once to remove the legend from a graph so there is a larger graph area to view. Click it again to return the legend.
- The **Print** button  on the right allows you to print the graph. You must have a default printer defined for the print operation to work.

In addition, you can right-click on any graph to display a menu that allows you to customize the graph. You can select more than one data point by holding down the [Shift] key and selecting each point with the mouse cursor. The table below explains each option on the menu.

Option	What it does
Display Selected Point Details	Launches a separate window that shows the current date, time and details for the data point on the graph that you selected
Clear Selected Points	Removes the data point you selected from the graph.
Clear All Points	Removes all data points from the graph
Space Between Bars	Inserts spaces between the bars in the graph. This option, and the option below, allow you to view the graph with spaces and without.
No Space Between Bars	Removes spaces between the bars in the graph

Option	What it does
Thin Bars	Makes the bars thinner than the default size
Standard Bars	The default size of the bars in all graphs
Wide Bars	Makes the bars thicker than the default size



Administrators Window

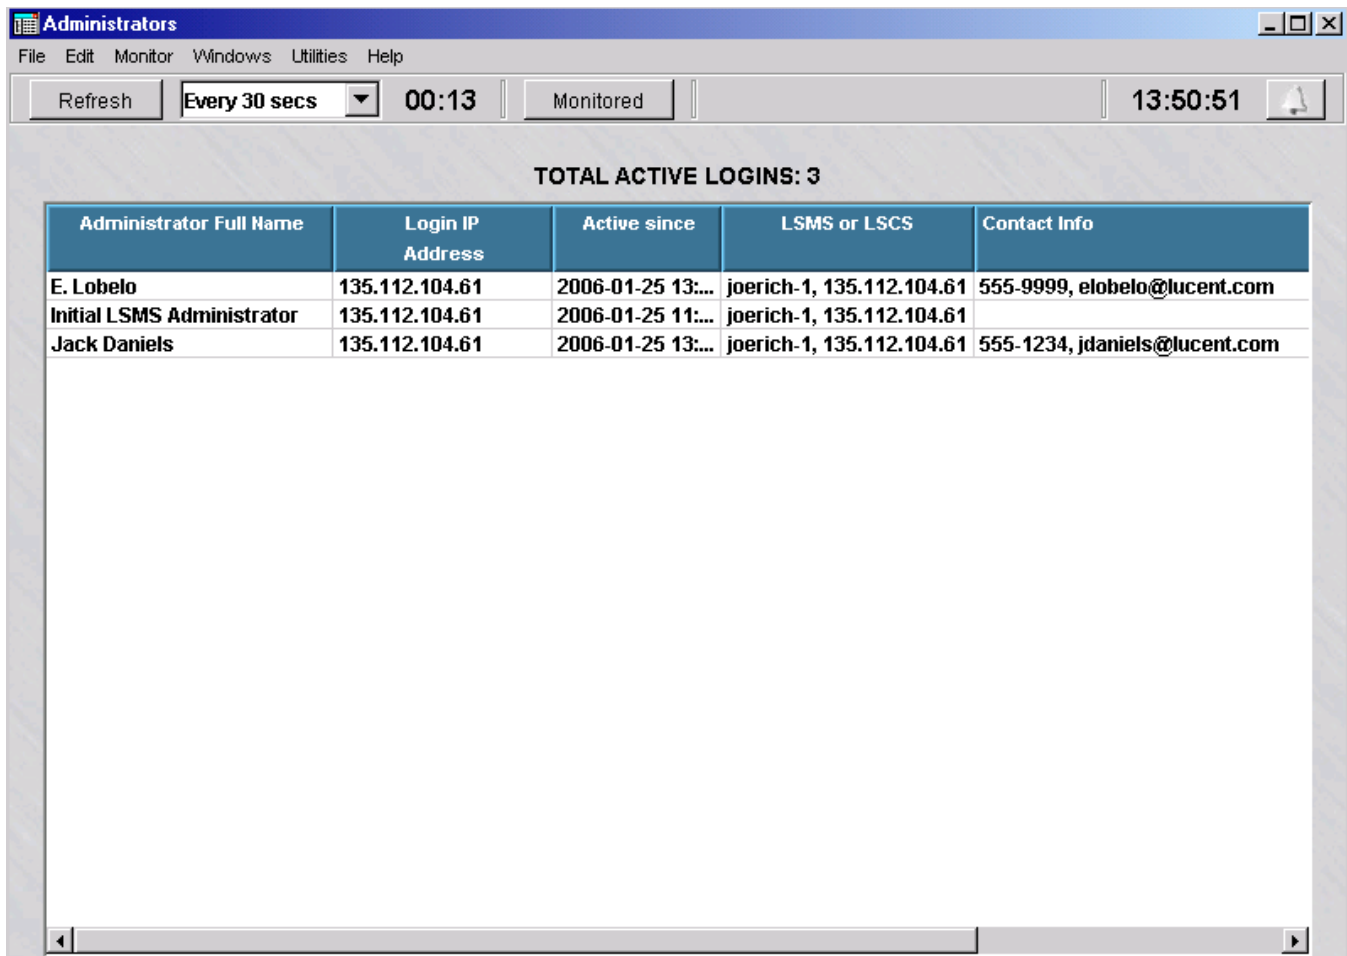
Overview

The Administrators window provides status information about each administrator who is currently logged in.

Window Components

To display the Administrators status window, open the **Monitor** Menu and select **Administrators**. [Figure 14-3, “Administrators Status Window” \(p. 14-10\)](#) shows a sample of the Administrators status window.

Figure 14-3 Administrators Status Window



The Administrators status window displays the status of all administrators currently logged into the LSMS you are logged into. This includes both LSMS and Group Administrators.

The total number of administrators currently logged in is given at the top of the table. For each administrator logged in, the Administrator table provides the following information:

Field	Description
Administrator	The full name of the administrator that was entered in the Full Name field when the administrator account was created
Login IP Address	The IP address of the administrator
Active Since	The date and time the administrator logged in
LSMS	The name and IP address of the LSMS this administrator is logged into
Contact Info	The telephone number and email address of the administrator, if this information was entered when the administrator account was created

LSMS/LSCS Window

The LSMS status window displays the status of the LSMS you are logged into and any associated Compute Servers (LSCSs) (if enabled). If this LSMS is part of a redundant pair, it also shows the status of the other LSMS in the pair. Details about each LSMS or Compute Server are presented on a separate line on the screen.

Figure 14-4, “LSMS/LSCS Status Window” (p. 14-12) shows a sample of the LSMS/Compute Server status window.

Figure 14-4 LSMS/LSCS Status Window

LSMS or LSCS Name	Type	IP Address	Status	Version	Associated With	Bricks Assigned	Assigned Homed	Not Assigned Homed
joerich-1	Primary	135.112.10...	Up	8.0.233		2	0	0
joesbrick-au...	Compute Server	135.112.10...	Lost		joerich-1	0	0	0
lscs-second...	Compute Server	135.112.10...	Lost		joerich-1	0	0	0
joerich-2	Secondary	135.112.10...	Lost			0	0	0

The total number of LSMSs and LSCSs, and the number currently up and lost, is given at the top of the table. For each LSMS or LSCS, the Administrator table provides the following information:

Field	Description
LSMS	The name of the LSMS or LSCS
Type	The type of monitored device. Possible values are LSMS or Compute Server.
IP Address	The IP address of the LSMS or Compute Server
Status	Up or lost

Field	Description
Version	The product version of the monitored device
Associated with	For a monitored LSCS, the name of the associated LSMS
Bricks Assigned	The number of Bricks assigned to the LSMS or LSCS
Assigned Homed	The number of Bricks for which this LSMS or LSCS is <i>priority 1</i> that are currently homed to this LSMS. This field is blank if the status of the LSMS or LSCS is lost.
Not Assigned Homed	The number of Bricks for which this LSMS or LSCS is <i>priority 2</i> that are currently homed to this LSMS or LSCS. This field is blank if the status of the LSMS or LSCS is lost. If you have a standalone LSMS, this field is always 0 .



Brick Status Windows

Overview

The Status Monitor provides ten distinct Brick Status windows. Seven of these windows are Brick Lists, which show different groupings of Bricks (such as all Bricks, only Bricks with a current status of *up*, or only Bricks in a certain folder), and information about each Brick in the list. The other three Brick Status windows are a Single Brick Status window, a Single Brick Ports window, and a single Brick Bandwidth Statistics window.

Brick Lists

To display a Brick List, open the Monitor menu, select **Brick Status**, and then select the list you want from the submenu. The table below indicates the Brick Lists that are provided and briefly describes each one:

Brick List	Description
All Bricks Assigned to	Lists all Bricks that are assigned to a specific administrator
All Bricks	Lists all Bricks over which you have device/view privileges, regardless of their current status
Monitored Bricks	Lists the specific Brick you have added to the Monitored Bricks list
Lost Bricks	Lists all Bricks with a current status of <i>lost</i>
Not Up Bricks	Lists all Bricks with a current status of <i>lost</i> or <i>unhealthy</i>
Up Bricks	Lists all Bricks with a current status of <i>up</i>
Bricks by Parent Folder	Lists all Bricks in the folder you select
Single Brick Status	Displays traffic and tunnel statistics for a selected Brick
Single Brick Ports	Displays data performance statistics for all configured ports of a selected Brick
Single Brick Bandwidth Statistics	Displays statistics about data packet traffic for the selected Brick <i>Note: The Enable Port Bandwidth Parameters checkbox must be checked on the Brick Ports Editor for one or more of the Brick ports to display this status information.</i>

The Bricks that are included in all of the Brick Lists except the Monitored Bricks List are determined by the current state of the Brick or parent folder. The Bricks in the Monitored Bricks List, however, are manually selected by the administrator. This is a convenience for administrators, who frequently find it useful to have a special list of Bricks that they want, for any of a number of reasons, to keep an eye on.

To create the list, select a Brick in any window that has the toolbar and click the **Monitored** checkbox in the toolbar. Repeat until all the Bricks you want are in the list.

To display the Monitored Bricks List, click the **Monitored** button in the toolbar of any Status Monitor window. To delete a Brick from the list, select the Brick and uncheck the **Monitored** checkbox.

You can also highlight a Brick in the list, right-click the Brick and select one of the following two options:

- Open Single Brick Window
- Open Single Brick Ports Window

These options display the same data that is displayed when you select Single Brick Window and Single Brick Ports Window from the Monitor menu.

Format and Contents

Figure 14-5, “Brick Lists (All Bricks)” (p. 14-16) shows a typical Brick List (All Bricks). All the Brick Lists share this same format.

