

Administrator Guide

Version 6.1





230 Third Ave, Waltham, MA 02451, USA T. 877 487 9400 F. 781 487 7900 www.Guardium.com

Copyright © 2007 Guardium. All rights reserved. Information in this document is subject to change without notice. Guardium, SQL Guard, Safeguarding Databases, SQL HealthGuard, SQL AuditGuard, SQL PolicyGuard, Guardium RemoteGuard and SQL Guard Security Suite are trademarks of Guardium, Inc. All other trademarks and trade names are the property of their respective companies.

Document Version 6.1 August 9, 2007

Contents

Contents.....	3
Preface	12
Chapter 1: Installation	13
Overview	13
Initial Configuration.....	13
IP Configuration.....	13
System Overview.....	14
Network Interfaces and Connectors.....	19
Step 1: Installation Preparation.....	20
Step 2: Physical Connectivity	20
Network Placement	20
Guidelines for Rack Mounting	22
Step 3: Initial System Configuration	23
Using the SQL Guard CLI	23
Logging in to SQL Guard.....	24
Optionally Reset the CLI Password.....	25
Configure Network Settings.....	25
Step 4: Install a Server Certificate (Optional)	30
4.1 Create a CSR	30
4.2 Submit the CSR to Your CA	31
4.3 Store the CA Certificate (Optional).....	31
4.4 Store the Server Certificate	32
Step 5: S-Tap and CAS Installation (Optional).....	33
Configure the Guardium Server to Control S-Taps.....	33
S-Tap Overview.....	34
Unix S-Tap Installation.....	35
Prepare to Install Unix S-Tap	35
Uninstall Previous Version of Unix S-Tap	40
Install Unix S-Tap on the Database Server.....	41
Complete the Unix S-Tap Configuration from the Guardium Administrator Portal	45
Windows S-Tap Installation	51
Prepare to Install Windows S-Tap.....	51
Uninstall Previous Version of Windows S-Tap	53
Install Windows S-Tap on the Database Server	54
Complete the Windows S-Tap Configuration from the Guardium Administrator Portal	56

Secondary SQL Guard Hosts for S-Tap	61
Viewing the S-Tap Events Panel	63
Prepare for Local Unix Clients to Use the TEE	63
Prepare for Local Unix DB2 Clients to Use the Tee.....	63
Prepare for Local Unix Informix Clients to Use the Tee.....	67
Prepare for Local Unix Oracle Clients to Use the Tee	70
Prepare for Local Unix Sybase Clients to Use the Tee.....	72
Installing CAS	73
Installing CAS on Database Servers	74
Activating CAS on SQL Guard Servers.....	76
Modifying CAS Options from the Administration Console.....	77
Setting Up and Maintaining Secondary Servers	82
Chapter 2: System Management	85
Web Browser Considerations	85
Logging in to SQL Guard	86
Changing Your Password	87
About the System Shared Secret	89
Changing the System Configuration.....	89
Working with Inspection Engines.....	93
About Kerberos-Encrypted Database User Names	93
Opening the Inspection Engine Configuration Panel	94
Selecting IP Addresses	95
Changing Settings that Apply to all Engines	96
Modifying Inspection Engine Configurations	98
Adding Inspection Engines.....	99
Starting and Stopping Inspection Engines	102
Removing Inspection Engines.....	103
Installing Policies	104
Changing the SQL Guard Portal.....	105
Changing the Portal User Authentication Method.....	105
RADIUS Authentication	106
LDAP Authentication	107
Setting Global Profile Defaults.....	108
Use Aliases Default	108
Alert Message Template	108
Application User Translation	111
Selective Audit Trail and Application User Translation	111
Configure Application User Detection	111
Populate Pre-Defined Application Groups	113
Regenerate Special Application Report Portlets	115
Configuring the Alerter	123
Stopping the Alerter.....	124

Configuring Anomaly Detection	125
Stopping Anomaly Detection	126
Session Inference	127
Stopping Session Inference	127
IP-to-Hostname Aliasing	128
IP-to-Hostname Aliasing	129
Viewing IP-to-Hostname Aliases	130
Upload Key File	132
Query Hint.....	132
Incident Generation	133
Aggregation.....	136
Overview.....	136
Exporting Data to an Aggregation Server	137
Stopping Export to an Aggregation Server.....	138
Importing Data on the Aggregation Server.....	139
Stopping Importing on an Aggregation Server	140
Archiving and Restoring.....	141
Archiving.....	141
EMC Centera Archive and Backup.....	144
TSM Archive and Backup.....	145
Stopping Archiving.....	147
Restoring	147
Exporting CSV Files.....	148
System Backup.....	150
Central Management	152
About Central Management	152
Users, Roles, and Groups under Central Management.....	154
Aliases and Groups under Central Management	154
Audit Processes under Central Management	155
Central Manager Reports Using Data from Managed Units	155
Non-Central Manager Tasks	155
Upgrade Considerations for Version 6.0 or Later.....	155
Implementing Central Management	156
If the Central Management Unit is Unavailable.....	159
Registering Units for Central Management	159
Synchronizing Portal User Accounts	163
Monitoring Managed Units	164
Installing Security Policies on Managed Units	169
Viewing Management Maps	170
Using S-Taps	171
Configuring S-Tap	172
Displaying S-Tap Information	172
Settings Available Only on the S-Tap Configuration File	181

Managing S-Taps	183
Adding or Modifying S-Tap Inspection Engines	185
Modifying or Removing S-Tap Inspection Engines	188
Applying Changes to S-Tap Configurations	188
Viewing the S-Tap Events Panel.....	189
S-Tap Error Messages	190
Reporting or Alerting on S-Tap Connectivity.....	191
Monitoring CAS Status	192
Stopping and Starting the CAS Agent.....	194
Exporting and Importing Definitions.....	196
Exporting SQL Guard Definitions	198
Importing SQL Guard Definitions	200
Custom Assessment Tests	201
Defining Custom Assessment Tests	201
Sample Custom Test Class.....	204
Installing and Managing Custom Assessment Tests	205
Custom Alerting	207
Alerting Overview	208
Using the Customer Defined Alerting Interface.....	210
Sample Custom Alerting Class.....	211
Managing Custom Alerting	212
Identifying Application Users and Events	215
Identification via Stored Procedures.....	215
Identification Using the Application Events API	216
Defining Custom Identification Procedures	218
Using the Task Scheduler.....	223
Defining a Schedule	223
Removing a Schedule	225
Pausing a Schedule.....	225
Monitoring Disk Space Use	225
Near Capacity Alert	225
Current Status Monitor.....	225
Using the Running Query Monitor	228
Monitoring via SNMP	228
SNMP Examples	229
Guardium SNMP OID	229
Chapter 3: User Management	231
About Users	231
User Account Security	231
Sample User Accounts	232
Managing User Accounts.....	232
Adding a New User	233

Adding Roles to a User	235
Editing User Information	236
Removing a User	236
Importing Users from an LDAP Server	238
Configuring LDAP User Import.....	238
Run an LDAP Import On Demand.....	240
Schedule LDAP Import.....	242
Chapter 4: Security Role Management	243
Default Roles	243
Users.....	244
About the admin User.....	245
Example User Accounts	246
Adding and Removing Roles	246
Adding a Role	247
Removing a Role	247
Editing Application Role Permissions	247
Application List	248
Chapter 5: Administrator Tools	255
Custom Tables, Domains and Queries.....	255
Creating Custom Tables.....	256
Uploading Data to Custom Tables	260
Scheduling Custom Data Uploads	262
Purging Data from Custom Tables	263
Defining Custom Domains.....	263
Working with Custom Queries	265
Value Change Auditing	266
Overview.....	266
Defining Audit Databases	267
Before Defining an Audit Database Under Informix or Sybase	267
Creating the Audit Database	269
Defining Monitoring Activities.....	274
Scheduling Value-Change Uploads.....	277
Maintaining Privileged Users Lists.....	279
Value-Change Reporting	281
Value Change Tracking Domain	281
Value Change Domain Default Reports	283
Chapter 6: Command Line Interface	285
Introduction	285
CLI Command Arguments.....	285
CLI Command Abbreviations	286

Accessing the CLI.....	286
Physical Console Access	286
Network SSH Access	287
CLI Command Categories	288
Documentation Conventions	288
Show Commands	289
show account Commands.....	289
show alerter Commands	289
show anomaly-detection Commands	291
show auth Command	291
show buffer Command	291
show build Command.....	291
show defrag Command	291
show fail-policy Command	291
show firewall Command	292
show gui port Command	292
show ignored port list Command.....	292
show inspection-engines Commands.....	292
show installed security policy Command	293
show license Command	293
show log Commands	293
show logging granularity Command	293
show maximum query duration Command.....	293
show network arp-table Command.....	293
show network interface all Command	293
show network interface inventory Command	294
show network interface port	294
show network macs Command	294
show network resolver Command	295
show network routes Commands	295
show password Commands	295
show product gid Command.....	296
show purge objects age Command	297
show remotelog Command	297
show security policies Command	297
show storage-system Command.....	298
show support state Command.....	298
show support-email Command.....	298
show storage-system Command.....	298
show system Commands	298
show system public key Commands	299
show throttle Command	299
show transfer-method Command	299

show unit type Command	300
Store Commands	300
store account Commands	300
store alerter Commands	301
store anomaly-detection Commands.....	303
store auth SQL_GUARD Command.....	303
store certificate Command	303
store defrag Commands.....	303
store fail-policy Command.....	304
store firewall Command.....	304
store full-bypass Command.....	305
store gui port Command.....	305
store ignored port list Command	305
store inspection-engine log sqlstrings Command	305
store installed security policy Command.....	305
store license Command.....	306
store local-stap Command	306
store log Commands	306
store logging granularity Command	307
store maximum query duration Command	307
store network interface Commands.....	307
store network resolver Command	308
store network routes Command	308
store password Commands.....	308
store product gid Command	309
store purge object Command	309
store remotelog Command	310
store stap certificate Command.....	311
store storage-system Command	311
store syslog-trap	311
store system Commands.....	312
store support state Command	313
store throttle Command.....	314
store transfer-method Command	314
store trusted certificate Command	314
store unit type <i>and</i> clear unit type Commands.....	314
store user password Command	315
Inspection Engine Control Commands	316
Operational Control Commands	318
? Command	318
aggregator backup keys file Command.....	318
aggregator clean shared-secret Command.....	318
aggregator debug start Command	318

aggregator debug stop Command.....	319
aggregator list failed imports Command.....	319
aggregator recover failed import Command.....	319
aggregator recover failed restore Command	319
aggregator restore keys file Command	320
backup and restore system Commands.....	320
commands Command	322
debug Command	322
diag Command	322
dsaa mail sender state Command.....	322
eject Command	322
export audit-data Command	322
export file Command	323
forward support email Command	324
help Command	324
import file Command	324
import tsm config Command	325
iptraf Command.....	325
license check Command	325
list audit-data Command	325
load-balancer Commands	326
ping Command	326
quit Command	326
remove audit-data Command.....	327
register / unregister for Central Management Commands	327
restart Commands	328
stop Commands	328
unlock admin Command.....	329
Certificate Commands	329
csr Command	329
store certificate Command	330
store trusted certificate Command	331
Diagnostics Command.....	331
Opening the Diagnostics Main Menu	331
1: Output Management.....	333
2: System Static Reports.....	337
3: System Interactive Queries	337
4: Perform Maintenance Actions	352
5: Exit to CLI	355
System Static Reports Overview.....	355
Appendix A: Time Zone List.....	361
Appendix B: Reinstalling SQL Guard	369

Reinstalling SQL Guard Software.....	370
--------------------------------------	-----

Preface

This document describes how to install and administer SQL Guard.

Target Audience

The SQL Guard Administrator Guide is written for a network-proficient administrator who must have knowledge of the network environment in which SQL Guard is installed, as well as a basic working knowledge of system configuration and management.

Related Documents

For a description of how to use the non-administrator functions of SQL Guard, see the *SQL Guard User Guide*.

Software Downloads from Adobe

SQL Guard creates reports and graphics that can be viewed using two software products available at no cost from Adobe, Inc. If you are viewing this document online using Adobe Acrobat Reader, you can click the buttons in the left-hand column to download the most recent versions of these products. If you are not online, or if the links do not work, use your browser to navigate to the Adobe home page: <http://www.adobe.com>, and search for the most recent versions of these products.

Product	Description
	Adobe Acrobat Reader – Use this product to view SQL Guard reports or the SQL Guard manuals online.
	Adobe SVG Viewer – Use this product to view SQL Guard Access Maps in your browser window.

Chapter 1: Installation

Overview

This chapter is organized as a series of installation steps that allows the administrator to completely configure the SQL Guard system. The initial configuration steps are performed using a local connection to the SQL Guard unit. The remaining configuration activities are performed over a network connection using a Web browser.

To successfully install your SQL Guard system, read through this overview and then follow the complete set of steps described in this chapter.

Once you have read through the chapter and gathered all necessary information, keep a copy of this information for easy reference in the event you need to reinstall the SQL Guard system later.

Initial Configuration

To initially configure the SQL Guard system:

1. Use the console to set the unit's IP configuration.
2. Place the SQL Guard system in its final network location and continue with the remaining configuration steps.

IP Configuration

To set the initial IP configuration for the unit, use the SQL Guard Command Line Interface (CLI), which is available from the serial port or on the system console.

To use a PC keyboard and monitor:

1. Attach a PC video monitor to one of the video connectors. There is one on the front of the unit and one on the back.
2. Attach a PC keyboard with a PS/2 style connector to the Keyboard connector on the back of the unit or attach a USB keyboard to a USB connector on the front or back of the unit.

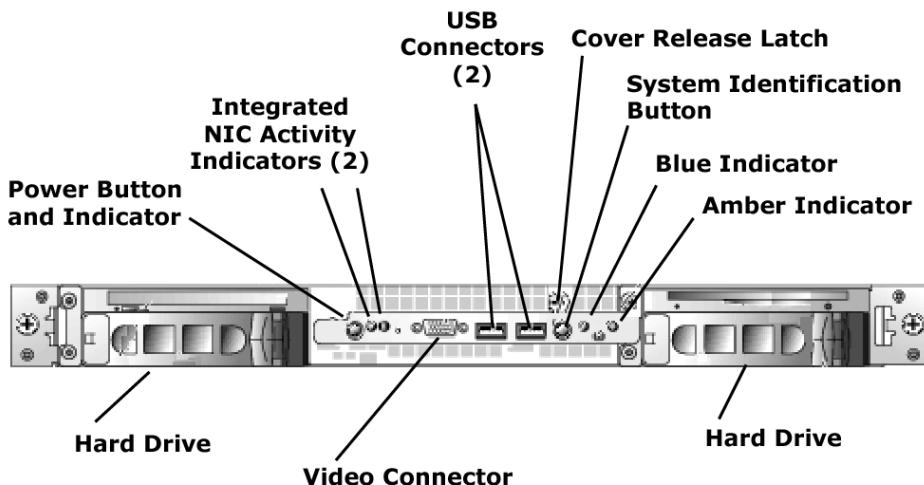
To use the serial port: Use a **NULL** modem cable to connect a terminal or another computer to the 9-pin serial port on the back of the unit. The terminal, or a terminal emulator on the attached computer, must be set to communicate as 19200-N-1 (19200 baud, no parity, 1 stop bit).

System Overview

The appearance of the SQL Guard System varies slightly depending on the model number and the options purchased.

System Front

There are several important items to note on the SQL Guard system front, each of which is described in the table below.



SQL Guard Unit – Front View

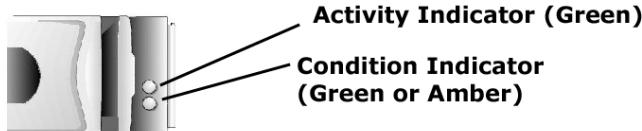
Note: Your system may not appear exactly the same as the one illustrated above.

Item	Description
USB Connectors (2)	You can connect a USB keyboard to one of these for initial installation or when using the CLI. You can also connect a USB or PS2 keyboard to the back of the unit. If your SQL Guard license is installed on a USB flash-memory card, you can insert that card here. (There are additional USB connectors on the back of the unit.)
Cover Release Latch	Use to remove the front cover.
System Identification	Use to locate a particular system in a rack. When pressed, the

Item	Description	
Button	blue indicator lights on both the front and back of the unit blink. When pressed a second time, the indicators stop blinking. There is also a system identification button on the back of the unit.	
Blue Indicator	Off	The system is off.
	Blue	The system is operating normally.
	Blinking	The system is identifying itself because the system identification button (see above) has been pressed.
Amber Indicator	Blinking	Indicates a fault with the system.
Hard Drive	Hard disk drive(s). There may be one on each side. See below for a description of the activity indicators.	
CD Drive	CD drive for installing upgrades or patches.	
Power Button and Indicator	Press to power the unit on or off. The indicator light may be:	
	Off	The system is off and AC power is not connected.
	Blinking	A blinking green light indicates that the power is connected, but the system is not powered on.
	On	A solid green light indicates that the system is powered on.
Integrated NIC Activity Indicators (2)	Activity indicators for the two <i>integrated</i> NICs, which may or may not be used, depending on your configuration. See Network Interfaces and Connectors below, for more information about the network connectors.	
Video Connector	Connect a PC monitor here for initial installation or when using the CLI. You can also connect a PC monitor to the back of the unit or you can connect a terminal or a PC to the serial port on the back of the unit.	

SCSI Hard-Drive Indicator Codes

The hard drives contain two indicator lights, on the lower right side:



SQL Guard Unit – Drive Indicator Lights

The **Activity Indicator** blinks when the drive is being accessed.

For non-RAID applications, the **Condition Indicator** is solid green when the unit is powered on. For RAID applications, see below.

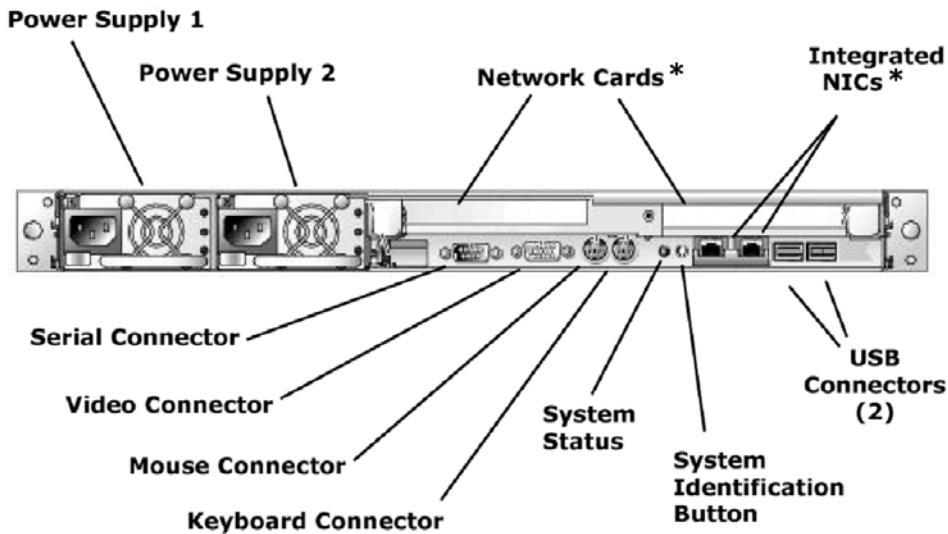
SCSI Hard-Drive RAID Indicator Codes

If RAID is activated, the two indicators on each of the hard-drive carriers provide information on the status of the SCSI hard drives. The following table lists the drive indicator patterns. Different patterns are displayed as events occur in the system. For example, if a drive fails, the *drive failed* pattern appears. After the drive is selected for removal, the *drive being prepared for removal* pattern appears, followed by the *drive ready for insertion or removal* pattern, and so forth.

Condition	Indicator Pattern
Identify drive	The condition indicator blinks green four times per second.
Drive being prepared for removal	The condition indicator blinks green two times per second.
Drive ready for insertion or removal	Both indicators are off.
Drive being prepared for operation	The condition indicator is solid green.
Drive predicted failure	The condition indicator slowly blinks green, amber, and off.
Drive failed	The condition indicator blinks amber four times per second.
Drive rebuilding	The condition indicator blinks green slowly.
Drive online	The condition indicator is solid green.

System Back

The appearance of the back of the SQL Guard system will vary slightly depending on the model and options purchased. The connectors described in the following table are present



SQL Guard Unit – Back View

Notes: Your system may not appear exactly the same as the one illustrated above.

Item	Description
Power Supply 1 Power Supply 2	Attach the supplied power cord to the Power Supply 1 socket. If the back-up power supply is installed, attach the second supplied power cord to the Power Supply 2 socket.
Network Cards* and Integrated NICs*	Two PCI slots on the back of the system may contain network cards, and one or two Integrated NICs may also be present. The use and location of all network cards is highly variable, depending on the options purchased. To connect network cables, refer to the network connection mapping document that shipped with the system. If you do not have this document, contact Guardium Support.

Item	Description
USB Connectors (2)	Connect a USB keyboard to either of the USB connectors for initial installation or when using the CLI. Typically, these are only used during the initial installation or for troubleshooting. If your SQL Guard license is installed on a USB flash card, you can insert that card in one of these connectors. There are two additional USB connectors on the front of the unit.
System Identification Button	Use this button to locate a particular system in a rack. When pressed, the system status indicator lights on both the front and the back of the unit blink blue. When pressed a second time, the status indicators stop blinking. There is also a system identification button on the front of the unit.
System Status	Blue Blinking blue indicates that a system identification button (on the front or back of the unit) has been pressed. Amber Blinking amber indicates the system needs attention due to a problem with power supplies, fans, system temperature, or hard drives. Off The system identification button has been pressed a second time.
Keyboard Connector	Connect a PS2 keyboard here for initial installation or when using the CLI. You can also connect a USB keyboard to the back (see below), or to the front of the unit (see above).
Mouse Connector	Not used.
Video Connector	Connect a PC monitor here for initial installation or when using the CLI. You can also connect a PC monitor to the front of the unit or you can connect a terminal or a PC to the serial port (see below) on the back of the unit.
Serial Connector	Use a NULL modem cable to connect a terminal or another computer to this 9-pin serial port; for use during initial installation or when using the CLI.

Network Interfaces and Connectors

Two PCI slots on the back of the system may contain network cards, and one or two Integrated NICs may also be present. The use and location of all network cards is highly variable, depending on the options purchased and the date the unit was built.

SQL Guard's use of network interfaces is described below. To connect the network cables, refer to the network connection mapping document that shipped with the system or with any upgrade to the unit that involved changing one or more network cards. If you do not have this document, contact Guardium Support.

ETH 0

Always use **ETH 0** to connect to the LAN over which users will access the SQL Guard user interface. The *primary* System IP Address is assigned to **ETH 0**, and the optional *secondary* System IP Address is always assigned to the *last* connector installed. You can assign the IP addresses using the CLI, as described later in this chapter. To enable the high-availability option, connect **ETH 3** to the same LAN as **ETH 0**, and use the [store network interface high-availability on](#) command (see Chapter 6).

SPAN Port Connections

Connect **ETH 1** to the first **SPAN** port. Optionally connect additional SPAN ports using the remaining connectors as necessary, in order.

TAP Connections

Beginning with **ETH 1 & ETH 2**, Use each pair of connectors in sequence, one per TAP.

Inline Installation (Database Firewall)

If installing SQL Guard inline to optionally provide database firewall functionality, use each *pair* of connectors beginning with **ETH 1 & ETH 2** to insert the SQL Guard unit between a group of one or more database servers and all of their clients.

Step 1: Installation Preparation

Before beginning the installation and configuration of SQL Guard, there are a number of steps that must be taken. Such steps include assuring that you have all of the parts of the SQL Guard system, gathering information about settings, and collecting necessary items such as hardware required to configure the SQL Guard system.

Prior to beginning SQL Guard configuration, select whether you will use a PC keyboard/monitor or configure through the serial port, as described above.

In addition to any hardware you may need that is specific for the selected configuration method, every installation and configuration requires the following:

Ethernet cable – To connect the unit to the network

Hardware shipped with the SQL Guard system, including:

- SQL Guard system
- Rack mount rails
- Documentation for rack mounting the system
- Metal front panel for the system
- Power cord
- Documentation for the Dell system
- Dell Driver CD

Step 2: Physical Connectivity

The SQL Guard system behaves as a network protocol analyzer. It must be connected to a switch, hub, or other local area network (LAN) device through which the database traffic flows. This equipment should be the switch or hub nearest to the database client application system or the switch or hub nearest to the database server itself.

Network Placement

For the most comprehensive monitoring of database communications, it is recommended that the SQL Guard system be located as close as possible to the protected resource: the database. If placed near the database client system, the SQL Guard system will see all traffic to or from that client and any of the databases with which it communicates. If placed near the database server, the SQL Guard system will see all traffic to or from any client to the database server. (Database servers monitored by the SQL Guard system can be Oracle, Sybase, Microsoft SQL 2000, Informix, and DB2.)

In order for the SQL Guard system to function properly, it must be able to collect the database communications that pass through the network segment on which it is connected. On a LAN that is implemented on a network hub, the SQL Guard system can view and collect network data packets. On a LAN that is implemented with network switches, viewing and collection of these data packets will not occur unless the switch is specifically configured to allow such actions.

If the SQL Guard system is placed on a switched network, that network switch must be configured to mirror all traffic *to and from* the databases to be monitored, to a port on which the SQL Guard unit will be connected. A network administrator will be able to perform this configuration. Consult your switch vendor's documentation on the exact method to perform this configuration. Some vendors call this mirroring feature Port Mirroring or Switched Port Analyzer (SPAN).

Note: In a Windows environment, if Kerberos authentication is used for MS SQL Server access, database user names are encrypted in the network traffic. In SQL Guard reports, these names appear as strings of hexadecimal characters. To have the SQL Guard server decode Kerberos-encrypted database user names automatically, mirror the Kerberos traffic to the SQL Guard server and enable the Kerberos decoding feature using the *store local-stap on* command, as described later under [Optionally Enable Automatic Decoding of Kerberos-Encrypted Database User Names](#).

If the SQL Guard database firewall will be enabled, the system must be installed *inline*, such that all client traffic to the protected servers passes through the SQL Guard unit.

The SQL Guard system provides administrative access from its first network interface card, whose connector is labeled ETH 0, and optionally from its *last* network interface card. The number of the last interface card varies, depending on what types of cards are installed (one-, two-, or four-port cards are available).

Database traffic is monitored either:

- Using SPAN ports connected in sequence to ETH 1, 2, 3, etc.

OR

- Using consecutive ETH connector pairs (1-2, 3-4, etc.), either to monitor traffic via network TAPs, or to install SQL Guard inline (in the latter case allowing it to be configured as a database firewall).

The network administrator:

- Provides an IP address for the ETH 0 connection to the desktop LAN, and optionally an IP address for a secondary management interface connection.

- If SQL Guard will be configured to function as a database firewall, centralizes all traffic to and from the databases to be protected, such that the SQL Guard unit can be inserted between all incoming and outgoing traffic for the resources to be protected.
- If SQL Guard will not be configured as a database firewall, configures one or more SPAN ports or network TAPs for use by SQL Guard.
- Provides the default router IP address.
- Provides DNS server IP addresses for from 1 to 3 DNS servers.
- Adds the new SQL Guard system to the company DNS server.
- If an NTP server will be used, provides its host name (you cannot specify an IP address for the NTP server).
- Provides SMTP configuration information (for email alerts): IP address, port, and if authentication is used, an SMTP user name and password.
- If SNMP will be used for alerts, provides SNMP configuration information: the IP address of the SNMP server and the trap community name to use.

The SQL Guard administrator:

- Coordinates with the network administrator to connect the desktop LAN to ETH 0, and to the optional secondary interface (if used).
- With the network administrator, connects the SPAN port(s), or uses one or more ETH pairs (1-2, 3-4, etc.) to either monitor traffic from network TAPs, or to insert the SQL Guard unit between one or more database servers and their clients. (The last case is required to allow SQL Guard to function as a database firewall.)
- Uses the SQL Guard administration console to ensure that the system and network settings are properly configured.

Guidelines for Rack Mounting

Different rails and rack mounting systems are available. See the separate document shipped with your unit for rack mounting instructions.

Step 3: Initial System Configuration

To configure the system initially, log into the unit and use the SQL Guard CLI. Later, you can use the Administration Console of the management interface to change most configuration settings.

Using the SQL Guard CLI

The CLI language is *not* case-sensitive.

All CLI examples are written in courier text. For example: `show system clock`

Notation for Command Arguments

Some command descriptions use delimiters to indicate which command arguments are mandatory and in which context. Each syntax description shows the dependencies between the command arguments by using special characters:

- The < and > symbols denote a required argument.
- The [and] symbols denote an optional argument.
- The | (vertical bar) symbol separates alternative choices when only one can be selected. For example: `store full-bypass <on | off>`

State Arguments

Commands that handle a “state” setting accept and use the following state arguments:

- on or off
- up or down
- enabled or disabled
- active or inactive
- 1 or 0

CLI Command Abbreviations

You can abbreviate commands and subcommands as long as you provide enough characters so the commands are not ambiguous.

For example: `show` can be shortened to: `sho`

In addition, there are several words that can be used as aliases for other CLI commands. For example:

```
exit = quit
passwd = password
cmds = commands
```

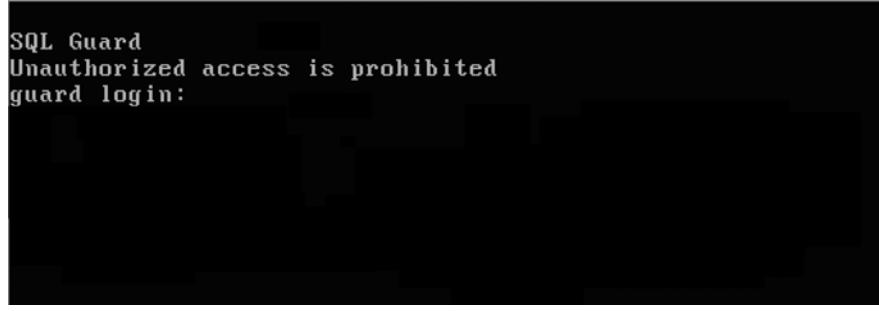
Logging in to SQL Guard

Once interactive administrative access is physically connected (via console or serial port), turn on the SQL Guard system.

If a serial terminal is connected, no text will be displayed until the system has completely finished its boot process. At that point, a login prompt is displayed.

If a PC keyboard and monitor are connected, a splash screen is displayed. The SQL Guard system then loads the operating system and displays various text messages as it progresses (*Setting clock, Loading default keymap, etc.*)

Once the system has finished booting, press the Enter key to obtain the SQL Guard *login* prompt.



```
SQL Guard
Unauthorized access is prohibited
guard login:
```

The SQL Guard login prompt

The administrative user for interactive, command line access is *cli*, with a default password assigned as noted in your installation package. After logging in, the following CLI prompt is displayed.



```
SQL Guard
Unauthorized access is prohibited
guard login: cli
Password:
Last login: Wed Oct 22 14:26:38 on tty1
guard.yourcompany.com>
guard.yourcompany.com>
guard.yourcompany.com> _
```

SQL Guard Command Line Interface (CLI) prompt

You can now start entering configuration settings.

Optionally Reset the CLI Password

To simplify the support process, we suggest that you keep the *cli* password assigned by Guardium. To change the *cli* password, use the **store user password** command. You will be prompted to enter the current password, and then the new password twice, as illustrated below. None of the password values you enter on the keyboard will display on the screen. The *cli* user password must:

- Be at least six characters in length.
- Contain at least one digit character (0-9).
- Contain at least one lowercase alphabetic character (a-z).
- Contain at least one uppercase alphabetic character (A-Z).

Note: The above rules differ from the more stringent rules applied to SQL Guard user passwords.

The **store user password** dialog should look like this, with the *guard.yourcompany.com* prompt replaced by the host and domain names configured for the SQL Guard server:

```
guard.yourcompany.com> store user password
Changing password for 'cli'.
Enter current password:
Enter new password:
Re-enter new password:
Ok
guard.yourcompany.com>
```

Note: There is no way to retrieve the CLI user password once it is set. If you lose this password, contact Guardium Technical Support to have it reset.

Configure Network Settings

For a complete list of commands and available through the CLI, see [Chapter 6: Command Line Interface](#).

Set the System IP Address

Users and remote components of SQL Guard access the system using one or two IP addresses. The *primary* IP address is for the ETH 0 connection, and is defined using the following two commands:

```
store network interface 1 ip <ip_address>
store network interface 1 mask <subnet_mask>
```

Optionally, a *secondary* IP address can be assigned to the *highest numbered* network connector (see the [Interfaces and Connectors](#) diagram), using similar commands:

```
store network interface 2 ip <ip_address>
store network interface 2 mask <subnet_mask>
```

Note: The remaining network connectors are used to monitor traffic, and do not require the assignment of an IP address.

Set the Default Router IP Address

Configure the default network router by entering the following command, replacing *default_router_ip* with the IP address of the device that routes for the management interface.

```
store network routes def <default_router_ip>
```

Set DNS Server IP Addresses

Set the IP address of one or more DNS servers to be used by the SQL Guard system to resolve host names and IP addresses. The first resolver is required, the others are optional. Replace *resolver_1_ip*, *resolver_2_ip*, and *resolver_3_ip* with the IP address for each DNS server used:

```
store network resolver 1 <resolver_1_ip>
store network resolver 2 <resolver_2_ip>
store network resolver 3 <resolver_3_ip>
```

Set Host and Domain Names

Configure the hostname and DNS domain name of the SQL Guard system.

Note: This name must match the hostname registered in the DNS system for the SQL Guard management interface's IP address. If not, administrative access will only be possible via IP address.

Replace *host_name* with the DNS host name for the unit and *domain_name* with the DNS domain name:

```
store system hostname <host_name>
store system domain <domain_name>
```

Set the Time Zone, Date, and Time

Configure the local time zone and current date and time. Replace *time_zone* and *date_time* with the proper values.

```
store system clock timezone <time_zone>
```

time_zone is the value for the unit's time zone, chosen from [Appendix A: Time Zone List](#).
For example: America/New_York

```
store system clock datetime <date_time>
```

date_time is in the format: YYYY-mm-dd hh:mm:ss

For example: 2004-07-11 10:40:00

Set the NTP Server Host Name and Activate It

If a network time protocol (NTP) server is available, configure the NTP settings and activate the use of that server. Otherwise, skip these two commands.

Replace *ntpserver_name* with the host name of the timeserver (you cannot specify an IP address here).

```
store system ntp server <ntpserver_name>
store system ntp state on
```

Optionally Enable Automatic Decoding of Kerberos-Encrypted Database User Names

In an MS SQL environment, database user names may be encrypted by Kerberos. These names will appear as strings of hexadecimal characters in SQL Guard reports. The SQL Guard server can decode these names automatically if it has access to the Kerberos traffic (as described previously) and the feature is enabled, as described below.

To enable the automatic decoding of Kerberos-encrypted database user names, enter the following commands:

```
store local-stap on
store unit type stap
```

Ignore any messages about restarting the inspection core or inspection engines. The correct settings will take effect when you restart the server after all initial settings have been configured (as described below).

Optionally Enable the Database Firewall

If the SQL Guard database firewall will be used, enable it using the following command:

```
store firewall on
```

Notes: The SQL Guard database firewall is not available on all SQL Guard server models. The server must be equipped with bypass network cards, which can be configured to allow network traffic to pass if the server is unavailable.

A unit configured as an aggregator cannot provide firewall protection.

Optionally Configure Port Forwarding Settings

The *store firewall on* command (see above) automatically adds inline as a unit type attribute for the system. If you are not using the firewall but you are using one or more pairs of interface cards to read and forward traffic (as opposed to using a SPAN or mirror port), enter the following command to enable inline mode:

```
store unit type inline
```

Note: This configuration is not recommended unless the server is equipped with bypass network cards, which allow traffic to pass if the unit is powered off or otherwise unavailable.

When inline access is enabled, you can control what happens to messages when the inspection engine (which tests policy rules against network traffic) is not running. By default, all messages will be forwarded. To block all messages when the inspection engine is down, use the following format of the store fail-policy command:

```
store fail-policy close
```

Validate All Settings

Before logging out of the CLI and progressing to the next configuration steps, it is important to validate the configured settings. Verify that each setting entered was entered and correctly, by entering the following sequence of show commands:

```
show network interface all  
show network routes defaultroute  
show network resolver all  
show system hostname  
show system domain  
show system clock timezone  
show system clock datetime  
show system ntp all  
show unit type  
show firewall
```

Reboot the System

Now that the basic system and network settings are configured, either stop the system and place it in its final network location or simply reboot the system if it is already in its final network location.

To reboot the system, enter the following command:

```
restart system
```

The system will shutdown and reboot immediately after the command is entered.

What to Do Next

If you will use a server certificate, continue with Step 4, below.

If you have purchased S-Tap for installation on database server systems, follow the procedure outlined in Step 5, below.

If you have *not* purchased S-Tap or after you have finished installing all S-Tap components, use the SQL Guard management console to begin setting up Inspection Engines and other SQL Guard components, as described in the remaining chapters of this document.

Step 4: Install a Server Certificate (Optional)

After you have configured the network settings and rebooted the system, you can obtain and store a server certificate following the process outlined below:

- 4.1 Use the CLI to create a Certificate Signing Request (CSR).
- 4.2 Submit the CSR to your Certificate Authority (CA) and obtain a server certificate in return.
- 4.3 If the server certificate returned by your CA *includes* the full trust path, skip ahead to step 4.4. Otherwise, store the CA certificate (and, if necessary, any intermediate certificates to the full trust path) on the SQL Guard unit. This must be done before storing the new server certificate.
- 4.4 Use the CLI to store the returned server certificate on the SQL Guard unit.

Each step is described in detail, below. Be aware that you perform the second step outside of the SQL Guard system, using whatever CA your company uses.

4.1 Create a CSR

Use the SQL Guard CLI to create a CSR. Be sure to enter all information correctly and do not enter this command until after your network settings have been configured. The generated CSR will be a PKCS7 file encoded in PEM (base64 ASCII text) format, so you can copy and paste it easily.

To create the CSR:

1. Log in to the SQL Guard unit as the *cli* user, as described previously
2. Enter the *csr* command:
`csr`
3. Reply to all prompts, which will be used in generating the request. Be aware that the common name (CN) is generated automatically from the host and domain name you assigned when configuring the unit:

```
What is the name of your organizational unit (OU=) ?  
What is the name of your organization (O=) ?  
What is the name of your city or locality (L=) ?  
What is the name of your state or province (ST=) ?  
What is the two-letter country code for this unit (C=) ?
```

After you respond to the last prompt, the system displays a description of the request, followed by the request itself, and followed finally by additional instructions.

For example:

```
This is the generated CSR:
Certificate Request:
Data:
Version: 0 (0x0)
Subject: C=US, ST=MA, L=Waltham, O=XYZCorp, OU=Accounting, CN=g2.xyz.com

-----BEGIN NEW CERTIFICATE REQUEST-----
MIICWjCCAhcCAQAwDElMAkGA1UEBhMCVVMxEDAOBgNVBAgTB1dhbHRoYW0xETAPBgNVBAoTCEdi
YXJkaXVtMRUwEYDVQQLEwxndWFyZG1lbSj5jb20xCTAHbGNVBAMTADCCAbgwggEsBgcqhkjOOAQB
MIIIBhkBgQD9f1OBHXUSKVLfSpwu70n9hG3UjzvRADDHj+A1tEmaUVdQCJR+1k9jVj6v8XluJD2
y5tbNeB04adNG/y2mC3a51QpaSfn+gEexAiwk+7qdf+8Yb+DtX58aophUPPBu9tPFhsMCNVQ/T
WhaRMvZ1864rYdcq7/IiAxmd0UgbXwIVAJdgU18VlwMspK5gqLrhAvwBz1AoGBAPfhoiXWmz3e
y7yrXDa4V7151K+7+jrggv1XTAs9B4JnUv1jrrUWU/mcQcQqYCSR2xi+hMKBYTt88JM0zIpUe8
FnqLWHYNKOCjhr4rs6Z1kW6jfwv6ITV18ftiegEkO8yk8b6oUZCJqIPf4VrlnwaSi2ZegHtVJWQB
TDv+z0kqA4GFAAKBgQCOnsEB4g4/limbHkuZ5YnLn9CMG3a2evEnqjXZts4itxeTYwPQvdkjdsMq
kaQ1BxmNUszOJ2rq5nC5Cg3X9spa+BzFr+PgR/Szka17nHcxkXJCjVjLk451L67K1LXv61TUfv/bU
PKm1GKDttP2ktG4dBFxFQdiCUEG0aaNFCYn6gAAMAsGBYqGSM44BAMFAAMwADA1ahUAhHTY5z9X
NiBauyAC9PS4Gr1eYakCFZ2kfxfjX1Bry51228XWMAUON95
-----END NEW CERTIFICATE REQUEST-----
```

Please copy and paste this output to a file, starting at the BEGIN and END lines, and use that file to work with your Certificate Authority in obtaining a certificate. I will be expecting the incoming certificate to be in PKCS#7 PEM format. Your CA will help you in receiving that format. Once you have it, please use the "store certificate" command to complete this operation.

4. Before continuing, check the Subject line to verify that you have entered your company information correctly. If you can submit a CSR online and obtain a server certificate quickly, remain logged in to SQL Guard. Otherwise, enter the *quit* command now to log out. Then log in again later after you have received the server certificate.

4.2 Submit the CSR to Your CA

When copying the CSR, be sure to select the entire request (shown highlighted above), *including the Begin and End request lines*. Most CAs provide online signing services, so you will be able to simply paste the CSR to a text box. If not, paste the CSR to a text file or into an email in the appropriate location.

Be sure to have the server certificate generated as a PKS7 file in PEM (base64 ASCII text) format, since you will need to copy and paste it into the SQL Guard CLI.

4.3 Store the CA Certificate (Optional)

Perform this step only if the server certificate returned from your CA does *not* include the full trust path.

Use the SQL Guard CLI to store the CA certificate and, if necessary, to store any intermediate certificates on the full trust path to the SQL Guard server. Certificates must be stored in hierarchical order, beginning with the CA certificate.

1. If you are not still logged in to the SQL Guard unit as the *cli* user, log in again as described previously.

2. Enter the *store trusted certificate* command:

```
store trusted certificate
```

The following prompt is displayed:

```
What is a one-word alias we can use to uniquely identify this  
certificate?
```

Enter a one-word name for the certificate and press Enter. The following instructions are displayed:

```
Please paste your CA certificate, in PEM format.  
Include the BEGIN and END lines, then press CTRL-D.
```

3. Copy the certificate, paste it to the command line, and press CRTL-D. You are informed of the success or failure of the store operation.
4. If there are intermediate certificates on the full trust path to the SQL Guard unit, repeat steps 2 and 3 above for each of those, in hierarchical order.