Figure 14-5 Brick Lists (All Bricks)

Name	IP Address	State	CPU%	Sess%	Brick Version	Failover Labels	Interface Status	Mgmt Server	Roaming
attila	10.1.1.1	✓ Up	3%	0%	6.0.308		Up:0 Down:0		
genghis	10.1.1.2	✓ Up	3%	0%	6.0.308		Up:0 Down:0		
tamerlane	10.1.1.3	✗ Lost	3%	0%	6.0.342		Up:0 Down:0		

For each Brick in the list, the following information is provided:

Field	Description
Name	The name of the Brick
IP Address	The IP address of the Brick
State	Up or Lost
CPU%	The percentage of the Brick CPU currently in use
Sess%	The percentage of the Brick session cache currently in use

Field	Description
Brick Version	<p>If the Brick is a standalone, this is the version of the Brick operating system.</p> <p>If the Brick is part of a failover pair, this is the version of the active Brick operating system.</p>
Failover Labels	<p>If this Brick is part of a failover pair, the label consists of the last two octets of each Brick's MAC address. The label also indicates which Brick is active and which is standby.</p> <p>If this Brick is not configured as part of a failover pair, the failover label is <i>Not Configured</i>.</p>
Port Status	The number of ports up and down
Mgmt Server	The name of the LSMS to which this Brick is currently homed
Roaming	<p>Indicates whether this Brick is "roaming" from its priority 1 LSMS:</p> <ul style="list-style-type: none"> • No means it is currently connected to its priority 1 LSMS • Yes means it is currently connected to its priority 2 LSMS

Single Brick Status Window

The Single Brick Status window provides information about a specific selected Brick. To display the Single Brick Status window, open the Monitor menu, select **Brick Status**, and then select **Single Brick Status** from the submenu.

Figure 14-6, “Single Brick Status Window” (p. 14-18) shows a typical Single Brick Status Window.

Figure 14-6 Single Brick Status Window

Item	Status
State	Lost since Fri Dec 31 10:06:51 EST 2004
Errors	
Failover Labels	
Brick Version	8.0.205
Folder	/system/Bricks/
Management Server	
Port Status	0 Up - 4 Down
Passed Throughput (Mbits/sec)	0.000
CPU Utilization	
Session Cache Utilization	
LAN-LAN Tunnels	0 (Up: 0 Down: 0 Manual: 0 Disabled: 0)
Brick Sessions	0 (Passed: 0 Dropped: 0 Proxy: 0 NAT: 0 Tunneled: 0)
Authenticated Users	0 (Firewall: 0 Client VPN: 0)
New Sessions by Protocol	0 (UDP: 0 TCP: 0 ICMP: 0 IPSEC: 0)
Packet Traffic	In: 0 Out: 0
Byte Traffic	In: 0 Out: 0
Model	Model 80 or 201
Roaming	

For the Brick shown, the following information is provided:

Field	Description
State	Up or Lost.
Errors	Dot3, collision or frame errors.
Failover Labels	If this Brick is part of a failover pair, the label consists of the last two octets of each Brick's MAC address. The label also indicates which Brick is active and which is standby. If this Brick is not configured as part of a failover pair, this field is blank.

Field	Description
Brick Version	If the Brick is a standalone Brick, this is the version of the Brick's operating system. If the Brick is part of a failover pair, this is the version of the active Brick's operating system.
Folder	The folder in which the Brick is found.
Management Server	The name of the LSMS to which this Brick is currently homed.
Port Status	The number of ports up and down.
CPU Utilization	The percentage of the Brick CPU currently in use.
Session Cache Utilization	The percentage of the Brick's session cache currently in use.
LAN-LAN Tunnels	The total number of LAN-LAN tunnels of which this Brick is an endpoint, divided into categories [IKEv1 Up, IKEv1 Down, IKEv2 Up, IKEv2 Down, Manual, Disabled].
Brick Sessions	The total number of sessions through this Brick, divided into categories (Passed, Dropped, Proxy, NAT, Tunneled).
Authenticated Users	The total number of authenticated users, divided into Firewall, IKEv1, and IKEv2 categories.
New Sessions by Protocol	The total number of new sessions, divided into categories by protocol (examples: UDP, TCP, ICMP, IPSec).
Packet Traffic	The total number of packets handled by the Brick in the last refresh cycle, divided into <i>In</i> and <i>Out</i> categories. This makes it possible to view the activity on a port, as well as the total of all the ports.
Byte Traffic	The total number of bytes handled by the Brick, divided into <i>In</i> and <i>Out</i> categories. This makes it possible to view the activity on a port, as well as the total of all the port.
Model	The Brick model. The alternatives are <i>Model 20 or 50</i> , <i>Model 80 150 or 201</i> , <i>Model 300</i> , <i>Model 500 or Model 1000</i> , or <i>Model 1100</i> .
Roaming	Indicates whether this Brick is "roaming" from its priority 1 LSMS: <ul style="list-style-type: none"> <i>No</i> means it is currently connected to its priority 1 LSMS <i>Yes</i> means it is currently connected to its priority 2 LSMS

The Single Brick Status window also has a Graph menu on the menubar at the top. This menu allows you to display graphs of the following:

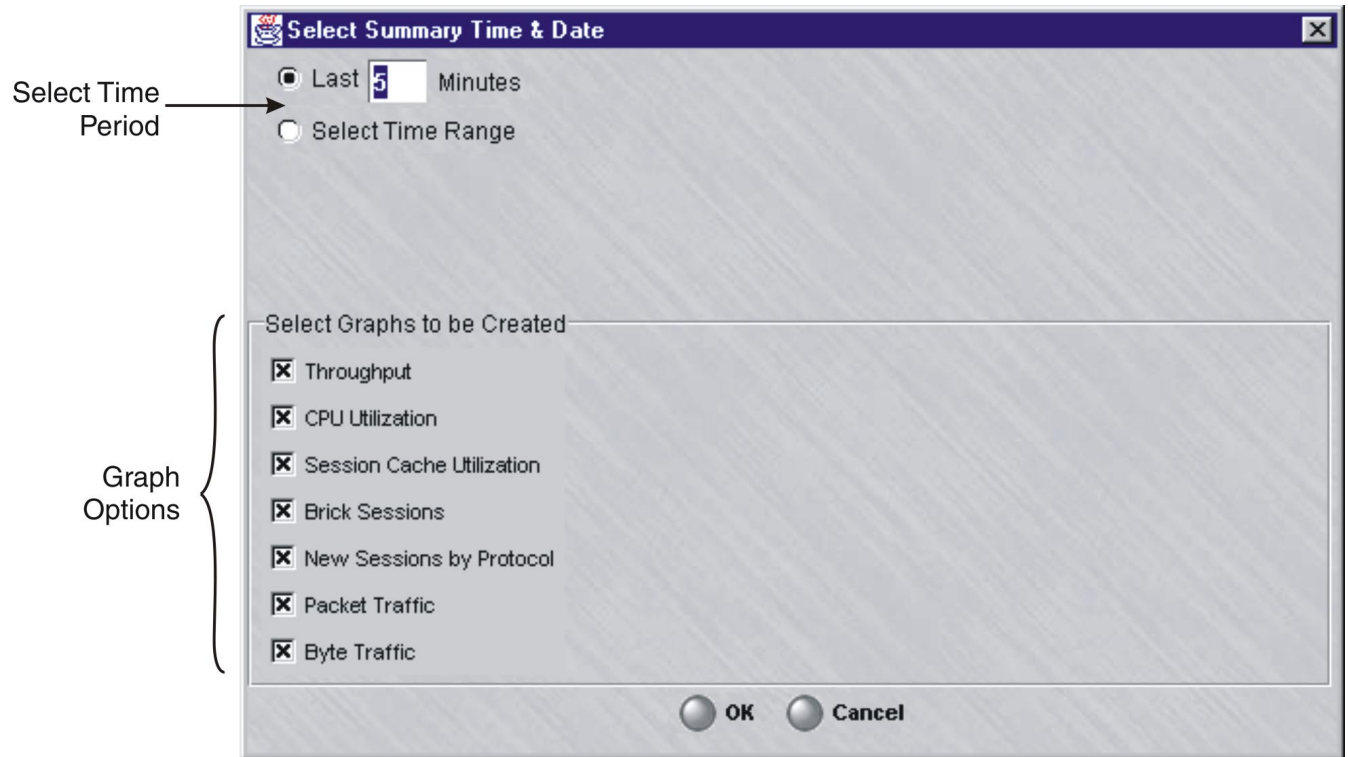
- Throughput
- CPU utilization
- Session cache utilization
- LAN-LAN tunnels (showing up, down manual and foreign tunnels)
- Brick sessions (showing passed, dropped, proxied, NATted and tunneled sessions)
- Authenticated users (showing firewall and Client VPN users)
- New sessions by protocol (showing UDP, TCP, ICMP and IPSec protocols)
- Packet traffic (showing in and out traffic)
- Byte traffic (showing in and out traffic)

If you want these graphs to reflect current Brick activity, select the appropriate graph from the Graph menu (or right-click in the Single Brick Status window and select the graph).

If you want the graph to display historical information, select **Summarize** from the Graph menu. The window shown in [Figure 14-7, “Select Summary Time and Date Window”](#) (p. 14-21) will appear. Indicate how many minutes previously you want the summary to cover (the default is 5), or click **Select Time Range** and enter a start date/time and end date/time.

Then, indicate which of the seven graphs you want created. By default, all checkboxes are checked, so all graphs will be created. When you click **OK**, a separate graph will be generated for each option checked.

Figure 14-7 Select Summary Time and Date Window



Single Brick Ports Window

The Single Brick Ports window provides information about the ports on a Brick™ you select. To display the Single Brick Ports window, open the Monitor menu, select **Brick Status**, and then select **Single Brick Ports**.

Figure 14-8, “Single Brick Ports Window” (p. 14-22) shows a typical Single Brick Ports window.

Figure 14-8 Single Brick Ports Window

The screenshot shows a window titled "Single Brick Ports - /system/Devices/Bricks/attila". The window has a menu bar (File, Edit, Monitor, Windows, Utilities, Graph, Help) and a toolbar with a Refresh button, a refresh interval dropdown set to "Every 30 secs", a timer showing "00:05", a "Monitored" button, a "Monitor" checkbox, and a clock showing "15:14:58". Below the toolbar, it displays "TOTAL INTERFACES: 4 [4 UP - 0 DOWN]". A table lists the statistics for each interface.

Interface	State	Bytes In	Bytes Out	Packets In	Packets Out	Error Packets In	Error Packets ...	Dot 3 Errors	Collision Errors	Frame Errors
ether0	✓ Up	4224	11404	66	110	0	0	0	0	0
ether1	✓ Up	2404	2404	31	31	0	0	0	0	0
ether2	✓ Up	2184	0	28	0	0	0	0	0	0
ether3	✓ Up	11514	7044	96	66	0	0	0	0	0

For each port, the window indicates whether it is up or down and displays the following information:

- Bytes in/out
- Packets in/out
- Error packets in/out
- Dot3, collision and frame errors.

The Single Brick Ports window also has a Graph menu on the menubar at the top. This menu allows you to display graphs of the following:

- Byte traffic
- Packet traffic
- Error traffic
- Errors.



If you want these graphs to reflect current port activity, select the appropriate graph from the Graph menu (or right-click in the Single Brick Status window and select the graph).

If you want the graph to display historical information, select **Summarize** from the Graph menu. A window similar to the one shown in [Figure 14-7, “Select Summary Time and Date Window”](#) (p. 14-21) will appear. Indicate how many minutes previously you want the summary to cover (the default is 5), or click **Select Time Range** and enter a start date/time and end date/time.

Then, indicate which of the four graphs you want created. By default, all checkboxes are checked, so all graphs will be created. When you click **OK**, a separate graph will be generated for each option checked.

At the bottom of each graph, there is a drop down list that lets you decide whether you want the graph to show all ports (the default), or one specific port.

In addition, each graph has two buttons to the right of the date:

- The **Legend** button  on the left acts as a toggle. Click it once to remove the legend from a graph so there is a larger graph area to view. Click it again to return the legend.
- The **Print** button  on the right allows you to print the graph. You must have a default printer defined for the print operation to work.

Single Brick Bandwidth Statistics Window

The Single Brick Bandwidth Statistics provides traffic statistics and performance information for each port of the selected Brick.

Important! The **Enable Port Bandwidth Parameters** checkbox must be checked on the Brick Ports Editor for each Brick port to display the port bandwidth statistics on this window. For instructions on how to configure Brick ports, refer to [Chapter 4, “Configuring Lucent VPN Firewall *Brick*® Device Ports”](#).

Figure 14-9, “Single Brick Zones Window” (p. 14-24) shows a sample Single Brick Bandwidth Zones window, displaying the bandwidth statistics for each Brick port.

Figure 14-9 Single Brick Zones Window

The screenshot shows a window titled "Single Brick Zones - /system/Devices/Bricks/150-twin". The window has a menu bar with "File", "Edit", "Monitor", "Windows", "Utilities", "Graph", and "Help". Below the menu bar is a toolbar with a "Refresh" button, a dropdown menu set to "Every 30 secs", a timer showing "00:18", a "Monitored" checkbox, and another "Monitor" checkbox. The main area contains a table with the following data:

Port	Zone	Passed Throughput In (Mbits/sec)	Passed Throughput ... (Mbits/sec)	Packets/sec In	Packets/sec Out	New Sessions/sec In	New Sessions/sec Out
ether0	pass-all	22.338	12.470	4969	4544	0	0
ether1	e1_vpn	2.419	17.967	3927	1657	390	16
ether3	e3_vpn	5.298	0.027	541	54	4	4
ether2	e2_vpn	14.248	0.212	1531	292	4	6
ether0	administrativezone	1.358	0.063	283	146	0	0

For each port, the window displays the following information by zone:

- Data packet throughput into the Brick (Mbits/sec)
- Data packet throughput out of the Brick (Mbits/sec)
- Packets per second into the Brick
- Packets per second out of the Brick
- New sessions per second into the Brick
- New sessions per second out of the Brick
- Megabit guarantee into the Brick
- Megabit limit into the Brick
- Session limit into the Brick
- Megabit guarantee out of the Brick
- Megabit limit out of the Brick
- Session limit out of the Brick

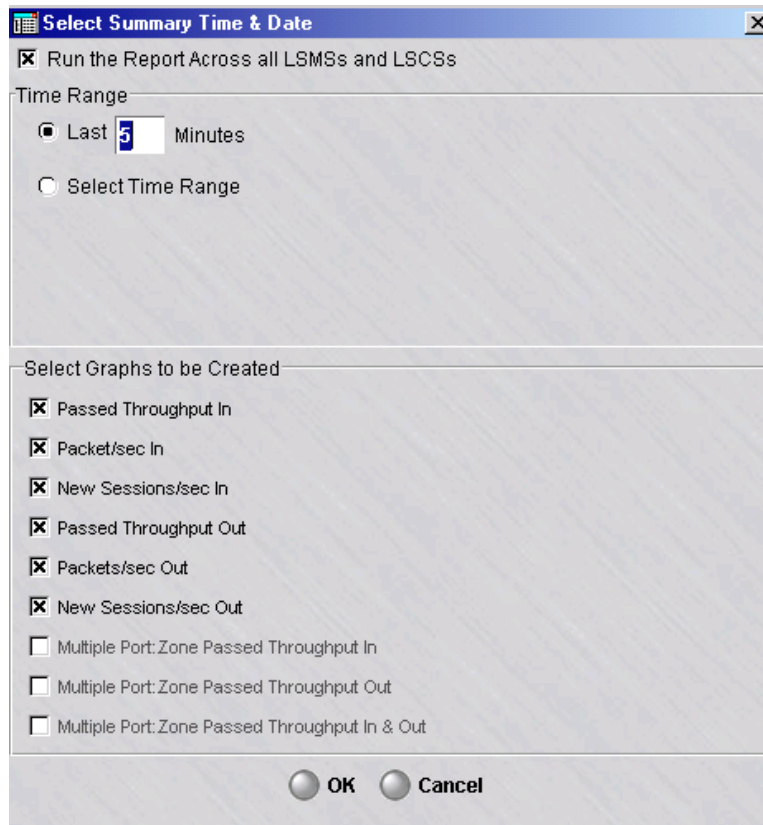
The Single Brick Zones window also has a Graph menu on its menu bar, to display a graphical representation of the data, according to user-specified criteria.

To select a graph that shows current Brick activity, select the appropriate graph from the Graph menu or right-click in the Single Brick Zones window and select the graph.

To display historical data graphically, select **Summarize** from the Graph menu.

The Select Summary Time & Date window is displayed (

Figure 14-10 Select Summary Time and Date Window (Brick Bandwidth Statistics)




Indicate the prior time period for the summary (default is five minutes), or click **Select Time Range** and enter a start and end date/time.

Indicate which graphs that you want to display, and click **OK**. A separate graph is generated for each option checked.

shows a sample graphical representation of the Single Brick Zones port bandwidth statistics.

Figure 14-11 Single Brick Zones Graphical Display



Console Alarms Window

Console alarms window display

The Console Alarms window displays all alarms that have been configured to send a message to the LSMS console. To display the Console Alarms window, open the Monitor menu and select **Console Alarms**.


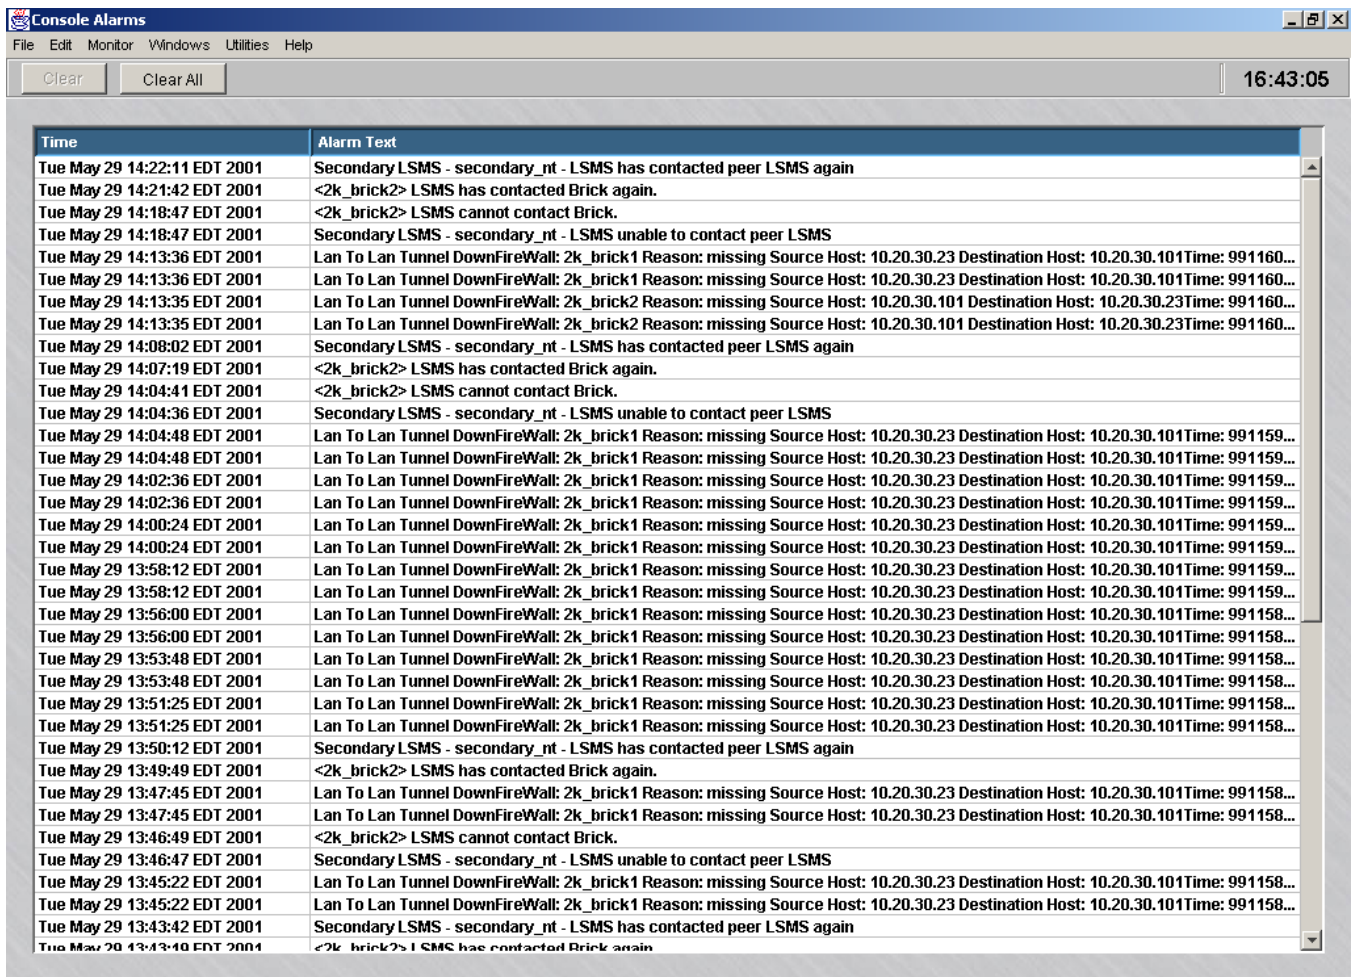
The yellow bell shown below will appear at the top of the LSMS window currently open when a new console alarm has been triggered. You can display the Console Alarms window by clicking this icon. 

Figure 14-12, “Console Alarms Window” (p. 14-27) shows a typical Console Alarms Window. For each alarm, the window shows the data and time, and the text of the console message.

Figure 14-12 Console Alarms Window



Appendix A: Administer a Lucent VPN Firewall *Brick*[®] Device Over the Internet from an Unregistered LSMS

Overview

Purpose

This chapter explains how to enable an LSMS with an unregistered IP address to administer a Brick over the Internet.

Contents

Background	A-2
Configure the Brick	A-3
Assign the Administrative Zone and Enter a VBA	A-4
Add NAT Rules to the administrativezone Ruleset	A-5
Activate the Remote Brick	A-8



Background

Remote administration of Bricks

There are a number of circumstances in which an LSMS with a private address might have to administer a Brick remotely over the Internet. For example, an LSMS may have been initially set up to manage several Bricks that were connected directly to its ports. In this case, the private address would not prevent the LSMS from communicating with these Bricks.

However, network growth could require the deployment of Bricks in other geographic locations. In this case, the LSMS would need to use the Internet to administer these Bricks, and the unregistered address would be a problem.

The solution is to employ the Brick network address translation (NAT) capability to convert the LSMS' private address to a registered address that can be used on the Internet.

To do this, you must first obtain a registered IP address for the LSMS to use. Once this is done you have to:

- Configure the remote Brick, making sure to enter the *registered* IP address as the LSMS IP address, not the private IP address you have been using
- Assign the *administrativezone* ruleset to the port connecting the Brick to the LSMS, and enter a virtual Brick address (VBA)
- Add three NAT rules to the *administrativezone* ruleset
- Modify the LSMS host group to include the registered address you will now be using for the LSMS
- Activate the remote Brick.