4.4 Store the Server Certificate

Use the SQL Guard CLI to store the server certificate:

1. If you are not still logged in to the SQL Guard unit as the *cli* user, log in again as described previously.
2. Enter the *store certificate console* command:

```
store certificate console
```

The following information and prompt is displayed:

```
Please paste your new server certificate, in PEM format.  
Include the BEGIN and END lines, then press CTRL-D.
```

3. Copy the server certificate, paste it to the command line, and press CRTL-D. You are informed of the success or failure of the store operation.
4. Use the following command to restart the GUI:

```
restart gui
```

Step 5: S-Tap and CAS Installation (Optional)

This section describes how to install S-Tap on a database server system. It is followed by extensive instructions on how to install CAS (see [Installing CAS](#), below). If you will be installing both components, be sure to read through both overview section below, and collect all required information or make any database configuration changes necessary before performing the S-Tap installation procedure.

Configure the Guardium Server to Control S-Taps

Any SQL Guard server that will function as an S-Tap and/or CAS host must be configured with an S-Tap *unit type*. To verify this, and to change the unit type setting if necessary, follow the procedure outlined below.

1. Log on to the *admin* portal of the Guardium server.
2. Verify that the menu on the Administration Console tab contains a Local Taps section, as illustrated to the right.



If the **menu contains a Local Taps section**, the unit type is set correctly and you should skip the remainder of this procedure.

If the **Local Taps section is missing from the menu**, continue with the remaining steps.

3. Log out of the Guardium server *admin* portal.
4. From an SSH client window, log in to the Guardium server command line interface, as the *cli* user.
5. Enter the following two commands:

```
store unit type stap  
restart inspection-core
```

For more detailed information about these commands, see the description of the *store unit type* command in Chapter 6 of the Administrator Guide.

6. Enter the **quit** command to log out of the CLI.
7. Log on to the Guardium Administrator portal and verify that the menu now appears.

S-Tap Overview

S-Tap is installed on a database server system. It monitors database traffic and forwards information about that traffic to a Guardium server, which can be deployed anywhere on the network. Only a minimal amount of configuration information is needed to install the S-Tap software on the database server, and once installed, it can be configured and controlled from a Guardium server.

S-Tap monitors database traffic using configurable *inspection engines*, which detect and report on traffic between specific clients and specific database servers. Functionally, the S-Tap inspection engine is nearly identical to the Guardium server inspection engine, which monitors database traffic on the network. But since it is installed on a database server system, S-Tap can see traffic that is *local* to that system in addition to being able to capture traffic that comes to this database server over the network.

As S-Tap collects data, it buffers the data and sends it to the Guardium server. The buffering of data allows S-Tap to continue to work when the Guardium server is not ready to receive data. If the Guardium server becomes unavailable for an extended period of time, S-Tap will fail over to a secondary server (if one or more secondary servers are defined in its configuration). It will continue to send data to the secondary server until either that server becomes unavailable, or until the S-Tap is restarted, at which point it will attempt to connect to its primary Guardium server first.

Note that because the traffic the Guardium server receives from S-Tap is a copy of the actual traffic, the Guardium server cannot provide database firewall protection for S-Tap traffic.

S-Tap on a database server must be at the same software release level as the Guardium server to which it sends data. This means that if you are upgrading from a previous version of the Guardium software, you must:

1. Stop and uninstall S-Tap on the database server.
2. Update the Guardium server to the new software release.
3. Install the new S-Tap software.

When you install S-Tap on a database server, you must set a small number of configuration properties, so that S-Tap can connect to a Guardium server. You can set the remaining properties from the Administration portal on the Guardium server. Although you can set all of the properties by editing the S-Tap configuration file on the database server, we recommend that you use the Guardium Administration portal for this purpose (where you are much less likely to introduce errors into the configuration).

The Unix and Windows installation procedures are described separately below.

Note: The S-Tap software is distributed on a separate CD. It is *not* contained on the standard SQL Guard Server installation disk.

Unix S-Tap Installation

To install and start using S-Tap on a Unix database server, you must complete the following tasks, each of which is described as a separate procedure, below:

- [Prepare to Install Unix S-Tap](#) – Gather information about the database server and network
- [Uninstall Previous Version of Unix S-Tap](#) – Required for upgrades only
- [Install Unix S-Tap on the Database Server](#)
- [Configure the Guardium Server to Control S-Taps](#) – You only need to do this once to control any number of S-Taps from a Guardium server
- [Complete the Unix S-Tap Configuration from the Guardium Administrator Portal](#)
- [Prepare Local Clients to Use the Tee](#) – Required only the Tee mechanism is being used (instead of the newer K-Tap)

Prepare to Install Unix S-Tap

Before starting the S-Tap installation procedure, read the following overview topics and then gather the information requested.

About Unix S-Tap Data Collection Mechanisms

Depending on how it is installed and configured, Unix S-Tap collects traffic using one of three mechanisms. To collect network traffic, it always uses PCap (see below). Regardless of the mechanism used, the traffic is filtered, so that only database related traffic for specific sets of client and server IP addresses is collected.

Mechanism	Description
PCap	<i>PCap</i> is a packet-capturing mechanism that listens to <i>network</i> traffic from and to a database server. PCap is used to monitor network traffic regardless of whether the K-Tap or Tee mechanism is used for local traffic. On Linux, PCap is also used to capture local TCP/IP traffic on the <i>lo</i> device.
K-Tap	<i>K-Tap</i> is the recommended mechanism to collect <i>local</i> traffic on a Unix database server. Unlike the Tee (see below), with K-Tap you do not need to change how database clients connect to the server. K-Tap is a kernel module that is installed into the operating system. Once installed, it can be enabled or disabled using a configuration file setting. When enabled, it observes all local access to a database server by hooking the mechanisms used to communicate between the database client and server. There are

several mechanisms used for each database type, and K-Tap hooks all of them. When K-Tap is disabled, the Tee can be used to monitor local traffic (see below). K-Tap and Tee are almost always mutually exclusive – to monitor local access you either use K-Tap or the Tee.

Tee

Tee is a proxy mechanism that reads and forwards traffic from local clients to a database server. As the Tee receives database traffic, it forwards one copy to the database server and one copy to S-Tap. When the Tee is used, database clients must connect to the Tee listening port instead of the database listening port. This means that you must either modify how the database client connects to the server, or how the database server accepts client connections. In either case, this is usually a minor configuration change to one or two files (depending on the database type) and the end result is that, as far as the clients are concerned, the Tee is the database, and as far as the database is concerned, the tee is the client. All this is transparent to both the clients and the server – but the configuration change is required to ensure that the connection is made through the Tee. When the Tee is used, database clients can bypass the Tee by connecting to the database listening port (instead of the Tee listening port), or by using named pipes, shared memory, or other inter-process connection mechanisms depending in the database type. We refer to any connections that are not made through the Tee listening port as *rogue* connections. When the Tee is used, you can enable an optional component called the Hunter to watch for, report upon, and optionally disable rogue connections. The Hunter runs at random intervals, so it may not detect all such connections, and while it can report on rogue connections and optionally disable them, it cannot audit what actions were performed by those connections. Another quirk of the Hunter is that when it wakes up to hunt for rogue connections, it can be CPU-intensive, so if you look at the Hunter process at that instant, it may appear to be consuming a lot of a server CPU resource. (The CPU use will drop quickly after the momentary spike.)

Note: To use the Hunter, version 5.8.0 or later of Perl must be installed in the /usr/bin/ directory.

Unix K-Tap Installation Decision

At installation time, you will choose whether or not to load the K-Tap kernel module to the server operating system (see the description of K-Tap, above). This is the only way to load that module. If you do not load K-Tap, and decide later that you want to use it (instead of the Tee), you will need to uninstall S-Tap, and then re-install it.

Unix S-Tap Installation Information

Verify the following items, or gather the information requested in this section of the document. When you have all of this information in hand, you will be ready to begin the Unix S-Tap installation process.

- Verify with your system administrator that:
 - The database server is patched to the latest patch level *recommended* by the database vendor.
 - The database server is configured to capture a core dump.
- (Linux only) Verify with Guardium Support that you have the correct version of S-Tap for your specific Linux OS type, level, Kernel version and 32/64 bit. To do this, execute the following commands and send the output to Guardium support (<mailto:support@guardium.com>):

```
> uname -a  
> lsmod  
> cat /etc/redhat-release  
> cat /etc/issue
```

- If installing CAS (the Change Audit System), verify that Java 1.4.2 or higher is installed on the database server. If it is not, you must install version 1.4.2 (or higher) of Java before installing CAS. To determine the version installed, use the **which** command to locate the **java** command directory. For example:

```
[root@yourserver ~]# which java  
/usr/local/j2sdk1.4.2_03/bin/java
```

If the **which java** command returns a symbolic link, use the **ls -ld <symbolic_link>** command to determine the real Java directory name.

If the **which java** command returns the message *command not found*, Java may be installed, but it has not been included in the PATH variable. If that happens, use the **find** command to locate the Java directory; for example:

```
[root@yourserver ~]# find . -name java  
./usr/bin/  
[root@yourserver ~]#
```

From the Java directory, run the **java -version** command to check the version number. For example:

```
[root@yourserver ~]# /usr/local/j2sdk1.4.2_03/bin/java -version  
java version "1.4.2_03"  
Java(TM) 2 Runtime Environment, Standard Edition (build 1.4.2_03-b02)  
Java HotSpot(TM) Client VM (build 1.4.2_03-b02, mixed mode)  
[root@yourserver ~]#
```

Note the JAVA_HOME directory, which is one directory above the **bin** directory – in the example above, **/usr/local/j2sdk1.4.2_03**, in the space provided below. During CAS installation, you will be prompted to supply this value.

JAVA_HOME	
------------------	--

- *Ignore this point if you will be using K-Tap to monitor local connections.* If you will be using the Tee to monitor local connections, and you will be using the optional Hunter component to detect (and optionally kill) processes that bypass the Tee listening port, verify that you have version 5.8.0 or later of Perl installed on the server in the /usr/bin directory (use the **/usr/bin/perl -v** command). If it is not installed, is installed in a different directory, or an older version is installed, install version 5.8.0 or later of Perl, now, in the /usr/bin directory.
- If there is a firewall between the database server and the Guardium server that will control it, you will need to verify that one or more of the following ports are open.

Port	Used for connection to Guardium server by...
16016 □ (open)	S-Tap
16017 □ (open)	Optional. CAS (Change Audit System)
16018 □ (open)	Optional. S-Tap, for secure TLS communications (see the following bullet).

Note: If you will also be installing S-Tap on Windows servers protected by the same firewall, there are additional ports that will need to be open for the Windows S-Tap (8075, 9501 and 60003). See the *Windows S-Tap Installation* topic for details.

- Decide if the Transport Layer Security (TLS) protocol will be used to ensure secure communications between S-Tap and the Guardium server, and if so, whether or not to fail over to an unencrypted connection if a TLS connection cannot be established. If TLS will be used, you will need to set the two following values in the S-Tap configuration file during the installation process:

Configuration File Property	Description
use_tls=1	TLS will be used. Default is 0 (do not use TLS).
failover_tls=1	The default, fail over to unencrypted communication if a TLS connection cannot be established. Set this value to zero only if you do not want fail over to an unencrypted connection.

- In the space provided below, note the IP addresses of the database server and the Guardium server that will control it. You will need to enter these values in the S-Tap configuration file during the installation process.

IP Address	Configuration File Property	Note the IP address here...
Database server	tap_ip=	
Guardium server	sqlguard_ip=	

- Ignore the following setting *except for* DB2 databases. For a DB2 database, determine the DB2 version and mark the checkbox for the value to use for the configuration file property below.

DB2 Version	S-Tap Configuration File Property
8.1 <input checked="" type="checkbox"/>	db2_fix_pack_adjustment=16
8.2 <input checked="" type="checkbox"/>	db2_fix_pack_adjustment=20

- Ignore the following setting *except for* DB2 databases. For all operating systems *except for* Linux, set the following attribute value to zero. For a Linux DB2 server, set this attribute according to the memory size used by db2bp. If you do not know this size, contact Guardium Support for assistance.

Memory Size Used by db2bp	S-Tap Configuration File Property
All non-Linux <input checked="" type="checkbox"/>	db2_shmem_client_position=0
Linux: 131072 <input checked="" type="checkbox"/>	db2_shmem_client_position=61440
Linux: 671744 <input checked="" type="checkbox"/>	db2_shmem_client_position=327680
Linux: 1064960 <input checked="" type="checkbox"/>	db2_shmem_client_position=524288

- To monitor local database connections, decide whether you will use the K-Tap data collection mechanism, or the Tee mechanism. If you will use the Tee, read *Appendix C: Prepare for Local Clients to Use the Tee* now, but do not perform those procedures until the S-Tap is installed and connected to the Guardium server. To start capturing data using the Tee, you will need to perform the appropriate procedure to prepare all local Unix clients to connect to the S-Tap Tee listening port, instead of the database sever listening port. Mark the local traffic data collection mechanism you will use, and note the properties you will have to set at installation time, below.

Local Traffic Data Collection Mechanism	S-Tap Configuration File Properties
---	-------------------------------------

K-Tap <input type="checkbox"/>	ktp_installed=1 tee_installed=0
Tee <input type="checkbox"/>	tee_installed=1 ktp_installed=0

- If you will use K-Tap to monitor *local* connections on an Oracle or DB2 database, you will need two additional pieces of information when you configure an inspection engine: the full path of the database installation directory, and the database executable file name. You can record those values in the space provided below. For Sybase and Informix, both of these attribute values must be NULL.

Database	S-Tap Configuration File Property	Note the Oracle or DB2 values here, or Enter NULL for both values for Sybase or Informix
Installation directory	db_install_dir=	
Executable file	db_exec_file=	

Uninstall Previous Version of Unix S-Tap

Skip this section if S-Tap is being installed for the first time. If S-Tap has been installed previously, there will be a directory named **/usr/local/guardium/guard_stap**. If you are uninstalling a previous version of S-Tap that used K-Tap, you will need to reboot the database server – *twice* – as described below. If K-Tap has been installed, you will have a device file named **/dev/guard_ktap**.

To uninstall Unix S-Tap:

- Log on to the database server system using the *root* account.
- Locate the **guard_stap** entry in the **etc/initab** file and remove it.
- Run the **init q** command, and then run **ps -ef | grep stap** to verify that S-Tap is no longer running.
- Copy the current S-Tap configuration file to a safe location (a “non-Guardium” directory). By default, the full path name is **/usr/local/guardium/guard_stap/guard_tap.ini**. You can use this file if you have to re-install the older version of the software, or you can refer to it when configuring the new software. But **do not** use the older configuration file directly with the newer version of the software – new properties may be missing, and the defaults taken may result in unexpected behavior when you start S-Tap.

5. If you are uninstalling a previous version of S-Tap that included K-Tap, reboot the database server now.

Do not skip this step.

6. Run the uninstall program, **uninstall.sh**. For example, if the default directory has been used:

```
[root@yourserver ~]# /usr/local/guardium/guard_stap/uninstall.sh
```

Note: *Do not run the uninstall program with S-Tap running. Be sure that you have stopped S-Tap as described in step 2 above.*

7. If your previous version of S-Tap included K-Tap, reboot the database server now (for the second time).

Do not skip this step.

What to do next:

1. If you have not yet done so, *before you install the new S-Tap on the database server*, upgrade the Guardium server to which the S-Tap sends data, to the new software release level.
2. Continue with the following topic, to install the new S-Tap software on the database server.

Install Unix S-Tap on the Database Server

Install Unix S-Tap (and optional CAS) software on the database server by running the installation script as described below. If any stage of the installation fails, undo all of the steps up to that point. Do not leave S-Tap partially installed.

1. Log on to the database server system using the *root* account.
2. Copy the appropriate S-Tap installer script from the CD (or network), to the **/tmp** or **/var/tmp** directory. The installer script name clearly identifies the operating system (for example, *aix53-sqltap_installer-nn_m.sh* or *sparc-solaris10-sqltap_installer-nn_m.sh*, where *nn_m* indicates the software release version number).

Some companies require the use of native installers to register packages on the system or perform other security or house-keeping functions. *If your company does not require the use of native installers, skip ahead to Step 3.*

If your company requires that you use native installers, see the appropriate section (depending on your operating system) in the *Using Native Installers* section, below, and then continue with Step 3.

Using Native Installers

Refer to the appropriate native installer section below if your company requires the use of operating system tools for installation. The S-Tap native installers generate the Guardium S-Tap script installation file on the target machine. Following this step, the user must run the .sh file the same way as always when installing S-Tap. There is a different native installer for each supported operating system. The main purpose of the native installers is to ensure that the S-Taps are registered within the operating system's "asset repository".

Solaris Native Installer

To install, use the following command, and then continue with Step 3 of the installation procedure:

```
pkgadd -d <filename>.pkg
```

Installation output should look something like this:

```
## Installing part 1 of 1.  
/var/tmp/sparc-solaris8-sqltap_installer-nn_m.sh  
WARNING: installing with  
default mode of 644  
verifying class ]  
## Executing postinstall script.  
sparc-solaris8-sqltap_installer-nn_m.sh available on /var/tmp
```

Installation of was successful.

To uninstall, use the following command:

```
pkgrm GrdTapIns?
```

Linux Native Installer

To install, use the following command, and then continue with Step 3 of the installation procedure:

```
rpm -i <filename>.rpm
```

For example:

```
rpm -i redhat-as-sqlstap-installer-n.m.i686.rpm
```

To uninstall, use the following command:

```
rpm -e <name>
```

For example:

```
rpm -e redhat-as-sqlstap-installer-n.m
```

To determine <name>, use the following command

```
rpm -qa
```

For example:

```
[root@uranus ~]# rpm -qa | grep tap  
redhat-as-sqlstap-installer-n.m
```

HPUX Native Installer

To install, use the following command, and then continue with Step 3 of the installation procedure:

```
swinstall -s /var/tmp/<filename>.depot @ <hostname>:/var/spool/sw
```

This is an interactive program. Follow the prompts and use the appropriate controls to install the appropriate S-Tap installation program, which will be located in /var/spool/sw/var/tmp

To uninstall, use the following command:

```
swremove @<hostname>:/var/spool/sw
```

AIX Native Installer

To install, use the following command, and then continue with Step 3 of the installation procedure:

```
/usr/lib/instd/sm_inst installp_cmd -a -Q -d '.' -f bff-file-name  
'-c' '-N' '-g' '-X' '-G'
```

To uninstall, use the following command:

```
/usr/lib/instd/sm_inst installp_cmd -u -f bff-file-name
```

3. Run the installer.
4. Respond to the legal notification and other prompts, as directed by the installer. If possible, we suggest that you accept all of the supplied defaults.

The installer opens the S-Tap configuration file for editing, under *vi*. Although you can modify all of the configuration file properties at this time, we suggest that you modify only the properties described below. As you read through the *Before You Start* topic above, you should have noted the values for all of these settings in that section of the document. Setting these properties will allow you to start the S-Tap and connect to the Guardium server. After that, you can make S-Tap configuration changes from the admin portal of the Guardium server.

5. Locate and set the following parameters, as described below:

- **tap_ip=NULL**

Set **tap_ip** to the IP address of the database server. For systems other than HPUX, you can enter the symbolic name of the sever instead of the IP address.

- **sqlguard_ip=NULL**

Set **sqlguard_ip** to the IP address of the primary Guardium server for this S-Tap. The *primary* server is the one that S-Tap will try to connect with each time that it restarts. You can identify secondary (failover) Guardium servers later, from the Guardium *admin* GUI.

- **tee_installed=0**
- **ktap_installed=1**

To use K-Tap, set **tee_installed=0** and **ktap_installed=1**. Or to use the Tee process instead of K-Tap, set the opposite values (**tee_installed=1** and **ktap_installed=0**).

- **use_tls=0**

To use Transport Layer Security (TLS), set **use_tls=1**. Note that there is a second TLS-related attribute, for which you will probably want to accept the default: **failover_tls=1** instructs S-Tap to fail over to an unencrypted connection if a secure TLS connection cannot be established. You should set **failover_tls=0** *only if* you do not want to fail over to an unencrypted connection if a secure TLS connection cannot be established.

- **db2_fix_pack_adjustment=16**

For DB2 databases only, set **db2_fix_pack_adjustment=16** for DB2 ver.8.1, or **20** for ver.8.2.

- **db2_shmem_client_position=0**

For DB2 databases only, set **db2_shmem_client_position=0** (zero, the default) for non-LINUX servers. For LINUX servers, set the value depending on the shared memory size used by **db2bp**:

Memory Size Used by db2bp	db2_shmem_client_position value
131072	61440
671744	327680
1064960	524288

Note: Contact Guardium support if you need help determining the shared memory size used by **db2bp**.

- **devices=none**

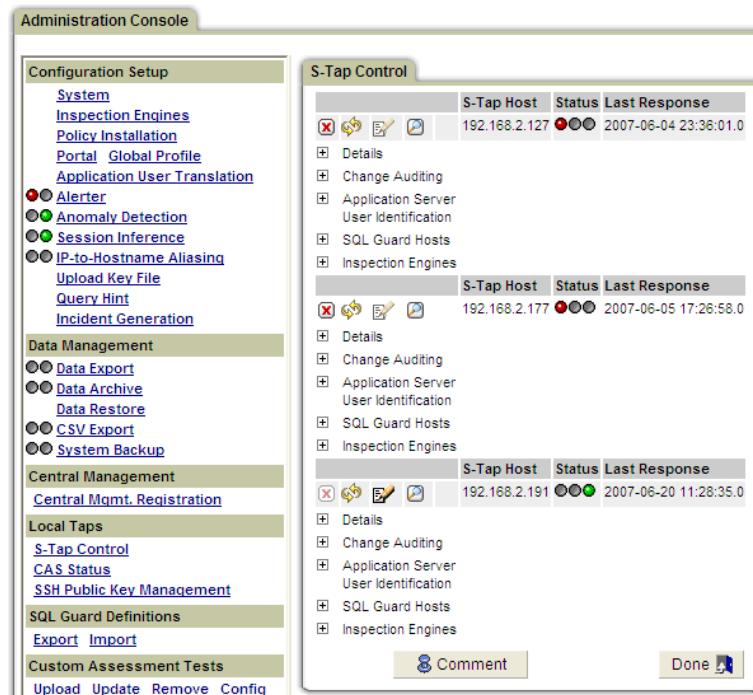
For all operating systems except for Linux, **devices=none** means that only local traffic will be monitored. To monitor local traffic only on a Linux system, set this value to **devices=lo** (as in “local loop back”). On any system, to monitor

both local and network traffic, set the **devices** parameter value to the device name on which the server IP address is defined. To find the device name on an HPUX system, use the **lanscan** command. On all other systems, use the **ifconfig -a** command.

6. Save the **guard_tap.ini** file and quit **vi** (use the **wq** command). The install program will check the parameter values you have set. If OK, it will continue to the next prompt. Otherwise, you will need to correct any erroneous parameters and then save the **guard_tap.ini** file again.
7. Respond to the CAS installation prompt.

Complete the Unix S-Tap Configuration from the Guardium Administrator Portal

1. Log on to the *admin* portal of the Guardium server configured to manage the S-Tap (the Guardium server whose IP address was specified in the **sqlguard_ip** parameter, as described previously).
2. Click **S-Tap Control** in the Local Taps section of the Administration Console, to open the S-Tap control panel. If there is no Local Taps section on the Administration Console menu, this Guardium server has not been configured as an S-Tap host. Before continuing, you must perform the procedure described at the start of this section: [Configure the Guardium Server to Control S-Taps](#).



3. Locate the S-Tap to be configured by finding the IP address or symbolic host name of the database server in the S-Tap Host column.
- Status**
4. Verify that the Status indicator is green (the rightmost light): . If it is not, the Guardium server and S-Tap are not connected.
 5. Click the (Edit) button for that S-Tap. If the Edit button is not active, this Guardium server is not the active host for this S-Tap. You will need to log on to the active host for this S-Tap to make any changes.
 6. Expand the Details section and make any necessary changes as described below:

The screenshot shows the 'Details' configuration pane for an S-Tap. It includes fields for 'Devices' (set to 'lan0'), 'Load Balancing' (unchecked), 'Messages' (with 'Remote' and 'Syslog' checked), 'Trace Files Dir' (set to '/usr/local/guardium/guard_stap/trace_files'), 'Alternate Ips' (empty), and 'App. Server User Identification' (unchecked).

- In the Devices box, enter any interfaces on which S-Tap should listen for remote connections to the database server.
 - Mark the Load Balancing box if S-Tap will balance traffic to all SQL Guard servers listed in the SQL Guard Hosts pane (see below). Load balancing is done by Client IP address and client port, since all traffic for a session between a specific client and a specific server must be viewed by the same Guardium server.
 - Mark the Remote box to send messages to the active SQL Guard host.
 - Mark the Syslog box to write messages to the syslog file on the database server.
 - In the Trace Files Dir box, enter the directory in which trace files will be created.
 - In the Alternate Ips box, enter all of the alternative or virtual IP addresses that are used to connect to this server. This is especially important when your server has multiple network cards with multiple IPs, or virtual IPs. S-Tap will only pass data to the Guardium server when the destination IP matches one of these listed IPs.
 - Mark the App. Server User Identification box if you are using this S-Tap for end-user identification and this S-Tap is being installed on the application server box rather than the database server. This type of installation is an advanced option not currently described in this document.
7. If the K-Tap has been installed, skip this step (there will be no Hunter section). Expand the Hunter section:

The screenshot shows the expanded 'Hunter' section of the configuration pane. It includes fields for 'Hunt' (empty), 'Sleep Time' (set to '1'), and 'DBs' (empty).

- Use the Hunt box to identify any processes to be killed, using the following syntax: `db_type:process[,db_type:process]`

Where the **db-type** can be DB2, Informix, Oracle or Sybase, and the processes may be any of the following:

SHM	Shared memory
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
FIFO	A named pipe IPC mechanism
PIPE	A simple (unnamed) pipe IPC
INET	Internet Protocol (HPUX)

These values are not case-sensitive, and each entry is separated from the next by a comma. Example: To kill Oracle Bequeath processes, which uses a simple pipe, you would enter: `oracle:pipe`

- In the Sleep Time box, enter the maximum number of seconds between the randomized starting time of the hunter's rogue process search routine. The start time is random to increase the difficulty of defeating it by running in fixed time slots or intervals. The recommended value for sleep time is anywhere between 60 and 300.
- In the DBs box, identify the databases to be reported on, separating each entry from the next with a comma. The allowed entries are: Informix, DB2, Sybase, and Oracle.

8. Optional. Expand the SQL Guard Hosts pane:

SQL Guard Hosts	
Active	SQL Guard Host
✓	192.168.3.101
	192.168.3.102
	192.168.3.222

If you have multiple SQL Guard servers, you can designate one or more of them as a secondary host for this S-Tap, in the SQL Guard Hosts pane. If the primary SQL Guard server becomes unavailable, S-Tap will fail over to a secondary host. See [Secondary SQL Guard Hosts for S-Tap](#) for instructions on how to define and configure secondary hosts.

9. On the add configuration panel, the Inspection Engines section is always expanded. This section displays differently depending on whether K-Tap or the Tee is being used. The K-Tap version of the panel is shown below. The only

difference is that if the Tee is being used, two boxes will replace the KTAP DB Real Port box, as shown below:

The screenshot shows a software interface for configuring inspection engines. It displays two separate configurations for the Oracle protocol. Each configuration includes fields for Protocol (Oracle), Port Range (1521-1521), KTAP DB Real Port (1521), IP (192.168.2.0), Mask (255.255.255.0), DB Install Dir (/home/oracle10), and Process Name (/home/oracle10/oracle/product/10.2.0/db_2). Below these, there is an 'Add Inspection Engine...' button, followed by another set of configuration fields for a new engine. This second set includes dropdown menus for Protocol and Port Range, and input fields for IP, Mask, Client Ip/Mask, Exclude Client Ip/Mask, DB Install Dir, and Process Name. At the bottom are 'Cancel' and 'Apply' buttons.

- To remove an existing inspection engine, click its Remove button (X).
- To begin defining a new inspection engine, select a protocol from the Protocol list in the Add pane (you cannot change an inspection engine's protocol once it has been defined).
- For the ports to be monitored, enter a beginning and ending port number in the Port Range boxes. Do not enter a large, all-inclusive range of ports here, as S-Tap may become bogged down attempting to parse traffic that is of no interest.
- If K-Tap is installed, enter the database listening port in the KTAP DB Real Port box, or if the Tee is being used, enter the Tee listening port in the Listen Port box, and the database listening port in the Real Port box.
- In the Client IP/Mask boxes, enter an IP mask and its corresponding subnet mask to select which clients to monitor. Click the + (Add Client Ip/Mask) button to add the entry. You can add multiple entries. To remove the last Client IP/Mask pair entered, click the - (Remove Client Ip/Mask) button.

Notes: Do *not* monitor traffic that is also monitored directly by the SQL Guard host for this S-Tap (for example through SPAN sessions). If that happens, that traffic will be ignored.

If the IP address is the same as the IP address for the database server, and a mask of 255.255.255.255 is used, all network traffic arriving at the server is filtered out.

- Optional. In the Exclude Client Ip/Mask box, enter an IP mask and its corresponding subnet mask to select which clients to exclude. Click the  (Add Client Ip/Mask) button to add the entry. You can add multiple entries. To remove the last Client IP/Mask pair entered, click the  (Remove Client Ip/Mask) button. This option allows you to configure the S-Tap to monitor all clients, except for a certain client or subnet (or a collection of these).
- For Sybase and Informix, the **DB Install Dir** and **Process Name** must be **NULL**. Oracle and DB2 require values to monitor local traffic. To obtain those values, log onto the server in command line mode and use **su - <account>**, on the database server, and then use **which oracle** or **which db2sysc** to display the **Process Name** value. The **DB install Dir** is derived from this value – the directories above which the executables reside.

For example, if **which oracle** returns /data1/home/oracle9/bin/oracle, then the **DB install Dir** would be set to /data1/home/oracle. If **which db2sysc** returns /export/home/db2inst1/sqllib/adm/db2sysc, then the **DB install Dir** would be set to /export/home/db2inst1.

10. Click the Add button if you are adding a new inspection engine.

11. Click Apply to save all changes.

In the S-Tap Control panel, you may see the status light  for the S-Tap turn yellow, indicating that configuration changes have been sent to the S-Tap, but the S-Tap has not yet restarted with the new configuration. If the light remains yellow for an extended period of time, you can assume that the S-Tap was unable to restart using the new configuration. When that happens, S-Tap attempts to restart using the last good configuration. When an error has occurred, you can review the errors by opening the S-Tap Events panel in a separate window by clicking the Show S-Tap Event Log button . In most cases the event log will contain error messages indicating what was wrong with the new configuration. See *Viewing the S-Tap Events Panel*, below. To

reload the last good configuration from the S-Tap host, click  (Refresh S-Tap Information) in the S-Tap Control panel.

This completes the Unix S-Tap installation procedure.

Windows S-Tap Installation

To install and start using S-Tap on a Windows database server, you must complete the following tasks, each of which is described as a separate procedure, below.

- [Prepare to Install Windows S-Tap](#) - Gather information about the database server and network
- [Uninstall Previous Version of Windows S-Tap](#) – Required for upgrades only
- [Install Windows S-Tap on the Database Server](#)
- [Configure the Guardium Server to Control S-Taps](#)– You only need to do this once to control any number of S-Taps from a Guardium server
- [Complete the Windows S-Tap Configuration from the Guardium Administrator Portal](#)

Prepare to Install Windows S-Tap

Before starting the S-Tap installation procedure, read the following overview topics and then gather the information requested.

About Windows S-Tap Drivers

Depending on how it is installed and configured, S-Tap collects local traffic using any of several drivers. In a default installation, all drivers are installed. Regardless of how it is captured, the traffic is filtered, so that only database related traffic for specific sets of client and server IP addresses is collected.

Data Collection Drivers

Driver	Monitors Database Access via...
LHmon	Local and remote TCP/IP traffic
Named Pipes	Named pipes
Shared Memory	Shared memory (except for DB2)
DB2 Shared Memory	Shared memory (DB2 only)

Kerberos Driver

The Kerberos driver is installed by default. It monitors Kerberos traffic, which allows Guardium to map and convert Kerberos-encoded user names to database user names. (The Kerberos encoded names look like clear-hex strings.)

Before starting the S-Tap installation procedure, verify the following items, or gather the information requested in this section of the document. When you have all of this information in hand, you will be ready to begin the installation process.

- Verify with your system administrator that:
 - The database server is patched to the latest patch level *recommended* by the database vendor.
 - The database server is configured to capture a core dump.
- If there is a firewall between the database server and the Guardium server that will control it, you will need to verify that the ports used by S-Tap are not being blocked by the firewall.

Port	Used for connection to Guardium server by...
8075 <input checked="" type="checkbox"/> (open)	S-Tap (always required)
60003 <input checked="" type="checkbox"/> (open)	Optional. CAS (Change Audit System)
9501 <input checked="" type="checkbox"/> (open)	Optional. S-Tap, for secure TLS communications (see the following bullet).

Note: If you will also be installing S-Tap on Unix servers protected by the same firewall, there are several additional ports (16016, 16017 and 16018) that may need to be open for the Unix S-Tap. See the *Unix S-Tap Guide* for more details.

- Decide if the Transport Layer Security (TLS) protocol will be used to ensure secure communications between S-Tap and the Guardium server, and if so, whether or not to fail over to an unencrypted connection if a TLS connection cannot be established. If TLS will be used, you will need to set the two following values in the S-Tap configuration file during the installation process:

S-Tap Configuration File Property	Description
<code>use_tls=1</code>	TLS will be used. Default is <code>0</code> (do not use TLS).
<code>failover_tls=1</code>	The default, fail over to unencrypted communication if a TLS connection cannot be established. Set this value to zero only if you do not want fail over to an unencrypted connection.

- Note below the IP addresses of the database server and the Guardium server that will control it.

IP Address	S-Tap Configuration File Property	Note the IP address here...
Database server		<code>tap_ip=</code>
Guardium server		<code>sqlguard_ip=</code>

- If you will monitor a DB2 database on this server, determine the DB2 version and mark the checkbox for the value to use for the configuration file property below. Ignore this setting if there is not a DB2 database on this server. Note that if you also have Unix based DB2 servers, the Unix values differ from the Windows values for each release of DB2.

DB2 Version	S-Tap Configuration File Property
8.1 <input checked="" type="checkbox"/>	<code>db2_fix_pack_adjustment=20</code>
8.2 <input checked="" type="checkbox"/>	<code>db2_fix_pack_adjustment=80</code>

- For DB2 databases on Windows servers, the following value must be **zero**.

db2bp Memory Size	S-Tap Configuration File Property
All Windows servers	<code>db2_shmem_client_position=0</code>

Uninstall Previous Version of Windows S-Tap

Skip this section if S-Tap is being installed for the first time. If you are uninstalling a previous version of S-Tap, you will need to reboot the database server.

To uninstall a Windows S-Tap:

1. Log on to the database server system using a system administrator account.
2. Copy the current S-Tap configuration file to a safe location (a “non-Guardium” directory). It is located in the `\Windows\System32` directory, and is named `guard_tap.ini`.
3. From the Services control panel, stop the GUARDIUM_STAP and GUARDIUM_TEE services if they are running. (Typically, the GUARDIUM_TEE service will *not* be running.)
4. From the Add/Remove Programs control panel, remove GUARDIUM_STAP.

5. Reboot the database server (***do not skip this step***).

What to do next:

1. If you have not yet done so, *before you install the new S-Tap on the database server*, upgrade the Guardium server to which the S-Tap sends data, to the new software release level.
2. Continue with the following topic, to install the new S-Tap software on the database server.

Install Windows S-Tap on the Database Server

Install S-Tap (and the optional CAS component) on the database server by following the procedure outlined below.

1. Log on to the host system using a system administrator account.

Notes: Read *all* bullets below before running the S-Tap installation wizard:

- If installing S-Tap on a Domain Controller to capture **Kerberos** traffic only, choose **Custom Installation**, and **clear** the checkboxes for all drivers *except* for the Kerberos driver, which is the only S-Tap driver you should install on the Domain Controller.
 - If you did not record the IP address of the database server or domain controller on which you are installing S-Tap in the previous section, obtain that address now.
 - If you did not record the IP address of the Guardium server that will control this S-Tap in the previous section, obtain that address now.
 - Decide if you want to install the CAS agent, so that you can respond appropriately when running the installation wizard.
 - Except as noted above, take the defaults for all other options suggested by the wizard. All of those items can be configured more easily from the Guardium administrator portal.
2. Insert the S-Tap installation disk in the CD drive and follow the installation instructions provided by the installation wizard.
 3. In this step, you must edit the **guard_tap.ini** file.

Perform this step only if:

- You will be using TLS for secure communications between S-Tap and the Guardium server

OR

- You will be monitoring a DB2 database on this server.
- a) Using a text editor, edit the **guard_tap.ini** file in the **\Windows\System32** directory. The file contains extensive comments. Each comment line begins with a semi-colon character (;) in the first column. The installation program stores a copy of the original, un-edited S-Tap configuration file (**guard_tap.ini**) in the installation directory. **Do not** edit that version of the file. The version of the configuration file used by S-Tap is stored in the Windows System32 directory.
- b) If you will be using TLS, locate the **use_tls** attribute, and set as follows:

use_tls=1

Note that there is a second TLS-related attribute, for which you will probably want to accept the default: **failover_tls=1** instructs S-Tap to fail over to an unencrypted connection if a secure TLS connection cannot be established. You should set **failover_tls=0** *only if* you do not want to fail over to an unencrypted connection when a secure TLS connection cannot be established.

- c) if you will be monitoring a DB2 database on this server Locate and set the **db2_fix_pack_adjustment** attribute value, as described earlier in the previous section (refer to that section for more detailed information).

DB2 Version	S-Tap Configuration File Property
8.1 <input type="checkbox"/>	db2_fix_pack_adjustment=20
8.2 <input type="checkbox"/>	db2_fix_pack_adjustment=80

- d) Save the **guard_tap.ini** file in **text** format in the **\Windows\System32** directory.

4. Restart S-Tap:

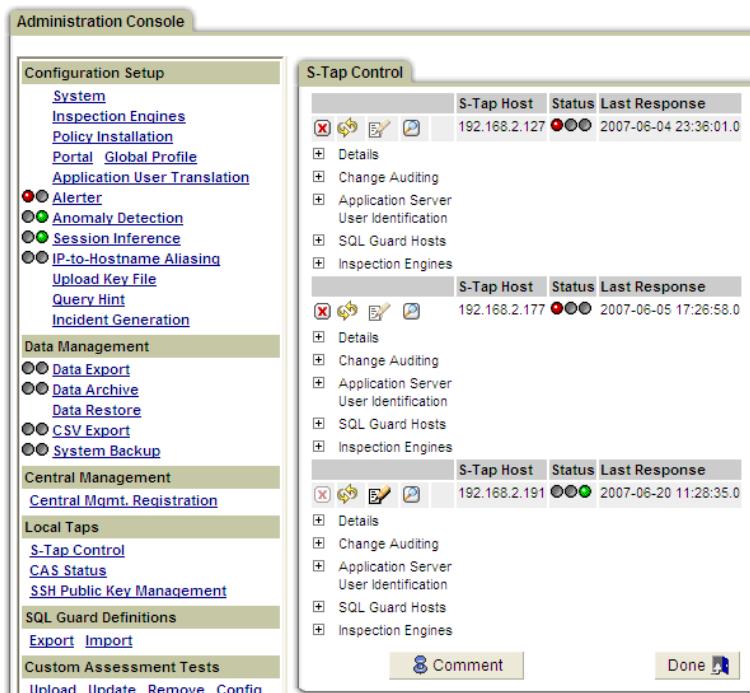
- Click **Start – Control Panel – Administrative Tools – Services** to open the Services utility and restart the SQL Guard service.
- After a few moments, check the Windows Event Log for any error messages.

Tip: If the following error message appears in the Event Log:
Configuration file has the error: No IP was found for host
???
you have edited the wrong version of the configuration file. You
must edit the version of **guard_tap.ini** that is stored in the
\Windows\System32 directory.

5. You must reboot *all* instances of the databases to be monitored, before any local traffic will be captured.

Complete the Windows S-Tap Configuration from the Guardium Administrator Portal

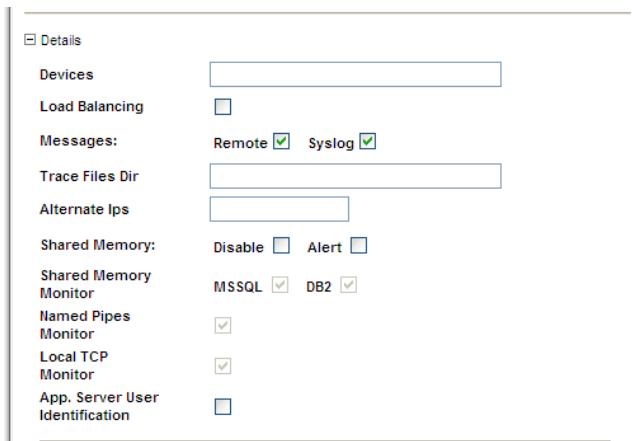
1. Log on to the *admin* portal of the Guardium server configured to manage the S-Tap (the Guardium server whose IP address was specified in the **sqlguard_ip** parameter, as described previously).
2. Click **S-Tap Control** in the Local Taps section of the Administration Console, to open the S-Tap control panel. If there is no Local Taps section on the Administration Console menu, this Guardium server has not been configured as an S-Tap host. Before continuing, you must perform the procedure described at the start of this section: [Configure the Guardium Server to Control S-Taps](#).



3. Locate the S-Tap to be configured by finding the IP address or symbolic host name of the database server in the S-Tap Host column.

Status

4. Verify that the Status indicator is green (the rightmost light):  . If it is not, the Guardium server and S-Tap are not connected.
5. Click the  (Edit) button for that S-Tap. If the Edit button is not active, this Guardium server is not the active host for this S-Tap. You will need to log on to the active host for this S-Tap to make any changes.
6. Expand the Details section and make any necessary changes as described below:



Devices	<input type="text"/>
Load Balancing	<input type="checkbox"/>
Messages:	Remote <input checked="" type="checkbox"/> Syslog <input checked="" type="checkbox"/>
Trace Files Dir	<input type="text"/>
Alternate Ips	<input type="text"/>
Shared Memory:	Disable <input type="checkbox"/> Alert <input type="checkbox"/>
Shared Memory Monitor	MSSQL <input checked="" type="checkbox"/> DB2 <input checked="" type="checkbox"/>
Named Pipes Monitor	<input checked="" type="checkbox"/>
Local TCP Monitor	<input checked="" type="checkbox"/>
App. Server User Identification	<input type="checkbox"/>

- Leave the Devices box blank.
- Mark the Load Balancing box if S-Tap will balance traffic to all SQL Guard servers listed in the SQL Guard Hosts pane (see below). Load balancing is done by Client IP address, since all traffic for a session between a specific client and a specific server must be viewed by the same Guardium server.
- Mark the Remote box to send messages to the active SQL Guard host.
- Mark the Syslog box to write messages to the syslog file on the database server.
- In the Trace Files Dir box, enter the directory in which trace files will be created.
- In the Alternate Ips box, enter all of the alternative or virtual IP addresses that are used to connect to this server.
- On the Shared Memory line, mark the Disable box if you want S-Tap to disable shared memory connections, and/or mark the Alert box if you

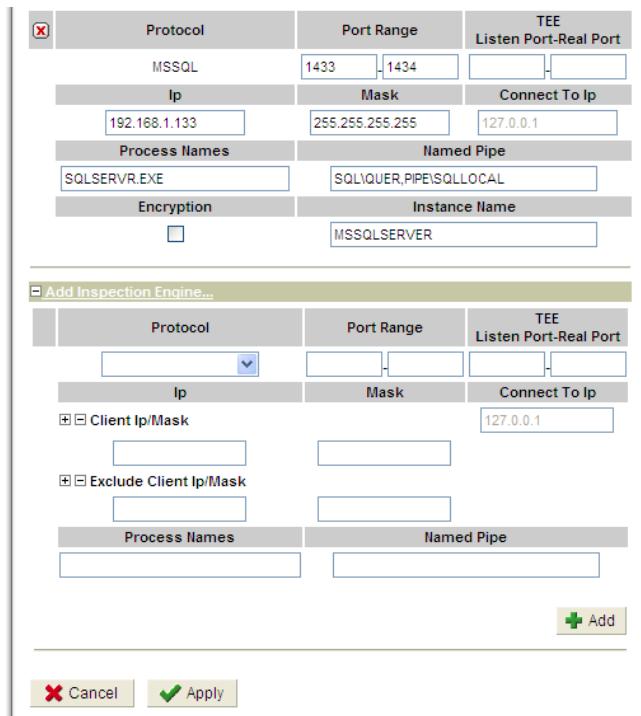
want S-Tap to issue an alert when shared memory connections are detected.

- The **Monitor** boxes (Shared Memory for MS SQL and DB2), Named Pipes and Local TCP indicate what drivers have been installed. These settings cannot be changed. To remove a driver, you must uninstall the product, re-boot, and then perform a **Custom** re-installation, omitting any drivers you do not want installed.
 - Mark the App. Server User Identification box if users will be identified by the Application Server.
7. Ignore the Hunter section. It applies for Unix S-Tap only.
8. Optional. Expand the SQL Guard Hosts Pane:

SQL Guard Hosts	
Active	SQL Guard Host
✓	192.168.3.101
	192.168.3.102
	192.168.3.222
	<input type="button" value="Add"/>

If you have multiple SQL Guard servers, you can designate one or more of them as a secondary host for this S-Tap, in the SQL Guard Hosts pane. If the primary SQL Guard server becomes unavailable, S-Tap will fail over to a secondary host. See [Secondary SQL Guard Hosts for S-Tap](#) for instructions on how to define and configure secondary hosts.

9. On the add configuration panel, the Inspection Engines section is always expanded. This section displays differently depending on whether you are looking at a Unix or Windows database server – a Windows example is illustrated below:



- To remove an existing inspection engine, click its Remove button (☒).
- To define a Kerberos inspection engine on a Domain Controller:
 - Expand the Add Inspection Engine panel as illustrated above.
 - Select **Kerberos** as the Protocol.
 - Enter **88** in both Port Range boxes.
 - Define one or more **Client IP/Mask** sets, and click Add.
- To catch Kerberos traffic using an inspection engine on an MS-SQL database server:
 - Expand the Add Inspection Engine panel as illustrated above.
 - Select **MSSQL** as the Protocol.
 - Enter **1433** and **1434**, respectively, in the Port Range boxes.
 - Enter **SQLSERVR.EXE** in the Process Names box
 - Enter **MSSQLSERVER** in the Instance Name box (when you select MSSQL as the Protocol, this name appears by default –

you will only need to change it if your server does not use the default).

- Define one or more **Client IP/Mask** sets, and click Add.
- To begin defining a new inspection engine, select a protocol from the Protocol list in the Add pane (you cannot change an inspection engine's protocol once it has been defined).
- For the ports to be monitored, enter a beginning and ending port number in the Port Range boxes. Do not enter a large, all-inclusive range of ports here, as S-Tap may become bogged down attempting to parse traffic that is of no interest.
- If the Tee mechanism is being used to monitor traffic (this is *never* recommended for Windows S-Tap), enter the Tee listening port in the Listen Port box, and the database listening port in the Real Port box.
- In the Client IP/Mask boxes, enter an IP mask and its corresponding subnet mask to select which clients to monitor. Click the  (Add Client Ip/Mask) button to add the entry. You can add multiple entries. To remove the last Client IP/Mask pair entered, click the  (Remove Client Ip/Mask) button.

Notes: Do *not* monitor traffic that is also monitored by the SQL Guard host for this S-Tap. If that happens, the SQL Guard unit receives duplicate packets, it is unable to reconstruct messages, and that traffic will be ignored.

If you are using the TEE and leave the Client IP/Mask boxes blank, all clients will be monitored.

If the IP address is the same as the IP address for the database server and a mask of 255.255.255.255 is used, all network traffic arriving at the server is monitored.

- Optional. In the Exclude Client Ip/Mask box, enter an IP mask and its corresponding subnet mask to select which clients to exclude. Click the  (Add Client Ip/Mask) button to add the entry. You can add multiple entries. To remove the last Client IP/Mask pair entered, click the  (Remove Client Ip/Mask) button.
- For Sybase and Informix, the **DB Install Dir** and **Process Name** must be **NULL**. Oracle and DB2 require values to monitor local traffic. To obtain those values, log onto the server in command line mode and use **su -**

<account>, on the database server, and then use **pwd** to display the **DB Install Dir** value, and **which oracle** or **which db2sysc** to display the **Process Name** value

10. Click the Add button if you are adding a new inspection engine.
11. Click Apply to save all changes.

In the S-Tap Control panel, you may see the status light  for the S-Tap turn yellow, indicating that configuration changes have been sent to the S-Tap, but the S-Tap has not yet restarted with the new configuration. If the light remains yellow for an extended period of time, you can assume that the S-Tap was unable to restart using the new configuration. When that happens, S-Tap attempts to restart using the last good configuration. When an error has occurred, you can open the S-Tap Events panel in a separate window by clicking the Show S-Tap Event Log button . In most cases the event log will contain error messages indicating what was wrong with the new configuration. See *Viewing the S-Tap Events Panel*, below. If the Events Panel does not provide enough information, see *Appendix D: Troubleshooting*.

To reload the last good configuration from the S-Tap host, click  (Refresh S-Tap Information) in the S-Tap Control panel.

This completes the Windows S-Tap installation procedure.

Secondary SQL Guard Hosts for S-Tap

If the Guardium server designated as the primary host for an S-Tap becomes unavailable, S-Tap can be configured to fail over to a secondary host. It remains connected to the secondary host until either that connection is lost or the S-Tap is restarted. Each time one of those events happens, S-Tap attempts to reconnect to its primary Guardium server.

Note that S-Tap restarts under slightly different conditions, depending on the database server operating system:

- On a Unix server, S-Tap restarts each time configuration changes are applied from the active host.
- On a Windows server, S-Tap restarts only when the server restarts, or when the S-Tap service is restarted from the (Windows) Services control panel.

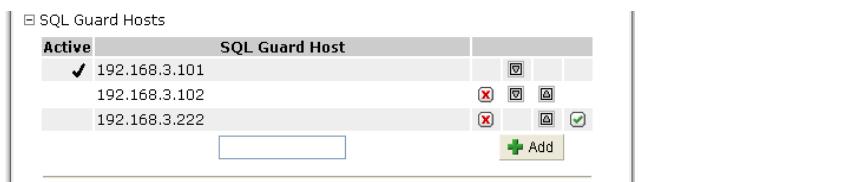
It thus follows that on a Windows database server, if you change the primary host for an S-Tap, you will need to restart the S-Tap service before S-Tap will start using the new primary host.

Before designating a Guardium server as a secondary host for S-Tap, you should verify the following points:

- The Guardium server must be configured to manage S-Taps. To check this and reconfigure the server if necessary, see [Configure the Guardium Server to Control S-Taps](#), at the start of this section.
- The Guardium server must have connectivity to the database server on which the S-Tap is installed. When multiple Guardium servers are used, they are often attached to disjointed branches of the network.
- The Guardium server must *not* have a security policy that will ignore session data from the database server on which the S-Tap is installed. In many cases, the Guardium security policy is built to focus on a very narrow subset of the observed database traffic, ignoring all other sessions. Either make sure that the secondary host will not ignore session data from the S-Tap, or modify the security policy on the Guardium server as necessary.

To define a secondary host for an S-Tap:

1. Log on to the *admin* portal for the *active* SQL Guard host for the S-Tap. The active host is the *only* host from which you can modify the S-Tap configuration.
2. Click **S-Tap Control** in the Local Taps section of the Administration Console, to open the S-Tap control panel.
3. Locate the database server on which the desired S-Tap is installed, and click  (Edit S-Tap Configuration) to open the S-Tap Configuration panel.
4. In the S-Tap Configuration panel, click the  button beside SQL Guard Hosts to expand that pane:



The first host listed is the *primary* host for the S-Tap. Following any outage or restart, S-Tap attempts to connect to the primary host first.

5. Enter the IP address of the secondary SQL Guard host in the text box.
6. Click the Add button.
7. Optional. To change the host designated as the primary host (the first in the list), or to change the order of secondary hosts, do one of the following:
 - Click the  or  buttons beside each host as necessary.

OR

- To designate the last host listed as the primary host, click the Set Primary button () in the right-most column for that host. This will move that host all the way to the top of the list.
8. Optional. To remove a secondary host, click its Remove button ().
 9. When you are done, click the Apply button at the bottom of the S-Tap Configuration panel.

Viewing the S-Tap Events Panel

You can use the S-Tap Events Panel to view the event messages output by S-Tap.

Note: If no messages display in the S-Tap Events Panel, the production of event messages may have been disabled in the configuration file for that S-Tap. If this is the case, you may be able to locate S-Tap event messages on the host system in the Event Log (Windows) or the *syslog* file (Unix/Linux).

To open the S-Tap Events panel for any S-Tap listed in the S-Tap Control panel, click the  (Show S-Tap Event Log) button for that S-Tap. The S-Tap Events Panel opens in a separate browser window.

Column or Control	Description
Event Type	Identifies a type of event: Success, Error Type, etc.
Event Description	Provides a short description of the event.
Timestamp	Provides the date and time that the event occurred.
Done	Click the Done button to close the window.

Prepare for Local Unix Clients to Use the TEE

This section does *not* apply if the K-TAP version has been installed. The K-TAP version does not use the TEE.

For Unix database servers, follow these procedures to re-direct local clients to the S-Tap Tee listening port.

Prepare for Local Unix DB2 Clients to Use the Tee

Do not perform this procedure until S-Tap has been installed on the DB2 server, and you are ready to start collecting data. For the local DB2 clients to use the Tee, you will create a database alias named *tee*, and the clients will change their login sequence to log into *tee* (instead of the DB2 server).

While performing this procedure, note the S-Tap configuration file properties you will have to set to start collecting data, in the table below.

S-Tap Configuration File Properties for DB2 Inspection Engine

Property=Value	Example Value
db_type=DB2	DB2
tee_listen_port=	12344
real_database_port=	50000

1. Log on to the database server system using an administrative account.
2. Locate the entry in the /etc/services file for the node name that clients use to connect to the database. Each entry in this file is in the following format:

```
node_name      port_number/protocol [aliases]
```

For example:

```
db2inst1      50000/tcp      # DB2 connection service port
```

Note the node name here (*db2inst1*, in the example above):

_____ , and record the port number (50000, in the example) in the table above as the **real_database_port** value.

3. Select an unused port number in the range of 1025-65535 for use by S-Tap. Search the services file for the selected port number to be certain that it is not used. In the table above, record the port number (12344, for example) as **tee_listen_port** value.
4. Enter the **db2** command to start the db2 command-line interface. To execute this command, you may need to add the command to the \$PATH, or switch users to a db2 user on the system.
5. Enter the **list node directory** command to list all nodes defined. A very simple example is illustrated below:

```
db2 => list node directory
Node Directory
Number of entries in the directory = 2
Node 1 entry:
  Node name          = GACCTEST
  Comment           =
  Directory entry type = LOCAL
  Protocol          = TCPIP
  Hostname          = merlin
  Service name      = 50000
```

```
Node 2 entry:  
Node name = LOCGOOSE  
Comment =  
Directory entry type = LOCAL  
Protocol = LOCAL  
Instance name = db2inst1
```

Note that the /etc/services entry that we looked at previously related the instance name *db2inst1* to the service name 50000.

6. Use the **catalog** command to create a node on the local server for the port to be assigned as the Tee listening port (12344 in our example). For example, to name the node LOCALTEE on the server named goose, we would enter the following command:

```
db2 => catalog tcpip node localtee remote goose server 12344  
DB20000I The CATALOG TCPIP NODE command completed successfully.  
DB21056W Directory changes may not be effective until the  
directory cache is refreshed.
```

7. Enter the **terminate** command to update the directory. (This closes the db2 utility.)

```
db2 => terminate  
DB20000I The TERMINATE command completed successfully.
```

8. Restart the db2 utility using the **db2** command, and then enter the **list node directory** command again to verify that the new node has been defined correctly. Continuing with our simple example, the new node now appears in the list:

```
db2 => list node directory  
Node Directory  
Number of entries in the directory = 3  
Node 1 entry:  
Node name = GACCTEST  
Comment =  
Directory entry type = LOCAL  
Protocol = TCPIP  
Hostname = merlin  
Service name = 50000  
Node 2 entry:  
Node name = LOCALTEE  
Comment =  
Directory entry type = LOCAL  
Protocol = TCPIP  
Hostname = goose  
Service name = 12344
```

```
Node 3 entry:  
  Node name          = LOCGOOSE  
  Comment            =  
  Directory entry type = LOCAL  
  Protocol           = LOCAL  
  Instance name      = db2inst1
```

9. Configure a database alias named *tee* for the database. In our example we will use a database named SAMPLE (replace this with the name of your database):

```
db2 => catalog database SAMPLE as tee at node localtee  
DB20000I  The CATALOG DATABASE command completed successfully.  
DB21056W  Directory changes may not be effective until the  
directory cache is  
refreshed.
```

10. Enter the **terminate** command to update the directory. (This closes the db2 utility.)

```
db2 => terminate  
DB20000I  The TERMINATE command completed successfully.
```

11. Restart the db2 utility using the **db2** command, and then enter the **list database directory** command to verify that the tee database alias has been defined correctly. Continuing with our simple example, the new database should appear in the list of databases (partial list only shown below):

```
db2 => list database directory  
System Database Directory  
Number of entries in the directory = 6  
Database 1 entry:  
. . .  
Database 3 entry:  
  Database alias          = DN0GOOSE  
  Database name            = SAMPLE  
  Node name                = DN0GOOSE  
  Database release level   = a.00  
  Comment                  =  
  Directory entry type     = Remote  
  Catalog database partition number = -1  
Database 4 entry:  
. . .  
Database 5 entry:
```

```
Database alias          = TEE
Database name           = SAMPLE
Node name               = LOCALTEE
Database release level = a.00
Comment                 =
Directory entry type   = Remote
Catalog database partition number = -1
```

12. Enter the **quit** command to close the db2 utility:

```
db2 => quit
```

Do not log out of the database server system yet. After configuring an S-Tap inspection engine, you will enter one or more SQL commands using the DB2 command-line SQL utility to verify the alias connection.

13. From the administrator portal of the Guardium server defined as the host for the S-Tap, use the S-Tap Configuration panel to define a DB2 inspection engine to listen on the selected Tee listening port, and forward messages to the real database port. Use the values you entered in the table above to set the appropriate value for each property. Be sure to set all other properties required for a DB2 inspection engine, as described elsewhere in this document.
14. Use the DB2 command-line SQL utility, db2sql92, to verify that the database connection through the local tee process works correctly. Log in to the database from the command line using a command like the following (where *db2inst1* is the database user name, *passwd* is the password, and *tee* is the database alias):

```
$ db2sql92 -a db2inst1/passwd -d tee
SQL authorization ID = DB2INST1
Local database alias  = TEE
Running in Embedded Dynamic mode.
-----
```
15. Enter a command that you know will create an SQL exception, and then quit the session. For example:
16. Now modify all client logins to log into the tee alias (instead of the DB2 server).

Prepare for Local Unix Informix Clients to Use the Tee

Do not perform this procedure until S-Tap has been installed on the Informix server, and you are ready to start collecting data. For the local Informix clients to use the Tee, you will create an *staptcp* service name in the /etc/services file, create an *stap_sqlhosts* file, and modify several environment variables such that local Informix clients will connect to the tee listening port instead of to the Informix server.

While completing the steps of this procedure, note the S-Tap configuration file properties you will have to set to start collecting data, in the table below:

S-Tap Configuration File Properties for Informix Inspection Engine

Property=Value	Example Value
db_type=Informix	Informix
tee_listen_port=	12344
real_database_port=	1400

1. Locate the sqlhosts file. The default file name is *sqlhosts*, and by default it is located in the *\$INFORMIXDIR/etc/* directory. The default filename and location can be overridden using the INFORMIXSQLHOSTS environment variable, which when present, defines the full path name for the file.
2. Make a copy of the sqlhosts file and name it *stap_sqlhosts*. You will modify the copy. *Do not modify the original sqlhosts file*. There are no naming requirements for the file. We will use the name *stap_sqlhosts* for the remainder of this section.
3. Open the *stap_sqlhosts* file in a text editor.
4. Locate the entry that local clients use to connect to the database. Each entry contains several positional parameters, in the following format:

dbservername netptype hostname servicename [options]

For example:

jumboinformix onsoctcp jumbo nettcp

The fourth parameter, *nettcp* in the example above, identifies a servicename that maps to a port number for this database server, in the services file.

5. Note the servicename specified: _____
You will use this name later to locate an entry in the services file.
6. Replace the servicename specified with a new servicename for S-Tap. There are no naming requirements. We will use *staptcp* for the remainder of this section. Continuing the example, the entry would be changed as follows:

jumboinformix onsoctcp jumbo staptcp

7. Save the *stap_sqlhosts* file.
8. Locate the services file. By default, it is in the /etc directory, but if Network Information Service (NIS) is used, you must edit the services file on the NIS server.
9. Make a backup copy of the file, and open the original for editing.

10. Locate the entry in the services file for the servicename that you replaced in the `stap_sqlhosts` file. Each entry in this file is in the following format:

```
servicename    port_number/protocol [aliases]
```

In our example services file, the `nettcp` entry is defined as follows:

```
nettcp          1400/tcp
```

In the table above, record the port number (1400, in the example) as the **real_database_port** value.

11. Select an unused port number in the range of 1025-65535 for use by S-Tap. Search the services file for the selected port number, to be certain that it is not used. In the table above, record the port number (12344, for example) as **tee_listen_port** value.
12. Add a line to the *services* file for S-Tap's listening port, `staptcp` in the example:

```
staptcp        12344/tcp
```

13. Save the services file.

14. Set the environment variable `INFORMIXSQLHOSTS`, to specify the full path name for the cloned version of the `sqlhosts` file that you created earlier. For example:

```
setenv INFORMIXSQLHOSTS $INFORMIXDIR/etc/stap_sqlhosts
```

15. When you are ready to start collecting data, from the administrator portal of the Guardium server defined as the host for the S-Tap, use the S-Tap Configuration panel to define an Informix inspection engine to listen on the selected Tee listening port, and forward messages to the real database port. Use the values you entered in the table above to set the appropriate value for each property.
16. You can use the `dbaccess` command to verify that client SQL requests are being seen by S-Tap. To use `dbaccess`, three environment variables must be set appropriately: `INFORMIXDIR`, `INFORMIXSERVER`, and `INFORMIXSQLHOSTS`. Verify that those variables are set correctly using the following command:

```
-bash-3.00# env | grep INFO
INFORMIXDIR=/data/informix
INFORMIXSERVER=jumboinformix
INFORMIXSQLHOSTS=/data/informix/etc/stap_sqlhosts
```

`INFORMIXSERVER` identifies the database server that we are trying to connect to (jumboinformix above).

`INFORMIXSQLHOSTS` identifies the `sqlhosts` file used to resolve connections to jumboinformix. During this resolution it will be either shared memory or a TCP connection. In our previous definition, it is a TCP connection with a

service name of *staptcp*. This will connect to the correct TCP port of 12344 which is resolved in the /etc/services file.

17. Enter the dbaccess command: **dbaccess**
18. Navigate to **connection – connect – Select Database Server** and select **jumboinformix**.
19. When prompted, enter an appropriate database user name and password.
20. Exit from the connection portion of the configuration and select **Query-language – select database**.
21. Select New and enter an SQL command, for example:
`select * from my_mistake`
22. Log in to a user portal on the Guardium server, and navigate to the Reports & Alerts – Report Templates – Exceptions tab, and select the SQL Errors report. You should be able to locate your SQL error near the “top” of the report, and thus verify that the tee is seeing the Informix traffic.

Prepare for Local Unix Oracle Clients to Use the Tee

Follow the procedure outlined below to modify the *tnsnames.ora* file, which maps service aliases to ports. Do not change this file until S-Tap has been installed and you are ready to start collecting data. Note the S-Tap configuration file properties you will have to set to start collecting data, in the table below.

S-Tap Configuration File Properties for Oracle Inspection Engine

Property=Value	Example Value
db_type=Oracle	Oracle
tee_listen_port=	12344
real_database_port=	1521

1. Make a backup copy of the *tnsnames.ora* file, which is located in the \$ORACLE_HOME/network/admin directory.
2. Open the *tnsnames.ora* file for editing in a text editor program.
3. Locate the entry in this file for the service alias used to access the database. An entry named **EAGLE10** on the **EAGLE** host is illustrated below:

```
EAGLE10 =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = eagle)(PORT = 1521)))
```

```
(CONNECT_DATA = (SERVICE_NAME = GUARD10))
)
```

4. In the table above, record the port number (1521, in the example) as the **real_database_port** value.
Do not change the above entry until you have verified that S-Tap is configured correctly.
5. Select an unused port number in the range of 1025-65535 for use by S-Tap. Search the file for the selected port number, to be certain that it is not used. In the table above, record the port number (12344, for example) as **tee_listen_port** value.
6. Create a duplicate entry for the service by copying and pasting the entry, and making the highlighted changes:

```
LOCALTEE =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = eagle)(PORT = 12344)))
(CONNECT_DATA = (SERVICE_NAME = GUARD10))
)
```

7. Save the tnsnames.ora file.
8. When you are ready to start collecting data, from the administrator portal of the Guardium server defined as the host for the S-Tap, use the S-Tap Configuration panel to define an Oracle inspection engine to listen on the selected Tee listening port, and forward messages to the real database port. Use the values you entered in the table above to set the appropriate value for each property.
9. Log on to the database server locally, using **sqlplus** to verify that S-Tap is configured properly and will see a local access. For example:

```
# sqlplus scott/tiger@LOCALTEE
```

Where **scott** is the database user name, **tiger** is the password, and **LOCALTEE** identifies the service.

10. Enter an invalid SQL command to create an SQL exception that will be easy to find. For example:

```
select * from my_mistake
```
11. Log in to a user portal on the Guardium server, and navigate to the Reports & Alerts – Report Templates – Exceptions tab, and select the SQL Errors report. You should be able to locate your SQL error near the “top” of the report, and thus verify that the tee is seeing the local Oracle traffic.

12. Reopen the tnsnames.ora file and replace the database service port number with the selected number. Continuing our example:

```
EAGLE10 =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = eagle)(PORT = 12344)))
    (CONNECT_DATA = (SERVICE_NAME = GUARD10))
  )
```

13. Save the tnsnames.ora file. All local clients connecting to EAGLE10 will now connect to port 12344 (the Tee listening port) instead of the actual database port.

Prepare for Local Unix Sybase Clients to Use the Tee

Follow the procedure outlined below to modify local *interfaces* file, which maps servers to ports. Do not change this file until S-Tap has been installed and you are ready to start collecting data. Note the S-Tap configuration file properties you will have to set to start collecting data, in the table below.

S-Tap Configuration File Properties for Sybase Inspection Engine

Property=Value	Example Value
db_type=Sybase	Sybase
tee_listen_port=	12344
real_database_port=	4100

1. Make a backup copy of the interfaces file, which is located in the \$SYBASE/directory.
2. Open the interfaces file for editing in a text editor program.
3. Locate the entry in this file whose name matches the Sybase server name. For example, a server named *parrot* might be defined as follows:

```
parrot
  master tcp ether parrot 4100
  query tcp ether PARROT 4100
```
4. In the table above, record the port number (4100, in the example) as the **real_database_port** value.
5. Select an unused port number in the range of 1025-65535 for use by S-Tap. Search the file for the selected port number, to be certain that it is not used. In the table above, record the port number (12344, for example) as **tee_listen_port** value.

6. Replace the port number with the selected number. For example:

```
parrot
    master tcp ether parrot 12344
    query tcp ether PARROT 12344
```

7. Save the interfaces file.
8. From GUI of the SQL Guard host defined for the S-Tap, use the S-Tap Configuration panel to define a Sybase inspection engine to listen on the selected Tee listening port, and forward messages to the real database port.

Installing CAS

Follow the procedures outlined below to install and configure CAS on database server systems. CAS monitoring and reporting topics are covered in the SQL Guard User Guide.

CAS Installation and Configuration Overview

CAS is a client-server application. You install a *CAS client* on a *CAS host*, which is usually a database server system, using the same utility that is used to install S-Tap. CAS shares configuration information with S-Tap, though each component runs independently of the other. Once the CAS client has been installed on the host, you configure the actual change auditing functions from the SQL Guard server portal.

Start Up and Failover Overview

When the CAS client starts on the host, it looks for a checkpoint file that it may have written to the system. This file tells CAS what it was doing the last time it was running. CAS then connects to its SQL Guard server. If it has found a checkpoint file, CAS will ask the SQL Guard server to verify its version of its monitoring assignment against what is stored in the SQL Guard database. While the CAS client and the SQL Guard server have been disconnected, there may have been changes to the assignment. When any differences are resolved, CAS will resume monitoring. If CAS does not find a checkpoint file, it will ask the SQL Guard server what it should do. If the SQL Guard server finds the CAS host in its database, then the associated template sets will be sent to the CAS client, expanded into monitored items, and monitoring will begin. If the SQL Guard server cannot find the CAS host in its database, it will add it to the database and send the default template set for the CAS host operating system.

If the CAS client loses its connection to the SQL Guard server or cannot make an initial connection, it opens a failover file and begins writing the messages that it would have sent to the server, to the failover file. When it reconnects, the CAS client shuts down and restarts, sending all messages stored in the failover file to the SQL Guard server, and deleting the file. If the CAS client was unable to make the initial connection, it will use the checkpoint file to determine what to monitor; otherwise it continues doing what it was

doing before communication failed. While it is in failover mode, CAS periodically checks to see if it can reconnect with the server. The number of times CAS will attempt to reconnect, and the average time interval between reconnect attempts are configurable.

If the reconnect attempt limit is met, the CAS client stops trying to reconnect, but continues to write data to a failover file. To cap disk space requirements on the database server, there are actually two failover files. CAS writes to one file until it reaches its maximum failover file size (which is configurable), and then switches to the other, overwriting any previous data on that file. The default failover file size is 50MB (for each of the files).

You can specify one or more secondary SQL Guard servers when configuring the CAS client. In failover mode, CAS only tries to reconnect to its primary server until the size of the failover file reaches a pre-configured limit. At that time, CAS begins trying to connect to any of the secondary servers, as well as its primary server (which is always the first server it tries to connect with during any reconnect attempt). While it is connected to a secondary server, CAS continues to try to reconnect to its primary server.

If secondary servers are used, you must be sure that the CAS configuration information on all secondary servers matches the CAS configuration defined on the primary server. This can be done by exporting all CAS definitions for a CAS host from the primary server, and importing those definitions on all secondary servers for that CAS host. See [Setting Up and Maintaining Secondary Servers](#) for more information.

As with S-Tap, CAS connectivity outages create exceptions on the SQL Guard server, so alerts can be issued within moments of detecting the outage.

Installing CAS on Database Servers

You install CAS on a database server system using the same script (for Unix) or wizard (for Windows) that you use to install S-Tap. You will need to provide the following information:

- The IP address of the database server system.
- The IP address of the SQL Guard server for CAS and S-Tap.
- For Unix servers only, the interface over which CAS will listen for SQL Guard server communications. Typical interface names, by operating system, are as follows:
 - AIX: Ew0
 - HPUX: 1An0
 - Solaris: hme0
 - Linux: eth1

- For Unix servers only, the directory location for version 1.4.2 or later of Java. In the unlikely event that it is not already installed, you will need to install Java version 1.4.2 or later before installing CAS. During the CAS installation, you will be prompted to supply the Java directory name. For example, if Java is located in the following directory: /usr/bin/jre1.4.2_12/bin/java, you would enter: /usr/bin/jre1.4.2_12/bin

Note: On Windows database servers, if Java has not been installed, it will be installed automatically.

Once it has been installed on the database server system, CAS installation parameters can be modified from the host SQL Guard server portal (see [Modifying CAS Options from the Administration Console](#)).

Once it has been installed, CAS will start trying to connect to SQL Guard. If it connects, it will run with a default configuration for the host operating system.

Before installing CAS, see the following topic on disk space requirements.

CAS Disk Space Requirements

The following table shows the amount of space required to install CAS. These numbers assume include the installation of Java version 1.4.202. If you have this version (or a later version) of Java installed, the space requirements will be less.

S-Tap and CAS Disk Space Requirements by Operating System (MB)

OS	HP-UX	Solaris	Linux	AIX	OSF1	Win32
CAS (including Java)	630	390	405	309	309	277

The following table describes files whose size may vary. As noted in the descriptions, all files are stored in one of three subdirectories (/bin, /conf, or /logs) of the following directories, for each operating system type:

- Unix:** /usr/local/guardium/guard_stap/cas
- Windows:** C:\Program Files\Guardium\GUARDIUM_STAP\cas

S-Tap and CAS File Descriptions

File(s)	Description
cas.log	This file in the logs subdirectory will remain very small under normal conditions. If Guardium Customer Service directs you to enable more detailed logging, this file may become very large, very quickly (for example, up to 100MB in a couple of days during

File(s)	Description
	testing). Each day a new file is created, with the old file renamed to indicate the date of the data it contains (cas.log.date).
fail_over_file and fail_over_file2	These files are stored in the bin subdirectory. If the SQL Guard server to which CAS sends data is unavailable, CAS writes that data to a failover file. That file can grow to a maximum size (50MB, by default), at which point it will be closed and a second will be created. If the second file reaches the maximum, it will be closed and the first file will be overwritten. The maximum failover file size can be set from the S-Tap configuration panel.

Activating CAS on SQL Guard Servers

For new SQL Guard servers beginning with version 6.0, if the CAS option has been purchased, Guardium activates the feature on the system before shipping. If your are adding the CAS feature to an existing system, or if a system has been rebuilt from scratch following a hardware or operating system failure or upgrade, you will need to enable the feature as described below, by performing one or both of the following steps.

Note: CAS must be activated on *any* SQL Guard Server that may act as a secondary (failover) SQL Guard server for CAS.

Step 1: Enable the Local Taps Menu

Perform this step only if there is no Local Taps section on the Administration Console menu of the Administrator portal.

1. From an SSH client command line window, log in to the SQL Guard server, as the CLI user.
2. Enter the following two commands:
`store unit type stap`
`restart inspection-core`
3. Enter the **quit** command to log out of the CLI session.
4. If you are still logged in to the SQL Guard administrator portal, log out (otherwise, you will still have the old configuration settings).
5. Log in to the SQL Guard administrator portal. The Local Taps section will now be present in the Administration Console.

Step 2: Install a New System License Key

Perform this step only if Guardium has provided you with a new license key to allow CAS functionality.

1. From an SSH client, log in to the SQL Guard server, as the CLI user.
2. If you are replacing an old license key with a new one, save a copy of the old license key as follows (otherwise, skip this step):
 - o Enter the **show license** command to display the current license key. The command and its output should look something like this:

```
suppl.guardium.com> show license
Host MAC: 00:C0:9F:41:CE:16
License:
hrKXsOL0i2Ri1+RMF1/uUD1exdbqxWQ9cEnwCBBloj2vv3hmzb7LKgy4jXH45n2
o8Qmw6s5e7PMLLb7GIRNE+8GKZkR/JQSB26xBvgpYyS5GtX/0mdoZNzXxy7z85P
uQH0EEHB0eGGcChO0MCDHAuq+YQXmrw6ReICR7kVTaHsg=
ok
```
 - o The actual license key begins following the colon after the word License, and always ends with (and includes) an equal sign. Copy and paste the old license key to a text file, and save the text file in a safe location.
3. Enter the **store license console** command. You will be prompted to paste the new license key:

```
suppl.guardium.com> store license console
Please paste the string received from Guardium Inc, press Enter.
```
4. Copy and paste the new license key at the cursor location, and then press Enter. Remember that the license key *includes* the trailing equal sign. A series of messages will display, ending with:

```
.... We recommend that the machine be rebooted at the earliest
opportunity in order to complete the license updating process.
ok
suppl.guardium.com>
```
5. Enter the **restart system** command to reboot the SQL Guard server:

```
suppl.guardium.com> restart system
Restarting system
ok
suppl.guardium.com>
```

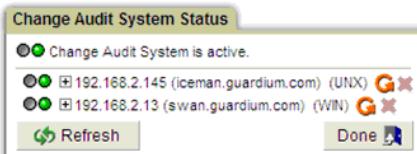
Your SSH session will be terminated as the system reboots.
6. Allow several minutes for the unit to reboot, and then log in to the SQL Guard Administrator portal.

Modifying CAS Options from the Administration Console

To change CAS installation options on the CAS host:

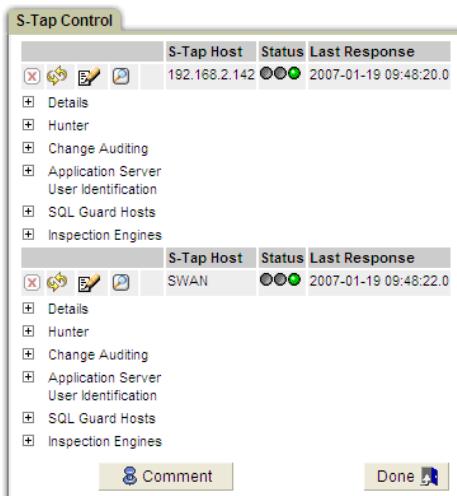
1. Log into the SQL Guard server administrator portal. The SQL Guard server you log into *must* be the *active host* for CAS on the database server.
2. First verify that CAS is active on the CAS host, as follows:

- From the Local Taps section of the Administration Console, select CAS Status to open the Change Audit System Status panel:



If the database server on which CAS was installed is *not* listed:

- You may be logged into the wrong SQL Guard server (i.e. not the one from which you had intended to manage CAS on this database server.)
 - The CAS configuration defined at installation time on the database server may specify the incorrect IP address for this SQL Guard server.
 - The CAS configuration defined at installation time on the database server may specify the incorrect IP address for the database server.
 - The communications path between the database server and the SQL Guard server may not be available.
 - The CAS executable may not have been started or may have been stopped on the database server.
 - If the CAS host *is* listed, verify that the associated green status light is illuminated. The green light is the rightmost of the two lights. If the green light is not illuminated, either CAS is not running on the database server, or communication with the database server has been lost.
3. S-Tap and CAS share configuration information on the database server, so you edit both configurations from the S-Tap Control panel. To open that panel, select **S-Tap Control** from the **Local Taps** section of the **Administration Console**.



4. Locate the CAS host in the S-Tap Host column. It may display as an IP address or as a DNS name. On this panel, you can ignore the status lights, since those apply to the S-Tap status on that host.
5. To the left of the Change Auditing label for the host you want to configure, click the (Show Change Auditing Configuration) button. For example, for the first host in the example above:

	S-Tap Host	Status	Last Response
	192.168.2.142		2007-01-23 09:34:39.0
Details			
Hunter			
Change Auditing	Task Baseline	task_baseline	
	Task Checkpoint	task_checkpoint	
	Client Baseline	client_baseline	
	Client Checkpoint	client_checkpoint	
	Checkpoint Period	60	
	Fail Over File	fail_over_file	
	Fail Over File Size Limit	50000	
	Max Reconnect Attempts	5000	
	Reconnect Interval	60	
	Raw Data Limit	1000	
	Md5 Size Limit	1000	
<hr/>			
Application Server			
User Identification			

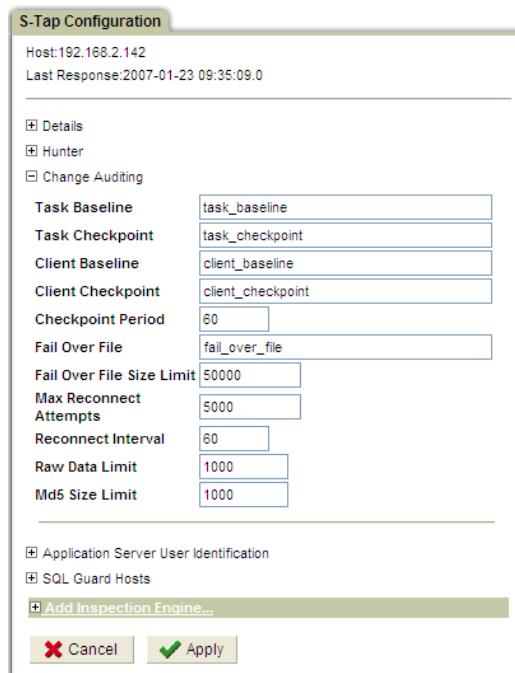
6. To edit the CAS configuration, click the (Edit S-Tap Configuration) button.

Note: If the Edit S-Tap Configuration button is disabled, check the Status light: if is red, the CAS host currently is not reachable;

if it is any other color, this SQL Guard server is not the active host for the S-Tap (and only the *active* host can modify the configuration).



7. To the left of the Change Auditing label, click the + (Show Change Auditing Configuration) button to expand the CAS configuration section of the S-Tap Configuration panel:



Note: When editing the Change Auditing pane, you have the opportunity to change several file names. *We strongly recommend that you keep the default names for these files.*

8. Referring to the table below, you can modify any of the CAS configuration parameters (but as noted above, we suggest that you do not change any of the default file names).

Change Auditing Pane Parameters

Parameter / Default	Description
Task Baseline task_baseline*	Reserved for future use.
Task Checkpoint task_checkpoint*	There will be a series of files beginning with this name, and ending with a uniqueness number. These files are used for restart processing.
Client Baseline client_baseline*	Reserved for future use.
Client Checkpoint client_checkpoint*	There will be a series of files beginning with this name, and ending with a uniqueness number. These files are used for restart processing.
Checkpoint Period 60 (sec)	Maximum number of seconds between checkpoints (60).
Failover File fail_over_file*	Name of the file to which data is written when the SQL Guard server cannot be reached. During this time, the file may grow to the maximum size specified (see Failover File Size Limit, below). When the limit is reached, a second file will be created, using the same name with the digit 2 appended to the end of the name. (This is the point at which CAS begins trying to connect to a secondary server.) If that file also reaches the maximum size, the first file will be overwritten, and if the first file fills again, the second file will be overwritten. Thus following an extended outage, you may lose data, but you will have an amount of data up to twice the Failover File Size Limit (see below).
Failover File Size Limit 50000 (KB)	Failover file maximum size, in KB (the default is 50 MB). There are two of these files, so the disk space requirement will be twice what you specify here. If you specify -1, there will be no limit on the file size.
Max Reconnect Attempts 5000	After losing a connection to the SQL Guard server, the maximum number of times CAS will attempt to reconnect. Set

Parameter / Default	Description
	this value to -1 to remove any maximum (CAS will attempt to reconnect indefinitely). The default is 5000 times, so using the default reconnect interval (see below), this is about 3.5 days. After the maximum has been met, CAS will continue to run, writing to the failover files, as described above, but it will not attempt to reconnect with a host.
Reconnect Interval 60 (seconds)	Number of seconds between reconnect attempts (60). See the description of the reconnection process, above.
Raw Data Limit 1000 (KB)	Maximum number of kilobytes written for an item when the <i>Keep data</i> checkbox is marked in the item template (1000). If you specify -1, there will be no limit.
Md5 Size Limit 1000 (KB)	Maximum size of a data item, above which the MD5 checksum calculation will not be performed (1000). If you specify -1, there will be no limit.

* Windows default file names are converted to all uppercase characters.

- When you are done making changes, click the Apply button at the bottom of the panel. When you click the Apply button (after making any changes):
 - The SQL Guard sever saves the new configuration and sends a copy to CAS on the database server.
 - CAS updates and begins using its new configuration.
 - CAS notifies the SQL Guard server that the new configuration is in use.

Note: Modifying the Change Auditing pane of the S-Tap Configuration panel *does not* affect the status of S-Tap, so you will not see any change in the status lights on the S-Tap Control panel. In contrast, when you apply updates to *any other pane* of the S-Tap Configuration panel, the S-Tap Control panel lights change – typically from green to red to yellow and back to green again – as the S-Tap receives, stops, and then restarts with the updated configuration.

Setting Up and Maintaining Secondary Servers

In the S-Tap/CAS configuration file on the database server system, one or more secondary SQL Guard servers can be defined. If the primary SQL Guard server becomes unavailable, CAS on that database server system will connect to a secondary SQL Guard server (as described previously, see [Start Up and Failover Overview](#)). If secondary servers are used, you must be sure that the CAS configuration information on all

secondary servers matches the CAS configuration defined on the primary server. This can be done from an administrator portal, by exporting all CAS definitions for a CAS host from the primary server, and then importing those definitions on all secondary servers for that CAS host.