□

Configure the Brick

Task

To configure the remote Brick, follow the steps below. (For a more detailed explanation of the configuration process, refer to [Chapter 3, “Configuring and Activating a Lucent VPN Firewall Brick® Device”](#).)

- 1 From the Navigator window, right-click the Bricks folder and select **New** from the pop-up menu. The Brick Editor will appear, with the Brick tab displayed.
- 2 Enter the Brick name, IP address and, if necessary, gateway IP address. The unregistered address you have been using will appear in the **LSMS IP Address** field if this is a standalone LSMS. Replace the unregistered address with the registered address.
- 3 Open the File menu and select **Save** to save the configuration.
- 4 Open the Brick Utilities men, and select **Make/Package Floppy**. Follow the on-screen instructions to make a floppy disk on the LSMS host, or to package the files for remote floppy creation. *However, do NOT load the floppy disk into the Brick until you are instructed to do so.*

END OF STEPS



Assign the Administrative Zone and Enter a VBA

When to use

The *administrativezone* ruleset is a pre-configured ruleset that allows the Bricks and LSMS to communicate, while protecting the LSMS from attack.

You will have to apply this ruleset to the port on the Brick that connects the Brick to the LSMS. You will also have to create a Virtual Brick address (VBA) and make the VBA the registered address you will now be using for the LSMS. Complete the following steps to do this:

- 1 With the Brick Editor displayed, click **Policy Assignment** to display the Policy Assignment tab.
- 2 Double-click the port connecting the Brick and LSMS. Since this Brick will be administered remotely, that should be the port connected to the router that Brick will use to communicate with the LSMS. The Brick Policy Assignment Editor is displayed.
- 3 In the **Zone Ruleset** field, select **administrativezone** from the drop-down list.
- 4 In the **Tunnel Endpoint Address/Virtual Brick Address** field, enter the registered IP address you will be using. If this is a standalone LSMS, this must be the same address you entered in the **LSMS IP Address** field when configuring the Brick.
- 5 Click **OK** to dismiss the Brick Policy Assignment Editor. You can ignore the other fields on the Brick Policy Assignment Editor for now.
- 6 Open the File menu and select **Save** to save the assignment.

END OF STEPS



Add NAT Rules to the *administrativezone* Ruleset

When to use

Once the *administrativezone* ruleset has been assigned to a port and the VBA has been entered, you have to enable and edit three rules in the ruleset to perform the required address translation:

- Rule 209
- Rule 210
- Rule 211

The first rule (209) performs destination address mapping. These rules will instruct the Brick to map all inbound sessions destined for the VBA (the registered IP address) to the unregistered IP address of the LSMS.

The second and third rules (210, 211) perform source address mapping. These rules instruct the Brick to map the source address of all outbound sessions (the unregistered IP address) to the VBA, so that the return sessions automatically have the VBA as their destination.

Task

To enable and edit these rules, follow the steps below. These rules will be made active in the *administrativezone* ruleset.

- 1 With the Navigator window displayed, click the appropriate Brick Zone Rulesets folder to display all existing Brick zone rulesets, and double-click **administrativezone**. The Brick Zone Ruleset Editor will appear, with the *administrativezone* rules displayed.
- 2 Double-click on Rule 209 in the rules table.

The Brick Zone Rule Editor is displayed with the Basic tab for Rule 209. Configure the rule as follows and make sure that it is active:

Direction	Source	Destination	Service	Action
In To Zone	brickRemoteAddresses	Virtual Brick Address	brick_to_SMS_Services	Pass

brickRemoteAddresses is a host group that can be selected from the drop-down list in the **Source** field. *Virtual Brick Address* is a keyword that can be selected from the drop-down list in the **Destination** field. *brick_to_SMS_Services* is a service group that can be selected from the drop-down list in the **Service or Group** field.

This rule will allow the Brick to send audit data to the LSMS and request policy downloads from the LSMS, using the registered address (the VBA) instead of the unregistered address that had been used previously.

-
- 3 Click **Address Translation** to display the Address Translation tab. Then, enter the LSMS' unregistered IP address in the **Destination Address Mapping** box (you may leave the mapping type as pool). Click **OK** to dismiss the Brick Zone Rule Editor and return to the Basic Tab of the Brick Zone Ruleset Editor.

-
- 4 Double-click on Rule 210.

The Brick Zone Rule Editor is displayed with the Basic tab for Rule 210. Configure the rule as follows and make sure that it is active:

Direction	Source	Destination	Service	Action
Out of Zone	LSMS	brickLocalAddresses	brick_from_SMS_Services	Pass

LSMS is a host group that can be selected from the drop-down list in the **Source** field. *brickLocalAddresses* is a host group that can be selected from the drop-down list in the **Destination** field. *brick_from_SMS_Services* is a service group that can be selected from the drop-down list in the **Service or Group** field.

This rule allows the LSMS to upload policy and configuration information to the LSMS, using the registered address (the VBA) instead of the unregistered address that had been used previously.

-
- 5 Double-click on Rule 211.

The Brick Zone Rule Editor is displayed with the Basic tab for Rule 211. Configure the rule as follows and make sure that it is active:

Direction	Source	Destination	Service	Action
In To Zone	brickRemoteAddresses	Virtual Brick Address	tcp/*/910	Pass

This is the same as the rule you just created, except that the service is different. The purpose of this rule is to allow policy download replies from the Brick to the LSMS when the Clear Cache option is invoked.

You also have to perform the same destination address mapping for this rule that you did for the Rule 209. Open the Address Translation tab and enter the LSMS unregistered IP address in the **Destination Address Mapping** box, as was done in Step 3.

-
- 6 Click **Address Translation** to display the Address Translation tab. Then, select the key word **Virtual Brick Address** from the drop-down list in the **Source Address Mapping** box (you may leave the mapping type as pool). Click **OK** to dismiss the Brick Zone Rule Editor and return to the Basic Tab of the Brick Zone Ruleset Editor.

-
- 7 Open the File menu and select **Save and Apply**, and then dismiss the Brick Zone Ruleset Editor.

END OF STEPS



Activate the Remote Brick

Task

You are now ready to activate the Brick, using the floppy you created on the LSMS host or on a remote host. To do this, follow the steps below:

- 1 Insert the floppy disk into the disk drive of the Brick.

- 2 Power up the Brick by toggling the Brick's power switch. The configuration information on the disk will be transferred to the Brick flash disk. The transfer process takes about 2.5 minutes.

- 3 When the transfer is complete, remove the floppy disk from the disk drive and power the Brick off and on to boot the Brick from its flash disk. The Brick is now ready to be deployed.

Important! If additional Bricks will be administered by this LSMS over the Internet, do the following:

1. Add the IP addresses of the new Bricks to the *BrickRemoteAddresses* host group.
2. Apply the policy to the new Brick.
3. Configure the new Brick, create a floppy, and activate the Brick using the floppy.

END OF STEPS



Appendix B: Sizing Guidelines

Overview

Purpose

To determine the proper size of the PC or host machine that is required to efficiently run the LSMS software, you need to understand the following:

- Resource utilization and behavior of the software
- Performance of the computer hardware
- Maximum user response time requirements
- How the system will be used.

Contents

Sizing Tool	B-2
Determine CPU Capacity	B-4
Memory Utilization	B-6
Disk Capacity for Log Files	B-7
Disk Configuration	B-8



Sizing Tool

Purpose of sizing tool

The LSMS includes a sizing tool that shows the current status of all LSMS processes. This is a very useful tool to help size the capacity of the hardware the LSMS is running on.

For example, if the amount of memory used for a particular service is nearing its maximum, you may need to go to the Tunable Parameters in the Configuration Assistant and increase the memory allocation for this service. For a more detailed discussion of the Configuration Assistant, see [Chapter 11, “Using the Configuration Assistant”](#) (in particular, ““TL1 Alarms” (p. 11-43)”).

Solaris Platform

On Solaris, there are two versions of this tool, a graphical version and a command line version. The only difference between the two is that the information is presented in graphical format on the graphical version.

The invocation method for each version is as follows:

- *Command Line*
To invoke the command line version, go to the directory `<ISMSHOME>` and enter `./lsmsStatus`
- *Graphical*
To invoke the graphical version, display the **Desktop** menu, select **Utilities**, and select the tool, or go to the directory `<ISMSHOME>` and enter `./StartLSMSStatus`

Windows Platform

On Windows, only the graphical version the tool is available. To invoke it, open the **START** menu, select **Utilities**, and select the tool.

Contents

The command line and graphical versions of the tool provide the following information:

Field	Description
LSMS Service	Name of the LSMS service
jre PID	Process ID of the service
Threads	Number of threads currently in the service

Field	Description
Active DB Conns	Current number of open database connections used by the service
Active Trans	Number of active database transactions in progress at the time the screen refreshed
Max Active Trans	Maximum number of active database transaction in progress at one time since the services were started
% CPU	Percent CPU utilization of this service
Maximum Heap	Maximum allocation for the heap memory of this service
Allocated Heap	Heap memory currently allocated by this service
Used Heap	Heap memory currently used by this service
Total CPU Time	Total CPU time used up by this service since it started
Last Started At	The last time this service was started (only available on graphical version)
Heap Snapshot At	Last time this service's allocated heap memory information was refreshed (only available on graphical version)

In addition, the tool also displays:

- Records logged by Lucent VPN Firewall *Brick*[®] devices to this LSMS
- Number of Bricks connected to this LSMS
- The time the LSMS Status window was invoked (this window can remain up even if you've stopped services)
- The last time that the number of Bricks connected to this LSMS changed.



Determine CPU Capacity

Overview

To determine the processor required by the machine running the LSMS software, analyze the following choices and choose the one that has the highest value.

- Windows — 400 MHz Pentium II processor or Solaris — 333 MHz Solaris Ultra Sparc
- Calculate the result of this formula: (2 MHz * the number of Bricks).
For example, if you have 350 Bricks, then a machine offering 700 MHz may best suit your needs. Acquire a machine that provides this processing power (or comes the closest) and is available on the market today.
- Calculate the result of this formula:
 - If you are using HTTP, then:
CPU Speed (MHz) = (L * 0.55) + (U * 50) + (V * 170*) + (A * 10)
 - If you are using HTTPS, then:
CPU Speed (MHz) = (L * 0.55) + (U * 133) + (V * 170*) + (A * 10)

* If using RADIUS, Local Password, or SecurID, the number is 170. If using a digital certificate, this number then becomes 250.

where:

L = Total network traffic through all Bricks to be logged in megabits/second—e.g., 200.

U = User authentications per second during "busy" time.

V = VPN negotiations per second during "busy" time.

A = Maximum simultaneous Administrators logged in.
--

Example

If your configuration includes 10 Bricks supporting 4 T3's and 12 T1's and each connection operates at an average of 50% utilization (full duplex), then the total network traffic is just under 200 Mbps.

User authentications and VPN negotiations include both initiated and failed. For this example, suppose you are using HTTPS and estimate 2 authentications per second and 1 VPN negotiation per second and you have two Administrators who can log in simultaneously.

In this example, the processor required for the machine running the LSMS software is 566 MHz.