Exporting CAS Hosts

1. From the Administration Console tab, under SQL Guard Definitions, click Export to open the Definitions Export panel.
2. Under Type, select CAS Hosts. A list of the CAS Hosts defined on this system will be displayed.
3. Select each CAS Host to be exported.
4. Click the Export button. A file named **exp_<date>_<time>.sql** will be saved on your system. This file will contain the definitions of all CAS hosts selected, and the definitions of any template sets used by those CAS hosts.

Importing CAS Hosts

1. From the Administration Console tab, under SQL Guard Definitions, click Import to open the Definitions Import panel.
2. Enter the name of the file containing the exported definitions or click the Browse button to select that file.
3. Click the Upload button. You are notified when the operation completes and the CAS host definitions contained in the file will be displayed.
 - o Click (Import this set of Definitions) to import the definitions.
 - o Click (Remove this set of Definitions without Importing) to remove the uploaded file without importing the definitions.

You are prompted to confirm either action.

Note: An import operation does not overwrite an existing definition.

If you attempt to import a definition with the same name as an existing definition, you are notified that the item was *not* replaced. If you *want* to overwrite an existing definition with an imported one, you must delete the existing definition before performing the import operation.

4. Click the Done button to close the panel when you have finished importing or removing all uploaded files.

Maintaining Secondary Servers for a CAS Host

After updating a CAS configuration on the primary server, you must update that configuration on all secondary servers. Since the import operation will not replace an existing definition, on each secondary server you must delete the old CAS host definition before importing the new one, as explained below.

Be sure to perform this procedure only while the selected CAS host is connected to its primary server.

1. **On the primary server**, from the SQL Guard administrator portal, export the definition of the CAS host (see Exporting CAS Hosts, above).
2. **On each secondary server**, from the SQL Guard administrator portal:
 - o Delete the old CAS host definition that you want to replace.
 - o Import the definitions that were exported from the primary server (see Importing CAS Hosts, above).

Chapter 2: System Management

Web Browser Considerations

Microsoft Internet Explorer is the only Web browser officially supported.

SVG Viewer Plugin

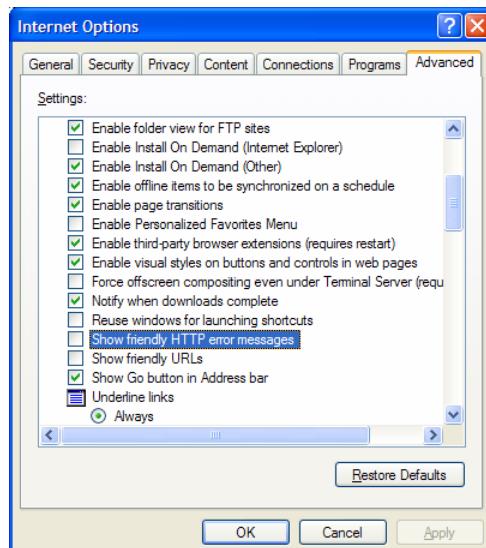
To view the Current Status Manager or management maps, you need the Adobe SVG Viewer plugin. See [Software Downloads from Adobe](#).

Caution: Disable “friendly” HTTP error messages

Microsoft Internet Explorer is typically configured by default to show “friendly HTTP error messages.” This means that instead of seeing any HTTP error messages produced by the SQL Guard management interface, you will receive friendly (but not very informative) messages output by Microsoft Internet Explorer.

To check or modify this setting:

1. Open Internet Explorer.
2. Select Internet Options on the Tools menu to open the Internet Options panel.
3. Click on the Advanced tab.
4. Scroll down in the Settings box to the *Show friendly HTTP error messages* setting. If it is checked, click on the checkbox to clear it.
5. If you cleared the checkbox, click the Apply button to save the setting.
6. Click the OK button to close the Internet Options panel.



Logging in to SQL Guard

To log in to the SQL Guard system:

1. Open the SQL Guard Login Page in your browser. If you have not saved the page address, type it in the Address box of your browser window, in the following format: **https://system:port**

Note: This address begins with the letters *https* (the *secure* hypertext transmission protocol, not the more common *http* protocol).

Substitute for the *system* and *port* components of the address as appropriate, using your installation values, where *system* is either the IP address or the network name for the system and *port* is the port assigned for the SQL Guard management interface. For example:

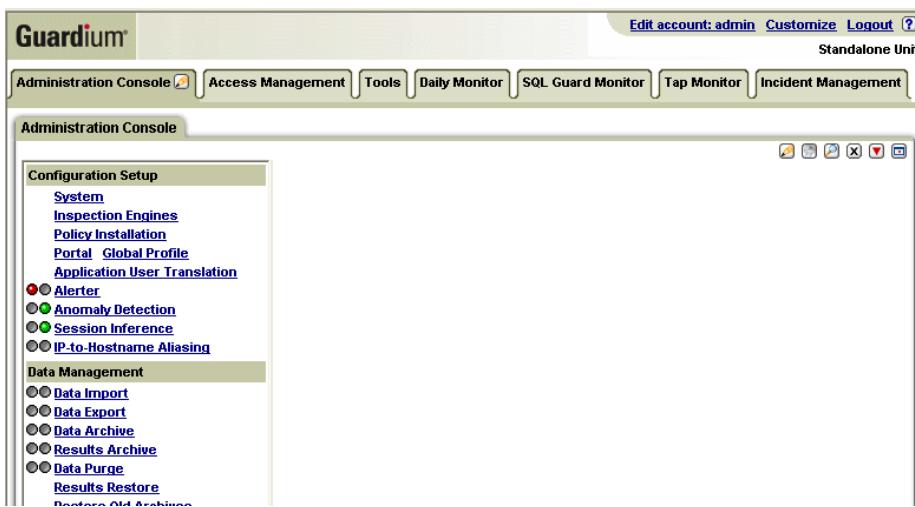
`https://192.168.3.47:8443`

`https://guard02:8443`



After typing the URL in the address box, press the Enter key to open SQL Guard Login Page:

2. Enter your assigned administrator user name and password, and click the Login button. (Note that each user can log into a SQL Guard unit from one IP address only. If the same user attempts to log in concurrently from two IP addresses, the second attempt will not be allowed.) This opens the SQL Guard Administrator portal, which by default contains the tabs illustrated below:



The Administration Console tab contains the bulk of the commands that are reserved exclusively for administrators, in a menu on the left-hand side of the panel (see above). Your default portal may contain a different set of tabs and menu items, depending on the options purchased and installed on the system you have logged into.

Changing Your Password

SQL Guard comes preconfigured with an initial *admin* password that you may want to change. To change your password:

1. Click the *Edit account* link in the upper, right-hand corner of the window to open the *Edit your account details* screen:

Edit your account details

Username:	admin
Old Password:	<input type="password"/>
Password:	<input type="password"/>
Password (confirm):	<input type="password"/>
First Name:	admin
Last Name:	admin
Email:	admin@guardium.com

In an effort to provide the highest level security, new passwords must be 8 or more characters in length and must include at least one uppercase letter, lowercase letter, digit, and special character. A special character is considered any of the following: @#\$%^&.:;+=_-

2. Enter your old password in Old Password box.
3. Enter your new password in the Password box. Then retype the new password in the Password (confirm) box. Unless you have disabled password validation (using the store password validation cli command).

When password validation is enabled, the password must be eight or more characters in length, and must include at least one uppercase alphabetic character (A-Z), one lowercase alphabetic character (a-z), one digit (0-9), and one special character from the table below.

Table of Special Characters for SQL Guard Passwords

Special Character	Character Name
@	Commercial at
#	Number sign
\$	Dollar sign
%	Percent sign
^	Circumflex accent (carat)
&	Ampersand
.	Full stop (Period)
;	Semicolon
!	Exclamation mark
-	Hyphen (minus)
+	Plus sign

=	Equals sign
_	Low line (underscore)

4. Enter your first and last names in the First Name and Last Name boxes.
5. Enter an email address in the Email box.
6. Click the Update Account button when you are done (or click the Cancel button to cancel the operation).

About the System Shared Secret

The SQL Guard administrator defines the System Shared Secret on the System Configuration panel, which is described in the following section. The system shared secret is used for archive/restore operations, and for Central Management and Aggregation operations. When used, its value must be the same for all units that will communicate. This value is null at installation time, and can change over time.

The system shared secret is used:

- When secure connections are being established between a Central Manager and a managed unit.
- When an aggregated unit signs and encrypts data for export to the aggregator.
- When any unit signs and encrypts data for archiving.
- When an aggregator imports data from an aggregated unit.
- When any unit restores archived data.

Depending on your company's security practices, you may be required to change the system shared secret from time to time. Because the shared secret can change, each system maintains a *shared secret keys* file, containing an historical record of all shared secrets defined on that system. This allows an exported (or archived) file from a system with an older shared secret to be imported (or restored) by a system on which that same shared secret has been replaced with a newer one.

Changing the System Configuration

To change system configuration information:

- Click System in the Configuration Setup section of the Administration Console menu to open the *System configuration* panel:

The screenshot shows the 'System Configuration' panel with the following fields:

- Unique Global Identifier:** 598182152
- System Shared Secret:** (empty)
- Retype Secret:** (empty)
- License Key:** (empty)
- Network Address:**
 - System Hostname:** g1
 - Domain:** guardium.com
 - System IP Address:** 192.168.3.101
 - SubNet Mask:** 255.255.255.0
 - Hardware (MAC) Address:** 00:04:23:A7:89:08
- Secondary Management Interface:** (empty)
- Routing:**
 - Default Route:** 192.168.3.1
- DNS:**
 - Primary Resolver:** 192.168.3.20 | [Test Connection](#)
 - Secondary Resolver:** (empty) | [Test Connection](#)
 - Tertiary Resolver:** (empty) | [Test Connection](#)

Buttons at the bottom include: **Stop**, **Restart**, and **Apply**.

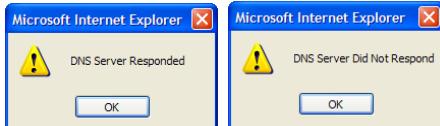
- Make any desired changes, referring to the *System Configuration Panel Reference* table below.
- Click the Apply button to save the updated system configuration when you are done making changes.

Note: The applied changes *do not* take effect until the unit is restarted. After applying configuration changes, click the Restart button to stop and restart the system (using the new configuration settings).

System Configuration Panel Reference

Field or Control	Description
Unique Global Identifier	This value is used for collation and aggregation of data. The default value is a unique value derived from the MAC address of the machine. It is strongly recommended that you do not change this value after the system begins monitoring operations.

Field or Control	Description
System Shared Secret	See About the System Shared Secret , above. <div style="border: 1px solid black; padding: 5px;"><i>Caution: When used, be sure to save the shared secret value in a safe location. If you lose the value, you will not be able to access archived data.</i></div>
Retype Secret	Any value you enter here does not display. Each character you type displays as an asterisk.
License Key	When entering or changing the system shared secret (see above), retype the new value a second time. This non-editable value is inserted in the configuration during installation. Do not modify this field unless you are instructed to do so by Guardium Support. If you install a new license key on a managed unit, when you click the Apply button you will receive a warning message that reads: "Warning: changing the license on a Central Management Unit requires refreshing all managed units." After you click OK to close the message window, you must click Apply a second time to install the new license key. You will know that the new license has been installed when you receive the message: "Data successfully saved."
System Hostname	The resolvable host name for the SQL Guard system. This name must match the DNS host name for the primary System IP Address (see below).
Domain	The name of the DNS domain on which the SQL Guard system resides.
System IP Address	The IP address for the network interface labeled ETH 0 on the back of the unit. Users and S-Taps connect to the SQL Guard server using the primary or optional secondary IP address.
SubNet Mask	The subnet mask for the primary System IP Address (above).
Hardware (MAC) Address	The MAC address for the primary network interface (above).
System IP Address (Secondary)	Optional. The IP address for the highest numbered network interface on the back of the unit.
SubNet Mask (Secondary)	Optional. The subnet mask for the secondary System IP Address (above).
Default Route	The IP address of the default router for the system.

Field or Control	Description
Primary Resolver	The IP address of from one to three DNS servers. You must specify the Primary Resolver. The others are optional.
Secondary Resolver	
Tertiary Resolver	
Test Connection	Click the Test Connection link to test the connection to the corresponding DNS server. This only tests that there is access to port 53 (DNS) on the specified host. It does <i>not</i> verify that this is a working DNS server. After a brief delay, you will receive one of the following notifications:
	 The image shows two side-by-side Microsoft Internet Explorer dialog boxes. Both have a yellow exclamation mark icon in the top-left corner. The left dialog box has the title "Microsoft Internet Explorer" and the message "DNS Server Responded" below it. The right dialog box also has the title "Microsoft Internet Explorer" and the message "DNS Server Did Not Respond" below it. Both dialogs have an "OK" button at the bottom.
Stop	Click the Stop button to shut the system down.
Restart	Click the Restart button to stop and then restart the system.
Apply	Click the Apply button to save the changes.

Working with Inspection Engines

An inspection engine monitors the traffic between a set of one or more servers and a set of one or more clients using a specific database protocol (Oracle or Sybase, for example). The inspection engine extracts SQL from network packets; compiles parse trees that identify sentences, requests, commands, objects, and fields; and logs detailed information about that traffic to an internal database.

You can configure and start or stop multiple inspection engines on the SQL Guard server.

Inspection engines cannot be defined or run on a Central Manager unit. However, you can start and stop inspection engines on *managed* units from the Central Manager control panel. See [Central Management](#) later in this chapter for more information.

Note: If this SQL Guard unit serves as an S-Tap host, be sure that it does not monitor the same traffic as the S-Tap. If that happens, the SQL Guard unit receives duplicate packets, it is unable to reconstruct messages, and that traffic is ignored.

About Kerberos-Encrypted Database User Names

This topic applies to MS SQL Server database access in Windows environments only. In these environments, database user names may be encrypted via Kerberos, in which case they appear in network traffic as strings of hexadecimal characters.

The SQL Guard server can be configured to automatically replace Kerberos-encrypted database user names with real database user names. It does this by monitoring Kerberos traffic, capturing and pairing real and encrypted names, and making the appropriate substitutions from that point forward. (SQL Guard does not actually decrypt the encrypted strings.)

Even without enabling the replacement of Kerberos-encrypted database user names, you may sometimes see real database user names in SQL Guard reports, when you know the names are Kerberos-encrypted. To understand how this happens, keep in mind that for a single session (a specific client-server connection from login to logout) SQL Guard is accumulating information over time. All of the information SQL Guard would like to have for that session is never contained in a single message. So over the life of a session, SQL Guard is constantly trying to locate missing items. If it encounters a real database user name during the course of the session, and all it has to that point is a Kerberos string, SQL Guard will replace the string with the real name.

There are two alternatives for handling Kerberos-encrypted names. The key to both approaches is that all Kerberos traffic must be sent to the SQL Guard server. The first approach is to SPAN or mirror all Kerberos traffic to the SQL Guard server, and the second approach is to install an S-Tap on the Kerberos domain controller, and configure the S-Tap

to send all Kerberos traffic to the SQL Guard server. Each approach is described separately, below.

SPAN Kerberos Traffic

To enable the automatic decoding of all Kerberos-encrypted database user names using SPAN-based Kerberos identification, you must perform the following setup:

1. Check with a network administrator to be certain that all traffic *from and to* the Windows domain controller on which the Kerberos ticket-granting authority runs will be *seen* by the SQL Guard server. This may require the reconfiguration of a SPAN or mirror port, or a network TAP from which SQL Guard views traffic.
2. Log in to the SQL Guard Server CLI, as the *cli* user.
3. Enter the following commands:

```
store local-stap on  
store unit type stap  
restart inspection-core  
restart inspection-engines
```

For more information about using these commands or the *cli*, see Chapter 6.

4. After entering the *cli* commands, check to make sure that the Local Taps section appears in the Administration Console menu.
5. Click S-Tap Control in that menu to open the S-Tap Control panel. An S-Tap Host entry for IP address 127.0.0.1 should appear.

Forward Kerberos Traffic Using an S-Tap

As an alternative to SPAN'ing Kerberos traffic to the SQL Guard server, you can install an S-Tap on the Kerberos ticket granting domain controller, and configure an inspection engine on that S-Tap to forward all Kerberos traffic. If multiple SQL Guard servers are defined for that S-Tap, the Kerberos traffic will be sent to *all* of them. For detailed information about installing S-Taps and configuring an S-Tap inspection engine for Kerberos traffic, see the following topics:

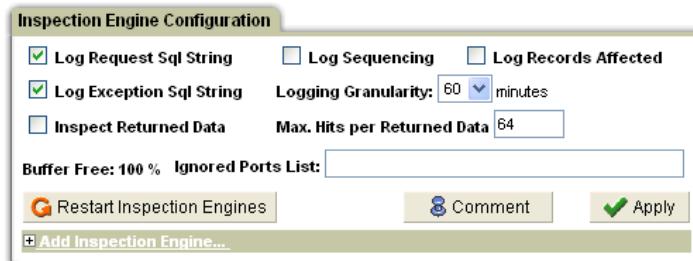
- [Installing S-Tap](#) in Chapter 1.
- [Adding or Modifying S-Tap Inspection Engines](#) later in this chapter.

Opening the Inspection Engine Configuration Panel

To work with inspection engine configurations:

1. Open the Administration Console panel (not shown).

2. In the Configuration Setup section of the Administration Console menu, click Inspection Engines to open the *Inspection Engine Configuration* panel:



The top pane of this panel contains information that applies to all inspection engines. Each pane below that describes a separate inspection engine (there are none defined in the example above). The order in which the inspection engines are listed determines when the various filtering mechanisms of each engine are applied. Once an inspection engine has been configured, you can move it up or down in the list of inspection engines.

After a discussion of how IP addresses are specified, the following sections describe how to define and manage inspection engines.

Selecting IP Addresses

Each inspection engine monitors traffic between one or more *client* and *server* addresses. You specify these addresses using an *IP address* and a *mask*. You can think of an *IP address* as a single location and a *mask* as a wild-card mechanism that allows you to define a range of IP addresses.

IP addresses have the format: **n.n.n.n**, where each **n** is an eight-bit number (called an *octet*) in the range 0-255.

For example, an IP address for your PC might be: 192.168.1.3. This address is used in the examples below. Since these are binary numbers, the last octet (3) can be represented as: 00000011.

The mask is specified in the same format as the IP address: **n.n.n.n**. However, a zero in any **bit** position of the mask serves as a wildcard. Thus, the mask 255.255.255.240 specifies all values from 0-15 in the last octet, since the value 240 in binary is 11110000; but only the values 192.168.1 in the first three octets. (Since 255 is 11111111 in binary – all 1s – no wildcarding is done for the first three octets).

Specifying binary masks can be a little confusing. However, for the sake of convenience, IP addresses are usually grouped in a hierarchical fashion, with all of the addresses in one category (desktop computers, for example) grouped together in one of the last two octets.

Therefore, in practice, the numbers you see most often in masks are either 255 (no wildcarding) or 0 (all).

Thus a mask 255.255.255.255 (which has no zero bits) identifies only the single address specified by *IP address* (192.168.1.3 in the example above).

Alternately, the mask 255.255.255.0, combined with the same IP address identifies all IP addresses beginning with 192.168.1.

Selecting All Addresses

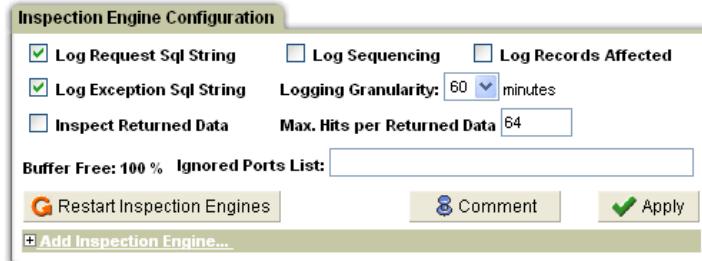
The IP address 0.0.0.0, which is sometimes used to indicate all IP addresses, is not allowed by SQL Guard. To select all IP addresses when using an IP address/mask combination, use any non-zero IP address followed by a mask containing all zeroes (e.g. 1.1.1.1/0.0.0.0).

Changing Settings that Apply to all Engines

The top pane of the *Inspection Engine Configuration* panel contains settings that apply to all inspection engines configured on the unit.

To view or change settings that apply to all inspection engines:

1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Inspection Engines to open the *Inspection Engine Configuration* panel.



See the following table for a description of the fields in the top pane. These fields apply to all inspection engines configured.

Inspection Engine Configuration Panel – Top Pane

Field or Control	Description
Log Request Sql String	If marked, each SQL request statement is logged in its sanitized format. Otherwise, no statements are logged.

Field or Control	Description
Log Sequencing	If marked, a record is made of the immediately previous SQL statement, as well as the current SQL statement, provided that the previous construct occurs within a short enough time period.
Log Records Affected	THIS IS A TECHNOLOGY PREVIEW – CONTACT GUARDIUM SUPPORT BEFORE USING THIS FEATURE.
Log Exception Sql String	If marked, when exceptions are logged, the entire SQL statement is logged.
Logging Granularity	The number of minutes (1, 2, 5, 10, 15, 30, or 60) in a logging unit. If requested in a report, SQL Guard summarizes request data at this granularity. For example, if the logging granularity is 60, a certain request occurred n times in a given hour. Although when exactly it occurred within the hour is not shown, if a rule in a policy is triggered by a request, a real time alert can indicate the exact time. When you define exception rules for a policy, those rules can also apply to the logging unit. For example, you might want to ignore 5 login failures per hour, but send an alert on the sixth login failure.
Inspect Returned Data	Mark to inspect data returned by SQL requests. If you want to include extrusion rules in your security policy, you must mark this checkbox.
Max. Hits per Returned Data	When returned data is being inspected, indicate how many hits (policy rule violations) are to be recorded.
Buffer Free n%	Display only. <i>n</i> is the percent of free buffer space available for the inspection engine process. This value is updated each time the window is refreshed. There is a single inspection engine process that drives all inspection engines. This is the buffer used by that process.

Field or Control	Description
Ignored Ports List	A list of ports to be ignored. Add values to this list if you know your database servers are processing non-database protocols and you want SQL Guard to not waste cycles analyzing non-database traffic. For example, if you know the host on which your database resides also runs an HTTP server on port 80, you can add 80 to the ignored ports list, ensuring that SQL Guard will not process these streams.
	Separate multiple values with commas, and use a hyphen to specify an inclusive range of ports. For example:
	101,105,110-223
Restart Inspection Engines	Click the Restart Inspection Engines button to stop and restart all inspection engines.
	Note: Any global changes made (and saved using the Apply button, below) do not take effect until you restart the inspection engines. However, individual inspection engine attributes, such as exclude, sequence order, etc., take effect immediately.
Comment	Click the Comment button to add comments to the Inspection Engine Configuration. See the <i>User Comments</i> topic in the <i>SQL Guard User Guide</i> for detailed instructions on how to use or view user comments.
Apply	Click the Apply button to save the configuration. See the note above regarding when changes are actually applied.

Modifying Inspection Engine Configurations

Once an inspection engine configuration has been added to the list of inspection engines, only the following settings for that inspection engine can be modified:

- Active on Startup
- Exclude From-IP

See [Adding Inspection Engines](#), below, for a description of these settings.

No other settings (for example, from-IP and to-IP address lists) can be modified once the configuration has been saved. If you configured an inspection engine incorrectly or need to make modifications due to changes in applications or network environment, add a new inspection engine and then remove the old one. To do this, see the following topics:

- [Adding Inspection Engines](#)

- [Removing Inspection Engines](#)

Once you have defined a new inspection engine, be sure to position it correctly in the list of inspection engines (using the up and down arrow keys on the pane title bar).

Adding Inspection Engines

To add an inspection engine:

1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Inspection Engines to open the *Inspection Engine Configuration* panel.
3. Click anywhere on the Add Inspection Engine pane title if the Add Inspection Engine pane is not expanded:



This opens the expanded Add Inspection Engine pane:

The screenshot shows the 'Add Inspection Engine...' pane expanded. It includes fields for 'Name' (text input), 'Protocol' (dropdown menu), 'Active on startup' (checkbox), 'Exclude From-IP' (checkbox), 'From-IP/Mask' (checkbox with four input fields), 'Ports' (checkbox with four input fields), 'To-IP/Mask' (checkbox with four input fields), and a large 'Add' button at the bottom right.

Refer to the following table for a description of the fields in the Add Inspection Engine pane.

Add Inspection Engine Pane

Field or Control	Description
Name	The new inspection engine name. It must be unique on the unit, and a meaningful name is recommended. We also recommend that you use only letters and numbers in the name, as the use of any special characters will prevent you from working with this inspection engine using the CLI.
Protocol	Select either the protocol to be monitored (Informix , DB2 , Sybase , MSSQL , Named Pipes , or Oracle) or the keyword Ignore . Select Ignore if you want all traffic between the specified client(s) and server(s) to be ignored.
Active on startup	If marked, the inspection engine is activated on system start-up.

Field or Control	Description
Exclude From-IP	If marked, the inspection engine monitors traffic from all clients <i>except</i> for those listed in the From-IP/Mask list (see below). To ignore a specific set of clients <i>without</i> including all others, define a separate inspection engine for those clients and use the <i>Ignore</i> protocol (see above).
From-IP/Mask	<p>A list of clients to be monitored (or excluded if the <i>Exclude From-IP</i> box is marked, as described above). The clients are identified by IP addresses and subnet masks.</p> <p>Click the plus button to add an IP address and subnet mask. Click the minus button to remove the last IP address and subnet mask (at the bottom of the list).</p>
Ports	<p>A single port or a range of ports over which traffic between the specified clients and database servers will be monitored. Most often, this should be a single port.</p> <hr/> <p>Warning: Do not enter a wide range of ports, just to be certain that you have included the right one! You may cause the inspection engine to bog down attempting to analyze traffic on ports that carry no database traffic or traffic that is of no interest for your environment.</p>
To-IP/Mask	<p>A list of database servers to be monitored, identified by IP addresses and subnet masks.</p> <p>Click the plus button to add an IP address and subnet mask. Click the minus button to remove the last IP address and subnet mask (at the bottom of the list).</p>
Add	<p>Click the Add button to save this inspection engine configuration.</p> <ol style="list-style-type: none"> 4. Click the Add button when you have supplied all of the information for this inspection engine. The new inspection engine configuration is displayed at the bottom of the list. 5. Optional: Remember that the filtering mechanisms defined in the inspection engine configurations are executed in the order in which the configurations are listed. If necessary, reposition the new inspection engine configuration, or any existing configurations, using the Up and/or Down arrow buttons in the configuration borders:



Click the Up button to move the configuration up one position in the list. Click the Down button to move the configuration down one position in the list.

6. Click the Start button (which will be replaced by a Stop button after the engine starts) to start the inspection engine just configured.

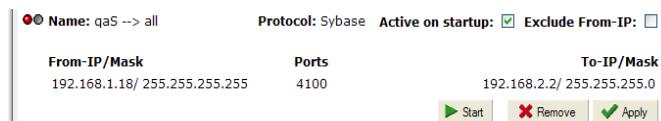
Starting and Stopping Inspection Engines

Once an inspection engine has been configured, you can start or stop it from the *Inspection Engine Configuration* panel.

Note: If you are using the optional Central Management feature to control multiple SQL Guard units from a central system, you can also start or stop inspection engines from the Central Management control panel. See [Central Management](#) later in this chapter for more information.

Starting Inspection Engines

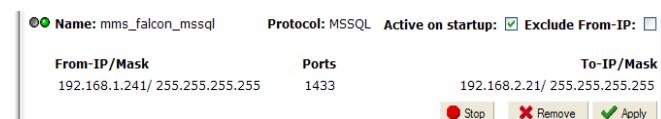
1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Inspection Engines to open the *Inspection Engine Configuration* panel.
3. Locate, in the list of Inspection Engines displayed, the inspection engine you want to start. A red light to the left of the Name field indicates that the inspection engine is not active. For example:



4. Click the Start button to start this inspection engine.

Stopping Inspection Engines

1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Inspection Engines to open the *Inspection Engine Configuration* panel.
3. Locate, in the list of Inspection Engines displayed, the inspection engine you want to stop. A green light to the left of the Name field indicates that the inspection engine is active. For example:

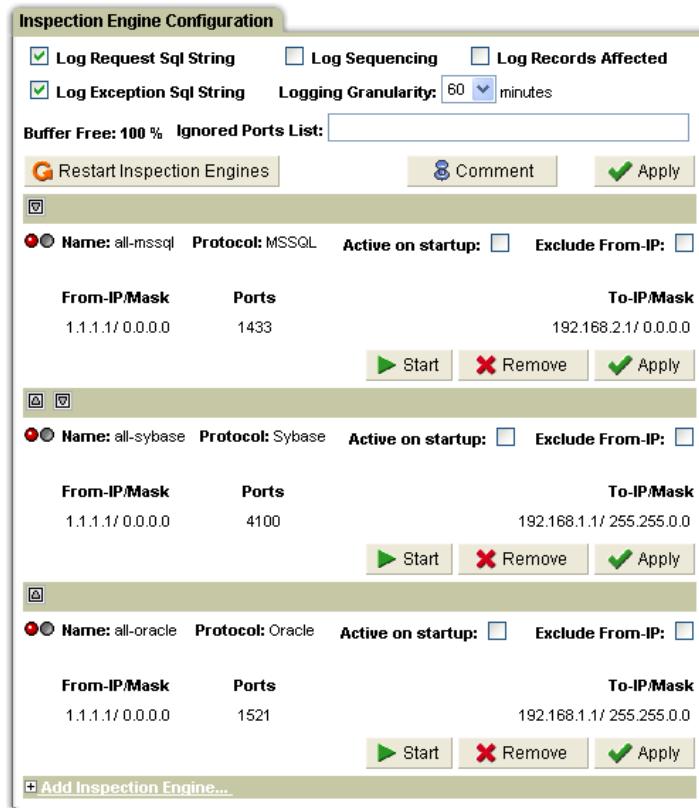


4. Click the Stop button to stop this inspection engine.

Removing Inspection Engines

To remove an inspection engine configuration:

1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Inspection Engines to open the *Inspection Engine Configuration* panel.



3. If the inspection engine to be removed is running, click that engine's Stop button to stop the engine before removing it.
4. Click the Remove button for the inspection engine to be removed. You are prompted to confirm the action.

Installing Policies

You can install a single policy on each SQL Guard server.

Note: Non-administrator users may also be authorized to install policies, from the Data Access Policy Application. See the *User Guide* for more information.

To install a policy on a standalone unit:

1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Policy Installation to open the policy installation panels:



The Currently Installed Policy panel describes the currently installed policy (none, in the example above).

3. Select, in the Policy Installer panel, the policy you want to install from the Policy Description list. Click the Install button. This installs the policy on the unit. However, the inspection engines will not use the policy selected until the engines are restarted.
4. Click the Refresh Inspection Engines button to start using the policy just installed. Otherwise, the policy just installed will not be used until the next time the inspection engines (or the unit) are restarted.
5. Optionally click the Comment button to add or view user comments. See the *User Comments* topic in the *User Guide* for detailed instructions on how to add or view user comments.

To install a policy on a managed unit, see [Installing Security Policies on Managed Units](#), under the *Central Management* topic, later in this chapter.

Changing the SQL Guard Portal

You can keep the SQL Guard Web server on its default port (8443) or reset the portal as described below. We strongly recommend that you use the default port number.

1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Portal to open the SQL Guard Portal panel:



3. Mark the Active on Startup checkbox (this should never be disabled).
4. Set the HTTPS Port to an integer value between 1025 and 65535.
5. Click the Apply button to save the value. (The Guardium security portal will not start listening on this port until it is restarted.)

OR

Click the Revert button to restore the value stored by the last Apply operation.

6. Click the Restart button to restart the unit if you have made and saved any changes. You can now connect to the unit on the newly assigned port.

Note: To re-connect to the unit once it has restarted with the new port number, you must change the URL used to open the SQL Guard Login Page on your browser.

Changing the Portal User Authentication Method

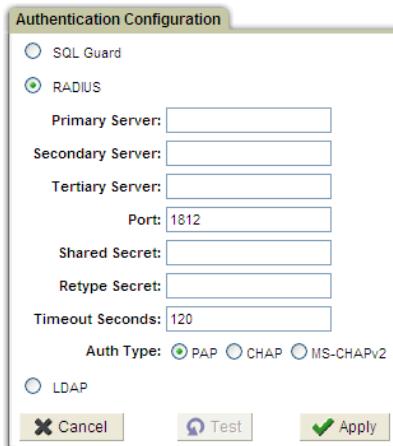
By default, SQL Guard users will be authenticated by SQL Guard. You can use the Authentication Configuration panel (illustrated below) to select an alternative authentication method.



If you select anything other than SQL Guard, only the special SQL Guard **admin** user will be able to log in locally using SQL Guard authentication. All other SQL Guard users will be authenticated via the selected authentication method. The alternative authentication methods are described in the following subsections:

RADIUS Authentication

When you select the RADIUS radio button on the Authentication Configuration panel, the following additional fields appear in the panel:



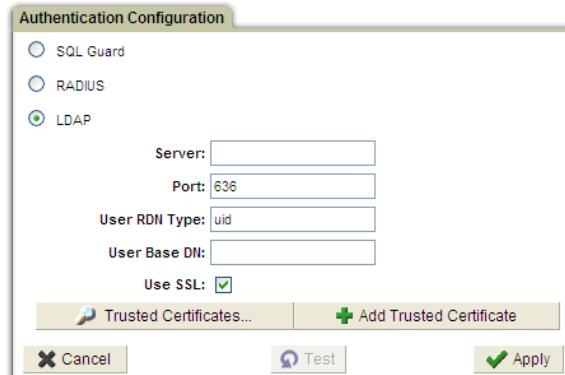
To configure RADIUS authentication, follow the procedure outlined below. If you need more detailed information, see the documentation for your RADIUS server:

1. Enter host name or IP address of the primary RADIUS server.
2. Optionally enter the host name or IP address of the secondary and tertiary RADIUS servers.
3. Enter the UDP port used (1812 or 1645) by your RADIUS server(s).
4. Enter the RADIUS server(s) shared secret twice.
5. Enter the timeout seconds (the default is 120).
6. Select the authentication scheme: PAP (password authentication protocol), CHAP (challenge-handshake authentication protocol), or MS-CHAPv2 (Microsoft version 2 of the challenge-handshake authentication protocol).

7. Optionally click the Test button to verify the configuration. You will be informed of the results of the test. The configuration will be tested whenever you click the Apply button to save changes (see below).
8. Click Apply. The SQL Guard server will attempt to authenticate a test user, and inform you of the results.

LDAP Authentication

When you select the LDAP radio button on the Authentication Configuration panel, the following additional fields appear in the panel:



To configure LDAP authentication, follow the procedure outlined below. If you need more detailed information, see the documentation for your LDAP server:

1. Enter the host name or IP address of the LDAP server.
2. Enter the port number (the default is 636 for LDAP over SSL).
3. Enter the user relative distinguished name type (uid by default).
4. Enter the user base distinguished name.
5. Mark or clear the Use SSL checkbox, as appropriate for your LDAP Server.
6. Optional. To inspect one or more trusted certificates, click Trusted Certificates and follow the instructions in that panel.
7. Optional. To add a trusted certificate, click Add Trusted Certificates and follow the instructions in that panel.
8. Optional. Click the Test button to verify the configuration. You will be informed of the results of the test. The configuration will be tested whenever you click the Apply button to save changes (see below).

9. Click Apply. The SQL Guard server will attempt to authenticate a test user, and inform you of the results.

Setting Global Profile Defaults

Use the Global Profile panel to set system-wide defaults, as explained below.



Use Aliases Default

Use Aliases in Reports unless otherwise specified

When marked, Aliases will be displayed by default for any report that does not explicitly have Show Aliases disabled in the Customize Portlet panel.

Alert Message Template

The Message Template box contains the template used to generate alert messages produced by rule actions. You can change the template by editing the text and clicking the Apply button. Any changes you make here will not take effect until the inspection engines are restarted. (To restart the inspection engines from the CLI, see the description of the restart inspection-engines command under the [restart Commands](#) topic in Chapter 6.)

The template contains placeholders for variables in the format **%%variableName**, where the variable name is terminated by a space character. If you modify the template, be sure to terminate each variable name that you use with a space character.

The default message template is illustrated below:

```
Alert based on rule ID %%ruleDescription
Category: %%category
Classification: %%classification
Severity %%severity
Rule # %%ruleID [%%ruleDescription ]
Request Info: [ Session start: %%sessionStart
Server Type: %%serverType
```

```

Client IP %%clientIP
ServerIP: %%serverIP
Client Port: %%clientPort
Server Port: %%serverPort
Net Protocol: %%netProtocol
DB Protocol: %%DBProtocol
DB Protocol Version: %%DBProtocolVersion
DB User: %%DBUser
Application User Name: %%AppUserName
Source Program: %%SourceProgram
Authorization Code: %%AuthorizationCode
Request Type: %%requestType
Last Error: %%lastError
SQL: %%SQLString

```

Each variable is described briefly, below.

Variable	Description
%%ruleDescription	The rule description from the policy rule definition
%%category	Category from the rule definition
%%classification	Classification from the rule definition
%%severity	Severity from the rule definition
%%ruleID	The rule number from the policy rule definition
%%sessionStart	Session start time (login time)
%%serverType	The database server type
%%clientIP	Client IP address
%%serverIP	Server IP address
%%clientPort	Client port number
%%serverPort	Server port number
%%netProtocol	Network protocol – for K-TAP on Oracle, this may display as either IPC or BEQ
%%DB Protocol	Database protocol
%%DBProtocolVersion	Database protocol version
%%DBUser	Database user name
%%AppUserName	Application user name
%%SourceProgram	Source program name

Variable	Description
<code>%%AuthorizationCode</code>	Authorization code
<code>%%requestType</code>	Request type
<code>%%lastError</code>	Last error description, which is only available when a SQL error request triggering an exception rule contains a <i>last error description</i> field
<code>%%SQLString</code>	SQL string (if any)

Note:

If upgrading to 6.0 or later of Guardium from a version prior to 6.0:
The default Message Template changed for version 6.0. The new message contains three new variables that were not available in previous releases: %%category, %%classification, and %%severity. During the upgrade process, the existing Message Template will not be updated to the new version (since you may have customized it), so if you want to include the new variables, you will need to update the Message Template manually after completing the upgrade procedure.

Application User Translation

Some applications manage a pool of database connections. In such three-tier architectures the pooled connections all log into a database using a single functional ID, and then manage all application users internally – when a user session needs access to the database it acquires a connection from the pool, uses it and then releases it back to the pool. When this happens, SQL Guard can see how the application interacts with the database, but it cannot attribute specific database actions to specific application users. For some widely used applications, SQL Guard has built-in support for identifying the end-user information from the application, and thus can relate database activity to the application end-users.

To use this facility, follow the procedure outlined below:

1. Define an Application User Translation configuration for the application.
2. Populate any pre-defined groups required for that application.
3. Regenerate any portlets for special reports for that application, and place the portlets on a page.

Each of these steps is described separately, below.

Selective Audit Trail and Application User Translation

If the installed security policy uses the selective audit trail feature to limit the amount of data logged, there are two important considerations:

- The policy will ignore all of the traffic that does not fit to the application user translation rule (for example, not from the application server).
- Only the SQL that matches the pattern for that security policy will be available for the special application user translation reports.

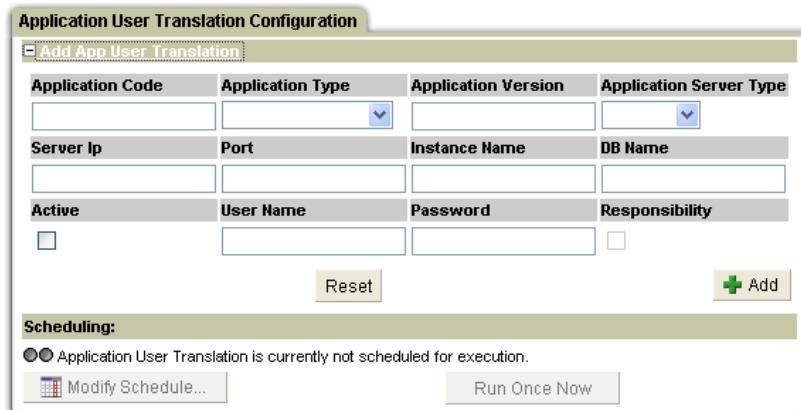
Configure Application User Detection

Follow the procedure outlined below to configure application user detection.

1. From the Configuration Setup section of the Administration Console menu, select Application User Translation to open the Application User Translation Configuration panel:



2. Click the  (Add) button to open the Add App User Translation panel:



The screenshot shows the 'Application User Translation Configuration' window. At the top, there's a toolbar with a green '+' icon labeled 'Add'. Below it is a table with several rows and columns for configuration. The first row has columns for 'Application Code', 'Application Type', 'Application Version', and 'Application Server Type'. The second row has columns for 'Server Ip', 'Port', 'Instance Name', and 'DB Name'. The third row has columns for 'Active', 'User Name', 'Password', and 'Responsibility'. There are also 'Reset' and 'Add' buttons. At the bottom, there's a section titled 'Scheduling:' with a note that says 'Application User Translation is currently not scheduled for execution.' and two buttons: 'Modify Schedule...' and 'Run Once Now'.

Use this panel to define each application for which you will want automatic user translation performed.

3. In the Application Code box, enter a unique code to identify the application.

Note: Under Central Management, you must use different application codes on different managed machines. This prevents aliases generated for the users from conflicting with each other. (Under Central Management, there is one set of aliases that is shared by all managed units.)

4. From the Application Type list, select the application type:

EBS
ICM
PeopleSoft
SAP
Siebel – Observed
Siebel – DB

5. In the Application Version box, enter the application version number (11, for example).
6. From the Application Server Type list, select the application server type. Only the server types available for the selected Application Type and Version (see above) will be displayed. For example, for **EBS** (Oracle E-Business Suite), Oracle will be the only server type in this list.
7. In the Server IP box, enter the IP address the application uses to connect to the database.

8. In the Port box, enter the port number the application uses to connect to the database.
9. In the Instance Name box, enter the instance name the application uses to connect to the database.
10. In the DB Name box, enter the database name for the application. (Required for some applications, not used for others.)
11. Mark the Active box to enable user translation. (Nothing will be translated until after the first import of user definitions – see below).
12. Enter a User Name for SQL Guard to use when accessing the database.
13. Enter a Password for SQL Guard to use when accessing the database.
14. Mark the Responsibility box if you want to associate responsibilities (Administration, for example) with user names. Or clear the Responsibility box to just record user names. When the box is cleared, all activities performed by a user will be grouped together, regardless of the responsibility at the time the activity occurred.
15. Click the Add button to save the Application User Translation definition.
16. Click Run Once Now to import the user definitions for this application (and any others defined).

Later, after verifying that the data import operation worked successfully, return to this panel and click the Modify Schedule button to define an import operation to run on a regular basis. You should schedule the importing of user definition data at whatever interval is suitable for your environment. The maximum time that a new application user name will not be available will the time between executions of the import operation. For instructions on how to use the scheduler, see [Using the Task Scheduler](#), later in this chapter.
17. From the Administration Console, select the Inspection Engines menu item, and click Restart Inspection Engines in the Inspection Engine Configuration panel.

Populate Pre-Defined Application Groups

This section applies only when Application User Translation is being used. When Application User Translation has been configured, you must populate at least two pre-defined groups. The following table identifies the groups you must populate for each application type:

Application	Pre-Defined Group	Group Type
EBS	EBS App Servers	Client IP
	EBS DB Servers	Server IP

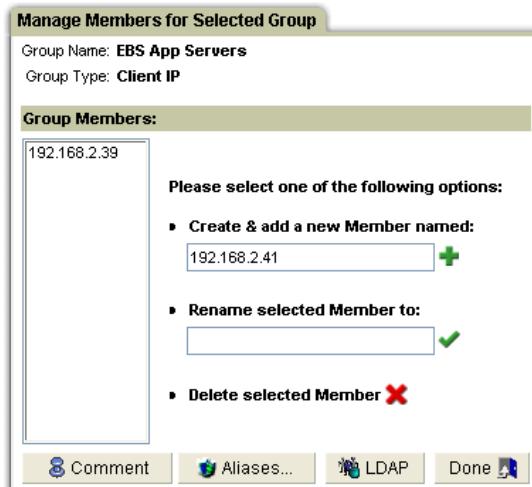
Application	Pre-Defined Group	Group Type
ICM	ICM App Servers	Client IP
	ICM DB Servers	Server IP
PeopleSoft	PSFT App Servers	Client IP
	PSFT DB Servers	Server IP
	PeopleSoft Objects	Objects
Siebel – DB	SIEBEL App Servers	Client IP
	SIEBEL DB Servers	Server IP
Siebel – Observed	SIEBEL App Servers	Client IP
	SIEBEL DB Servers	Server IP
SAP	SAP App Servers	Client IP
	SAP DB Servers	Server IP

To populate a pre-defined group:

1. Navigate to the Group Builder (Tools tab – Config & Control tab – Group Builder menu selection).
2. In the Modify Existing Groups panel, select the appropriate group from the list above (EBS App Servers, in the example below). from the list and



3. Click the Modify button to open the Manage Members for selected Group panel:



4. In the *Create and add a new Member named* box, enter a group member name (in the example above, the IP address for an EBS application server), and then click the **+** (Add) button.

Repeat this step for each server.

5. Click the Done button to return to the Modify Existing Groups panel (illustrated previously, above).

Repeat this procedure for each pre-defined group, for each application for which you have configured application user translation.

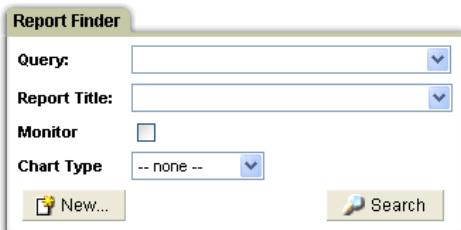
Regenerate Special Application Report Portlets

This section applies only when Application User Translation is being used. For some application types, one or more special report portlets must be regenerated. For example, there are two pre-defined EBS reports, and two pre-defined PeopleSoft reports. These reports cannot be modified. After populating the pre-defined EBS or PeopleSoft groups, as described above, follow the procedure outlined below to regenerate the EBS or PeopleSoft portlets and place them on a page.

The examples in this section are for the EBS portlets, but the procedure is identical for other application types.

1. Optional. Log in to the SQL Guard management interface using an appropriate user account for viewing the application reports.

2. Navigate to the Report Finder panel as follows:



If logged in as *admin*: Tools tab – Report Building tab – Report Builder menu selection.

If logged in under a user account: Reports & Alerts tab – Custom Reports tab – Report builder button.

3. Click the Search button to open the Report Search Results panel:



4. Select a report portlet for the application type (EBS Application Access, in the example above), and click the Regenerate Portlet button.

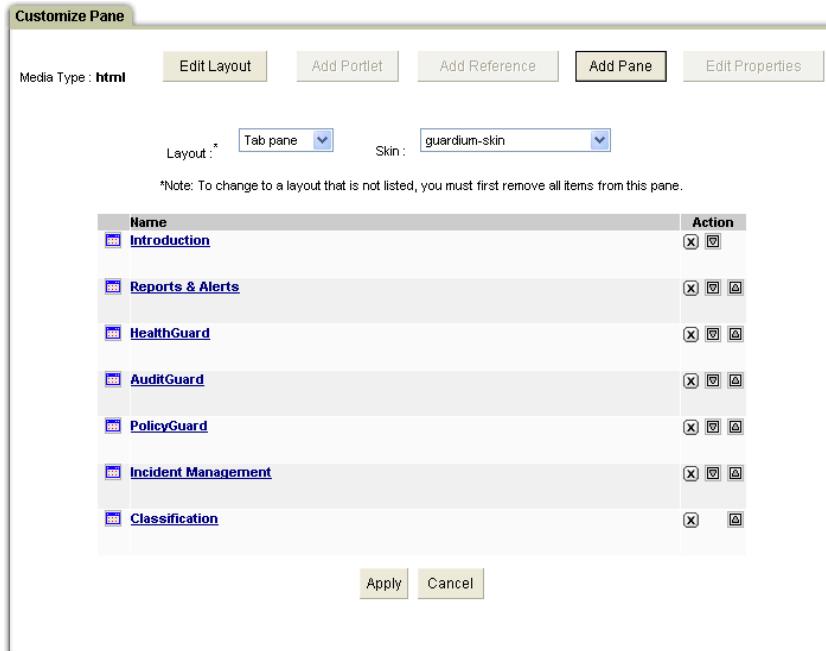
You will be informed that the portlet has been regenerated:



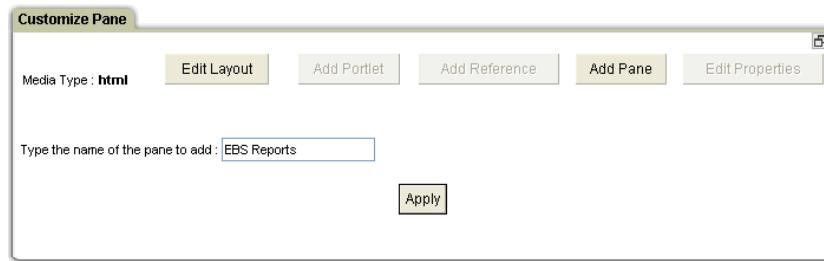
5. Repeat the above step for the EBS Processes Database Access, or the PSFT Processes Database Access report.

Now add a tab to your layout, and include the two regenerated portlets on that tab.

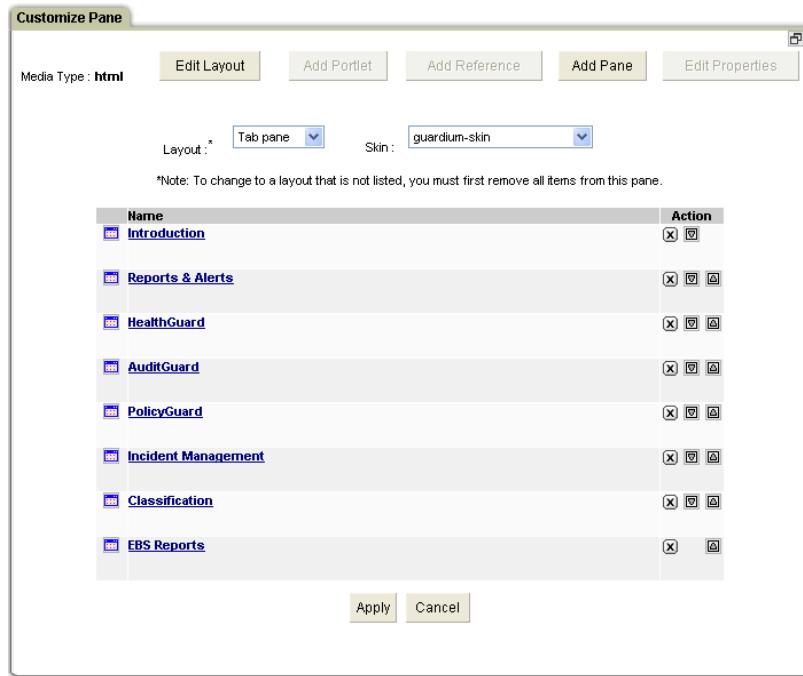
6. Click the **Customize** link at the top of the SQL Guard window, to open the Customize Pane (a standard user tabbed pane layout is illustrated below and is used for the remainder of this section):



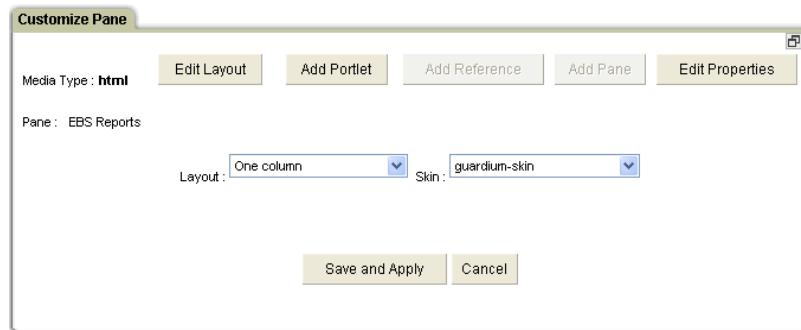
7. Click the Add Pane button to define a new tab:



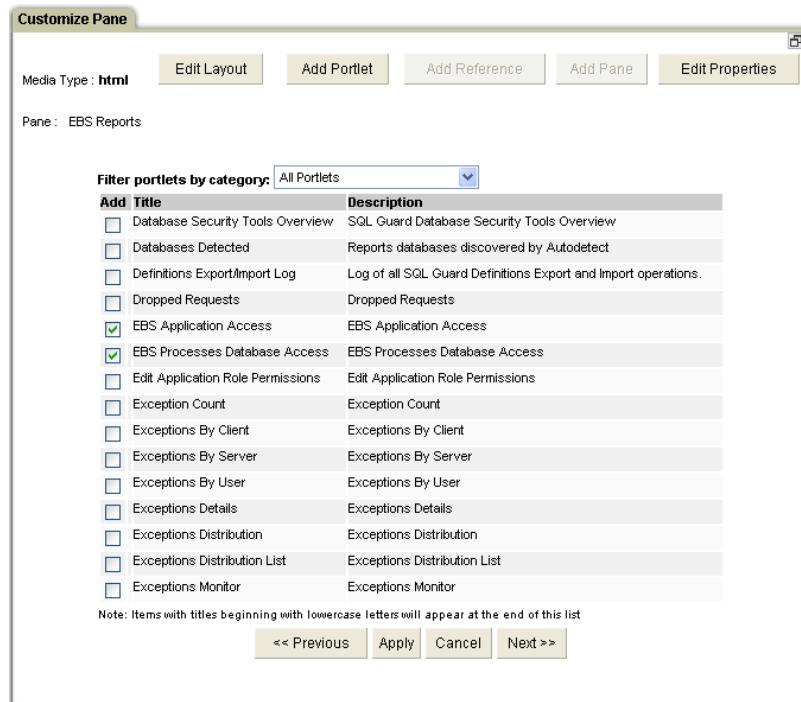
8. Enter a name for the tab, *EBS Reports* in the example above, and click *Apply*:



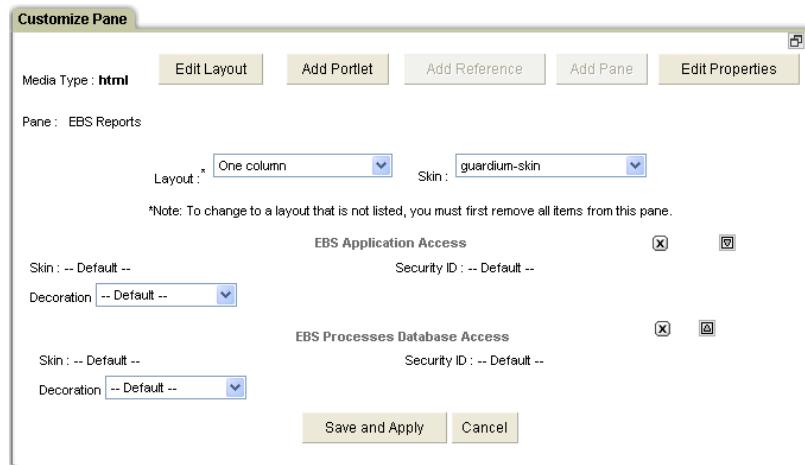
9. The new tab appears as the last tab in the list. Click on the new tab name to edit that pane:



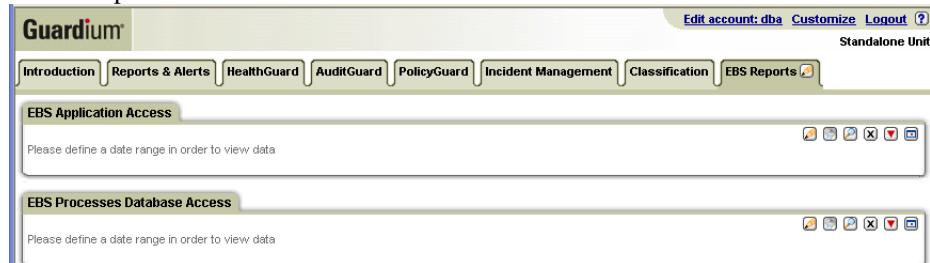
10. Click the Add Portlet button, and click the Next button until you locate the reports you want (EBS in the example). For example:



11. Mark the checkboxes beside each desired report, and then click Apply:



12. Click Save and Apply to save the new pane layout.
13. Click Apply (not shown) to save the new tab layout.
14. The new tab will appear at the end of the first row of tabs. Click on the new tab name to open the tab



15. Now set the date range and Show Aliases run-time parameters for each report as follows: Click the customize button at the right side of the portlet panel. The

Customize Portlet panel for the first report is shown below:

Customize Portlet

Report: **EBS Application Access** Based on Query: **EBS Application Access**

Title: EBS Application Access

Skin: -- Default --

Run Time Parameters

GROUPING_SUB_TYPE: Choose A Group Type Or Sub-Type to Group By

Choose Grouping Type

QUERY_FROM_DATE: Enter Period From

QUERY_TO_DATE: Enter Period To

REMOTE_SOURCE: Remote Data Source

SHOW_ALIASES: Show Aliases

Presenation Parameters

fetchSize: 20

refreshRate: 0

Update Cancel

The screenshot displays a 'Customize Portlet' dialog box. At the top, it shows the report name 'EBS Application Access' and its base query. Below this, there are fields for 'Title' (set to 'EBS Application Access') and 'Skin' (set to 'Default'). A section titled 'Run Time Parameters' contains fields for 'GROUPING_SUB_TYPE' (with a dropdown menu 'Choose A Group Type Or Sub-Type to Group By'), 'QUERY_FROM_DATE' (labeled 'Enter Period From' with a date input field), 'QUERY_TO_DATE' (labeled 'Enter Period To' with a date input field), 'REMOTE_SOURCE' (set to 'none'), and 'SHOW_ALIASES' (with radio buttons for 'On', 'Off', and 'default', where 'On' is selected). Another section titled 'Presentation Parameters' includes 'fetchSize' (set to 20) and 'refreshRate' (set to 0). At the bottom are 'Update' and 'Cancel' buttons.

16. In the QUERY_FROM_DATE field, enter the starting point for the report (now -1 day in the example above).
17. In the QUERY_TO_DATE field, enter the endpoint for the report (now in the example).
18. In the SHOW_ALIASES row, select the On button. You should always leave the Show Aliases button enabled for these reports.
19. Click the Update button. The output of the first report will be displayed.
20. Repeat steps 15-19 above to set the runtime parameters for the second report.

The output for each report is described in the following section.

About the Special Application Reports

As mentioned earlier, if the installed security policy is logging a limited amount of traffic due to the use of the selective audit trail feature, only the SQL that matches the pattern for that security policy will be available for these (and all other) reports.

There are pre-configured reports for most of the special applications. You cannot modify these reports, but you can clone the queries they are based upon and generate your own reports. You must regenerate the portlets for these reports after populating any pre-defined groups for that application, as described previously.

The following table lists the reports available for each application type, and the reporting domain used.

Incident Generation Processes Panel Reference Table

Application	Default Reports	Domain
EBS	EBS Application Access	Access Tracking
	EBS Processes Database Access	Access Tracking
ICM	ICM Application Access	Application Tracking
PeopleSoft	PSFT Application Access	Access Tracking
	PSFT Processes Database Access	Access Tracking
Siebel	Siebel Application Access	Application Tracking
SAP	SAP Application Access	Application Tracking

For each of these reports, you can access a number of drill-down reports, as would be the case for any report portlet. Drill-down capabilities and reports are described in the *SQL Guard User Guide*.

Configuring the Alerter

Use the Alerter panel to define your SMTP and SNMP setup to SQL Guard, so that the SQL Guard system can send notifications using one or both of these methods.

Note: Notifications of any type cannot be sent until the Alerter configuration has been completed with the *Active on Startup* box marked (see below).

To configure the Alerter:

1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Alerter to open the Alerter panel:



3. Mark the Active On Startup box to start the Alerter on startup of the SQL Guard system.
4. In the Polling Interval box, enter the number of seconds between checks for outbound messages.

In the SMTP panel, configure SQL Guard for email notifications:

5. Enter the IP address for the SMTP gateway in the IP Address box.

6. Enter the SMTP port number (it is almost always 25) in the Port box.
7. Optional: Click the Test Connection hyperlink to verify the SMTP address and port. This only tests that there is access to specified host and port. It does *not* verify that this is a working SMTP server. A dialog box is displayed, informing you of the success or failure of the operation.

Note: If this SMTP server uses authentication, you must supply a valid User Name and Password for that mail server in the following two fields. Otherwise, those fields can be left blank.

8. Enter a valid user name for your mail server in the User Name box if your SMTP server uses authentication.
9. Enter the password for the above user in the Password box if your SMTP server uses authentication. Re-enter it in the Re-enter Password box.
10. In the Return Email Address box, enter the return address for email sent by the system. This address is usually an administrative account that is checked often.
11. Select Auth in the Authentication Method if your SMTP server uses authentication. Otherwise, select None. When Auth is selected, you must specify the user name and password to be used for authentication.

In the SNMP panel, configure SQL Guard for SNMP notifications:

12. In the IP Address box, enter the IP address to which the SNMP trap will be sent.
 13. Optional: Click the Test Connection hyperlink to verify the SNMP address and port (22). This only tests that there is access to specified host and port. It does *not* verify that this is a working SNMP server.
A dialog box is displayed, informing you of the success or failure of the operation.
 14. In the “Trap” Community box, enter the community name for the trap. Retype the community in the Retype Community box.
 15. Click the Apply button to store the values in the configuration database.
-
- Note:** The Alerter will not begin using a new configuration until it is restarted.
-
16. Click Restart to restart the Alerter with the new configuration.,

Stopping the Alerter

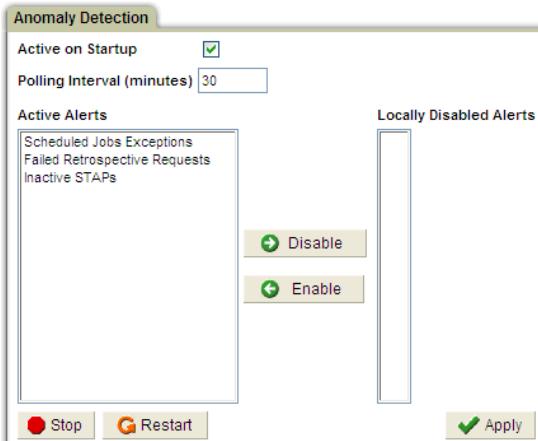
To stop the alerter, open the Alerter panel as described above, and click the Stop button.

Configuring Anomaly Detection

Both the Alerter and Anomaly Detection must be configured with the *Active On Startup* box marked, before the SQL Guard statistical alerts option can be enabled. The Anomaly Detection polling interval controls the frequency with which SQL Guard checks log data for anomalies. (Anomaly Detection is required for statistical alerts, but not for policy-triggered run-time alerts.)

To configure the Anomaly Detection polling interval:

1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Anomaly Detection to open the Anomaly Detection panel:



3. Mark the Active on Startup box to start Anomaly Detection on startup of the SQL Guard server.
4. Enter the number of minutes between checks of the log data.
5. To disable or enable an alert that is marked Active (in the alert definition), select it in the appropriate list and click the Disable or Enable button.
6. Click the Apply button to store the interval in the configuration database.

Note: Anomaly Detection will not begin using a new configuration until it is restarted.

7. Click Restart to restart Anomaly Detection with the new interval.

Stopping Anomaly Detection

To stop Anomaly Detection, open the Anomaly Detection panel as described above and click the Stop button.

Session Inference

The Session Inference process checks for sessions that have not been closed but are also not active for a specified period of time.

To configure the Session Inference options:

1. Open the Administration Console panel (not shown).
2. In the Configuration Setup section of the Administration Console menu, click Session Inference to open the Session Inference panel:



3. Mark the Active On Startup box to start Session Inference on startup of the SQL Guard system.
4. Use the default (120 minutes) or enter the desired value in the Polling Interval box.
5. Use the default (720 minutes) or enter the desired value in the Max Inactive Period box.
6. Click the Apply button to store the values in the configuration database.

Note: Session Inference will not begin using a new configuration until it is restarted.

7. Click Restart to restart Session Inference with the new configuration.,

Stopping Session Inference

To stop Session Inference, open the Session Inference panel as described above and click the Stop button.

IP-to-Hostname Aliasing

About Aliases

An *alias* is a synonym that is substituted for a stored value. An alias is most commonly used to provide a meaningful name that substitutes for a data value. For example, *QA Sybase Server* could be defined as an alias for the server IP address of *192.168.18.1*.

Once an alias has been defined:

- The alias can be displayed instead of the data value in report results.
- The alias can be used instead of the data value to formulate queries and to enter parameter values.

Defining Aliases

There are two ways to define aliases that are generally available to all users:

- **Alias Builder** – The Alias Builder is a tool that you can use to define aliases manually for the values of various entities (IP addresses, users, fields, etc.).
- **Alias Quick Definition** – You can open the Alias Quick Definition window either from a report or from the Group Builder. When opened from a report, you can define an alias for any report element for which aliases are allowed on the selected row/drill down. When opened from the Group Builder you can define an alias for any member of the group being defined or for the group name itself.

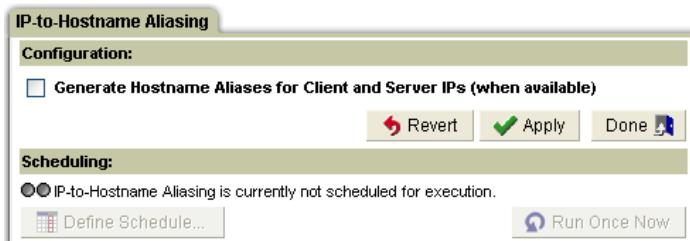
Since the above methods can be used by most users (in addition to administrators), they are described in the *SQL Guard User Guide*.

IP-to-Hostname Aliasing

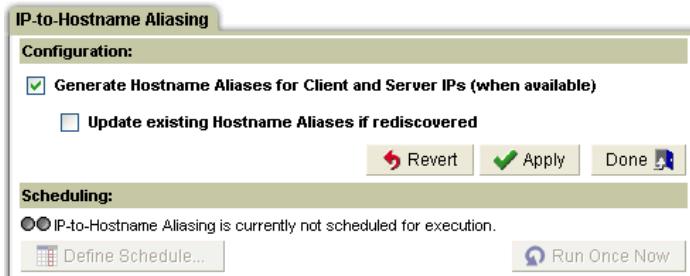
In addition to the alias definition methods described above, **IP-to-Hostname Aliasing** allows SQL Guard administrators to use the DNS table to generate hostname aliases for SQL Guard client and server IP addresses. This task, which can be scheduled to run on a regular basis, is performed from the Administration Console – Configuration Setup menu.

To generate hostname aliases:

1. Select IP-to-Hostname Aliasing from the Administration Console menu to open the IP-to-Hostname Aliasing panel:



2. Mark the *Generate Hostname Aliases...* checkbox to enable hostname aliasing. A second checkbox is displayed:



3. Mark the second checkbox only if you want to update a previously defined alias in the event that the hostname for that IP address changes. Marking this checkbox also causes any aliases that are not the same as the hostname to be overwritten by the hostname. For an example, see the note below.

Note: You may prefer to manually define aliases for some client or server IP addresses. For example, the hostname for server IP address *1.2.3.4* might be *dbserver04.guardium.com*, but that server might be known in the company as *QA Sybase Server*. If you define the alias manually and mark the second checkbox, the alias is overwritten by the hostname.

4. Click the Apply button to save the IP-to-Hostname Aliasing configuration.
5. Click Run Once Now to generate the aliases immediately,

Note: To schedule the IP-to-Hostname Aliasing function to run on a regular basis, see: [Using the Task Scheduler](#).

Viewing IP-to-Hostname Aliases

To view the aliases that have been generated (or defined manually), follow the procedure outlined below:

1. Select **Tools – Config & Control – Alias Builder** to open the Alias Finder:



2. Select from the drop-down list either Client IP or Server IP, to display the client or server IP aliases (respectively).
3. Click Search to open the client or server IP aliases in the Alias Builder panel (Server IP was selected in the example below).

Alias Builder

Group Type: Server IP

	Value	Alias
<input checked="" type="checkbox"/>	192.168.1.18	qaserver.guardium.com
<input checked="" type="checkbox"/>	192.168.1.8	midgard.guardium.com
<input checked="" type="checkbox"/>	192.168.2.100	ostrich.guardium.com
<input checked="" type="checkbox"/>	192.168.2.12	eagle.guardium.com
<input checked="" type="checkbox"/>	192.168.2.17	hamesh.guardium.com
<input checked="" type="checkbox"/>	192.168.2.22	seagull.guardium.com
<input checked="" type="checkbox"/>	192.168.2.247	terminator.guardium.com
<input checked="" type="checkbox"/>	192.168.2.45	robin.guardium.com
<input checked="" type="checkbox"/>	192.168.2.48	crane.guardium.com
<input checked="" type="checkbox"/>	192.168.2.55	phoenix.guardium.com
<input checked="" type="checkbox"/>	192.168.2.64	viper.guardium.com
<input checked="" type="checkbox"/>	192.168.2.75	blackbird.guardium.com
<input checked="" type="checkbox"/>	192.168.2.82	winx64.guardium.com

Add

	Value	Alias
Add		

Done

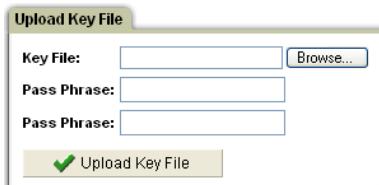
Upload Key File

Key files are required when you have a Microsoft SQL Server environment that is configured using the *force protocol encryption* option, or if you have a Microsoft SQL Server 2005 environment and are using encrypted login sessions with SQL Server mixed authentication. In these cases SQL Guard must decrypt the database streams, and thus requires the key file for the server.

Since a single SQL Guard unit may be monitoring multiple SQL Server instances, you may need to upload multiple key files.

To upload a key file to the SQL Guard server:

1. Open the Administration Console panel (not shown).
2. In the Upload Key file panel of the menu, click Upload to open the Upload Key File panel:



3. Click the Browse button to locate the key file you want to upload. The key file name must be the fully qualified domain name of the SQL Server. The class file cannot be renamed – it must be created with that name.
4. Enter the pass phrase in the Pass Phrase box, and re-enter the phrase in the Pass Phrase Confirm box.
5. Click the Upload Key File button. You will be informed of the results of the operation.

Query Hint

This feature is password protected and can be used only as directed by Guardium Support.



Incident Generation

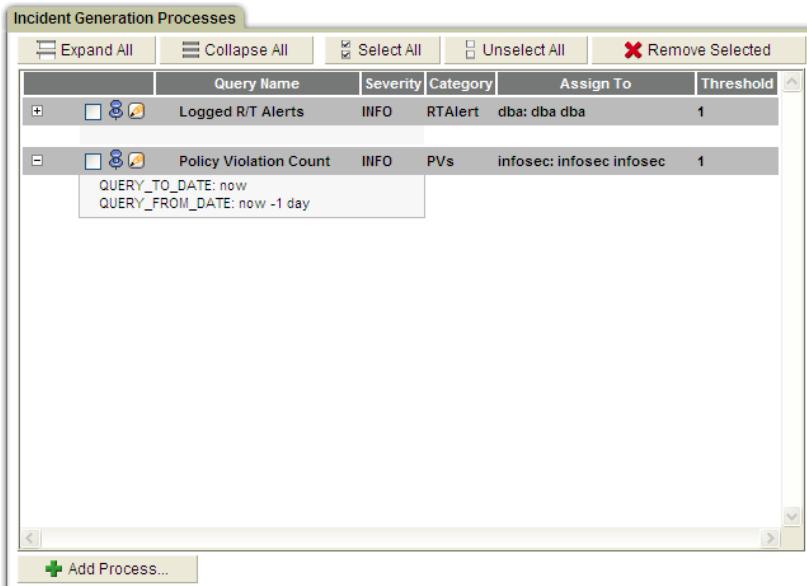
This section describes how to define and schedule incident generation processes. These activities are by default restricted to *admin* users. Once an incident has been generated, both administrators and all other users work with incidents on the Incident Management tab, which is included on both the admin and user portals. See the User Guide for detailed instructions on how to use the Incident Management functions (to open incidents, assign incidents to users, send notifications, and so forth).

An *incident generation process* queries the policy violations domain, which is a log of all policy violations recorded. Before defining an incident generation process, you will need to define a query in that domain, or locate a suitable pre-defined query.

Note: Beginning with version 6.0 of SQL Guard, a threshold alert can be configured to log a policy violation, which means that you can generate incidents based on both threshold and policy alerts.

To work with incident generation processes:

1. Select **Administration Console – Incident Generation** to open the Incident Generation Processes panel:

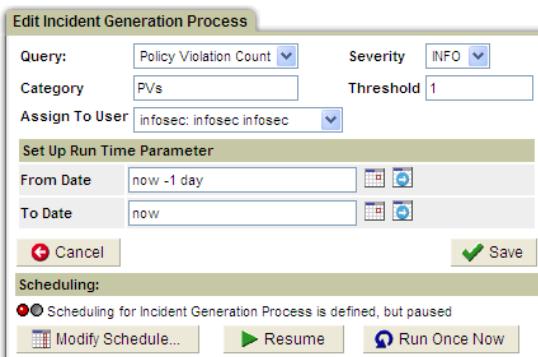


In the example above, two incident generation processes have been defined. The following table describes the controls on this panel

Incident Generation Processes Panel Reference Table

Control	Description
Expand All	Click to display all runtime parameters for all processes.
Collapse All	Hide all runtime parameters for all processes.
Select All	Mark all of the selection checkboxes.
Unselect All	Clear all of the selection checkboxes.
Remove Selected	Remove the selected processes.
[+] or [-]	Click to display or hide all runtime parameters for a process.
[checkmark]	Mark to select the associated process.
[pencil]	Click to edit the associated process definition.
[pushpin]	Click to add a user comment to the process definition. If the process already contains comments, a slip of paper displays beneath the push-pin:  .

2. Click Add Process to open the Edit Incident Generation Process panel:



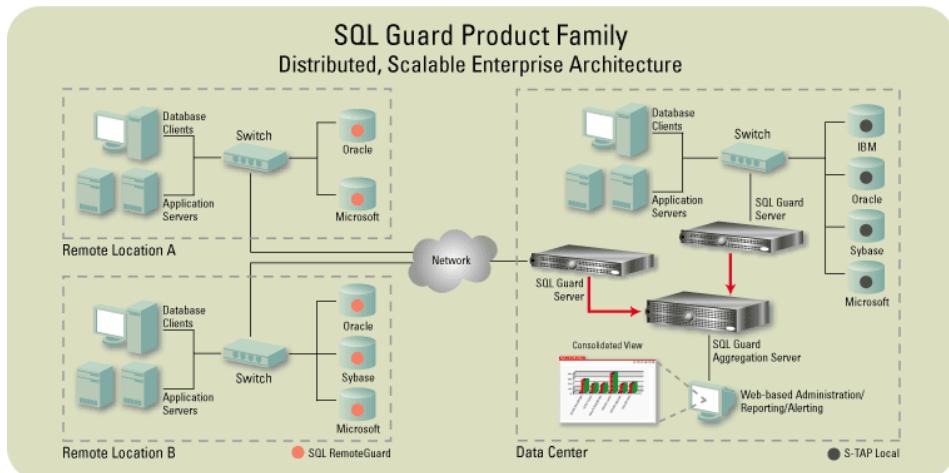
3. Select a query from the list. Only queries from the policy violations domain are available for this purpose.
4. Select a Severity for the incident.
5. Optionally enter a Category for the incident.
6. Select the user to whom the incident will be assigned.
7. Enter From and To Dates for the query.
8. Click Save to save the process definition.

9. To run the query now, click Run Once Now.
10. To schedule the query, click Modify Schedule to open the general-purpose scheduling utility. See [Using the Task Scheduler](#) for a description of how to use the scheduler.

Aggregation

SQL Guard aggregation allows you to collect and merge information from multiple SQL Guard Servers to a single SQL Guard Aggregation Server.

Overview



If you are running SQL Guard in an enterprise deployment, you may have multiple SQL Guard servers monitoring different environments (different geographic locations or business units, for example). It may be useful to collect all data in a central location to facilitate an enterprise view of database usage. You can accomplish this by exporting data from a number of SQL Guard servers to another SQL Guard server that has been configured as an *aggregation* server. In such a deployment, you typically run all reports, assessments, audit processes, and so forth, on the aggregation server to achieve an enterprise view.

Hierarchical Aggregation

SQL Guard supports hierarchical aggregation, where multiple aggregation units merge upwards to a higher-level, central aggregation server. This is useful for multi-level views. For example, you may need to deploy one aggregation server for North America aggregating multiple units, another aggregation server for Asia aggregating multiple units, and a central, global aggregation server merging the contents of the North America and Asia aggregation servers into a single corporate view.

To consolidate data, all aggregated SQL Guard servers export data to the aggregation server on a scheduled basis. The aggregation server imports that data into a single database on the

aggregation server, so that reports run on the aggregation server are based on the data consolidated from all of the aggregated SQL Guard servers.

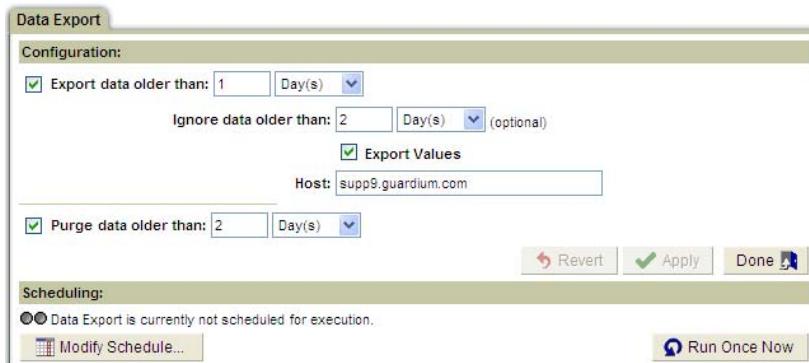
Aggregating, Archiving, and Purging Operations

Scheduled *export* operations send data from SQL Guard *collector* units to a SQL Guard aggregation server. On its own schedule, the aggregation server executes an *import* operation to complete the aggregation process. On either or both units, *archive* and *purge* operations are scheduled to back up and purge data on a regular basis (both to free up space and to speed up access operations on the internal database). The export, archive, and purge functions typically do not operate on the same data. For example, you may want to export and archive all information older than one day and purge all information older than one month, thereby always leaving one month of data on the sending unit. The archive and purge operations are described in subsequent sections.

Exporting Data to an Aggregation Server

To export data to an aggregation server, follow the procedure below. You can define a single export configuration for each SQL Guard unit.

1. Open the Administration Console panel (not shown).
2. In the Data Management section of the Administration Console menu, click Data Export to open the Data Export panel:



3. Mark the Export checkbox.
4. In the boxes following *Export data older than*, specify a starting day for the export operation as a number of days, weeks, or months prior to the current day, which is day *zero*. These are calendar measurements, so if today is April 24, all data captured on April 23 is one day old, regardless of the time when the operation is performed. To archive data starting with yesterday's data, enter the value 1.

5. Optionally, use the boxes following *Ignore data older than* to control how many days of data will be archived. Any value specified here must be greater than the *Export data older than* value, so you always export at least two days of data.

Note: If you leave the *Ignore data older than* row blank, you export data for *all* days older than the value specified in the *Export data older than* row. This means that if you export daily and purge data older than 30 days, you will export each day of data 30 times before it is purged on the 31st day.

6. To remove all data values from the exported data, clear the Export Values box. (It is marked by default.) This may be a requirement if the collector resides in a country that prohibits the export of data, and the aggregation server resides in another country.
7. In the Host box, enter the IP address or DNS host name of the aggregation server to which this system's encrypted data files will be sent.

Note: *This unit and the aggregation server to which it is sending data must have the same System Shared Secret.* If not, the export operation works, but the aggregation server that receives the data is not able to decrypt the exported file. For more information about the System Shared Secret, see [Changing the System Configuration](#) at the start of this chapter.

8. Click the Apply button to save the export and purge configuration for this unit.

When you click the Apply button, the system attempts to verify that the specified aggregator host will accept data from this unit. If the operation fails, the following message is displayed and the configuration will not be saved:

A test data file could not be sent to this host. Please confirm the hostname or IP address is entered correctly and the host is online.

If the Apply operation succeeds, the buttons in the Scheduling panel become active.

9. To run the operation once, click the Run Once Now button.
10. To schedule this operation to run on a regular basis, click the Modify Schedule button. The general-purpose task scheduler is opened. For details on using the general-purpose task scheduler, see [Using the Task Scheduler](#) in Chapter 2.

Stopping Export to an Aggregation Server

To stop the export of data to an aggregation server:

1. Open the Administration Console panel (not shown).

2. In the Data Management section of the Administration Console menu, click Data Export to open the Data Export panel (see above).
3. Clear the Export checkbox.
4. Click Apply.

Importing Data on the Aggregation Server

SQL Guard *collector* units export encrypted data files to another SQL Guard server configured as an *aggregation* server. The encrypted data files reside in a special location on the aggregation server until the aggregation server executes an *import* operation to decrypt and load all data to its own internal database. To avoid the possibility of importing files that have not completely arrived, the aggregation server will *not* import files that have changed in the last two minutes.

Follow the procedure outlined below to define the *Data Import* operation on an aggregation server. You can define only a single Data Import configuration on each unit.

1. Open the Administration Console panel (not shown).
2. In the Data Management section of the Administration Console menu, click Data Import to open the Data Import panel.
3. Mark the Import checkbox. This causes the appearance of an additional non-modifiable field indicating the location of the data files to be imported:



4. Click Apply to save the configuration. The Apply button is only available when you toggle the *Import data from* checkbox on or off.
5. Click the Run Once Now button to run the operation once.
6. Click the Modify Schedule button to schedule the operation to run on a regular basis. The general-purpose task scheduler is opened. For instructions on how to use the general-purpose task scheduler, see [Using the Task Scheduler](#) at the end of Chapter 2.

Note: *This aggregation server and all units exporting data to it must have the same System Shared Secret.* If not, the export operations will still work, but the aggregation server will not be able to decrypt the files of exported data. For more information about the System Shared Secret, see [Changing the System Configuration](#) at the beginning of this chapter.

Stopping Importing on an Aggregation Server

To stop importing data sent from other SQL Guard units:

1. Open the Administration Console panel (not shown).
2. In the Data Management section of the Administration Console menu, click Data Import to open the Data Import panel (see above).
3. Clear the *Import data from* checkbox.
4. Click Apply.

Note: Stopping importing does not stop other SQL Guard units from exporting data to this system. To stop that, you must stop the Export operation on each sending unit.

Archiving and Restoring

The archive and purge process frees space and preserves SQL Guard information for future use. You should periodically archive and purge data from standalone units and from aggregation units that do not themselves have an aggregation export schedule. On SQL Guard units that export data to aggregation servers, data is typically archived only from the highest-level aggregation server, although it is possible to archive from any and all units.

SQL Guard's archive function creates signed, encrypted files that cannot be tampered with. It may be necessary to run reports or investigations on this data at some point. For example, some regulatory environments may require that you keep this information for three, five, or even seven years in a form that can be queried within 24-hours. This functionality is supported by the SQL Guard *restore* capability, which allows you to restore archived data to the unit.

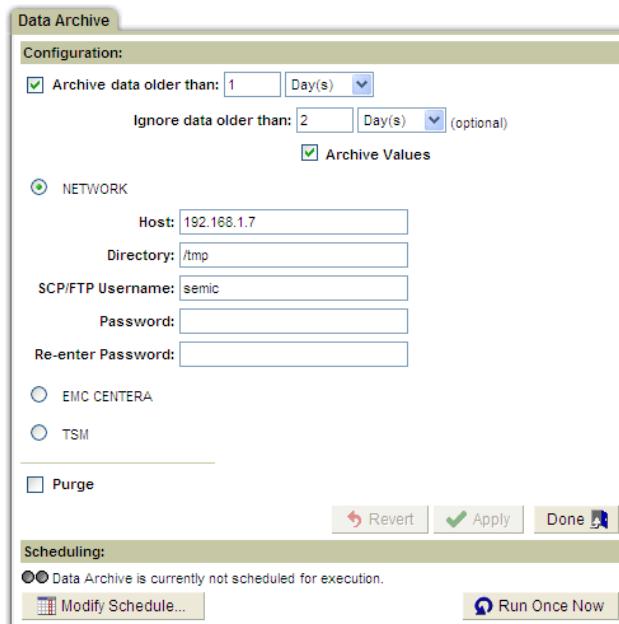
The following sections describe how to define and schedule archiving and how to restore from an archive.

Note: The archive and restore operations depend on the file names generated during the archiving process. **DO NOT** change the names of archived files.