Since 566 MHz exceeds a 400 MHz Pentium II processor, acquire a machine that provides this processing power (or comes the closest) and that is available on the market today.

$$\text{processor} = (200 * 0.55) + (2 * 133) + (1 * 170) + (2 * 10) \text{ 566 MHz} = 110 + 266 + 170 + 20$$

Summary

The above guidelines should be sufficient for most installations.

Naturally, they will depend on the amount of administrator activity and network traffic mix.

Important! *Logging DROPS only*

Logging dropped traffic only dramatically reduces the amount of log traffic. Many Administrators find this sufficient.



Memory Utilization

Overview

To determine the amount of RAM required by the machine running the LSMS software, analyze the following choices and choose the one that has the highest value.

- 256 MB
- Calculate the result of this formula:
Physical Memory (MB) = $(170 + F/4) + (L * .05) + R/25 + (\# \text{ tunnels} * .06)$

where:

F = Number of Brick(s) in the network.
L = Network traffic through a Brick to be logged in megabits/second — Example: 200.
R = Number of routers.

Example

If your configuration includes 10 Bricks supporting 4 T3's and 12 T1's and each connection operates at an average of 50% utilization (full duplex), then the total network traffic is just under 200 Mbps.

In this case, the memory required by the machine running the LSMS software is **182.5 MB** = $172.5 + 10$.

If you are using Internet Explorer as your browser, assume you will need approximately 32 MB more of memory.

Important! *Virtual Memory*

Virtual memory should be at least twice the size of physical memory.



Disk Capacity for Log Files

Use the following guidelines to determine how much disk space you need to allocate for the audit logs:

- Each log record occupies 100 bytes of disk space.
- Each megabit of traffic that is audited generates 10 session records per second. If the traffic travels through a VPN tunnel, an additional five records per second are audited. If only drops are audited, the session records will drop significantly. In this case, you may see only about 1/2 session records per second per megabit.
- Each user authentication request generates 1 to 2 records.
- Each VPN session generates 2 to 6 records.
- Proactive monitoring traffic amounts to 20 records per Brick per minute, or about 2KB per Brick per minute.
- Allow an extra 10 percent to minimize fragmentation and an extra 10 percent for administrative events.



Disk Configuration

Use the following chart for guidelines on how to configure your disk given a logging records-per-second range:

Records per Second	Windows	Solaris (May Require a tune of the file system)
< 2,000	Single drive	Single drive
2,000 - 5,000	At least one drive with a minimum 4K cluster size.	Single drive
5,000 - 10,000	One drive with a 32K or higher cluster size or multiple drives with striping. The LSMS should be installed on a separate partition from the boot drive (i.e., not on C:\).	2-3 drives with striping
10,000 - 15,000	2-3 drives with striping and 32K cluster size	3-4 drives with striping
15,000 - 20,000	4 or more drives with striping and a 64K cluster size. Consider a high performance disk array.	4 or more drives with striping
20,000 - 30,000	While this data rate can be supported with a high performance disk subsystem, it must be done with extreme care. An inadequate disk subsystem can put the LSMS in a state in which it cannot catch up with the incoming data.	

General disk configuration guidelines include:

- When the FTP scheduler is used at a high data rate, it should be scheduled either:
 - for times when no one is expected to be logged in
 - or
 - to run very frequently (every 1 to 2 minutes) so that the data is still in the memory cache.
- The disk drive should be a minimum of 4 GB, 7200 RPM, and <10 msec seek time. For rates above 5000 records per second, faster disk drives are highly desirable (e.g., 10,000 RPM or better).
- Use defragmentation with care. Some tools may interfere with the normal operation of the LSMS by running at too high a priority level.

□

Appendix C: Changing the IP Address of the LSMS

Overview

Purpose

This chapter explains how to change the IP address of the LSMS on both the Windows and Solaris platforms. A change in the IP address of the LSMS may be necessary if you are rearranging the elements of your network.

A utility called "changeIP" was created to simplify the process. The procedure is the same for both Solaris and Windows LSMS. However, the steps differ slightly, depending on whether the LSMS is a Primary-only LSMS or a Primary/Secondary pair.

Follow the procedures exactly as described in this Appendix, and you will not have to reboot Lucent VPN Firewall *Brick*[®] devices or interrupt service to your user community.

Contents

Procedure (Primary-only LSMS)	C-2
Procedure (Primary/Secondary Redundant Pair)	C-4
After the Update	C-6



Procedure (Primary-only LSMS)

Task

To change the IP address of an LSMS, you must log into the LSMS Navigator *directly from the LSMS host for which the IP address will be changed*. Once you have logged in, use the following procedure to change the IP address:

- 1 Update the "LSMS" host group** — From the Navigator, open the **Host Groups** subfolder under the **Policies** folder. Locate the **LSMS** host group and double-click on it to open it for edit. Click on the "+" radio button, enter the new IP address to which the LSMS will be changed, and click **OK**. Save and Apply this change.

The LSMS host group should now contain two addresses: the current IP address and the proposed new IP address.

- 2 Update the LSMS information for each Brick** — Open the **Bricks** folder and double-click on a Brick to open it for edit. In the **LSMS IP Address** field, replace the current IP address with the new IP address for the LSMS. Save and Apply the change. Repeat this step for each Brick controlled by this LSMS until all Bricks have been modified.

Important! *If multiple Bricks are configured to this LSMS, be sure to update the Brick directly in front of the LSMS last.*

- 3 Exit the LSMS Navigator.**
-

- 4 Stop LSMS Services.**
-

- 5 Change the IP address of the LSMS host machine.**
-

- 6 Execute the changelP utility on the LSMS** — Open a DOS command window (*Windows*[™]) or a terminal window (*Solaris*[®]) and connect to the LSMS installation directory. In a *Windows*[™] environment, this might typically be *c:\users\isms\lmf*. In a *Solaris*[®] environment, this might typically be */opt/isms/lmf*. Run the changelP utility as follows, where *x.x.x.x* is the new IP address of the LSMS:

Windows[®]: *local\bin\changeIP x.x.x.x*

Solaris[®]: */local/bin/changeIP x.x.x.x*

7 Start LSMS Services.

8 *Log into the **LSMS** Navigator and use **Monitor > Brick Status > All Bricks** to verify that all Bricks are active.*

END OF STEPS



Procedure (Primary/Secondary Redundant Pair)

Task

To change the IP address of an LSMS, you must log into the LSMS Navigator *directly from the LSMS host for which the IP address will be changed*. Once you have logged in, use the following procedure to change the IP address:

- 1 Update the "LSMS" host group** — From the Navigator, open the **LSMS/LSCS** folder under the **Policies** and double-click on the LSMS (Primary or Secondary) for which the IP address is to be changed to open it for edit. In the **IP Address - IP Public** field, enter the new IP address. Use **FILE**, **SAVE** and **CLOSE** to save this entry.

- 2 Update the LSMS information for each Brick** — Open the **Bricks** folder and double-click on a Brick to open it for edit. In the **Home LSMS/LSCS Priority** section of the Brick Editor, locate the LSMS for which the address is to be changed and double-click on this entry to open it for edit. Double-click the **LSMS/LSCS IP Address** field and in the ensuing dialog window, click the drop-down arrow and select the new IP address. Click **OK**, then **Save and Apply** the change to the Brick. Repeat this step for all Bricks that are configured for this LSMS.

Important! *If multiple Bricks are configured to this LSMS, be sure to update the Brick directly in front of the LSMS last.*

- 3 Exit the LSMS Navigator.**

- 4 Stop LSMS Services** — When the Services are stopped, any Bricks "homed" to the LSMS for which the address is being changed will lose contact with this LSMS and "rehome" to the other LSMS in the redundant pair.

- 5 Change the IP address of the LSMS host machine.**

- 6 Execute the changelP utility on the LSMS** — Open a DOS command window (*Windows*[™]) or a terminal window (*Solaris*[®]) and connect to the LSMS installation directory. In a *Windows*[™] environment, this might typically be *c:\users\isms\lmf*. In a *Solaris*[®] environment, this might typically be */opt/isms/lmf*. Run the changeIP utility as follows, where *x.x.x.x* is the new IP address of the LSMS:

Windows[™]: *local\bin\changeIP x.x.x.x*

Solaris™: /local/bin/changeIP x.x.x.x

Status messages are displayed while the utility is updating the database with the new IP address.

-
- 7 Start LSMS Services** — If the LSMS for which the address has been changed is the higher-priority LSMS in the redundant pair, the Brick will now “rehome” to this LSMS when the Services are started.

-
- 8 Log into the LSMS Navigator and use Monitor > Brick Status > All Bricks to verify that all Bricks are active.**

-
- 9 Run the dbsetup utility** — dbsetup must be executed on the redundant LSMS (the LSMS for which the address was not changed).

On the redundant LSMS:

1. Exit the LSMS Navigator (if running).
2. Stop LSMS Services.
3. Open a DOS window and connect to the LSMS installation directory (such as *c:\users\isms*). On a *Solaris®* host, open a terminal window and connect to the LSMS installation directory (such as */opt/isms*).
4. Execute dbsetup as follows:
local\bin\dbsetup (*Windows™*)
/local/bin/dbsetup (*Solaris®*)
5. When dbsetup completes, start LSMS Services.
6. Log into the LSMS and use **Monitor > LSMS/LSCS** to verify that the Primary and Secondary LSMS are connected

.....

END OF STEPS



After the Update

Post-update activities

Once you have updated the Bricks and the LSMS database and the LSMS services have been restarted, what else needs to be done?

The remaining things to be done are:

- **Primary LSMS** — Once services have been restarted, the Bricks that are normally "homed" to the primary LSMS should start "rehomeing" from the secondary LSMS to the primary. This will depend on the settings on the Brick under "LSMS Rehome Options" under the Brick tab on the Brick Editor.
In addition, the database on the primary LSMS can no longer synchronize with the database on the secondary LSMS. On the secondary LSMS, from the LSMS installation directory, you must run a `dbsetup`. This will reinitialize the database on the secondary and copy the database from the primary to the secondary.
For more information on "dbsetup", refer to *Chapter 6. Database Utilities* in the *LSMS Tools and Troubleshooting Guide*.
- **Secondary LSMS** — Once services have been restarted, the Bricks that are normally "homed" to the secondary LSMS should start "rehomeing" from the primary LSMS to the secondary. This will depend on the settings on the Brick under "LSMS Rehome Options" under the Brick tab on the Brick Editor.
The database on the secondary will automatically update the database on the primary with the change to its LSMS IP address.

□

Appendix D: Support for Non-IP Protocols

Overview

Purpose

The Lucent VPN Firewall *Brick*[®] device is originally designed to work with IP packets. Information within the IP packet is used to make access control decisions, such as whether or not to allow this packet to pass through the Brick.

In addition to handling IP packets, a Brick can be also be configured to pass packets from non-IP protocols, for example, AppleTalk, Novell IPC, and Decnet/LAT. For example, you may have a Brick between an Apple client and server that need to talk to one another.

Non-IP protocol packets

Important! CAUTION

Filtering is not performed on these individual packets. Either they ALL are entirely allowed or denied.

To pass non-IP protocol packets, you need to edit a file on the LSMS host. The file then needs to be applied to the Brick that will be passing the non-IP packets.