Archiving

Archive data files can be sent to an SCP or FTP host on the network, or to an EMC Centera or TSM storage system (if configured). You can define a single archiving configuration for each unit To archive data to another host on the network and optionally purge data from the unit, follow the procedure outlined below.

1. Open the Administration Console panel.
2. In the Data Management section of the Administration Console menu, click Data Archive to open the Data Archive panel (not shown).
3. Mark the *Archive* checkbox. This displays additional fields, as illustrated below:



4. In the boxes following *Archive data older than*, specify a starting day for the archive operation as a number of days, weeks, or months prior to the current day, which is *day zero*. These are calendar measurements, so if today is April 24, all data captured on April 23 is one day old, regardless of the time when the operation is performed. To archive data starting with yesterday's data, enter the value 1.
5. Optionally, use the boxes following *Ignore data older than* to control how many days of data will be archived. Any value specified here must be greater than the *Archive data older than* value, so you always archive at least two days of data.

Note: If you leave the *Ignore data older than* row blank, you archive data for *all* days older than the value specified in the *Archive data older than* row. This means that if you archive daily and purge data older than 30 days, you archive each day of data 30 times (before it is purged on the 31st day).

Depending on the archive options configured for your system (using the **store storage-system** CLI command), you may have EMC Centera or TSM options on your panel, as illustrated above. If you select one of those archive destinations, see the appropriate topic:

- o [EMC Centera Archive and Backup](#)
- o [TSM Archive and Backup](#)

Steps 6-9 apply only when System is selected as the archive destination.

6. In the Host box, enter the IP address or DNS host name of the host to receive the archived data.
7. In the Directory box, identify the directory in which the data is to be stored. How you specify this depends on whether the file transfer method used is FTP or SCP. If you are unsure which file transfer method has been configured, use the [show transfer-method](#) CLI command (described in Chapter 6).

For FTP: Specify the directory relative to the FTP account home directory.

For SCP: Specify the directory as an absolute path.

8. In the Username box, enter the user name to use for logging onto the host machine. This user must have write/execute permissions for the directory specified in the Directory box (above).
9. In the Password box, enter the password for the above user, then enter it again in the Re-enter Password box.
10. Mark the Purge checkbox to purge data, whether or not it is archived. When this box is marked, the *Purge data older than* fields display.

* **IMPORTANT:** The Purge configuration is used by both Data Archive and Data Export. Changes made here will apply to any executions of Data Export and vice-versa. In the event that purging is activated and both Data Export and Data Archive run on the same day, the first operation that runs will likely purge any old data before the second operation's execution. For this reason, any time that Data Export and Data Archive are both configured, the purge age must be greater than both the age at which to export and the age at which to archive.

11. If purging data, use the *Purge data older than* fields to specify a starting day for the purge operation as a number of days, weeks, or months prior to the current day, which is day zero. All data from the specified day and all older days will be purged, except as noted below. Any value specified for the starting purge date must be greater than the value specified for the *Archive data older than* value. In addition, if data exporting is active (see [Exporting Data to an Aggregation Server](#), above), the starting purge date specified here must be greater than the *Export data older than* value.

Notes: There is no warning when you purge data that has not been archived or exported by a previous operation.

The purge operation **does not** purge restored data whose age is within the *do not purge restored data* timeframe specified on a restore operation. For more information, see [Restoring Archived Data](#), below.

12. Click Apply to verify and save the configuration changes.

When you click the Apply button, the system attempts to verify the specified Host, Directory, Username, and Password by sending a test data file to that location. If the operation fails, the following message is displayed and the configuration will not be saved:

A test data file could not be sent to this host with the parameters given. Please confirm the hostname or IP address is entered correctly, the host is online, the target directory exists and can be written to by the user given, and the password given is correct for that user.

If the Apply operation succeeds, the buttons in the Scheduling panel will become active.

13. Click the Run Once Now button to run the operation once.
14. Click the Modify Schedule button to schedule the operation to run on a regular basis. The general-purpose task scheduler is opened.

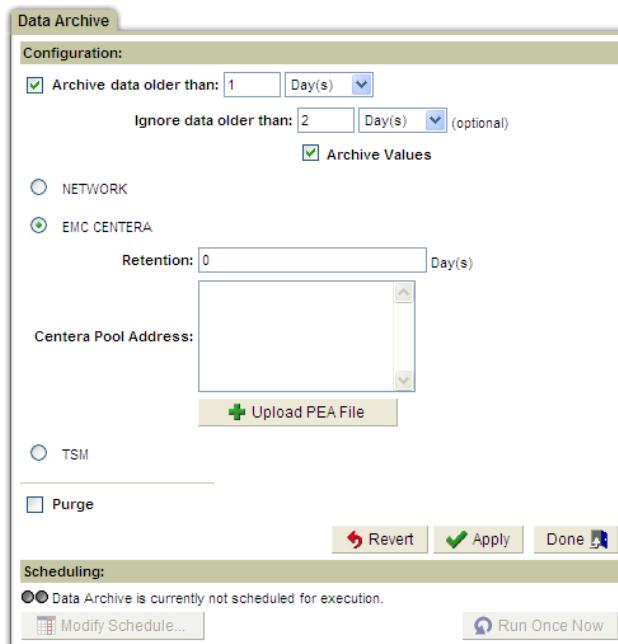
Note: For instructions on how to use the general-purpose task scheduler, see [Using the Task Scheduler](#) in Chapter 2.

EMC Centera Archive and Backup

When you select EMC Centera as an archive or backup destination, the EMC Centera portion of the archive or backup configuration panel expands, as illustrated below (it is the same for both operations – only one version of the panel is illustrated).

To use EMC Centera:

1. Open the Data Archive or System Backup panel. Initially, the Network radio button is selected by default, and the Network backup parameters are displayed (not shown here).
2. Select the EMC Centera radio button. The EMC Centera parameters will be displayed on the panel, as illustrated below for the Data Archive panel (they are identical for the System Backup panel):



3. In the Retention box, enter the number of days to retain the data. The maximum is 24855 (68 years). If you want to save it for longer, you can restore the data later and save it again.
4. In the Centera Pool Address box, enter the Centera Pool Connection String; for example:
10.2.3.4,10.6.7.8?/var/centera/profile1_rwe.pea
5. Click the Upload PEA file button to upload a Centera PEA file to be used for the connection string.
6. Click the Apply button to save the configuration. The system will attempt to verify the Centera address by opening a pool using the connection string specified. If the operation fails, you will be informed and the configuration will not be saved.
7. Return to the general instructions for Archiving or System Backup.

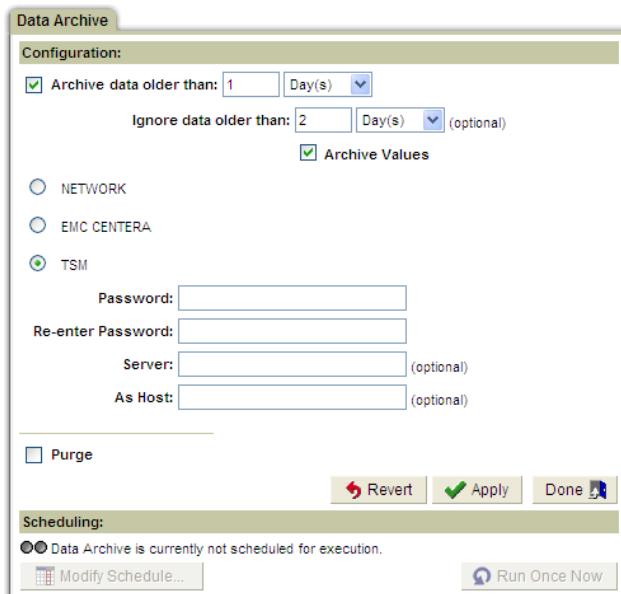
TSM Archive and Backup

When you select TSM as an archive or backup destination, the TSM portion of the archive or backup configuration panel expands, as illustrated below (it is the same for both operations – only one version of the panel is shown).

Before setting TSM as an archive or backup destination, the SQL Guard system must be registered with the TSM server as a client node. A TSM client system options file (*dsm.sys*) must be created (on your PC, for example) and uploaded to SQL Guard. Depending on how that file is defined, you may also need to upload a *dsm.opt* file. For help creating a *dsm.sys* file for use by SQL Guard, consult with your company's TSM administrator. To upload a TSM configuration file, see the [import tsm config](#) CLI command in Chapter 6.

To use TSM:

1. Open the Data Archive or System Backup panel. Initially, the Network radio button is selected by default, and the Network backup parameters are displayed (not shown here).
2. Select the TSM radio button. The TSM parameters will be displayed on the panel, as illustrated below for the Data Archive panel (they are identical for the System Backup panel):



3. In the Password box, enter the TSM password that this SQL Guard unit uses to request TSM services, and re-enter it in the Re-enter Password box.
4. Optionally enter a Server name matching a *servername* entry in your *dsm.sys* file.
5. Optionally enter an As Host name.
6. Click the Apply button to save the configuration. When you click the Apply button, the system attempts to verify the TSM destination by sending a test file to

the server using the **dsmc archive** command. If the operation fails, you will be informed and the configuration will not be saved

7. Return to the general instructions for Archiving or System Backup.

Stopping Archiving

To stop the archiving of data:

1. Open the Administration Console panel (not shown).
2. In the Data Management section of the Administration Console menu, click Data Archive to open the Data Archive panel (see above).
3. Clear the *Archive data older than...* checkbox.
4. Click Apply.

Restoring

As described previously, archives are written to an SCP or FTP host, or to a Centera or TSM storage system (see [Archiving Data](#), above). To restore archives, you must copy the appropriate file(s) back to the SQL Guard unit on which the data is to be restored. There is a separate file for each day of data. Depending on how your archive/purge operation is configured, you may have multiple copies of data archived for the same day. Archive and export data file names have the same format:

`<daysequence>-<hostname.domain>-w<run_datestamp>-d<data_date>.dbdump.enc`

For example:

`732423-g1.guardium.com-w20050425.040042-d2005-04-22.dbdump.enc`

The date of the data contained on the file (with the date being written in *yyyy-mm-dd* format) is *data_date*. The date is located near the end of the file name (just before *.dbdump.enc*). Do not confuse the *data_date* with the *run_datestamp*, which appears earlier in the file name, and is the date that the data was archived or exported.

Unless you are restoring data from the first archive created during the month, you will need to restore multiple days of data. That is because when restoring data, SQL Guard needs to have all of the information that it had when the data being restored was archived. After the archive was created, some of that information may have been purged due to a lack of use. All information needed for a restore operation is archived automatically, the first time that data is archived each month. So, when restoring data, you must always start with the first day of data archived for the month containing the data you want to restore.

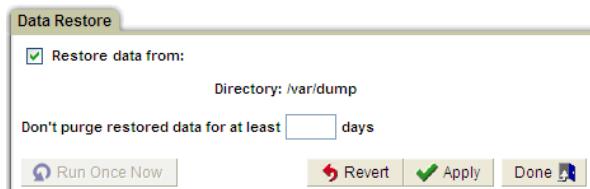
To restore archives:

1. From the *cli*, use a separate *import file* command to copy each archived data file to be restored to the SQL Guard unit.

Note: The archive and restore operations depend on the file names generated during the archiving process. **DO NOT** change the names of archived files. If a generated file name is changed, the restore operation will not work.

For more information on using the *import file* command, see the description of the [import file Command](#) in Chapter 6.

2. Open the Administration Console panel (not shown).
3. In the Data Management section of the Administration Console menu, click Data Restore to open the Data Restore panel:



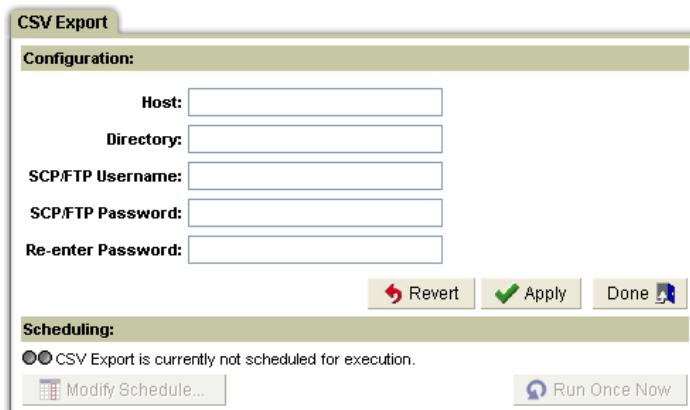
4. Mark the *Restore data from* box to activate the data restore function.
5. In the Directory box, enter the directory name to which the archived data files have been copied. The *import file* command copies all files with an *enc* extension to the /var/dump directory.
6. Optionally specify in the *Don't purge restored data for at least n days* box a number of days to “protect” restored data from any purge operations. Otherwise, the restored data may be purged the next time a purge operation runs.
7. Click Apply to save the configuration.
8. Click Run Once Now to run the restore operation.

Exporting CSV Files

Prior to release 5.0, the data archiving operation transferred CSV files (as well as archived data) to a single destination system and directory. Beginning with release 5.0, the transfer of CSV files and archived data has been separated.

To export CSV files, follow the procedure outlined below.

1. Open the Administration Console panel (not shown).
2. Click CSV Export in the Data Management section of the Administration Console menu to open the CSV Export panel:



3. In the Host box, enter the IP address or DNS host name of the host to receive the CSV files.
4. In the Directory box, identify the directory in which the data is to be stored. How you specify this depends on whether the file transfer method used is FTP or SCP. If you are unsure which file transfer method has been configured, use the [show transfer-method](#) CLI command (described in Chapter 6).

For FTP: Specify the directory relative to the FTP account home directory.

For SCP: Specify the directory as an absolute path.

5. In the Username box, enter the user name to use for logging onto the host machine. This user must have write/execute permissions for the directory specified in the Directory box (above).
6. In the Password box, enter the password for the above user, and enter the password again in the Re-enter Password box.
7. Click the Apply button to save the configuration.

When you click the Apply button, the system attempts to verify the specified Host, Directory, Username, and Password by sending a test data file to that location. If the operation fails, the following message is displayed and the configuration will not be saved:

A test data file could not be sent to this host with the parameters given. Please confirm the hostname or IP address is entered correctly, the host is online, the target directory exists and can be written to by the user given, and the password given is correct for that user.

If the Apply operation succeeds, the buttons in the Scheduling panel will become active.

8. Click the Run Once Now button to run the operation once.

9. Click the Modify Schedule button to schedule this operation to run on a regular basis. The general-purpose task scheduler is opened. For details on using the general-purpose task scheduler, see [Using the Task Scheduler](#) in Chapter 2.

To verify that CSV files have been exported, check the Aggregation/Archive report on the SQL Guard Activity Monitor tab. There should be a Send activity for each CSV file exported.

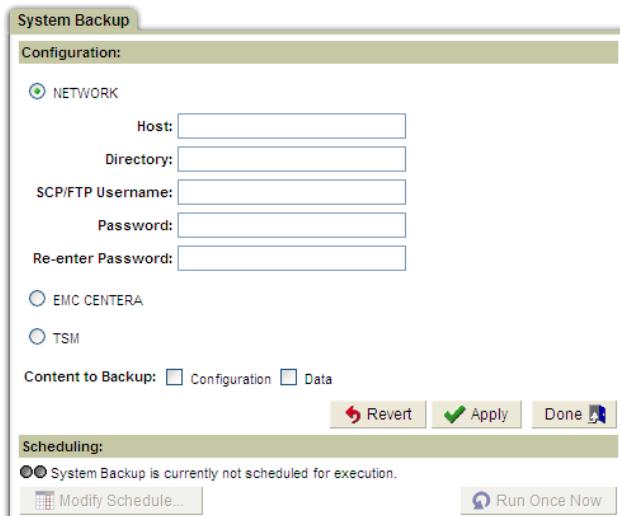
System Backup

The System Backup command allows you to define a backup operation that can be run on demand or on a scheduled basis. All SQL Guard configuration information and data is written to a single encrypted file and sent to the specified destination, using either SCP or FTP, depending on the [transfer-method](#) CLI command setting. Unlike the CLI [backup system](#) command, this backup operation does *not* stop web services or inspection engines.

To restore backed up information, use the **restore system** CLI command. For more information about restoring data, see [backup and restore Commands](#) in Chapter 6.

To backup system information:

1. Open the Administration Console panel (not shown).
2. Click System Backup in the Data Management section of the Administration Console menu to open the System Backup panel:



Depending on the archive options configured for your system (using the **store system-storage** CLI command), you may have EMC Centera or TSM options on

your panel, as illustrated above. If you select one of those archive destinations, see the appropriate topic:

- [EMC Centera Archive and Backup](#)
- [TSM Archive and Backup](#)

Steps 3-6 apply for a Network backup only.

3. In the Host box, enter the IP address or DNS host name of the host to receive the system backup file.
4. In the Directory box, identify the directory in which the data is to be stored. How you specify this depends on whether the file transfer method used is FTP or SCP. If you are unsure which file transfer method has been configured, use the [show transfer-method](#) CLI command (described in Chapter 6).

For FTP: Specify the directory relative to the FTP account home directory.

For SCP: Specify the directory as an absolute path.

5. In the Username box, enter the user name to use for logging onto the host machine. This user must have write/execute permissions for the directory specified in the Directory box (above).
6. In the Password box, enter the password for the above user, and enter the password again in the Re-enter Password box.
7. In the Content to Backup row, mark one or both checkboxes:
 - Mark the Configuration checkbox to back up all definitions.
 - Mark the Data checkbox to back up all data.
8. Click the Apply button to save the configuration.

When you click the Apply button, the system attempts to verify the specified Host, Directory, Username, and Password by sending a test data file to that location. If the operation fails, the following message is displayed and the configuration will not be saved:

A test data file could not be sent to this host with the parameters given. Please confirm the hostname or IP address is entered correctly, the host is online, the target directory exists and can be written to by the user given, and the password given is correct for that user.

If the Apply operation succeeds, the buttons in the Scheduling panel will become active.

9. Click the Run Once Now button to run the operation once.

10. Click the Modify Schedule button to schedule this operation to run on a regular basis. The general-purpose task scheduler is opened. For details on using the general-purpose task scheduler, see [Using the Task Scheduler](#) in Chapter 2.

Central Management

Before using central management, be sure that you understand what happens where in a central management configuration. This is especially important when placing existing, standalone systems under central management.

About Central Management

In a central management configuration, one SQL Guard unit is designated as the Central Manager. That unit can be used to monitor and control other SQL Guard units, which are referred to as *managed* units. Unmanaged units are referred to as *standalone* units.

The following table identifies which components are taken from which location in a central management environment.

SQL Guard Component Sources

Central Manager	Managed Unit
Users*	System Configuration
Security Roles*	Inspection Engines
Application Role Permissions	Alerter (configuration)
Queries	Anomaly Detection
Reports	Session Inference
Time Periods	IP-to-Hostname Aliasing
Alerts	System Backup
Security Assessments	Aggregation / Archiving
Audit Processes* (definitions)	Custom Assessment Tests
Audit Process Results*	Custom Alerting
To-Do Lists*	Custom Identification Procedures
Privacy Sets	Exported CSV Output
Baselines	Schedules
Policies	DB Auto-discovery Configurations
Groups*	
Aliases	

* These elements are exported from the Central Manager to all managed units on a scheduled basis, as described later.

From the Central Manager, the SQL Guard administrator can:

- Register SQL Guard units for management
- Monitor managed units (unit availability, inspection engine status, etc.)
- View system log files (syslogs) of managed units
- View reports using data on managed units
- View main statistics for managed units
- Install SQL Guard security policies on managed units
- Restart managed units

- Manage SQL Guard inspection engines on managed units
- Maintain the complete set of Users, Security Roles, Groups, and Application Role Permissions used on all managed systems

Note: Application Role Permissions can also be changed by the SQL Guard administrator from any managed unit. When this happens, the permissions are changed for all managed units.

Users, Roles, and Groups under Central Management

The Central Manager controls the definition of users, roles, and groups for all managed systems by exporting the Central Manager's *complete set* of SQL Guard user, security role, and group definitions on a scheduled basis or on demand (see [Synchronizing Portal User Accounts](#)). The managed units update their internal databases on an hourly basis, which means that there may be a delay of up to an hour between the time that a managed unit receives updates and the time that the managed unit applies those updates.

New users must log onto the Central Manager before logging onto a managed unit.

Notes: If you have SQL Guard users or security roles defined on an existing standalone unit that is about to be registered for central management, those definitions will not be available after the system is registered, unless those users and security roles have also been defined on the Central Manager.

You cannot administer users or security roles on a managed unit. Those definitions can only be administered when logged on to the Central Manager.

When a unit is unregistered for central management, the users and security roles that were backed up when the unit was registered are restored.

When installing an Accelerator add-in product (PCI, SOX, etc.), in a Central Manager environment, install it first on the Central Manager and then on the managed unit. Add any roles and users as required for the Accelerator on the Central Manager (and those will be synchronized with the managed unit from there). See your Accelerator documentation for more information.

Aliases and Groups under Central Management

On all processes that automatically generate aliases or groups, for example: import user groups from LDAP, group generation from queries, alias generation from queries,

classifier, etc. if the same group or alias is automatically generated on more than one managed machine, (managed by the same manager) then it may conflict with an existing group or alias, which will not be replaced.

Audit Processes under Central Management

All Audit Process definitions are stored on the Central Manager, but all Audit Process Results and User To-Do Lists are stored on the managed units. Audit Process definitions are exported from the Central Manager to the managed units as part of the user synchronization process (see [Synchronizing Portal User Accounts](#)). When audit process results have been produced, the results will be available to users, but on managed units, there may be a delay of up to an hour before reports or monitors such as Outstanding Audit Process Reviews are updated.

Central Manager Reports Using Data from Managed Units

From the Central Manager, reports and audit processes can use data from a managed unit. The managed unit is selected as a run-time parameter, and is referred to as a remote datasource. When an audit process references a remote datasource, that audit process can be run from the Central Manager only, so it will not appear in a list of audit processes displayed on a managed unit.

If a report on a pane of the Central Manager portal contains data from a remote datasource, and the managed unit becomes unavailable (due to a network outage, for example), the pane on which the report resides cannot be refreshed, which means that other reports on the same pane may not be displayed, even though data for those reports may be available. For this reason, when using remote datasources for a report, it is best to use a menu layout, with one report per menu entry, so that the unavailability of one remote source does not prevent any other reports from being displayed.

Non-Central Manager Tasks

When a server is configured as a Central Manager, you must be aware of the tasks that cannot be performed on that unit, but rather must be performed on other (non-Central Manager) units. This includes the following:

- Inspection engines cannot be defined on the Central Manager.
- Load Balancing cannot be performed from the Central Manager.

Upgrade Considerations for Version 6.0 or Later

If you have upgraded your Central Manager to version 6.0, but have not yet upgraded all managed units, there are some functions that will not work properly until the managed unit has been upgraded to the same version. All known issues at the time this document was prepared are described below. See the SQL Guard Version 6.0 *Upgrade Guide* and *Release*

Notes for more information. All of these issues apply to the case where the Central Manager has been upgraded to 6.0, and the managed unit is at release 5.0 or later.

- Before any pre-6.0 managed unit can work with a 6.0 Central Manager, the managed unit must have a pre-upgrade patch applied. This is not the full 6.0 upgrade, but rather just what needs to be updated for the managed unit to function with the 6.0 Central Manager.
- Until the managed unit is upgraded to 6.0, all new reports, alerts, policies and audit processes must be defined on the Central Manager.
- For an audit process defined on the Central Manager, you cannot use a pre-6.0 managed unit as a remote data source (the remote data source feature is new with version 6.0).
- From the Central Manager, you cannot view the installed policy on a managed unit that has not been upgraded to version 6.0.
- For a pre-6.0 managed unit managed by a 6.0 Central Manager, you cannot view the Policy Violations report on the managed unit.
- For a pre-6.0 managed unit managed by a 6.0 Central Manager, you cannot view the Logged Alerts on the managed unit.
- You cannot create a new audit process on a pre-6.0 managed unit managed by a 6.0 Central Manager.
- You cannot view the Aggregation/Archive Log on a pre-6.0 managed unit managed by a 6.0 Central Manager.
- You cannot create a new statistical alert on a pre-6.0 managed unit managed by a 6.0 Central Manager.
- You cannot open the Access Trace query builder on a pre-6.0 managed unit managed by a 6.0 Central Manager (this component was obsolete in version 5, but was still present).

Implementing Central Management

In a new SQL Guard installation, implementing central management is a straight-forward process. In an existing SQL Guard environment, conversion to central management can be more complicated if you want to preserve components (reports, policies, etc.) that have been defined on standalone units. The following sections provide general guidelines for implementing central management in both situations.

Implementing Central Management in a New Installation

In a new installation, follow the procedure outlined below:

1. Select and make note of the System Shared Secret that will be used by the Central Manager and all managed units. See the [System Configuration Panel Reference](#) for more information about the System Shared Secret.
2. Install the Central Manager unit. See [Chapter 1: Installation](#).
3. From a command-line window, log into the unit as the *cli* user and use the [store unit type](#) command to set the *manager* attribute for the Central Manager.
4. Repeat for each managed unit:
 - o Install the managed unit. See [Chapter 1: Installation](#).
 - o Register the managed unit for central management. See [Registering Units for Central Management](#), below.

Note: To avoid confusion, do not define anything (reports, users, policies, etc.) on a managed unit until *after* it has been registered for central management.

Implementing Central Management in an Existing Installation

In an existing SQL Guard environment, refer to the procedure outlined below to develop a plan for implementing central management. If you are converting an existing SQL Guard unit to a Central Manager, keep in mind that a Central Manager can *not* monitor network traffic (i.e., inspection engines cannot be defined on a Central Manager).

1. Select a System Shared Secret to be used by the Central Manager and all managed units. See the [System Configuration Panel Reference](#) for more information about the System Shared Secret.
2. Install the Central Manager unit or designate one of the existing systems as the Central Manager. In either case, use the [store unit type](#) command to set the *manager* attribute for the Central Manager.
3. Any definitions from the standalone unit that you want to have available in the central management environment will have to be exported before the standalone unit is registered for management. Later, those definitions will be imported on the Central Manager. **BEFORE** exporting or importing any definitions, follow the procedure outlined below for each standalone unit that is to become a managed unit, and read through the introductory information under [Exporting and Importing Definitions](#).
 - o Decide which users, security roles, queries, reports, groups, time periods, alerts, security assessments, audit processes, privacy sets, baselines, policies,

and aliases from the standalone system you want to have available after the system becomes a managed unit. For the remainder of this discussion, ignore any components on the standalone system you do not want to have available.

- Compare the security roles and groups defined on the standalone unit with those defined on the Central Manager. Under central management, a single version of these definitions applies to all units.
- If a security role with the same name exists on both systems and it is used for different purposes, add a new role on the Central Manager and assign the new role to the appropriate definitions after they are imported.
- If the same group name exists on the standalone unit and the Central Manager but it has different members, create a new duplicate group on the standalone system, taking care to select a group name that does not exist on the Central Manager. In all of the definitions to be exported, change the old group name references to new group name references.
- Note all security roles assigned to all definitions that will be exported from the standalone system. When definitions are imported, they are imported **WITHOUT** roles, so you have to add them manually.
- Check the application role permissions on each system. If any security roles assigned to an application on the standalone unit are missing from the Central Manager, add them to the Central Manager.
- Export all queries, reports, groups, time periods, alerts, security assessments, audit processes, privacy sets, baselines, policies, and aliases from the standalone system that you want to have available after the system becomes a managed unit. (See [Exporting and Importing Definitions](#), later in this chapter.) **Do not** export users or security roles. If you are unsure about a definition, export it in a separate export operation so that you can decide later whether or not to import that definition to the Central Manager. Once you register for central management, none of the old definitions from the standalone unit are available.
- On the standalone unit, if there are any audit process *results* that you want to view in the future, create PDF versions of those results and store them in an appropriate location. Under central management, only the audit results produced under central management are available.
- On the standalone unit, instruct all users to remove all portlets containing custom report, and to *not* create any new reports until the conversion to central management is complete.
- On the Central Manager, manually add all users from the standalone unit.

- On the standalone unit, delete all user definitions except for the *admin* user (which cannot be deleted).
- Register the standalone unit for central management. See [Registering Units for Central Management](#), below.
- On the Central Manager, import all definitions exported from the standalone system. Check to make sure that references to included items (receivers in alert notifications, for example) are correct.
- Re-assign security roles, as necessary, to all imported definitions.
- Inform users of the managed unit that they must use the Report Builder application to re-generate the portlets for any custom reports they want to display in their layouts.

If the Central Management Unit is Unavailable

If the Central Manager is unavailable to a managed unit (due to a network or system failure, for example), a message is displayed prominently in the SQL Guard management interface window when you log into the managed unit:



You will be able to perform a very limited number of functions on that system, since most functions rely on the internal database stored on the Central Manager.

Registering Units for Central Management

You can register SQL Guard units for central management either from the Central Manager or from the unit itself. Regardless of how the registration is done, the Central Manager and all managed units must have the same System Shared Secret. For more information, see the description of the System Shared Secret under [Changing the System Configuration](#) above.

If the unit to be managed is already registered for central management with another manager, un-register that unit from that manager before registering it with the new manager.

Each procedure for registering is described separately below. Before registering a unit, be sure that you understand exactly what happens to that unit when it is registered and unregistered for central management (see below).

About Central Manager Licenses

The Central Manager license limits the number of units that can be managed. If you attempt to register more units than are permitted, the operation will not be allowed.

When a unit is unregistered, you should *always* perform that function from the Central Manager. This is the only way that the Central Manager reduces its count of managed units. You can unregister from the managed unit, but that capability is intended for emergency use only (for example, if the Central Manager becomes unavailable). If you unregister only from the managed unit, the Central Manager will still count that unit as a managed unit for licensing purposes, and you may not be able to register another unit with the Central Manager.

What Happens During Registration and Unregistration

When you *register* a unit for central management, the system makes a pre-registration backup of the registered unit's configuration. The backup includes all definition data on that machine: queries, reports, users, etc. – everything except actual logging data.

While registered and under central management, local definitions of users and roles are used to control access, but these definitions cannot be modified while logged on to the managed unit. All other definitions and components are taken directly from the Central Manager, except for custom assessment tests, custom alerting classes, and custom identification procedures, as noted previously.

If a security policy is installed on the managed unit, it is stored in the SQL Guard database on that unit, but the definition of that security policy is not available, except on the Central Manager.

When you *unregister* a unit from central management, the unregister process restores the configuration for that unit from the pre-registration backup. This means that any changes made to the configuration of this unit from the Central Manager (the definition of new users or the installation of a security policy, for example) will be overwritten by the pre-registration configuration during the unregister process.

Caution: When unregistering a unit, if the pre-registration backup was created under a previous release of the Guardium software, restoring that configuration without first applying a patch to bring it to the current software release level will disable the unit, potentially causing the loss of all data stored there. Accordingly, ***do not unregister a unit*** until you have verified that the pre-registration configuration is at the current software release level. If you are unsure about how to verify this, contact Guardium Support ***before unregistering the unit***.

Registering a Unit from the Central Manager

To register a unit for central management from the Central Manager, follow the procedure outlined below. The unit to be managed does not have to be online when it is registered (see the last step of the procedure for more information), but as mentioned earlier it must have the same System Shared Secret as the Central Manager.

- Select Central Management from Central Management section of the Administration Console menu to open the Central Management panel:

	Unit Name	Online	Inspection Engines	Installed Security Policy	SQL Guard Model	SQL Guard Version	Last Ping Time
<input type="checkbox"/>	g1.guardium.com			text7 installed at 2007-04-05 08:33:05	G2000	6.0	4/6/07 10:14 AM
<input type="checkbox"/>	x4.guardium.com			LogFullDetailsEveryThing2	G2000	6.0	4/6/07 10:14 AM
<input type="checkbox"/>	g2.guardium.com				G2000	6.0	4/6/07 10:14 AM

Buttons at the bottom: Select All, Unselect All, Unregister, Install Policy, Back, Register New..., Show Distributed Map, Distributed Monitor, Done.

- Click the Register New button to open the Unit Registration panel.
- Enter in the Unit IP box the IP address of the SQL Guard unit to be managed.

Note: If the unit you specify is already managed by another Central Manager, you will get an error message and the registration will fail. (You can unregister that unit from the other Central Manager, or directly from that unit.)

- Enter the port number configured for the SQL Guard software on that unit.
- Click the Save button. What happens next depends on whether or not the specified unit can be accessed by the Central Manager.
 - If the specified unit can be accessed by the Central Manager, the registration will occur immediately.
 - If the specified unit *cannot* be accessed, the Central Manager will list the unit as a managed (but offline) unit and it will continue attempting to access and register that unit for a maximum of seven days.

Regardless of what happens, you will be informed of the result of the operation.

If the Registered Unit Status Remains Offline

If you know the unit just registered is online and accessible from the Central Manager, but its status in the Central Management panel remains *offline*:

- Verify that the unit to be managed is online, accessible, and operational by using a browser window to log in to the SQL Guard system on that unit.
- In the Central Management panel, click the refresh button for the unit:

- Check that you have entered the correct IP address for the unit.
- Check that the unit has the same shared secret as the Central Manager.

Unregistering a Unit from the Central Manager

To unregister a managed unit from the Central Manager:

1. Select Central Management from the Administration Console menu to open the Central Management panel:



2. Mark the checkbox for the managed unit you want to unregister.
3. Click the Unregister button. You will be prompted to confirm the action.

Registering from a Managed Unit

You can register a unit either from the Central Manager or the managed unit. On a managed unit, you can also use the CLI *register* command to register the unit (see [register / unregister commands](#) in Chapter 6).

To register for central management from a managed unit:

1. Select *Central Mgmt. Registration* from the Administration Console menu to open the Central Management Registration panel:



2. In the *Central Management Host Ip* box, enter the IP Address of the SQL Guard unit from which this unit will be managed.

3. In the Port box, enter the port number for the Central Management unit.
4. Click the Register button.

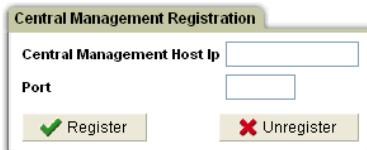
Note: The central management unit must be online and accessible by this unit when you register for central management. In contrast, when you register units for management from the central management unit, you can register units that are not currently accessible.

Unregistering from a Managed Unit

When a unit is unregistered, you should *always* perform that function from the Central Manager. This is the *only* way that the Central Manager decrements its count of managed units. You can unregister from the managed unit, but this capability is provided for emergency use only, for example if the Central Manager becomes unavailable. If you unregister only from the managed unit, the Central Manager will still count that unit as a managed unit for licensing purposes, and you may not be able to register another unit with the Central Manager.

To unregister from a managed unit:

1. Select *Central Mgmt. Registration* from the Administration Console menu to open the Central Management Registration panel:



2. Click the Unregister button. You will be prompted to confirm the action.

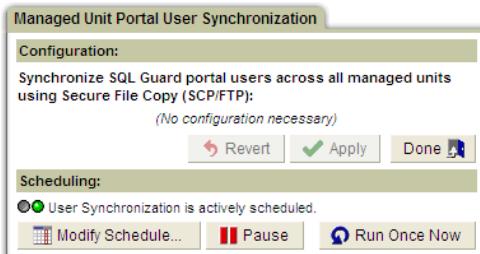
On a managed unit, you can also use the CLI *unregister* command to unregister the unit (see [register / unregister commands](#) in Chapter 6).

Synchronizing Portal User Accounts

As mentioned earlier, the Central Manager controls the definition of Users, Security Roles, and Groups for all managed units. It does this by making an encrypted and signed copy of its complete set User, Security Role, and Group definitions, and transmitting that information to all managed units. The managed units then update their internal databases on an hourly basis, which means that there may be a delay of up to an hour between the time that the managed unit receives updates and the time that the managed unit applies those updates.

To manage portal user synchronization:

1. Select *Portal User Sync.* from the Administration Console menu to open the Managed Unit Portal User Synchronization panel:



2. Do one of the following:
 - Click **Modify Schedule** to change the user synchronization task schedule using the standard task scheduler (see [Using the Task Scheduler](#), later in this chapter).
 - If the task is actively scheduled, click **Pause** to stop further scheduled executions.
 - If the task is paused, click **Resume** to start running the task again (according to the defined schedule).
 - Click **Run Once Now** to run the synchronization task immediately.

Note: The task being scheduled or “run once now” refers to the collection of data and its transmission to the managed units only – the managed units may not use that data to update their user tables until up to one hour after it has been received.

Monitoring Managed Units

To monitor managed units:

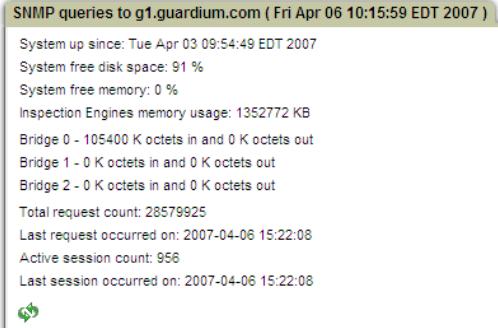
1. Select Central Management from the Administration Console menu to open the Central Management panel:



Each component of the Central Management panel is described in the table below.

Central Management Panel

Control	Description
<input type="checkbox"/>	Mark this box to select the unit for an unregister or policy installation operation (see the button descriptions below).
	Refresh Unit Info – Refreshes all information displayed in the expanded view of that unit by issuing new requests to that unit.
	Reboot Unit – Reboots the unit at the operating-system level. By default, the SQL Guard portal is started at startup.
	Restart Unit Portal – Restarts the SQL Guard application portal on the managed unit. You can then log in to that unit to perform SQL Guard tasks that must be performed on that unit (defining or removing inspection engines, for example).
	View Unit SNMP Attributes – Opens the SNMP Viewer panel (illustrated below) in a separate window.

Control	Description
	 <p>SNMP queries to g1.guardium.com (Fri Apr 06 10:15:59 EDT 2007)</p> <p>System up since: Tue Apr 03 09:54:49 EDT 2007 System free disk space: 91 % System free memory: 0 % Inspection Engines memory usage: 1352772 KB Bridge 0 - 105400 K octets in and 0 K octets out Bridge 1 - 0 K octets in and 0 K octets out Bridge 2 - 0 K octets in and 0 K octets out Total request count: 28579925 Last request occurred on: 2007-04-06 15:22:08 Active session count: 956 Last session occurred on: 2007-04-06 15:22:08</p>

Click the  (Refresh) button in the lower left corner of the panel to refresh the data in the window.



View Unit Syslog – Opens the Syslog Viewer (illustrated below) in a separate window, displaying the last 64KB of syslog messages.

Syslog of g1.guardium.com
 <pre> Apr 6 01:32:39 g1 sshd[32767]: Did not receive identification string from UNKNOWN Apr 6 05:36:54 g1 sshd[2490]: Connection closed by 192.168.1.8 Apr 6 01:37:45 g1 sshd[2973]: Did not receive identification string from UNKNOWN Apr 6 01:37:45 g1 sshd[2974]: Did not receive identification string from UNKNOWN Apr 6 01:37:45 g1 sshd[2976]: Did not receive identification string from UNKNOWN Apr 6 01:37:45 g1 sshd[2978]: Did not receive identification string from UNKNOWN Apr 6 01:37:45 g1 sshd[2977]: Did not receive identification string from UNKNOWN Apr 6 01:38:32 g1 sshd[3402]: Did not receive identification string from 192.168.1.8 Apr 6 01:42:50 g1 sshd[5726]: Did not receive identification string from UNKNOWN Apr 6 01:42:50 g1 sshd[5727]: Did not receive identification string from UNKNOWN Apr 6 01:42:50 g1 sshd[5728]: Did not receive identification string from UNKNOWN Apr 6 01:42:50 g1 sshd[5730]: Did not receive identification string from UNKNOWN Apr 6 01:42:50 g1 sshd[5729]: Did not receive identification string from UNKNOWN Apr 6 05:47:16 g1 sshd[8021]: Connection closed by 192.168.1.8 Apr 6 01:47:56 g1 sshd[8410]: Did not receive identification string from UNKNOWN Apr 6 01:47:56 g1 sshd[8411]: Did not receive identification string from UNKNOWN Apr 6 01:47:56 g1 sshd[8413]: Did not receive identification string from UNKNOWN Apr 6 01:47:56 g1 sshd[8414]: Did not receive identification string from UNKNOWN Apr 6 01:47:56 g1 sshd[8412]: Did not receive identification string from UNKNOWN Apr 6 01:48:58 g1 sshd[8928]: Did not receive identification string from 192.168.1.8 </pre>

Click the  (Refresh) button in the lower left corner of the panel to refresh the data in the window.



Shortcut to Unit Portal – Opens the SQL Guard login page for the managed unit, in a separate browser window.

Control	Description
Unit Name	The host name of the managed unit. If you hold the mouse pointer over the unit name, its IP address displays as a tool tip. Note: If the hostname changes on the unit, the Central Manager will no longer see that unit when automatically refreshing the Online status. If you suspect the hostname has changed, use the Refresh button on the toolbar (described above) to obtain the changed hostname and update the displayed current Online status and other information for that unit.
Online	Indicates whether or not the unit is online. If the green indicator is lit, the unit is online; if the red indicator is lit, the unit is offline. The Central Manager refreshes this status at the refresh interval specified in the central management configuration (one minute by default). Note: If an error occurred connecting to a unit, the error description can be viewed as a tool tip when you hover the mouse indicator over that unit's record in the management table.

Control	Description
Inspection Engines	Click the (plus) button to expand the list of inspection engines; click the (minus) button to hide the list of inspection engines. The information displayed for each inspection engine is described below. Note: This information is fetched from the managed unit when the Refresh button is pressed, not on every ping.
Name	The name of the inspection engine.
Protocol	The protocol monitored by the inspection engine: Oracle, MSSQL, or Sybase, Informix, or DB2.
Active on Startup	Indicates if the inspection engine starts on system startup.
Exclude From-IP	Indicates if the list of from-IP addresses is to be excluded (i.e., not examined).
From-IP/Mask	A list of the IP addresses and subnet masks of the clients whose database traffic to the To-IP/Mask addresses (see below) the inspection engine monitors.
Ports	The ports on which database clients and servers communicate; can be a single port, a list of ports, or a range of ports.
To-IP/Mask	A list of IP addresses and subnet masks of servers whose traffic from the corresponding client machine (see From-IP/Mask, above) is monitored.
Installed Security Policy	The name of the security policy installed on the managed unit. This field is updated on every ping.
SqlGuard Model	The SQL Guard model number of the managed unit.
SqlGuard Version	The SQL Guard version number of the managed unit.
Last Ping Time	The last time that the unit was pinged by the Central Manager to determine the managed unit's online/offline status.
Select All	Selects all managed units.
Unselect All	Unselects all managed units.
Unregister	Registers all selected units. See the topic above for more information on the implications of unregistering a unit.
Install Policy	Opens the Install Security Policy panel, to install a security policy on all selected units. See Installing Security Policies on Managed Units , below.