Contents

Ethertype and DSAP Files	D-2
Procedure for Passing Non-IP Packets	D-3



Ethertype and DSAP Files

Non-IP protocol file editing

To support the passing of non-IP protocols, you need to edit one or both of these files:

- *ethertype*
The *ethertype* file is used to support EV2 ethertype protocols. These protocols are represented by a four-digit hexadecimal number.
- *dsap*
The *dsap* file supports 802.2 DSAP protocols. These protocols are represented by a two-digit hexadecimal number.

File Location

For each Brick in your network, the *ethertype* and *dsap* files reside in a folder under the root installation directory.

- On Windows, typically the folder is:
c:\users\isms\lmf\firewalls\<Brick_name>
- On Solaris, typically the folder is:
/opt/isms/lmf/firewalls/<Brick_name>

Obtain Hexadecimal Number

Before editing the *ethertype* or *dsap* file, you need to identify the hexadecimal number that represents the protocol you need to support. You can either:

- Access this URL to obtain a current list of ethertypes:
http://www.iana.org
The reference for this site is RFC 1700, "Assigned Numbers", J. Reynolds, J. Postel, October 1994.
- Attach a packet sniffer to your network and examine the packets that are generated when attempting to pass packets of the protocol.



Procedure for Passing Non-IP Packets

Task

To pass non-IP protocol packets (for example, AppleTalk packets), you must:

- 1 Edit a file named *ethertype* on the LSMS. Save and exit the file.
 - 2 Apply the changes to the Brick. Refer to [Chapter 3, “Configuring and Activating a Lucent VPN Firewall Brick® Device”](#) for details on applying changes to a Brick.
-

END OF STEPS

Edit ethertype File on LSMS

Use the following procedure to edit a file so that the Brick will pass AppleTalk packets:

- 1 For each Brick that needs to pass non-IP protocol packets, you need to navigate to a folder under the installation root directory folder.
 - On Windows, typically the folder is:
c:\users\isms\lms\firewalls\<Brick_name>
 - On Solaris, typically the folder is:
/opt/isms/lms/firewalls/<Brick_name>
 - 2 Open the *ethertype* file with an ASCII editor on the platform you are using, for example, Notepad on Windows, vi on Solaris.
 - 3 To support AppleTalk packets, enter the following two four-digit hexadecimal numbers on separate lines in the *ethertype* file:
-

809B

80F3

Important! DSAP Values

Since AppleTalk generates AA DSAP types, you do not need to edit the dsap file. This applies to all non-IP protocols that generate AA types for ethertypes 0x5FE or greater. For these ethertypes or for any other DSAP other than AA, you must explicitly enter them in the dsap file in the format dsap, ethertype1, ethertype2, etc.

.....
4 Save and exit the file.
.....

5 Repeat Steps 1 - 4 for each Brick that needs to pass AppleTalk packets.

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

Apply Changes to the Brick

Perform the steps below to load the new configuration information to a Brick:

.....
1 Log onto the LSMS host.

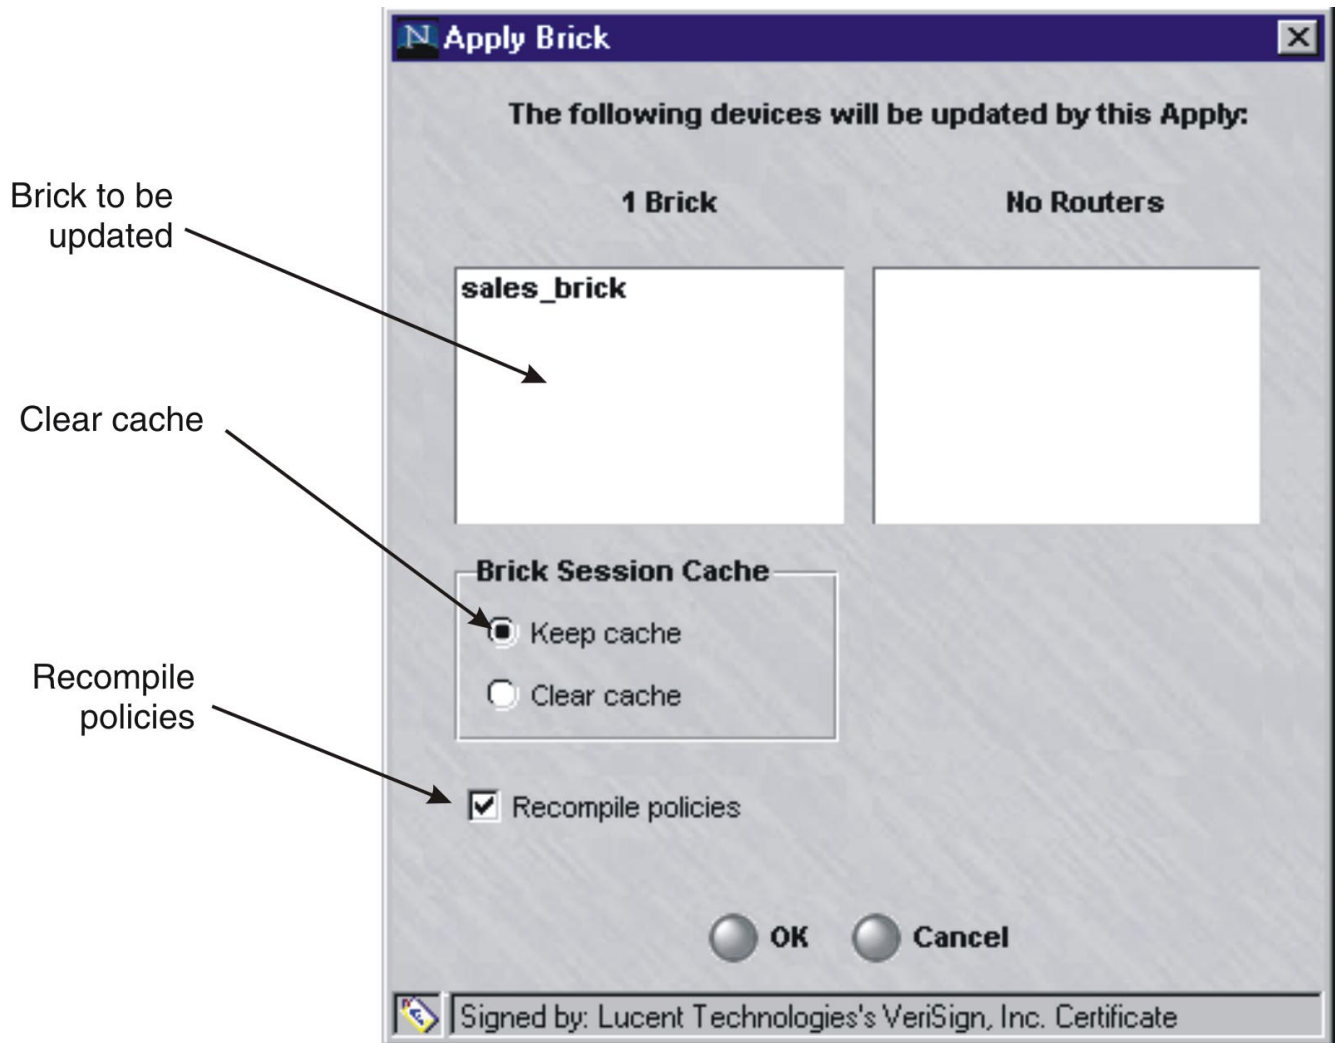
For details, refer to *Chapter 2. Getting Started*.

.....
2 If the Brick is currently displayed in the Brick Editor, display the Utilities menu and select **Brick ► Apply**.

If the Bricks are displayed in the Navigator window, right-click the Brick you want and click **Apply** from the pop-up menu.

The Apply Brick window will appear. It is shown in Figure D-1. Note that the Brick you are applying (updating) appears in the left panel on top.

Figure D-1 Apply Brick Window



- 3 When performing the apply, you have the option of keeping or clearing the Brick session cache. The default is to keep the cache. To clear the cache, click the **Clear Cache** radio button.

If you decide to clear the cache, you will terminate any client tunnels currently established to this Brick. You will have to contact the client users and instruct them to re-enable their tunnels.

You could also disrupt some sessions in progress, for example, FTP sessions or sessions allowed by rules with dependency masks.

-
- 4 You have the option of recompiling the policies associated with this Brick before applying the Brick. The default is to recompile the policies. If you do not want to recompile them, uncheck the **Recompile Policies** checkbox.

If you recompile the policies, all changes to the policies will be applied to the Brick. Therefore, if you want to update the Brick configuration, but *not apply any policy changes at this time*, uncheck the **Recompile Policies** checkbox.

- 5 When you are ready to begin the apply, click **OK** to dismiss the Apply Brick window. The apply will take place.
-

- 6 Repeat steps 2 - 5 for each Brick that needs to pass non-IP protocol packets.

END OF STEPS



Appendix E: New Feature Setup

Overview

Purpose

By default, when you purchase LSMS R9.1, your installation key(s) will provide you with the ability to manage up to 25 managed Lucent VPN Firewall *Brick*® devices as well as up to 100 IPsec Client users.

If you wish to expand the number of Bricks or IPsec users to be managed through your LSMS, you must purchase a separate options license key. These options license keys will also only be valid for specific LSMS installations and not valid for other LSMS installations. For redundant scenarios, two option license keys are needed for each LSMS in order to enable purchased options.

The purpose of "New Feature Setup" is to allow an administrator to install a new key and provide the additional LSMS management capacity.

Contents

What Do You Have Now?	E-2
Using New Feature Setup	E-4



What Do You Have Now?