Control	Description
Back	Closes the panel and returns to the Administrator Console.
Register Now	Opens the Unit Registration panel to register a new unit for management.
Show Distributed Map	Displays a map of the Central Manager unit and all managed units. See Viewing Management Maps , below.
Distributed Monitor	Opens the Distributed Monitoring of Managed Nodes report in a separate window (not shown).
Done	Closes the panel and returns to the Administrator Console.

Installing Security Policies on Managed Units

To install a security policy on a managed unit:

1. Select Central Management from the Administration Console menu to open the Install Security Policy panel:



2. Select each unit on which you want to install the same security policy. To select a unit, mark the checkbox in the first column of the row for that unit.
3. Click the Install Policy button to open the Install Security Policy panel:



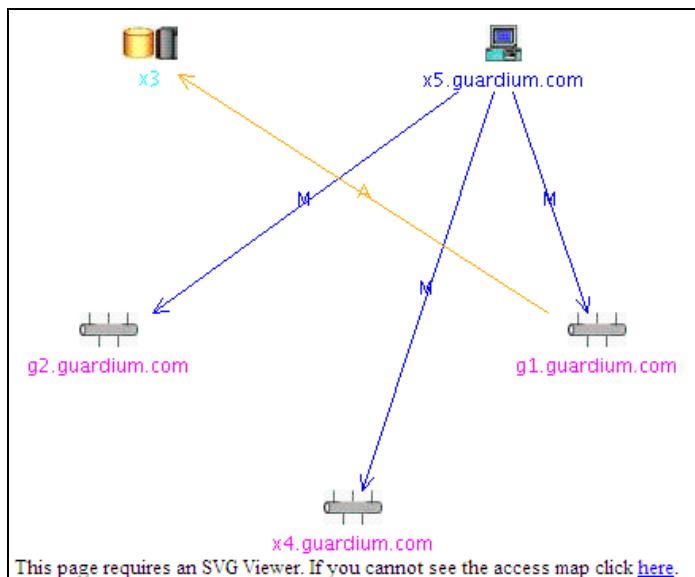
4. From the Policy list, select the policy you want to install.
5. Click the Install Policy button. You will be informed of the success (or failure) of each policy installation. If a selected unit is not available (it may be offline or a link may be down), the Central Manager will inform you of that fact. It will continue attempting to install the new policy for a maximum of seven days (as long as that unit remains registered for central management).

Viewing Management Maps

To view management maps, you need the Adobe SVG Viewer. See [Software Downloads from Adobe](#). If a map does not display as expected, or does not display at all, you may need to update your version of the SVG Viewer. Use the link above to check your SVG Viewer version.

To view a map showing all managed units:

1. Select Central Management from the Administration Console menu.
2. Click the Show Distributed Map button to display a map of the central manager unit and all managed units:



The following table describes the symbols used in the map.

Distributed Map Symbols

Symbol	Description
--------	-------------

Symbol	Description
 guy	The Central Manager unit, labeled with its hostname.
 unit_4	A managed unit, labeled with its hostname.
 topAggregator	An aggregator unit, labeled with its hostname.
 M	A blue arrow labeled with the letter M connects the Central Manager unit with all managed units (which are not also aggregation units).
 A	Yellow arrows labeled with the letter A connect aggregation units with the units being aggregated (unless the unit is also a managed unit). The arrow indicates the direction of aggregation.
 A/M	Green arrows labeled with the letters A/M relate managed aggregation units to the Central Manager unit. The arrows indicate the direction of aggregation (and may be included on both ends if the Central Manager unit is also an aggregation unit).

Using S-Taps

The following section describes how to manage S-Taps using the SQL Guard management interface.

- For an S-Tap overview, see [S-Tap Overview](#) in Chapter 1.
- For information about installing S-Tap, see [Step 4: S-Tap Installation](#), also in Chapter 1.
- For information about the report templates pertaining to S-Tap, see the *SQL Guard User Guide*.

Configuring S-Tap

Users configure the S-Tap application to:

- Detect traffic between specific database servers and clients, over specific ports.
- Forward that traffic to a specific SQL Guard server.
- Designate one or more alternative SQL Guard servers to receive that information, in the event that the primary SQL Guard server is not available.

Configuration properties can be changed in two ways:

- Locally on the database server, using a text editor to modify the S-Tap configuration file. Although it is not recommended, any property in the configuration file can be edited using this method.
- From the Administration Console when logged in to the *active* SQL Guard host for the S-Tap, when the S-Tap is online. If the active SQL Guard host is managed by the Central Manager, you still must log into the active SQL Guard host directly, as S-Tap control on a managed unit is not available from the Central Manager. See below for further information.

Managing S-Tap from the Administration Console

From the S-Tap Control Panel of the Administration Console, you can display configuration information for up to 64 S-Taps. If the SQL Guard unit you are logged into is currently the *active* host for an S-Tap, and that S-Tap is online, you can also start and stop, or modify the configuration of that S-Tap.

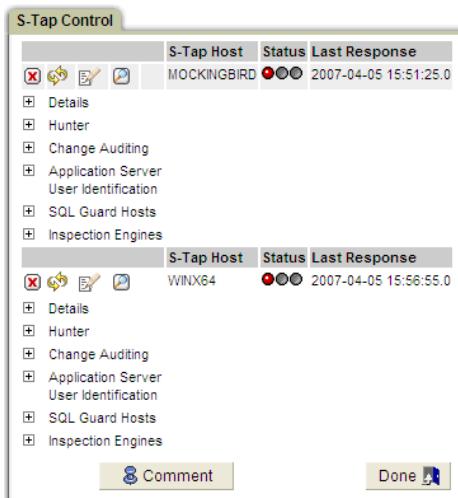
The database access information forwarded by S-Tap is essentially the same information that the SQL Guard server captures, so all of the monitoring, alerting, and reporting capabilities available in SQL Guard are available for S-Tap data.

The remainder of this section contains the following topics:

- **Displaying S-Tap Information**
Describes how to view information about S-Taps.
- **Managing S-Taps**
Describes how add, modify, or remove inspection engines on the S-Tap using the S-Tap Configuration panel.

Displaying S-Tap Information

Click S-Tap Control in the Local Taps section of the Administration Console menu (not shown) to open the S-Tap Control panel:



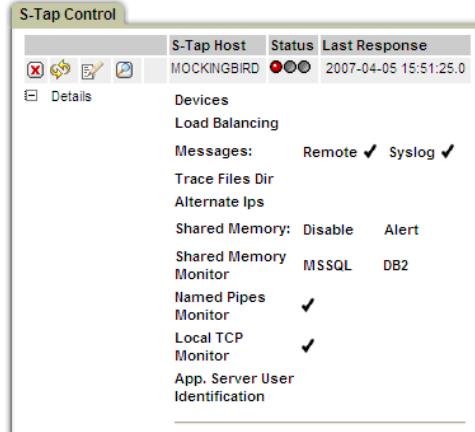
If there is no Local Taps section in the Administration Console menu, the SQL Guard unit you are logged into has not been configured with the S-Tap unit type. (See the [store unit type](#) command description in Chapter 6.)

The following table describes the columns and controls of the S-Tap Control Panel.

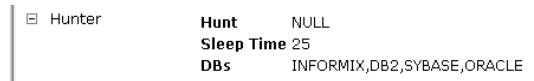
S-Tap Control Panel Reference

Column or Control	Description
(Remove)	Click to remove this S-Tap from the list of S-Taps displayed. Note: Clicking this button does not result in the S-Tap ceasing to send information, nor does it remove this SQL Guard host from the list of SQL Guard hosts stored in the configuration file for that S-Tap. This button is useful to “clean up” your display of information when you know that an S-Tap has become inactive or when this SQL Guard unit is no longer listed as a host in that S-Tap’s configuration file. In either of those cases, the S-Tap is displayed indefinitely with an <i>offline</i> status (since all this SQL Guard unit “knows” is that the S-Tap is unavailable).

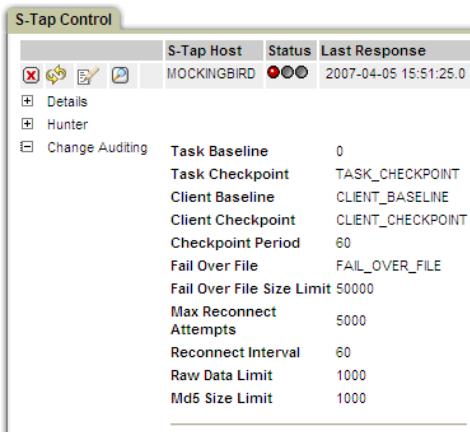
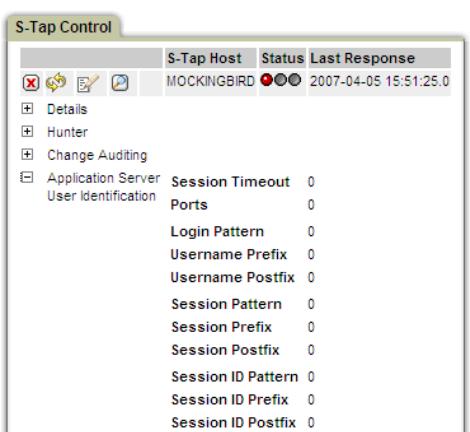
Column or Control	Description
 (Refresh S-Tap Information)	Click to refresh the display of information for this S-Tap. The S-Tap Control panel refreshes the status displayed for each S-Tap every 30 seconds. All other S-Tap information is refreshed only when the S-Tap starts, when its configuration changes, or when this button is clicked.
 (Edit S-Tap Configuration)	Enabled only for the <i>active</i> SQL Guard host (see below), and only when the S-Tap Status is <i>Green</i> (meaning it is online). Click to open the S-Tap Configuration panel in a separate window. See Managing S-Tap Inspection Engines , below.
 (Show S-Tap Event Log)	Click to open the S-Tap Events panel in a separate browser window. See Viewing the S-Tap Events Panel below.
S-Tap Host	Identifies the host on which the S-Tap is installed.
	<p>Note: An S-Tap Host with an IP address of 127.0.0.1 indicates that the local-stap property has been enabled. This feature provides automatic decoding of Kerberos-encrypted database user names in an MS SQL Server (Windows) environment. For more information, see About Kerberos-Encrypted Database User Names.</p>
Status 	One of the three lights will be illuminated:
	Green (Online) – The S-Tap is functioning normally.
	Red (Offline) – The S-Tap is not responding.
	Yellow (Not Synchronized) – Configuration changes have been sent to the S-Tap, but the S-Tap has not yet acknowledged that the changes were applied. If the light remains yellow for an extended period of time, you can assume that the S-Tap was unable to restart using the new configuration. When that happens, S-Tap attempts to restart using the last good configuration. When an error has occurred, you can open the S-Tap Events panel in a separate window by clicking the Show S-Tap Event Log button (see above). In many cases the event log will contain error messages indicating what was wrong with the new configuration.
	To reload the last good configuration from the S-Tap host, click the Refresh S-Tap information button (see above).

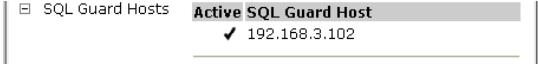
Column or Control	Description
	<p>Note: If you have trouble determining the color of the light, hold the mouse pointer over the set of lights to display the current status (Offline, Not Synchronized, or Online) as a tool tip.</p>
Last Response	Date and time of the last heartbeat from the S-Tap.
Details Pane	 <p>The screenshot shows the 'S-Tap Control' window with the 'Details' pane expanded. The top bar displays 'S-Tap Host: MOCKINGBIRD', 'Status: Offline', and 'Last Response: 2007-04-05 15:51:25.0'. The 'Details' pane lists various configuration settings:</p> <ul style="list-style-type: none"> Devices: (checkboxes for Remote and Syslog are checked) Load Balancing: (checkbox is checked) Messages: (checkboxes for Remote and Syslog are checked) Trace Files Dir: (checkbox is checked) Alternate Ips: (checkbox is checked) Shared Memory: (checkboxes for Disable and Alert are checked) Shared Memory Monitor: (checkboxes for MSSQL and DB2 are checked) Named Pipes Monitor: (checkbox is checked) Local TCP Monitor: (checkbox is checked) App. Server User Identification: (checkbox is checked)
Devices	Identifies the interfaces on which the S-Tap listens for SQL Guard server communications. May be blank for Windows servers.
Load Balancing	A checkmark indicates that this S-Tap balances traffic to all servers listed in the SQL Guard Hosts pane (see below). Load balancing is by Client IP address, since all traffic for a session between a specific client and a specific server must be viewed by the same SQL Guard unit
Messages	<p>Remote. When marked, messages are sent to the active SQL Guard host.</p> <p>Syslog. When marked, messages are sent to the syslog file.</p>
Trace Files Dir	The directory in which trace files are stored.
Alternate IPs	Additional IP address for the database server system on which the S-Tap is installed.

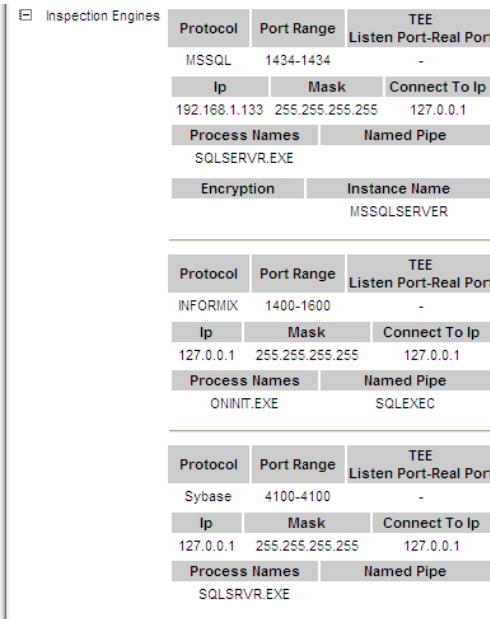
Column or Control	Description
Shared Memory	Windows only. Disable. When marked, shared memory connections are disabled whenever they are encountered. Alert. When marked, alert messages will be sent when shared memory connections are detected.
Shared Memory Monitor	Windows only. MSSQL. When marked, MSSQL shared memory connections will be monitored (and reported). DB2. When marked, DB2 shared memory connections will be monitored (and reported).
Named Pipes Monitor	Windows only. When marked, named pipes connections will be monitored (and reported).
Local TCP Monitor	Windows only. When marked, local TCP connections will be monitored (and reported).
App. Server User Identification	Indicates if application server user identification is in use.
Hunter Pane	The hunter feature is used only on Unix database servers, and only when the TEE is enabled, as opposed to the K-TAP (kernel TAP). The hunter can report on, and optionally kill, rogue processes.



Column or Control	Description												
Hunt	<p>NULL indicates that no processes will be killed. Otherwise, each process of the specified type for the specified database will be killed when detected. These are listed in the form:</p> <p><code>db_type:process[,db_type:process]</code></p> <p>With each db-type entry separated from the next by a comma. The database types are FTP, DB2, Informix, Oracle and Sybase; and the processes are any of the following:</p> <table><tbody><tr><td>SHM</td><td>Shared memory</td></tr><tr><td>IPv4</td><td>Internet Protocol version 4</td></tr><tr><td>IPv6</td><td>Internet Protocol version 6</td></tr><tr><td>FIFO</td><td>A named pipe IPC mechanism</td></tr><tr><td>PIPE</td><td>A simple (unnamed) pipe IPC</td></tr><tr><td>INET</td><td>Internet Protocol (HPUX)</td></tr></tbody></table> <p>These values are not case-sensitive.</p> <p>Example: To kill Oracle Bequeath processes, which uses a simple pipe, you would enter: oracle:pipe</p>	SHM	Shared memory	IPv4	Internet Protocol version 4	IPv6	Internet Protocol version 6	FIFO	A named pipe IPC mechanism	PIPE	A simple (unnamed) pipe IPC	INET	Internet Protocol (HPUX)
SHM	Shared memory												
IPv4	Internet Protocol version 4												
IPv6	Internet Protocol version 6												
FIFO	A named pipe IPC mechanism												
PIPE	A simple (unnamed) pipe IPC												
INET	Internet Protocol (HPUX)												
Sleep Time	The maximum number of seconds between the randomized starting time of the hunter's search routine. The start time is random to increase the difficulty of defeating it by running in fixed time slots or intervals.												
DBs	Identifies the types of databases to be reported on by the hunter: Informix, DB2, Sybase, Oracle.												

Column or Control	Description
Change Auditing Pane	Applies only when a Change Auditing System client has been installed on the database server:  The screenshot shows the S-Tap Control interface with the 'Change Auditing' section expanded. It lists various configuration parameters: Task Baseline (0), Task Checkpoint (TASK_CHECKPOINT), Client Baseline (CLIENT_BASELINE), Client Checkpoint (CLIENT_CHECKPOINT), Checkpoint Period (60), Fail Over File (FAIL_OVER_FILE), Fail Over File Size Limit (50000), Max Reconnect Attempts (5000), Reconnect Interval (60), Raw Data Limit (1000), and Md5 Size Limit (1000). <p>See the CAS configuration topic for a description of these settings.</p>
Application Server User Identification	Applies only when Application Server User Identification is installed:  The screenshot shows the S-Tap Control interface with the 'Application Server User Identification' section expanded. It lists various session-related parameters: Session Timeout (0), Ports (0), Login Pattern (0), Username Prefix (0), Username Postfix (0), Session Pattern (0), Session Prefix (0), Session Postfix (0), Session ID Pattern (0), Session ID Prefix (0), and Session ID Postfix (0). <p>See the Application Server User Identification topic for a description of these settings.</p>

Column or Control	Description
SQL Guard Hosts Pane	 <p>This pane lists all Guardium servers defined as hosts for the S-Tap. The first listed is the <i>primary</i> host, which is the first host the S-Tap attempts to connect with each time it restarts, or when the connection with the current host is lost. The <i>active</i> host is indicated by a checkmark. The <i>active</i> host is the one to which the S-Tap is currently sending data, and it is the only host from which you can edit the S-Tap configuration. Additional hosts are <i>secondary</i> hosts. If S-Tap loses its connection to the active host, and it cannot re-connect to the primary host, it will attempt to connect to a secondary host (in the order listed).</p>
Active	A checkmark in the Active column identifies the active host for the S-Tap.
SQL Guard Host	This column lists the IP address or system name for each SQL Guard unit that can function as the active host for this S-Tap.

Column or Control	Description																																																												
Inspection Engines	The layout of the Inspection Engines panel varies depending on the server operating system, the database protocol, and for Unix systems, whether the KTAP or TEE is installed. Several typical examples are illustrated below:																																																												
	 <p>The screenshot shows three separate sections of the Inspection Engines panel, each for a different database protocol:</p> <ul style="list-style-type: none"> MSSQL Example: <table border="1"> <thead> <tr> <th>Protocol</th> <th>Port Range</th> <th>TEE</th> </tr> </thead> <tbody> <tr> <td>MSSQL</td> <td>1434-1434</td> <td>-</td> </tr> <tr> <td>Ip</td> <td>Mask</td> <td>Connect To Ip</td> </tr> <tr> <td>192.168.1.133</td> <td>255.255.255.255</td> <td>127.0.0.1</td> </tr> <tr> <td>Process Names</td> <td>Named Pipe</td> <td></td> </tr> <tr> <td>SQLSRVR.EXE</td> <td></td> <td></td> </tr> <tr> <td>Encryption</td> <td>Instance Name</td> <td></td> </tr> <tr> <td></td> <td>MSSQLSERVER</td> <td></td> </tr> </tbody> </table> INFORMIX Example: <table border="1"> <thead> <tr> <th>Protocol</th> <th>Port Range</th> <th>TEE</th> </tr> </thead> <tbody> <tr> <td>INFORMIX</td> <td>1400-1600</td> <td>-</td> </tr> <tr> <td>Ip</td> <td>Mask</td> <td>Connect To Ip</td> </tr> <tr> <td>127.0.0.1</td> <td>255.255.255.255</td> <td>127.0.0.1</td> </tr> <tr> <td>Process Names</td> <td>Named Pipe</td> <td></td> </tr> <tr> <td>ONINIT.EXE</td> <td></td> <td>SQLEXEC</td> </tr> </tbody> </table> Sybase Example: <table border="1"> <thead> <tr> <th>Protocol</th> <th>Port Range</th> <th>TEE</th> </tr> </thead> <tbody> <tr> <td>Sybase</td> <td>4100-4100</td> <td>-</td> </tr> <tr> <td>Ip</td> <td>Mask</td> <td>Connect To Ip</td> </tr> <tr> <td>127.0.0.1</td> <td>255.255.255.255</td> <td>127.0.0.1</td> </tr> <tr> <td>Process Names</td> <td>Named Pipe</td> <td></td> </tr> <tr> <td>SQLSRVR.EXE</td> <td></td> <td></td> </tr> </tbody> </table> 	Protocol	Port Range	TEE	MSSQL	1434-1434	-	Ip	Mask	Connect To Ip	192.168.1.133	255.255.255.255	127.0.0.1	Process Names	Named Pipe		SQLSRVR.EXE			Encryption	Instance Name			MSSQLSERVER		Protocol	Port Range	TEE	INFORMIX	1400-1600	-	Ip	Mask	Connect To Ip	127.0.0.1	255.255.255.255	127.0.0.1	Process Names	Named Pipe		ONINIT.EXE		SQLEXEC	Protocol	Port Range	TEE	Sybase	4100-4100	-	Ip	Mask	Connect To Ip	127.0.0.1	255.255.255.255	127.0.0.1	Process Names	Named Pipe		SQLSRVR.EXE		
Protocol	Port Range	TEE																																																											
MSSQL	1434-1434	-																																																											
Ip	Mask	Connect To Ip																																																											
192.168.1.133	255.255.255.255	127.0.0.1																																																											
Process Names	Named Pipe																																																												
SQLSRVR.EXE																																																													
Encryption	Instance Name																																																												
	MSSQLSERVER																																																												
Protocol	Port Range	TEE																																																											
INFORMIX	1400-1600	-																																																											
Ip	Mask	Connect To Ip																																																											
127.0.0.1	255.255.255.255	127.0.0.1																																																											
Process Names	Named Pipe																																																												
ONINIT.EXE		SQLEXEC																																																											
Protocol	Port Range	TEE																																																											
Sybase	4100-4100	-																																																											
Ip	Mask	Connect To Ip																																																											
127.0.0.1	255.255.255.255	127.0.0.1																																																											
Process Names	Named Pipe																																																												
SQLSRVR.EXE																																																													

Protocol	The type of database server being monitored (FTP, DB2, Informix, Oracle, Sybase, MSSQL, MSSQL-Named Pipes, or Kerberos).
Port Range	The range of ports monitored for this database server. There is usually only a single port in the range. For a Kerberos inspection engine, this value should always display as 88-88.
TEE Listen Port - Real Port	Not used for Windows. Required for Unix. The TEE Listen Port is the port on which S-Tap listens for and accepts local database traffic. The Real Port is the port onto which S-Tap forwards traffic. By inserting itself between the two ports, the inspection engine in effect creates a "T."

Column or Control	Description						
IP	A list of Client IP addresses used with the corresponding mask (see below) to determine which clients to monitor. If the IP address is the same as the IP address for the database server, and a mask of 255.255.255.255 is used, all network traffic arriving at the server is monitored. If an IP address of 0.0.0.0 is used, there are special uses for the mask value, as described below (see Mask).						
Mask	A mask used with the corresponding IP address (see above) to determine which clients to monitor. If an IP address of 0.0.0.0 is used, there are special uses for the following Mask values: <table border="1"> <thead> <tr> <th>Mask</th><th>Description</th></tr> </thead> <tbody> <tr> <td>0.0.0.0</td><td>All traffic will be collected.</td></tr> <tr> <td>255.255.255.255</td><td>Nothing will be collected.</td></tr> </tbody> </table>	Mask	Description	0.0.0.0	All traffic will be collected.	255.255.255.255	Nothing will be collected.
Mask	Description						
0.0.0.0	All traffic will be collected.						
255.255.255.255	Nothing will be collected.						
Connect To Ip	The IP address for S-Tap to use to connect to the database. Some databases accept local connection only on the “real” IP of the machine, and not on the default (127.0.0.1).						
Process Names	Windows only. A list of database server process names, used only for Oracle, or for MS SQL Server when named pipes are used. For Oracle, the list is usually <i>oracle.exe,tnslsnr.exe</i> . For MS SQL Server, the list is usually <i>sqlservr.exe</i> .						
Named Pipe	Windows only. Specifies the name of a named pipe. If a named pipe is used, but nothing is specified here, S-Tap retrieves the named pipe name from the registry.						
Instance Name	Database instance name (MS SQL Server).						

Settings Available Only on the S-Tap Configuration File

The settings described in this section can be changed only by editing the configuration file on the database server (*guard_tap.ini* – see the Installation instructions in Chapter 1 for information about how to locate and edit this file). These settings cannot be accessed from the GUI.

MS SQL Server 2005 Using Encryption

In the inspection engine section for any MS SQL Server using encryption, use the **INSTANCE_NAME** parameter to specify the service name for MS SQL Server, which is **MSSQLSERVER** by default. If you have used another service name, use that name instead:

```
INSTANCE_NAME=MSSQLSERVER
```

The INSTANCE_NAME parameter specifies the service name for MS SQL Server, which is MSSQLSERVER by default. If you have used another name, substitute that name here.

For example:

```
[DB_MSSQL1]
PORT_RANGE_START=1434
PORT_RANGE_END=1434
TAP_DB_PROCESS_NAMES=SQLSERVR.EXE
NAMED_PIPE=SQL\QUER,PIPE\SQLLOCAL
DB_TYPE=MSSQL
INSTANCE_NAME=MSSQLSERVER
NETWORKS=127.0.0.1/255.255.255.255
```

MS SQL Server Using Kerberos Authentication

In the inspection engine section for any MS SQL Server using Kerberos authentication, use the INSTANCE_NAME parameter to specify the service name for MS SQL Server, which is MSSQLSERVER by default. If you have used another service name, use that name instead:

```
INSTANCE_NAME=MSSQLSERVER
```

For example:

```
[DB_MSSQL1]
PORT_RANGE_START=1434
PORT_RANGE_END=1434
TAP_DB_PROCESS_NAMES=SQLSERVR.EXE
NAMED_PIPE=SQL\QUER,PIPE\SQLLOCAL
DB_TYPE=MSSQL
INSTANCE_NAME=MSSQLSERVER
NETWORKS=127.0.0.1/255.255.255.255
```

DB2 Using Shared Memory

In the inspection engine section for any DB2 server using shared memory, set the DB2_FIX_PACK_ADJUSTMENT parameter to one of the following values:

16 for DB2 v8.1 or v8.2 on Unix K-Tap

20 for DB2 v8.1 on Windows

80 for DB2 v8.2 on Windows

For example, for DB2 v8.2 on Windows:

```
[DB_DB22]
PORT_RANGE_START=50000
PORT_RANGE_END=50000
DB_TYPE=DB2
DB2_FIX_PACK_ADJUSTMENT=80
```

NETWORKS=192.168.1.0/255.255.255.0

Managing S-Taps

You can manage S-Taps only when the SQL Guard system you are logged into is the *active* SQL Guard host for that S-Tap, and only when that S-Tap is online. When these conditions are met, the  (Edit S-Tap Configuration) button is enabled.

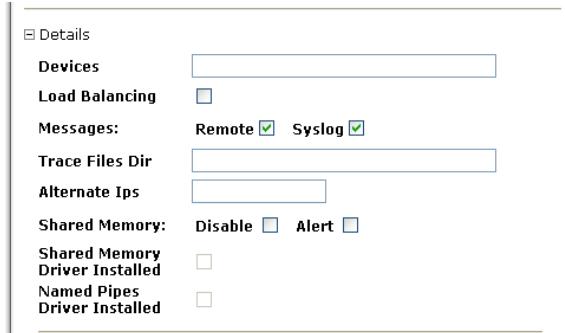
To manage S-Taps, follow the procedure outlined below:

1. Open the S-Tap Control panel as described previously.
2. Click  (Edit S-Tap Configuration) for the S-Tap you want to modify, which opens the S-Tap Configuration panel for that S-Tap.

The editing procedure for each major pane of the S-Tap Configuration panel is described separately below.

Details Pane Changes

1. Expand the Details pane:



Details

Devices:

Load Balancing:

Messages: Remote Syslog

Trace Files Dir:

Alternate Ips:

Shared Memory: Disable Alert

Shared Memory Driver Installed:

Named Pipes Driver Installed:

2. In the Devices box, enter the interfaces on which the S-Tap should listen for SQL Guard server communications. It can be left blank for Windows, but must be entered for Unix.
3. Mark the Load Balancing box if S-Tap will balance traffic to all SQL Guard servers listed in the SQL Guard Hosts pane (see below). Load balancing is by Client IP address, since all traffic for a session between a specific client and a specific server must be viewed by the same SQL Guard unit.
4. Mark the Remote box to send messages to the active SQL Guard host.
5. Mark the Syslog box to write messages to the syslog file on the database server.

6. In the Trace Files Dir box, enter the directory in which trace files are to be installed.
7. In the Alternate Ips box, enter any additional IP addresses for the database server on which the S-Tap is installed.
8. For Windows servers only, mark the Disable box to have S-Tap disable shared memory connections whenever they are discovered.
9. For Windows servers only, mark the Alert box to have S-Tap send a system alert message whenever shared memory connections are discovered.
10. For Windows servers only, mark the Shared Memory Driver Installed box to indicate that the shared memory driver has been installed.
11. For Windows servers only, mark the Named Pipes Driver Installed box to indicate that the Named Pipes Driver has been installed.
12. Optionally click the Hide Details button () to close the Details pane when you have finished making all changes.

Hunter Pane Changes

1. Expand the Hunter pane:



2. Use the Hunt box to identify any processes to be killed, using the following syntax: `db_type:process[,db_type:process]`

Where the **db-type** can be FTP, DB2, Informix, Oracle or Sybase, and the processes may be any of the following:

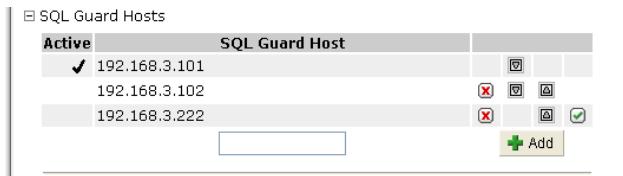
SHM	Shared memory
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
FIFO	A named pipe IPC mechanism
PIPE	A simple (unnamed) pipe IPC
INET	Internet Protocol (HPUX)

These values are not case-sensitive, and each entry is separated from the next by a comma. Example: To kill Oracle Bequeath processes, which uses a simple pipe, you would enter: **oracle:pipe**

3. In the Sleep Time box, enter the maximum number of seconds between the randomized starting time of the hunter's rogue process search routine. The start time is random to increase the difficulty of defeating it by running in fixed time slots or intervals.
4. In the DBs box, identify the databases to be reported on, separating each entry from the next with a comma. The allowed entries are: Informix, DB2, Sybase, and Oracle.
5. Optionally click the Hide Hunter button () to close the pane when you have finished making all changes.

SQL Guard Hosts Pane Changes

The SQL Guard Hosts pane is shown expanded, below:



Active	SQL Guard Host	
<input checked="" type="checkbox"/>	192.168.3.101	 
	192.168.3.102	 
	192.168.3.222	  

For detailed instructions on changing the list of SQL Guard hosts for an S-Tap, see [Secondary SQL Guard Hosts for S-Tap](#), in Chapter 1.

Inspection Engines Pane Changes

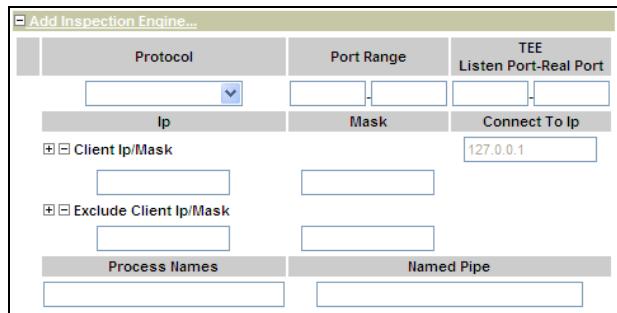
Refer to the following topics (below) to make changes to the S-Tap inspection engine configuration:

- o Adding S-Tap Inspection Engines
- o Modifying or Removing S-Tap Inspection Engines

Adding or Modifying S-Tap Inspection Engines

To add an inspection engine to an S-Tap configuration:

1. See [Managing S-Tap Inspection Engines](#) (above) to open the S-Tap Configuration panel.
2. Click the Add Inspection Engines title bar to expand the pane:



3. From the Protocol list, select the protocol to be monitored: FTP, Informix, DB2, Sybase, MSSQL, Named Pipes, Windows File Share, Oracle, or Kerberos. Only one can be selected. To monitor multiple database servers on the same host, define multiple inspection engines.

Named Pipes and Kerberos have special uses, as described below.

Kerberos: For users of Microsoft SQL Server, where the domain controller is a Windows 2003 server using Kerberos authentication, an S-Tap inspection engine can be configured to forward Kerberos traffic to SQL Guard. The SQL Guard server will then substitute real database user names for the Kerberos-encrypted database user names. When a Kerberos inspection engine is defined, the Kerberos traffic will be forwarded to *all* SQL Guard Hosts (not just the active host).

When the Kerberos protocol is selected, set the Port Range (described below) to the single value of 88, and leave all other fields empty.

As an alternative to forwarding Kerberos traffic via S-Tap, if the Kerberos traffic can be mirrored to the SQL Guard server, the SQL Guard server can be configured to perform the decryption using the mirrored traffic. For information about how to configure the SQL Guard server to translate Kerberos user names, see [About Kerberos-Encrypted Database User Names](#) in Chapter 2.

Named Pipes: On Windows systems, this protocol is used to access Microsoft SQL Server, Informix, or Oracle databases using the named pipes mechanism. When selected, enter the name of the named pipe in the Named Pipe box, only if the default named pipe name for the database is not used.

4. In the Port Range boxes, enter the single range of ports to be monitored, from the lowest number to the highest number. For Kerberos, enter port **88**. Do not enter an all-inclusive range of ports here, as the S-Tap may become bogged down attempting to interpret traffic that is of no interest.
5. Optional for Windows, required for Unix/Linux based servers only: In the TEE Listen and Real Port boxes, enter the port numbers between which S-Tap inserts itself,

creating a “Tee” (i.e., it accepts database traffic on the Listen Port and forwards it onto the Real Port).

6. In the Client IP/Mask boxes, enter an IP mask and its corresponding subnet mask to select which clients to monitor. Click the  (Add Client Ip/Mask) button to add the entry. You can add multiple entries.

Notes: Do *not* monitor traffic that is also monitored by the SQL Guard host for this S-Tap. If that happens, the SQL Guard unit receives duplicate packets, it is unable to reconstruct messages, and that traffic will be ignored.

If you are using the TEE and leave the Client IP/Mask boxes blank, all clients will be monitored.

If the IP address is the same as the IP address for the database server and a mask of 255.255.255.255 is used, all network traffic arriving at the server is monitored.

To remove the last Client IP/Mask pair entered, click the  (Remove Client Ip/Mask) button.

7. Optional. In the Exclude Client Ip/Mask box, enter an IP mask and its corresponding subnet mask to select which clients to exclude. Click the  (Add Client Ip/Mask) button to add the entry. You can add multiple entries. To remove the last Client IP/Mask pair entered, click the  (Remove Client Ip/Mask) button.
8. Windows only, for Oracle or MS SQL Server databases only. In the Process Names box, enter a list of database server process names, with each name separated from the next by a comma, and no blank spaces between entries. For Oracle, the list is usually *oracle.exe,tnslsnr.exe*. For MS SQL Server, the list is usually *sqlservr.exe*.
9. In the Named Pipe box, if the protocol is Named Pipes if named pipes are used and you are not using the default named pipes name for the database type, enter the Named Pipe name here
10. Click the Add button when you are done entering all information for this inspection engine.
11. Optional: Add more inspection engines by repeating the steps above.
12. Click the Apply button when you are done adding or modifying information. For a description of how the change is actually made on the S-Tap, see [Applying Changes to S-Tap Configurations](#), below.

What happens next: The inspection engine status light will turn yellow, because the S-Tap configuration defined on the SQL Guard server no longer matches the

configuration defined on the database server. The SQL Guard server will send an updated configuration to the database server, and the S-Tap will stop and attempt to restart using the new configuration. If the S-Tap cannot restart with the new configuration, it will attempt to restart with the last good configuration (which it saves).

Modifying or Removing S-Tap Inspection Engines

To modify or remove an inspection engine definition in an S-Tap configuration:

1. See [Managing S-Tap Inspection Engines](#) (above) to open the S-Tap Configuration panel for the S-Tap whose configuration you want to modify:
2. To remove an inspection engine, click its  (Remove) button.

OR

To modify an inspection engine definition, type over any values that you want to replace. See [Adding S-Tap Inspection Engines](#) (above) for detailed information on the use of each field.

Note: You can modify existing information but you cannot add or remove Client IP/Mask pairs. To make that change, you have to remove the inspection engine and re-add it using the correct Client IP/Mask information.

3. Click the Apply button when you are finished making all changes. For a description of how the change is actually made on the S-Tap, see [Applying Changes to S-Tap Configurations](#), below.

Applying Changes to S-Tap Configurations

When you click the Apply button after modifying an S-Tap configuration, the following sequence of events occurs:

1. SQL Guard sends the updated configuration to the S-Tap host.
2. SQL Guard changes the Status light for that S-Tap from green to yellow.
3. S-Tap attempts to use the new configuration. If it cannot, it reverts to the last good configuration.
4. S-Tap signals the SQL Guard host that it is back online, and indicates whether or not it is using the new configuration.
5. If S-Tap is using the new configuration, SQL Guard changes the Status light to green. If S-Tap cannot use the new configuration, SQL Guard leaves the Status

light set to yellow. When this happens, check the S-Tap Events Panel for error messages (see below).

Viewing the S-Tap Events Panel

You can use the S-Tap Events Panel to view the event messages output by S-Tap.

Note: If no messages display in the S-Tap Events Panel, the production of event messages may have been disabled in the configuration file for that S-Tap. If this is the case, you may be able to locate S-Tap event messages on the host system in the Event Log (Windows) or the *syslog* file (Unix/Linux).

To open the S-Tap Events panel for any S-Tap listed in the S-Tap Control panel, click the  (Show S-Tap Event Log) button for that S-Tap. The S-Tap Events Panel opens in a separate browser window:

S-Tap Events		
Event Type	Event Description	Timestamp
SUCCESS	Heartbeat was received from 192.168.3.104	2007-04-05 16:09:42.0
SUCCESS	Guardium STAP registered with 192.168.3.104	2007-04-05 16:09:42.0
INFORMATION_TYPE	Guardium_STAP Started	2007-04-05 16:09:42.0
SUCCESS	Heartbeat was received from 192.168.3.104	2007-04-04 13:39:58.0
SUCCESS	Guardium STAP registered with 192.168.3.104	2007-04-04 13:39:54.0
INFORMATION_TYPE	Guardium_STAP Started	2007-04-04 13:39:54.0
SUCCESS	Heartbeat was received from 192.168.3.104	2007-04-04 13:36:17.0
SUCCESS	Guardium STAP registered with 192.168.3.104	2007-04-04 13:36:15.0
INFORMATION_TYPE	Guardium_STAP Started	2007-04-04 13:36:15.0
SUCCESS	Heartbeat was received from 192.168.3.104	2007-04-04 01:50:07.0
SUCCESS	Guardium STAP registered with 192.168.3.104	2007-04-04 01:50:05.0
INFORMATION_TYPE	Guardium_STAP Started	2007-04-04 01:50:05.0
SUCCESS	Heartbeat was received from 192.168.3.104	2007-04-02 14:08:32.0
SUCCESS	Guardium STAP registered with 192.168.3.104	2007-04-02 14:08:29.0
INFORMATION_TYPE	Guardium_STAP Started	2007-04-02 14:08:29.0

Done 

Column or Control	Description
Event Type	Identifies a type of event: Success, Error Type, etc.
Event Description	Provides a short description of the event.
Timestamp	Provides the date and time that the event occurred.
Done	Click the Done button to close the window.

S-Tap Error Messages

The following table describes the error messages produced by S-Tap, in alphabetical sequence. Items shown in *italicized text* are variable.

Message	Description
Cant read ini file /usr/local/guardium/guard_stap/guard_tap.ini: Cannot resolve hostname xxx for the IP address parameter sqlguard_ip in section SQLGUARD_x. Reverting to /usr/local/guardium/guard_stap/guard_tap.ini.bak	The S-Tap configuration file (guard_tap.ini) has been corrupted, which is most likely to happen when it has been edited manually. When this happens, S-Tap attempts to restart from the last known good backup file (if one is available).
guard_tap[n]: bind: Address already in use <i>[DB server name or IP]</i> guard_tap[n]: Cant bind listening socket for tee: Address already in use	A port that an S-Tap TEE is trying to use is already in use. For example, if you configure a TEE to listen on port 4100, and Sybase is already listening on that port, you will receive this message. This message is issued only when the S-Tap TEE process is starting.
guard_tap[n]: connect: Network is unreachable	The standard message received when trying to reach a host that is not accessible. In 99% of the cases this means that the SQL Guard server down (is not answering ARP requests).
guard_tap[n]: Delayed server connection error: Connection refused	The SQL Guard server is refusing a connection request from this S-Tap. That SQL Guard server either has no inspection engine running (not likely), or it is not configured to accept S-Tap connections (check the unit_type setting for that SQL Guard unit).
guard_tap[n]: Deleting connection on unknown pid:n	Not an error message; please disregard.
guard_tap[n]: Got a connection from a remote machine, ignoring	S-Tap has received a connection request (to a TEE port) from an unrelated application at a remote host, and is ignoring that request.
guard_tap[n]: Got new configuration	The SQL Guard administrator has updated the configuration while logged into the SQL Guard server, and the updated configuration file has been received by the S-Tap.
guard_tap[n]: Guard Tee is accepting	Normal TEE process start-up message

Message	Description
connections on port 12346	(appears only when the TEE is installed).
guard_tap[n]: Guardium TAP starting	Normal S-Tap process start-up message.
guard_tap[n]: read from socket: Connection reset by peer	The database server or database client is down. For example, someone ran an Oracle <i>sqlplus</i> session and used <i>ctrl-C</i> to exit (in this case, it does not indicate a problem.)
guard_tap[n]: Server wasn't heard from for 180 sec, closing and re-opening	S-Tap has not received a heartbeat signal from the SQL Guard server for 180 seconds. It will attempt to reconnect with the server. No data is lost (unless the server does not respond to the connection request).
guard_tap[n]: SQLguard socket read: Connection reset by peer	The SQL Guard server closed the connection to the S-Tap. This happens when the SQL Guard server restarts, or when the SQL Guard server inspection engine automatically goes down and comes up again (in which case, it does not indicate a problem).
guard_tap[n]: waitpid: No child processes	Not an error message; please disregard.
hunter[n]: killed n2-	The S-Tap hunter process has killed an unauthorized connection identified by n2.