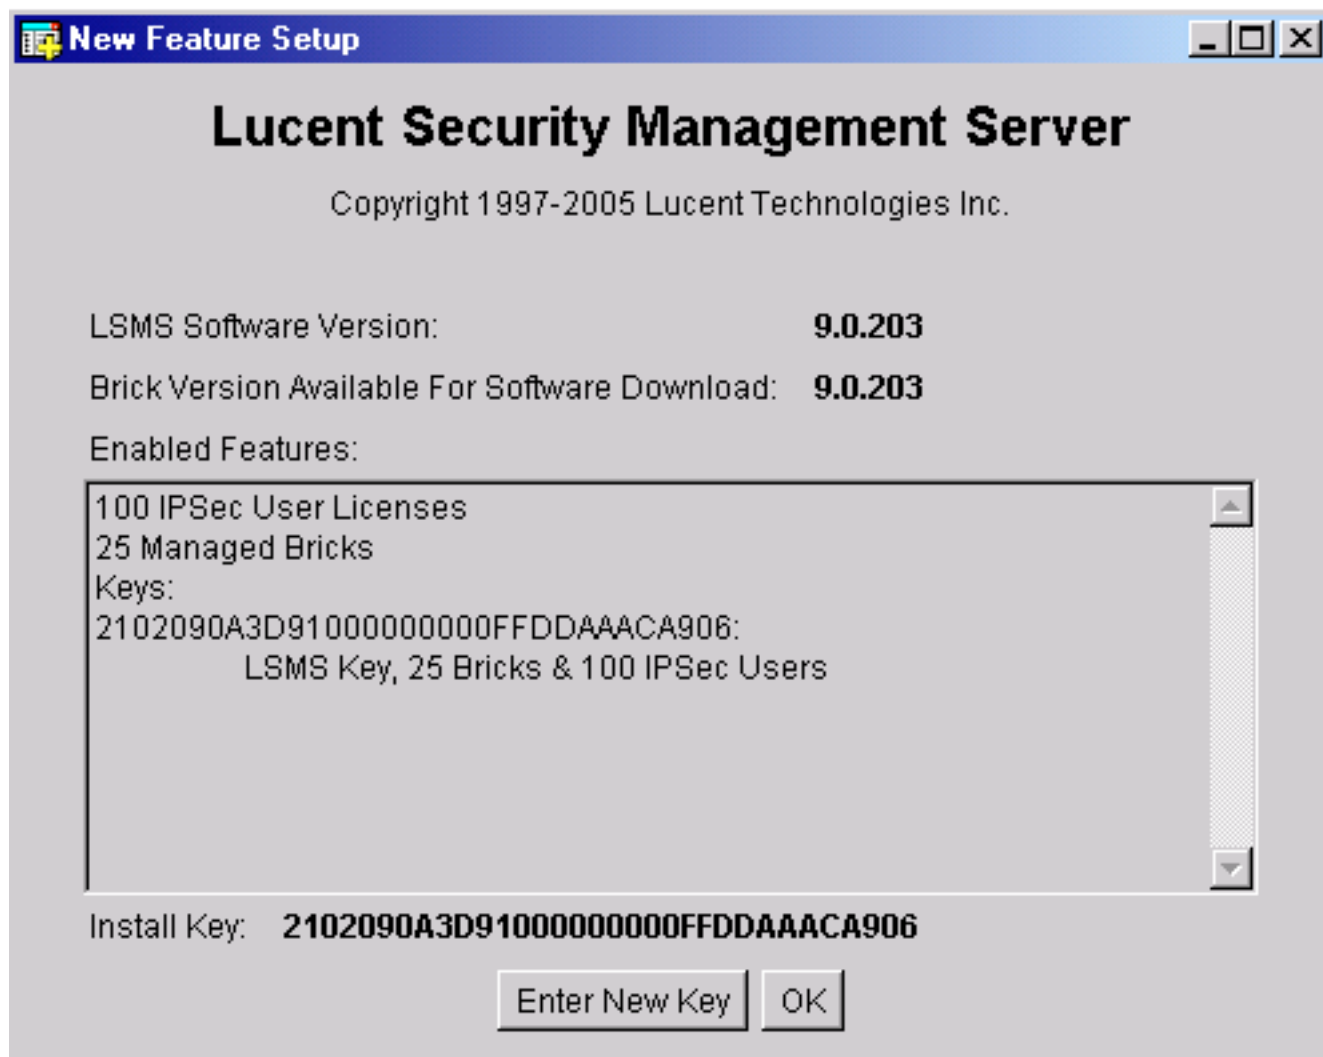
Determining current feature setup

To determine how many licenses that you currently have, open the LSMS Navigator or LSMS Remote Navigator. Then, choose Help from the menu bar and click About from the list of options. In the Enabled Feature Options section, you will see how many IPSec Client users and Bricks that your LSMS can currently manage.

Figure E-1 shows a sample of the New Feature Setup window. In the sample shown, the LSMS can manage one hundred IPSec Client users and twenty five bricks, but not Model 1100 Bricks. If the Premium Brick Management key had already been added, a separate line item would list "Premium Brick Management" in this window. Additional

application information such as currently installed license keys and associated features is also shown in the window.

Figure E-1 New Feature Setup Window



□

Using New Feature Setup

Task

Once you have purchased the additional key, follow the instructions below:

- 1 Go to
http://www.lucent.com/security
and click **IP Service Product Registration and Support**. Then, click **Register Now**.

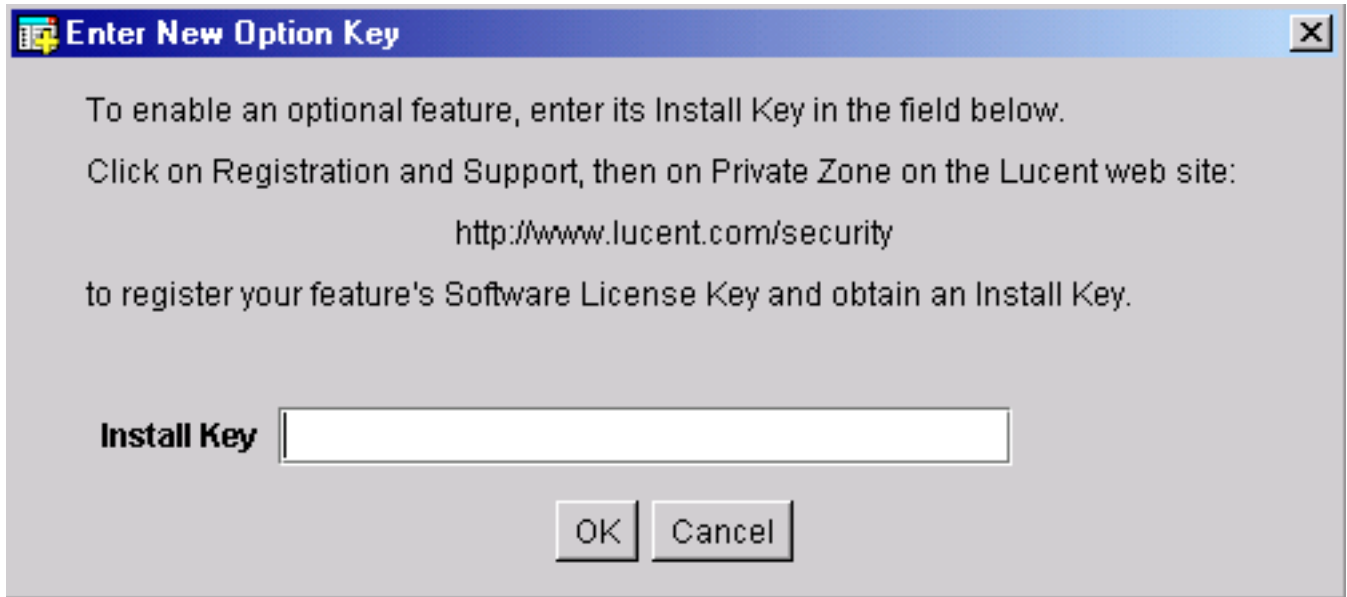
- 2 Proceed with the section titled **Registered Customers of Lucent Internet Security Products** to obtain your new installation key.

- 3 After obtaining the new installation key, return to the LSMS console and bring up **New Feature Setup**.
 - For Windows users, click **Start => Programs => Lucent Security Management Server => Utilities => New Feature Setup**. The Login New Feature Setup window is displayed. Enter your login and password.
 - For Solaris users, "cd" to your LSMS installation directory (default choice is */opt/isms/lmf*), and type:
./newFeatureSetup
The New Feature Setup window is displayed ([Figure E-1, "New Feature Setup Window"](#) (p. E-3) shows a sample screen).

- 4 Click on the **Enter New Key** button

The Enter New Option Key window is displayed (xxx shows a sample screen).

Figure E-2 New Option Key Window



5 In the **Install Key** field, type in your new key.

6 Click the **OK** button.

The new license limits are displayed on the New Feature Setup window. For the new key to take effect, you need to restart the LSMS services.

- For Windows users, click on **Start => Programs => Lucent Security Management Server =>** and select **Restart Services**.
- For Solaris users, "`cd`" to your LSMS installation directory (default choice is `/opt/isms/lmf`), and type:
`.restartServices`

Note that restarting LSMS services will automatically logout all existing administrator sessions. Also, if your entered options key is not accepted, it's usually due to the key not being intended for the LSMS on which it is being installed. Please refer to the registration website and ensure that that options key is matched for the LSMS under question.

END OF STEPS



Index

- A** Activate
 - Brick, [3-34](#)
- Activate a Brick, [3-42](#)
- Add
 - Static route, [4-22](#)
- Admin ID, [1-3](#), [1-7](#), [10-14](#)
- Administer a brick over the Internet from an unregistered LSMS, [A-1](#)
- Administer a Brick over the Internet from an unregistered LSMS, [A-8](#)
 - Activate the remote brick, [A-8](#)
 - Add NAT rules to the administrativezone ruleset, [A-5](#)
 - Assign the Administrative Zone, [A-4](#)
 - Configure the Brick, [A-3](#)
 - Enter a VBA, [A-4](#)
- Administrator account
 - Create, [8-10](#)
 - Delete, [8-24](#)
 - Edit, [8-21](#)
 - Explained, [8-1](#)
 - Maintain, [8-21](#)
- Administrators & LSMS window, [14-10](#), [14-14](#)
 - LSMS Table, [14-11](#)
- Alarms parameters, [11-9](#)
- Allow MAC addresses to move, [3-19](#)
- AppleTalk, [D-1](#)
- Apply, [1-20](#), [1-20](#), [1-20](#), [1-21](#)
 - Brick, [5-6](#)
 - How to apply, [1-21](#)
 - What to apply, [1-20](#)
 - When to apply, [1-20](#)
- Assign
 - Groups, [8-17](#)
- Assign multiple rulesets to a port, [4-17](#)
- Assign the same VBA to multiple ports, [4-18](#)
- Audience
 - Audience:Network Administrators, [xix](#)
-
- B** Backing up and restoring data, [12-1](#), [12-11](#)
 - Automatic backup, [12-2](#)
 - Manual backup, [12-3](#), [12-3](#)
 - Restore scenarios, [12-10](#)
 - Restoring data, [12-7](#)
- Brick
 - rehomeing options, [2-7](#)
- Brick devices
 - supported, [xxii](#)
- Brick failover, [3-22](#)
 - manually initiate, [3-32](#)
 - primary Brick and, [3-24](#)
 - set up, [3-26](#)
- Brick serial port
 - activate login banner, [4-28](#)
- Brick Status windows, [14-14](#), [14-27](#)
 - Brick lists, [14-14](#)
 - Single Brick Bandwidth Statistics window, [14-23](#)
 - Single Brick Ports window, [14-21](#)
 - Single Brick Status window, [14-17](#)
- Bricks
 - activate, [3-34](#)
 - Apply changes, [5-6](#)
 - boot, [3-41](#)
 - Bridge or router, [3-2](#)
 - Configuration options, [3-17](#)

- Configure a physical port, [4-2](#)
 - Configuring, [3-1](#), [4-1](#), [5-1](#), [5-28](#)
 - Delete, [5-9](#)
 - Directly connected to the LSMS, [3-2](#)
 - Download software, [5-22](#)
 - failover, [3-22](#)
 - Firewall, [3-4](#)
 - heartbeat messages and, [3-22](#)
 - Intelligent cache management, [4-1](#)
 - Make floppy, [3-37](#)
 - Modify, [5-5](#)
 - Move, [5-10](#)
 - Reboot, [5-11](#)
 - Refresh the MAC Table, [5-14](#)
 - Tunnel endpoint, [3-5](#)
 - view snapshot of configuration, [5-3](#)
 - Bricks host group, [8-3](#)
 - Buttons, [1-14](#)
 - Editing buttons, [1-14](#)
 - Tunnel buttons, [1-15](#)
-
- C** Change the IP address of the LSMS, [C-1](#), [E-1](#)
 - Compute servers
 - redundancy and, [2-2](#)
 - Concurrency control, [1-22](#)
 - enable, [1-23](#), [1-25](#)
 - Force a logout of administrator, [1-27](#)
 - force logout of an administrator, [1-24](#)
 - lock status timeout, [1-23](#)
 - Configuration Assistant, [11-1](#), [11-32](#), [11-45](#), [11-54](#)
 - Alarms parameters, [11-9](#)
 - Direct paging parameters, [11-13](#)
 - Log files parameters, [11-20](#)
 - Log transfer parameters, [11-24](#)
 - LSMS web server parameters, [11-29](#)
 - SNMP agent parameters, [11-34](#)
 - Software download parameters, [11-37](#)
 - Start Configuration Assistant from the LSMS host, [11-3](#)
 - strong passwords feature, [11-51](#)
 - User authentication parameters, [11-48](#)
 - View parameters, [11-8](#)
 - VPN debugging parameters, [11-53](#)
 - Configuration options, [3-17](#)
 - Allow MAC addresses to move, [3-19](#)
 - Halt all traffic if audit fails, [3-18](#)
 - NOC Gateway, [3-19](#)
 - Route multicast packets to first matching zone, [3-19](#)
 - Configure a physical port, [4-2](#), [4-8](#)
 - Configuring, [3-1](#)
 - Configuring Bricks, [4-1](#)
 - Configuring bricks, [5-1](#)
 - Configuring Bricks, [5-28](#)
 - Configuring bricks
 - Configurations options, [3-17](#)
 - Configuring Bricks
 - Primary LSMS, [3-8](#)
 - Redundant LSMS, [3-14](#), [3-16](#)
 - Console Alarms window, [14-27](#)
 - Contents panel, [1-9](#)
 - CPU capacity, [B-4](#)
 - Create
 - Administrator accounts, [8-10](#)
 - Groups, [8-5](#)
-
- D** Decnet/LAT, [D-1](#)
 - Delete
 - Administrator account, [8-24](#)
 - Brick, [5-9](#)
 - Group, [8-8](#)
 - Policy assignment, [4-19](#)
 - Static route, [4-27](#)
 - Direct paging parameters, [11-13](#)
 - Disk capacity for log files, [B-7](#)
 - Download software to a Brick, [5-22](#)
 - dsap file, [D-2](#)
-
- E** Edit
 - Administrator account, [8-21](#)
 - Group, [8-7](#)
 - Editing buttons, [1-14](#)

- Enable
 - strong passwords feature, [11-51](#)
 - ethertype file, [D-2](#)
-
- F** Failover
 - Brick heartbeat messages and, [3-22](#)
 - Folders, [1-17](#)
 - Of a group, [8-2](#)
 - Folders panel, [1-7](#)
 - Force a logout of administrator, [1-27](#)
 - Force logout of an administrator, [1-24](#)
-
- G** Group
 - Create, [8-5](#)
 - Defined, [8-2](#)
 - Delete, [8-8](#)
 - Edit, [8-7](#)
 - Folders, [8-2](#)
 - Subfolders, [8-2](#)
 - Group administrator, [8-3](#), [8-9](#)
 - Privileges, [8-3](#), [8-17](#)
 - Group Administrators, [1-7](#)
 - Groups, [1-17](#)
 - Assign privileges, [8-17](#)
 - Maintain, [8-7](#)
 - New, [8-3](#)
-
- H** Halt all traffic if audit fails, [3-18](#)
 - Halt Traffic if Log Full checkbox, [11-23](#)
-
- Heartbeat messages
 - redundant LSMS configuration, [2-6](#)
- Heartbeats
 - LSMS Failover, [2-6](#)
- Host group
 - Bricks, [8-3](#)
-
- I** Intelligent cache management, [4-1](#)
 - Set the threshold levels, [5-29](#)
-
- L** Log files
 - Disk capacity, [B-7](#)
 - Log files parameters, [11-20](#)
 - Log in, [1-2](#), [1-2](#), [1-3](#), [1-3](#), [10-15](#)
 - Log in from a remote host, [1-3](#)
 - Log in from the LSMS host, [1-2](#)
 - Status monitor only login, [1-3](#), [10-15](#)
 - Log off, [1-2](#)
 - Log transfer parameters, [11-24](#)
 - Login banner
 - Brick serial port, [4-28](#)
 - LSMS administrator, [8-3](#), [8-9](#)
 - LSMS Administrators, [1-7](#)
 - LSMS host, [1-2](#), [1-3](#), [10-14](#)
 - LSMS Messenger, [8-25](#)
 - LSMS Navigator, [1-2](#)
 - LSMS Remote Navigator, [1-3](#)
 - LSMS web server parameters, [11-29](#)
-
- Lucent Netcare Professional Services, [xxii](#)
-
- M** Maintain
 - Administrator accounts, [8-21](#)
 - Groups, [8-7](#)
 - Make floppy, [3-37](#)
 - LSMS host, [3-37](#)
 - Memory utilization, [B-6](#)
 - Menu bar, [1-10](#)
 - Messages
 - send to administrators, [8-25](#)
 - mgmt-tunnel ruleset, [8-3](#)
 - Modem port, [11-13](#)
 - Modify
 - Brick, [5-5](#)
 - Policy assignment, [4-18](#)
 - Static route, [4-25](#)
 - Mouse actions, [1-14](#)
 - Move
 - Brick, [5-10](#)
-
- N** Navigator window, [1-7](#)
 - Contents panel, [1-9](#)
 - Folders panel, [1-7](#)
 - Network administrators, [xix](#)
 - New feature setup, [2-3](#)
 - New groups, [8-3](#)
 - NOC Gateway, [3-19](#)
 - nocgwzone, [8-3](#)
 - Non-IP protocols, [D-1](#)
 - Non-IP Protocols, [D-4](#)

- Non-IP protocols
 - Edit the ethertype and dsap files, [D-2](#)
 - Novell IPC, [D-1](#)

 - P** Parameters
 - Alarms, [11-9](#)
 - Direct paging, [11-13](#)
 - Log files, [11-20](#)
 - Log transfer, [11-24](#)
 - LSMS web server, [11-29](#)
 - Reports, [11-32](#)
 - SNMP agent, [11-34](#)
 - Software download, [11-37](#)
 - Tunable, [11-45](#)
 - User authentication, [11-48](#)
 - VPN debugging, [11-53](#)
 - Ports, [4-2](#), [4-8](#), [4-17](#), [4-18](#), [4-18](#), [4-19](#), [4-19](#)
 - Assign a security policy, [4-8](#)
 - Assign multiple rulesets, [4-17](#)
 - Assign the same VBA to multiple ports, [4-18](#)
 - configure, [4-2](#)
 - Delete a policy assignment, [4-19](#)
 - Modify a policy assignment, [4-18](#)
 - Re-order the policy assignment entries, [4-19](#)
 - Primary Brick
 - failover and, [3-24](#)
 - Primary LSMS
 - installation, [2-3](#)
 - Privileges
 - Of a group administrator, [8-17](#)

 - R** Reboot a Brick, [5-11](#)
 - Redundancy
 - Compute servers and, [2-2](#)
 - load sharing, [2-6](#)
 - Logging and, [2-9](#)
 - LSMS, [2-1](#)
 - monitoring and, [2-9](#)
 - Redundant LSMS, [3-14](#)
 - Refresh the MAC Table, [5-14](#)
 - Remote administration, [10-14](#), [10-14](#)
 - Create the host group, [10-11](#)
 - Create the rules, [10-12](#)
 - Prepare the remote host, [10-11](#)
 - Remote login See Remote administration, [10-10](#)
 - Reports parameters
 - Reports parameters, [11-32](#)
 - Restoring data, [12-1](#)
 - Route multicast packets to first matching zone, [3-19](#)
 - ruleset
 - mgmt-tunnel, [8-3](#)

 - S** Sarbanes-Oxley (SOX) password compliance
 - enable password restrictions for, [11-51](#)
 - Secondary LSMS
 - installation, [2-3](#)
- See Backing up and restoring data, [12-1](#)
- Set
- Alarms parameters, [11-9](#)
 - Direct paging parameters, [11-13](#)
 - Log files parameters, [11-20](#)
 - Log transfer parameters, [11-24](#)
 - LSMS web server parameters, [11-29](#)
 - Reports parameters, [11-32](#)
 - SNMP agent parameters, [11-34](#)
 - Software download parameters, [11-37](#)
 - System wide parameters, [11-1](#)
 - Tunable parameters, [11-45](#)
 - User authentication parameters, [11-48](#)
 - VPN debugging parameters, [11-53](#)
- Sizing guidelines, [B-1](#), [B-8](#)
- Determine CPU capacity, [B-4](#)
 - Determine memory utilization, [B-6](#)
 - Disk capacity for log files, [B-7](#)
 - Sizing tool, [B-2](#)
- Sizing tool, [B-2](#)
- SMTP host, [11-9](#)
- SNMPagent parameters, [11-34](#)
- Software download parameter, [11-37](#)

- SOX password compliance
 - See: Sarbanes-Oxley (SOX) password compliance
 - Static routes, [4-20](#)
 - activate, [4-26](#)
 - add, [4-22](#)
 - Add, [4-22](#)
 - cost-based selection of, [4-20](#)
 - deactivate, [4-26](#)
 - delete, [4-27](#)
 - Delete, [4-27](#)
 - modify, [4-25](#)
 - Modify, [4-25](#)
 - Status Monitor, [14-1](#), [14-27](#)
 - Administrators & LSMS window, [14-10](#)
 - Brick states, [14-3](#)
 - Brick Status windows, [14-14](#)
 - Console Alarms window, [14-27](#)
 - Display the Status Monitor, [14-2](#)
 - Status Monitor data, [14-3](#)
 - Status Overview window, [14-6](#)
 - Toolbar, [14-4](#)
 - Status monitor only login, [1-3](#), [10-15](#)
 - Status Overview window, [14-6](#), [14-10](#)
 - Brick buttons, [14-6](#)
 - Brick graphs, [14-7](#)
 - Status folder drop-down, [14-7](#)
 - Strong passwords feature
 - enabling, [11-51](#)
 - Subfolders
 - Of a group, [8-2](#)
 - Syslog port, [11-10](#)
 - System group, [1-17](#)
 - Defined, [8-1](#)
 - System wide parameters, [11-1](#)
-
- T** Technical support, [xxii](#)
 - Trigger Alarm Code checkbox, [11-10](#)
 - Tunabale parameters
 - Tunable parameters, [11-45](#)
 - Tunnel buttons, [1-15](#)
 - Tunnel endpoint, [3-5](#)
-
- U** User authentication parameters, [11-48](#)
-
- V** VLANs, [6-1](#), [6-23](#)
 - Always Show VLAN Information checkbox, [6-7](#), [6-7](#)
 - Assign a policy to a port, [6-13](#)
 - Associate a network with a VLAN, [6-17](#)
 - Configure Brick physical ports and, [6-7](#)
 - Default VLAN ID, [6-10](#)
 - Receive format, [6-10](#)
 - Transmit format, [6-11](#)
 - VLAN domain, [6-10](#)
 - VLAN membership, [6-10](#)
 - What is a VLAN?, [6-2](#)
 - Why build VLANs?, [6-4](#)
 - Zone VLAN ID field, [6-13](#)
 - VPN debugging parameters, [11-53](#)