Reporting or Alerting on S-Tap Connectivity

SQL Guard logs two types of exceptions relating to S-Tap connectivity:

- *STAP Connectivity timeout* indicates that an S-Tap has not been heard from within the timeout interval.
- *STAP Connectivity reconnect* indicates that the S-Tap has re-connected.

To create reports or alerts based on either of these exception types, see the following sections of the SQL Guard User Guide:

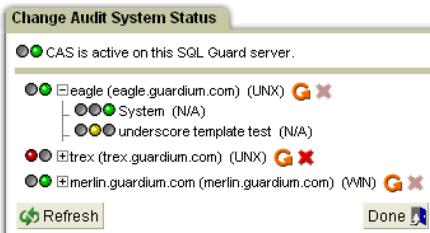
- Chapter 7: Building Queries and Reports
- Chapter 8: Statistical Alerts

In either case, the query domain you use is the Exceptions domain. The Exception Description attribute is equal to one of the two exception type values shown above.

Monitoring CAS Status

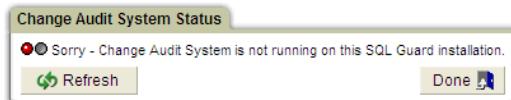
By default, the functions described in this section are available to the *admin user*, and users with the *admin role*. Open the Administrator portal and locate the Local Taps section of the Administration Console. If there is no Local Taps section, the **unit type** setting for this SQL Guard server needs to be changed. See [Activating CAS on the SQL Guard Server](#) for instructions on how to enable the Local Taps menu.

To monitor CAS status, select **CAS Status** in the **Local Taps** section of the **Administration Console** to open the Change Audit System Status panel:



For each database server where CAS is installed and running, and where this SQL Guard server is configured as the *active SQL Guard host*, this panel displays the CAS status, and the status of each CAS instance configured for that database server.

Note: If the Change Audit System Status panel looks like this:



The installed SQL Guard server system license key does not include CAS. See [Activating CAS on the SQL Guard Server](#) for instructions on how to install a new license key.

Regarding the sets of status lights on the Change Audit System Status panel: when you hover the mouse over a set of status lights, a pop-up text box displays the current status. If you have trouble distinguishing the colors on your monitor, for all status light sets, the leftmost light is always red, the rightmost light is green, and on sets of three lights, the middle one is yellow.



Change Audit System Status Panel

Component	Description
(top row – CAS on SQL Guard Server)	Red: CAS is not running on this SQL Guard Server. Green: CAS is active on this SQL Guard server.
(monitored system)	For each CAS host where this SQL Guard server is the active SQL Guard server, the status lights indicate whether CAS is connected: Red: Host and/or the CAS agent is offline or unreachable. Green: Host and CAS agent are online.
(monitored system)	<i>Reset the CAS agent on this monitored system.</i> This stops and restarts the CAS agent on the database server.
(monitored system)	<i>Remove this monitored system from CAS.</i> Click this button to remove all CAS information for this monitored system from the SQL Guard internal database. This button is disabled if the CAS agent is running on this system. You must stop the CAS agent to use this button. See Stopping and Starting the CAS Agent .
(CAS instances)	Indicates the status of a CAS instance on the monitored system. <i>If the owning monitored system status is red (indicating that the CAS agent is offline), ignore this set of status lights.</i> <ul style="list-style-type: none"> Red The instance is disabled. Green The instance is enabled and online, and its configuration is synchronized with the SQL Guard server configuration. Yellow The instance is enabled, but the instance configuration on the SQL Guard server does not match the instance configuration on the monitored system (it has been updated on the SQL Guard server, but that update has not been applied on the monitored system).

Component	Description
Refresh	Rechecks the status of all servers in the list. This button <i>does not</i> stop and/or restart CAS on a database server – it only checks the connection between CAS on the SQL Guard server and CAS on each database server.

Stopping and Starting the CAS Agent

There are several situations where you may need to stop or start the CAS agent on a monitored system. Follow the procedures outlined below.

Note: If all you want to do is stop and restart the CAS agent, you can do that from the Administrator Console of the SQL Guard server, using the  (Reset the CAS agent on this monitored system) button on the Change Audit System Status panel.

Stopping CAS on a Unix Host

1. Edit the file /etc/inittab.
2. Find the line:

```
cas:2345:respawn:/usr/local/guardium/guard_stap/cas/bin/run_wrapper.s
h /usr/local/guardium/guard_stap/cas/bin
```

near the end of the file.
3. Comment out the line by inserting the # (pound sign) character in the first character position.
4. Save the file.
5. Issue the command:

```
init -q
```
6. Issue the command:

```
ps -ef | grep wrapper
```
7. Note the PID of each of the processes listed.
8. For each of the processes listed, issue the following command:

```
kill -9 <pid>
```
9. In the Change Audit System Status panel of the SQLGuard administrator portal, the status light for this CAS host should be red, and the  button should be enabled (which allows you to remove data from this CAS host from the SQL Guard server internal database).

Starting CAS on a Unix Host

Use this procedure to restart the CAS agent only when it has been stopped by editing the /etc/inittab file as described above.

1. Edit the file /etc/inittab.

2. Find the line:

```
#cas:2345:respawn:/usr/local/guardium/guard_stap/cas/bin/run_wrapper.  
sh /usr/local/guardium/guard_stap/cas/bin
```

near the end of the file.

3. Uncomment the line by removing the # (pound sign) character in the first character position.

4. Save the file.

5. Issue the command:

```
init -q.
```

The CAS agent will restart and connect to the server.

Starting and Stopping CAS on a Windows Host

On Windows CAS runs as a System Service.

1. In the Services panel, highlight the Change Audit System Client item.
2. Select either Start or Stop from the Action menu.

Exporting and Importing Definitions

If you have multiple Guardium systems with identical or similar requirements, you can define the components you need on one system and export those definitions to other systems, provided those systems are on the same Guardium software release level.

You can export one type of definition (reports, for example) at a time. Each element exported can cause other referenced definitions to be exported as well. For example, a report is always based on a query, and it can also reference other items, such as IP address groups or time periods. All referenced definitions (except for security roles) are exported along with the report definition. However, only one copy of a definition is exported if that definition is referenced in multiple exported items.

Note: When you export graphical reports, the presentation parameter settings (colors, fonts, titles, etc.) are not exported. When imported, these reports will use the default presentation parameter settings for the importing SQL Guard system.

Importing Groups

When importing a group that already exists, members may be added, but no members will be deleted.

Importing Aliases

When importing aliases, new aliases may be added, but no aliases will be deleted.

Definitions that can be exported and imported

- Access Map
- Alert
- Alias
- Audit Process
- Auto-discovery Process
- CAS Hosts
- CAS Template Sets
- Classification Process
- Classifier Policy
- Custom Class Connection Permission
- Custom Domain
- Custom Table
- Datasource
- Group
- Period (time period)
- Policy (but *not* an included Baseline)
- Privacy Set
- Query
- Report
- Role
- Security Assessment
- User

Definitions that cannot be exported and imported

- Baseline or Baseline included in a Policy
- Custom Alerting Class
- Custom Assessment Test
- Custom Identification Procedure

Ownership, Role and User Considerations

Special considerations apply to the handling of ownership, users, and roles when exporting and importing definitions. For detailed information about managing users and roles, see [Chapter 3: User Management](#) and [Chapter 4: Security Role Management](#).

Ownership of Imported Definitions

When a definition is created, the SQL Guard user who creates it is saved as the owner of that definition. The significance of this is that if no security roles are assigned to that definition, only the owner and the privileged SQL Guard *admin* user have access to it.

When a definition is imported, the owner is always changed to *admin*.

Roles for Imported Definitions

References to security roles are removed from exported definitions. So any imported definitions will have no roles assigned.

Users for Imported Definitions

A reference to a user in an exported definition causes the user definition to be exported.

When definitions are imported, the referenced user definitions are imported *only if they do not already exist on the importing system*. In other words, existing user definitions are never overwritten. This has several implications, as described in the Duplicate Role and User Implications topic, below.

In addition, imported user definitions are disabled. This means that imported users can receive email notifications sent from the importing system, but they are not able to log into that system, unless the SQL Guard administrator enables that account.

Duplicate Group and User Implications

As mentioned earlier, if a group referenced by an exported definition already exists on the importing system, the definition of that group from the exporting system will not be imported. This may create some confusion if the group is not used for the same purposes on both systems.

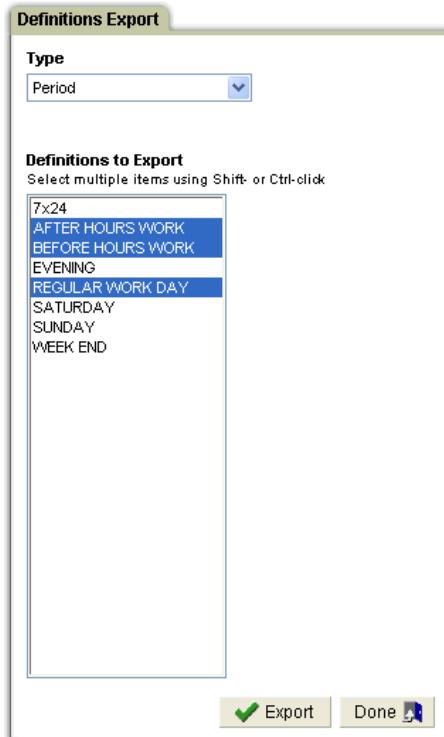
If a user definition already exists on the importing system, it may not be for the same person defined on the exporting system. For example, assume that on the exporting system the user *jdoe* with the email address *john_doe@aaa.com* is a recipient of output from an exported alert. Assume also that on the importing system, the *jdoe* user already exists for a person with the email address *jane_doe@zzz.com*. The exported user definition is not imported, and when the imported alert is triggered, email is sent to the *jane_doe@zzz.com* address. In either case, when security roles or user definitions are not imported, check the

definitions on both systems to see if there are differences. If so, make the appropriate adjustments to those definitions.

Exporting SQL Guard Definitions

To export SQL Guard definitions for use on other systems:

1. Open the Administration Console panel (not shown).
2. In the SQL Guard Definitions section of the Administration Console menu, click Export to open the Definitions Export panel:

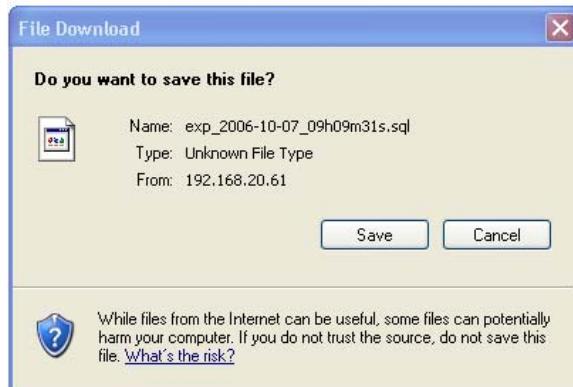


3. From the Type list, select the type of definition to export. The Definitions to Export box is populated with definitions of the selected type. In the example to the right, Period was selected as the definition type and all periods defined are listed in the Definitions to Export box.
4. Select all of the definitions of this type to be exported.

- *To select multiple contiguous definitions:*
Click the mouse on the first definition to copy, hold down the Shift key, and click the mouse on the last definition to copy.
- *To select multiple non-contiguous definitions:* Hold down the Ctrl key and click the mouse to select multiple non-contiguous definitions.

Note: Do *not* export a Policy definition whose name contains one or more quote characters. That definition can be exported, but it *cannot* be imported. To export such a definition, make a clone of it, naming the clone without using any quote characters, and export the clone.

5. Click the Export button. Depending on your browser security settings, you may receive a warning message asking if you want to save the file or to open it using an editor:



6. Click the Save button if you receive a warning like the one mentioned above.
7. A standard (for your operating system) Save dialog box is displayed. Supply a file name and directory for the exported definition file (not shown). Click Save or OK to save the file and close the dialog box.
8. Click the Done button to close the Definitions Export panel when you finish exporting all types.

Note: Some types of definitions include other types. For example, a report is always based on a query and could optionally include other SQL Guard items such as time periods or IP address groups. When you export a report, you also export the query it uses, as well as any other definitions included in that report (time periods or IP address groups, for example).

Importing SQL Guard Definitions

To import SQL Guard definitions:

1. Open the Administration Console panel (not shown).
2. In the SQL Guard Definitions section of the Administration Console menu, click Import to open the Definitions Import panel.



3. Enter the name of the file containing the exported definitions or click the Browse button to select that file.
4. Click the Upload button. You are notified when the operation completes and the definitions contained in the file are displayed, as illustrated below:



5. Upload additional files by repeating the previous two steps or import or remove uploaded definitions as follows:
 - o Click (Import this set of Definitions) to import the definitions.
 - o Click (Remove this set of Definitions without Importing) to remove the uploaded file without importing the definitions.

You are prompted to confirm either action.

-
- Note:** An import operation does not overwrite an existing definition. If you attempt to import a definition with the same name as an existing definition, you are notified that the item was *not* replaced. If you *want* to overwrite an existing definition with an imported one, you must delete the existing definition before performing the import operation.
-
6. Click the Done button to close the panel when you have finished importing or removing all uploaded files.
-
- Note:** Exported Policy definitions with names containing one or more quote characters *cannot* be imported. See [Exporting SQL Guard Definitions](#), above.
-

Custom Assessment Tests

A *custom assessment test* is a user-written security assessment test implemented as a Java class. For general information about using security assessment tests and the ready-to-go assessment tests provided with the system, see the *SQL Guard User Guide*. For instructions on how to write a custom assessment test, see [Defining Custom Assessment Tests](#), below. For instructions on how to upload and maintain custom tests on the SQL Guard system, see [Installing and Managing Custom Assessment Tests](#), later in this chapter.

-
- Note:** Before creating a custom assessment test, contact Guardium Support to obtain the necessary interfaces file.
-

Defining Custom Assessment Tests

This section assumes that you are an experienced Java language programmer. A custom assessment test produces a numeric score and a text recommendation based on that score. The test can be based on data contained in any SQL Guard report or on data accessed outside of the SQL Guard system.

Regardless of the origin of the data, a custom assessment test class must implement the **CustomerDefinedTestIfc** interface. This interface contains two assessment methods: one to return the test score and one to return the text recommendation based on that score. These methods are described below (see [Using the Customer Defined Test Interface](#)).

When calling either of the customer defined test interface methods, SQL Guard passes an interface to the SQL Guard report results header class, **ReportResultHeaderIfc**. This interface allows access to all of the report data. Therefore, if you need data from within the SQL Guard report (for example, to calculate the test score or to list detailed information in the recommendation text), use the methods of the report results header class to access the appropriate data. The methods available in the report results class are described below (see [Using the Report Result Header Interface](#)).

Using the Customer Defined Test Interface

The customer defined test interface, `CustomerDefinedTestIfc`, is included in the `com.guardium.assessment.tests` package. The SQL Guard report results header that this interface references is defined in `com.guardium.datamodel.results`. To use the interface, include both of the following statements in your program:

```
package com.guardium.assessment.tests  
import com.guardium.datamodel.results.ReportResultHeaderIfc;
```

The customer defined test interface is defined as illustrated on the next page. Following that, the two methods are described in detail.

```
public interface CustomerDefinedTestIfc {  
    public double getTestScore(ReportResultHeaderIfc repResultHeader,  
                             double totalNumberOfRequests);  
    public String getRecommendationText (double score,  
                                         ReportResultHeaderIfc repResultHeader);  
}
```

When a security assessment task runs and if that task includes a custom assessment test, SQL Guard invokes both methods of the custom test: first the `getTestScore` method and then the `getRecommendationText` method. Each of these methods is described below.

getTestScore Method

Syntax `public double getTestScore(ReportResultHeaderIfc repResultHeader,
 double totalNumberOfRequests)`

Parameters `ReportResultHeaderIfc` is the SQL Guard report result header interface, described below.

`totalNumberOfRequests` is the total number of requests sent during the assessed period.

Returns The calculated test score, as a function of the report result and the total number of requests during the tested period.

getRecommendationText Method

Syntax `public String getRecommendationText (double score,
 ReportResultHeaderIfc repResultHeader);`

Parameters The calculated score (returned by `getTestScore`, above).

`ReportResultHeaderIfc` is the SQL Guard report result header interface, described below.

Returns A text recommendation to be displayed in the results viewer.

Using the Report Result Header Interface

As mentioned previously, the methods of the SQL Guard report result header interface provide access to SQL Guard report data, specifically:

- The total number of rows in the report table
- The overall value for the report (usually the sum of the last column of the report)
- The string value of a particular cell in the report table
- The total number of columns in the report table

getTotalRowCount Method

Syntax public int getTotalRowCount();

Parameters None.

Returns The total number of rows in the report results table.

getOverallValue Method

Syntax public double getOverallValue();

Parameters None.

Returns The total number of rows in the report results table.

getCellValue Method

Syntax public String getCellValue(int row, int column);

Parameters The row number (beginning with 1).

The column number (beginning with 1).

Returns The string value of the report results table cell identified by the passed row and column parameters.

getNumberOfColumns Method

Syntax public int getNumberOfColumns();

Parameters None.

Returns The total number of columns in the report results table.

Sample Custom Test Class

The following program illustrates the definition of a custom test class. This program checks the report results of a custom defined security assessment report listing all predefined Oracle user accounts used during the test period. If the report has no output, it returns a score of 10 and a short message saying that all is well. If there is report output, the program returns a score of zero and a recommendation suggesting that the use of these accounts (which it lists at the end of the recommendation) should be eliminated.

To calculate the test score, the program checks for any output. If there is output, the program extracts the name of each user account used from the passed report results table.

```
/*
 * Sample Custom Test Class
 *
 */
package com.guardium.assessment.tests;
import com.guardium.datamodel.results.ReportResultHeaderIfc;

public class CustTestOracleUsers implements CustomerDefinedTestIfc {
    /*
     * Calculate and return the test score.
     * If any predefined Oracle users, return 0.
     * If no predefined Oracle users, return 10
     */
    public double getTestScore(
        ReportResultHeaderIfc repResultHeader,
        double totalNumberOfRequests) {
        double totalAccess = repResultHeader.getOverallValue();
        if ( totalAccess > 0 )
            return 0;
        else
            return 10;
    }
    /*
     * Return a recommendation for each possible score.
     * If predefined Oracle user accounts were used, add a list
     * of them to the end of the message.
     */
    public String getRecommendationText(double score,
        ReportResultHeaderIfc repResultHeader) {
        String retStr = "";
        if (score == 10)
        {
            retStr = "No Oracle Predefined user accounts were used, ";
            retStr = retStr.concat("indicating a controlled environment.");
        }
        else
        {
            retStr = "Predefined Oracle user accounts were used. ";
            retStr = retStr.concat("You should consider eliminating use of ");
            retStr = retStr.concat("the following accounts: ");
            for (int i = 1 ; i <= repResultHeader.getTotalRowCount(); i++) {
```

```
        String us = repResultHeader.getCellValue(i,1);
        if (i > 1)
            retStr = retStr.concat(" , ");
        retStr = retStr.concat(us);
    }
    retStr = retStr.concat(".");
}
return retStr;
}

}
```

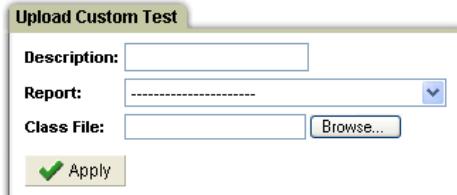
Installing and Managing Custom Assessment Tests

The Custom Test commands on the Administrator Console panel provide the ability to upload a custom test class to SQL Guard, update it, or remove it. Each command is described below.

Uploading Custom Tests

Once a custom class has been compiled, you can upload it to the SQL Guard system. To upload a custom test class:

1. Open the Administration Console panel (not shown).
2. In the Custom Assessment Tests section of the Administration Console menu, click Upload to open the Upload Custom Test panel:

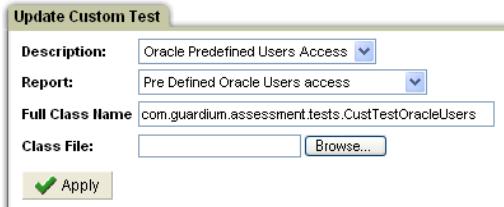


3. Enter a unique description for the custom test class in the Description box.
4. Optional: Select a report from the Report list.
5. Click the Browse button to open an operating-system dependent file navigation window, from which you can select the class file to be uploaded. The full path name of the class file will be inserted into the Class File box.
6. Click the Apply button after selecting the file (or typing the full path name in the Class File box). This action uploads the class and closes the panel.

Updating Custom Tests

To update a custom test class:

1. Open the Administration Console panel (not shown).
2. In the Custom Assessment Tests section of the Administration Console menu, click Update to open the Update Custom Test panel:

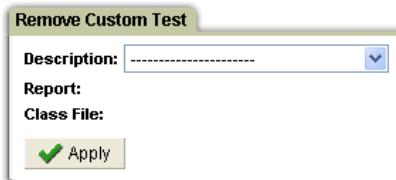


3. Select, in the Description box, the class to update. The full class name will display in the Full Class Name box.
4. Optional: Select a report from the Report list.
5. Click the Browse button to open an operating-system dependent file navigation window, from which you can select the class file to replace the existing class file. The full path name of the class file will be inserted into the Class File box.
6. Click the Apply button after you have selected the file (or typed the full path name in the Class File box). This action updates the class and closes the panel.

Removing Custom Tests

To remove a custom test class:

1. Open the Administration Console panel (not shown).
2. In the Custom Assessment Tests section of the Administration Console menu, click Remove to open the Remove Custom Test panel:



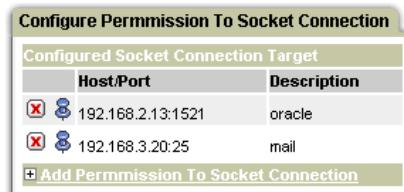
3. Select, in the Description box, the class to be removed.
4. Click Apply to remove the class. You are prompted to confirm the action.

Configuring Permission to Socket Connection

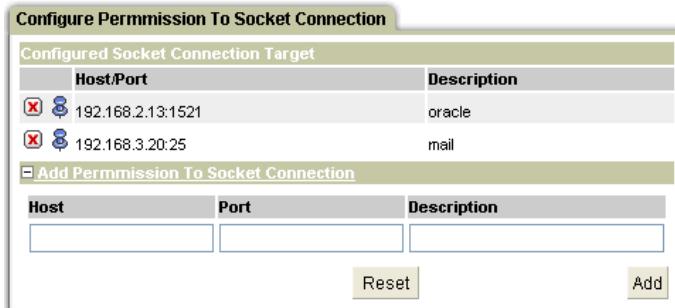
This procedure is the same for Custom Assesment Tests and Custom Alerting.

To configure permissions for socket connections that are used by custom tests:

1. Select **Administration Console – [Custom Assessment Tests or Custom Alerting] – Config** to open the Configure Permission To Socket Connection panel:



2. Click Add permission To Socket Connection to expand the Add pane:



3. Enter the IP address or host name for the host.
4. Enter a port number for the socket connection.
5. Enter a description.
6. Click Add.

Custom Alerting

SQL Guard provides for the distribution of alert messages via e-mail, SNMP or user-written Java classes. The last option is referred to as *custom alerting*. When an alert is triggered, a custom alerting class might update a Web page, for example, or it could take any other action appropriate for the situation.

Note: Before creating a custom alerting class, contact Guardium Support to obtain the necessary interfaces file.

Alerting Overview

An *alert* is a message indicating that an exception or policy rule violation was detected. Alerts are triggered in two ways:

- A *statistical alert* is triggered by a query that looks back over a specified time period to determine if the query condition has been satisfied. The SQL Guard Anomaly Detection component runs statistical alerts on a scheduled basis.
- A *real-time alert* is triggered by a security policy rule. The SQL Guard Inspection Engine component runs the security policy as it collects and analyzes database traffic in real time.

Regardless of how they are triggered, SQL Guard logs all alerts the same way: the alert information is logged in the SQL Guard internal database. The amount and type of information logged depends on the specific alert type. The SQL Guard Alerter component, which also runs on a scheduled basis, processes each new alert, passing the logged information for each alert to any combination of three notification mechanisms:

- SMTP – The SMTP (outgoing e-mail) server. The Alerter passes standard email messages to the SMTP server for which it has been configured.
- SNMP – The SNMP (network information and control) server. When SNMP is selected for an alert notification, the Alerter passes all alert messages of that type to the single trap community for which the Alerter has been configured.
- Custom – A user written Java class to handle alerts. The Alerter passes an alert message and timestamp to the custom alerting class. There can be multiple custom alerting classes and one custom alerting class can be an extension of another custom alerting class.

Alerting tasks can be performed by both SQL Guard administrators and users, as described below.

Alerting Tasks for SQL Guard Administrators

SQLGuard administrators perform the following tasks, all of which relate to the handling of alerts, and all of which are *optional*. SQL Guard administrators:

- Configure the Alerter to deliver messages to a single SMTP server and/or a single SNMP Server (see [Configuring the Alerter](#))
- Upload and manage custom alerting classes on the SQL Guard server (see [Managing Custom Alerting](#))
- Start and stop the Alerter, which delivers the alert messages to SMTP, SNMP or Custom alerting classes (see [Configuring the Alerter](#))

- Start and stop Anomaly Detection, which runs the statistical alerts (see [Configuring Anomaly Detection](#))
- Configure, start, and stop the Inspection Engine, which runs the installed security policy as database traffic is detected (see [Working with Inspection Engines](#))

Alerting Tasks for SQL Guard Users

Authorized SQL Guard users can perform the tasks described below, which relate to the handling of alerts. SQL Guard administrators may also be authorized to perform any of these tasks.

- Define statistical alerts (see *Chapter 8: Building Statistical Alerts* in the *SQL Guard User Guide*).

Note: When defining a statistical alert, be sure to mark the Active box. Also, to start logging statistical alerts, the SQL Guard Anomaly Detection component must be started (see [Configuring Anomaly Detection](#)).

- Define and install the security policy, which may contain real-time alerts (see *Chapter 11: Building Policies*, in the *SQL Guard User Guide*).

Note: To start logging real-time alerts, the policy containing the alert actions must be installed and the inspection must be restarted (see [Working with Inspection Engines](#)).

- Write custom alerting classes (see [Using the Customer Defined Alerting Interface](#)).

Using the Customer Defined Alerting Interface

The custom alerting class must be in the `com.guardium.custom` package and must implement the `com.guardium.custom.CustomerDefinedAlertingIfc` interface:

```
package com.guardium.custom  
public class YourClassNameHere implements CustomerDefinedAlertingIfc {
```

The interface contains the five methods described below.

processAlert Method

Description Process a single alert message.
Syntax `public void processAlert (String message, Date timeStamp)`
Parameters A String containing the message generated by the alert.
A `java.util.Date` for the time the alert message was created.

getMessage Method

Description Return the alert message.
Syntax `public String getMessage ()`
Returns A String containing the alert message.

getTimeStamp Method

Description Return the timestamp associated with the alert message.
Syntax `public Date getTimeStamp ()`
Returns A `java.util.Date` for the time the alert message was created.

setMessage Method

Description Set the alert message.
Syntax `public void setMessage (String inMessage)`
Parameter A String containing the alert message.

setTimeStamp Method

Description Set the timestamp associated with the alert message.
Syntax `public void setTimeStamp (Date inDate)`
Parameter A `java.util.Date` for the time the alert message was created.

Sample Custom Alerting Class

The following sample program implements the five methods described in the previous section. For the processAlert method, this program simply writes the alert message and timestamp to the system console.

```
/*
 * Sample Custom Alerting Class
 *
 */
package com.guardium.custom;
import java.text.DateFormat;
import java.util.Date;

public class HandleAlerts implements CustomerDefinedAlertingIfc {
    private String message = "";
    private Date timeStamp = null;
    public void processAlert(String message, Date timeStamp){
        setMessage(message);
        setTimeStamp(timeStamp);
        System.out.println(getMessage() + " on " +
                           DateFormat.getDateInstance().format(getTimeStamp()));
    }
    public void setMessage(String inMessage){
        message = inMessage;
    }
    public String getMessage(){
        return message;
    }
    public void setTimeStamp(Date inDate){
        timeStamp = inDate;
    }
    public Date getTimeStamp(){
        return timeStamp;
    }
}
```

Testing the Custom Alerting Class

After compiling the class, follow the procedure outlined below to test it.

1. Upload the custom class to the SQL Guard server.
2. Define a statistical or real-time alert to use the custom alerting class. Regardless of which alert type generates the alert, testing is easier if you assign a second notification type (email, for example) against which you can compare the custom alerting results.

3. Check the SQL Guard environment by doing one of the following:

For a statistical alert:

- Check that the Anomaly Detection polling interval is suitable for testing purposes and that Anomaly Detection has been started.
- **Note:** If the polling interval is too long (it may be 30 minutes or more), you may have a long wait before the query runs.
- Check that the Alerter polling interval is suitable for testing purposes and that the Alerter has been started.
- Check that the alert to be tested has been marked *active*.

For a real-time alert:

- Check that policy containing the rule with the custom alert action is the installed policy.
- Verify that the inspection engine was restarted after the updated policy was installed.
- Check that the Alerter polling interval is suitable for testing purposes and that it has been started.

4. Take whatever action is necessary to trigger the alert (generate a number of login failures, for example).

Managing Custom Alerting

The Custom Alerting commands on the Administrator Console panel provide the ability to upload, update, or remove custom alerting classes on the SQL Guard server. Each command is described below.

Uploading Custom Alerting

To upload a custom alerting class:

1. Open the Administration Console panel (not shown).
2. In the Custom Alerting section of the Administration Console menu, click Upload to open the Upload Custom Alerting panel:

The image shows two overlapping panels. The top panel is titled 'Upload Custom Alerting' and has fields for 'Description' (empty) and 'Class File' (empty with a 'Browse...' button). The bottom panel is titled 'Update Custom Alerting' and has fields for 'Description' (containing 'gregtest 2.8'), 'Full Class Name' (containing 'com.guardium.custom.CASecureTest08 FileAccess'), and 'Class File' (empty with a 'Browse...' button). Both panels have an 'Apply' button with a checkmark icon. Below the 'Update' panel, there is an orange message: 'In use by: [Statistical Alert] An alert to test system security, [Statistical Alert] java custom class alerts, [Statistical Alert] java custom classes'.

3. Enter a unique description for the custom alerting class in the Description box.
4. Click the Browse button to open an operating-system dependent file navigation window, from which you can select the class file to be uploaded. The full path name of the class file is inserted into the Class File box.
5. Click the Apply button after you have selected the file (or typed the full path name in the Class File box). This action uploads the class and closes the panel.

Updating Custom Alerting

To update a custom alerting class:

1. Open the Administration Console panel (not shown).
2. In the Custom Alerting section of the Administration Console menu, click Update to open the Update Custom Alerting panel:

The image shows the 'Update Custom Alerting' panel. It has fields for 'Description' (containing 'gregtest 2.8'), 'Full Class Name' (containing 'com.guardium.custom.CASecureTest08 FileAccess'), and 'Class File' (empty with a 'Browse...' button). Below the fields is an orange message: 'In use by: [Statistical Alert] An alert to test system security, [Statistical Alert] java custom class alerts, [Statistical Alert] java custom classes'. At the bottom is an 'Apply' button with a checkmark icon.

3. In the Description box, select the class you want to update. All alerts using the class will be displayed, as illustrated above. The full class name displays in the Full Class Name field.
4. Click the Browse button to open an operating-system dependent file navigation window, from which you can select the class file to replace the existing class file. The full path name of the class file is displayed in the Class File box.

5. Click the Apply button after you have selected the file (or typed the full path name in the Class File box). This action updates the class and closes the panel.

Removing Custom Alerting

1. Open the Administration Console panel (not shown).
2. In the Custom Alerting section of the Administration Console menu, click Remove to open the Remove Custom Test panel:



3. Select, in the description box, the class to be removed.
4. Click Apply to remove the class. You are prompted to confirm the action.

Configuring Permission to Socket Connection for Custom Alerting

See [Configuring Permission to Socket Connection](#) which describes this procedure for both Custom Alerting and Custom Assessment Tests.

Identifying Application Users and Events

Some database applications are designed to use or share a small number of database user accounts. These applications manage their users independently of the database management system, which means that when observing database traffic from outside of the application, it can be difficult to determine the application user who is controlling a database connection at any given point in time. However, when questionable database activities occur, you need to relate specific actions to specific individuals, rather than to an account shared by groups of individuals. In other words, you must know the *application* user, not just the *database* user.

When tracking usage internally, many applications use database stored procedures to identify the application user. Should this be the case, user information can be extracted easily by SQL Guard (see [Identification Via Stored Procedures](#), below).

In other situations, the application user may not be identified by a stored procedure call. In those instances, you can use the SQL Guard Application Events API. The Application Events API allows you to signal SQL Guard when an application user takes or relinquishes control of a connection or when any other event of interest occurs (see [Identification Using the Application Events API](#), below).

It may be necessary to use both methods to identify users, depending on the application. Regardless of the method used, when creating SQL Guard queries and reports, you can access the application user name and event attributes (event type, value, date, etc.) from either the Access Tracking domain or the Policy Violations domain (described in more detail later).

Identification via Stored Procedures

As mentioned above, in many existing applications, all of the information needed to identify an application user can be obtained from existing database traffic, from stored procedure calls. This can be achieved once SQL Guard knows what calls to watch for and which parameters contain the user name or other information of interest.

In the simplest case, an application might have a single stored procedure that sets a number of property values, one of which is the user name. A call to set the user name might look like this:

```
set_application_property('user_name', 'JohnDoe');
```

In a custom procedure definition (described later), you can tell SQL Guard to watch for a stored procedure named *set_application_property* with a first parameter value of *user_name*. When this is found, you can set the application user to the value of the second parameter in the call (*JohnDoe*, in the example above).

There may be multiple stored procedures for an application: one to start an application user session, one to end a session, and others to signal key events particular to that application. SQL Guard's custom identification procedure mechanism can be used to track *any* application events you want to monitor.

Since each of your applications may have a different way of identifying users, you may have to define separate custom identification procedures for each application. For instructions on how to set the application user based on information contained in existing stored procedures, see [Defining Custom Identification Procedures](#), below.

Identification Using the Application Events API

In some cases, the application user may not be identified in a stored procedure call or it may be impossible to extract that information from existing procedure calls. When this happens, you can use the SQL Guard Application Events API. The Application Events API provides simple “no-op” calls that can be issued from within your application to signal SQL Guard when a user acquires or releases a connection or when any other event of interest occurs.

The syntax for each type of call is described below.

Note: If your SQL Guard security policy has Selective Audit Trail enabled, the Application Events API commands used to set and clear the application user and/or application events will be ignored by default, and the application user names and/or application events will not be logged. To log these items so that they will be available for reports or exceptions, you should include a policy rule to recognize the appropriate commands, specifying the Audit Only rule action.

Setting the Application User via GuardAppUser

Use this call to indicate that a new application user has taken control of the connection. The supplied application user name will be available in the Application User attribute of the Access Period entity (illustrated to the right). For this session, from this point on, SQL Guard will attribute all activity on the connection to this application user, until SQL Guard receives either another *GuardAppUser* call or a *GuardAppUserReleased* call (which clears the application user name, as described below).

If you also want to signal when other events occur (you can define event types as needed), use the *GuardAppEvent* call, described in the following section.

Syntax: `SELECT 'GuardAppUser:user_name' FROM location`

 Access Period
 Period Start
 Period Start Date
 Period Start Weekday
 Period Start Time
 Timestamp
 Period End
 Period End Date
 Period End Weekday
 Period End Time
 Application User
 Average Execution Time
 Application Event Id
Total Records Affected
Total Records Affected (Desc)

user_name is a string containing the application user name. This string will be available as the Application User attribute value in the Access Period entity.

FROM location is used only for Oracle, DB2, or Informix. (Omit for other database types.) It must be entered exactly as follows:

Oracle: FROM DUAL

DB2: FROM SYSIBM.SYSDUMMY1

Informix: FROM SYSTABLES

Use the GuardAppUserReleased call to signal that the current user has relinquished control of the connection. SQL Guard will clear the application user name, which will remain empty until SQL Guard receives another GuardAppUser call.

Syntax: SELECT 'GuardAppUserReleased' FROM location

FROM location is used only for Oracle, DB2, or Informix. (Omit for other database types.) It must be entered exactly as follows:

Oracle: FROM DUAL

DB2: FROM SYSIBM.SYSDUMMY1

Informix: FROM SYSTABLES

Setting an Application Event via GuardAppEvent

This call provides a more generic method of signaling the occurrence of application events. You can define your own event types and provide text, numeric, or date values to be stored with the event. You may want to use this call together with the GuardAppUser call described above. SQL Guard will attribute all activity on the connection to this application event, until SQL Guard receives either another GuardAppEvent:Start command or a GuardAppEvent:Released command.

Syntax: SELECT 'GuardAppEvent:Start|Released',
'GuardAppEventType:type',
'GuardAppEventUserName:string',
'GuardAppEventStrValue:string',
'GuardAppEventNumValue:number',
'GuardAppEventDateValue:date' FROM location

Start | Released use the keyword Start to indicate that the event is taking control of the connection or Released to indicate that the event has relinquished control of the connection.

type identifies the event type. It can be any string value, for example: Login, Logout, Credit, Debit, etc.

string is any string value to be set for this event. For example, for a Login event you might provide an account name.

number is any numeric value to be set for this event. For example, for a Credit event you might supply the transaction amount.

date is a user-supplied date and optional time for this event. It must be in the format **yyyy-mm-dd hh:mm:ss**, where the time portion (hh:mm:ss) is optional. It may be the current date and time or it may be taken from a transaction being tracked.

FROM location is used only for Oracle, DB2, or Informix. (Omit for other database types.) It must be entered exactly as follows:

Oracle: FROM DUAL

DB2: FROM SYSIBM.SYSDUMMY1

Informix: FROM SYSTABLES

The GuardAppEvent call populates an Application Events entity containing the attributes illustrated to the left. When creating SQL Guard queries and reports, you can access the Application Events entity from either the Access Tracking domain or the Policy Violations domain.

If any Application Events entity attributes have not been set using the GuardAppEvent call, those values will be empty.

Regarding the two date attributes, Event Date is set using the GuardAppEvent call (or from a custom identification procedure as described in the following section), and Timestamp is the time that SQL Guard stores the instance of the Application Event entity.

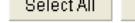
Defining Custom Identification Procedures

A *custom identification procedure* definition tells SQL Guard how to extract application user or event information from an existing stored procedure.

To define a custom identification procedure:

1. Open the Administration Console panel (not shown).
2. In the Custom Identification Procedures section of the Administration Console menu, click Manage Custom ID Procedures to open the Custom Identification Procedures panel:

Custom Identification Procedures									
Select	Custom Map Name	Procedure Name	Action	DB Username	Server Type	Server IP	Database Name	More Info.	
<input type="checkbox"/>	setAppEventMap	setAppEvent	Set	sa	MS SQL SERVER	ANY			
<input type="checkbox"/>	setAppUserMap	setAppUser	Set	sa	MS SQL SERVER	ANY			
<input type="checkbox"/>	userMojganStart	startUserMojgan	Set	sa	MS SQL SERVER	ANY			



To view an existing map, hold the mouse pointer over the More Info column Question Mark icon for the row containing the map you want to view:

Custom Identification Procedures									
Select	Custom Map Name	Procedure Name	Action	DB Username	Server Type	Server IP	Database Name	More Info.	
<input type="checkbox"/>	setAppEventMap	setAppEvent	Set	sa	MS SQL SERVER	ANY			
<input checked="" type="checkbox"/>	setAppUserMap	setAppUser	Set	sa	MS SQL SERVER	ANY			
<input type="checkbox"/>	setCCEventMap	setCCEvent	Set	sa	Custom Map Name: setAppUserMap Condition1 Location: 0 Condition1 Value: Condition2 Location: 0 Condition2 Value: Application Username Position: 2 Event String Value Position: 0 Event Number Value Position: 0 Event Type Position: 0 Event Date Position: 0				
<input type="checkbox"/>	userMojganStart	startUserMojgan	Set	sa					





- Click on the Add Mapping pane title to expand the Add Mapping pane:

The screenshot shows a configuration dialog titled "Add Mapping...". It is organized into several sections:

- Custom Map Name:** A text input field.
- Procedure Name:** A text input field.
- Action:** A dropdown menu set to "Set".
- Condition1 Location:** A text input field containing "0".
- Condition1 Value:** A text input field.
- Condition2 Location:** A text input field containing "0".
- Condition2 Value:** A text input field.
- Parameter Position:** A section containing five pairs of input fields:
 - Application Username Position: "0"
 - Event String Value Position: "0"
 - Event Number Value Position: "0"
 - Event Type Position: "0"
 - Event Date Position: "0"
- Server Information:** A section containing four pairs of input fields:
 - Server Type: A dropdown menu set to "-----".
 - DB Username: A text input field.
 - Server IP: A text input field.
 - Server IP Group: A dropdown menu set to "-----".
 - Database Name: A text input field.
 - Server Net Mask: A text input field.
- Groups...**: A button with a icon showing a group of servers.
- Add**: A green button with a plus sign icon.

4. In the Custom Map Name box, enter the name to be used for this mapping.
5. In the Procedure Name box, enter the name of the database procedure that will supply information.
6. Select Set or Clear from the Action list to indicate whether the procedure call will set or clear application values. The Event Type Position field has a special use when the Clear action is selected. (See Step 8, below.)
7. If application information can be obtained from an existing stored procedure call, but only under one or two conditions, use a Condition Location box to specify which stored procedure call parameter is to be tested, and the corresponding Condition Value box to specify the value that must be matched to set application information from one or more of the other parameters.

For example, assume that a stored procedure named *set_context* is used by an application to set a number of values, one of which is the user name. The procedure is passed through three parameters: an application name, a property name, and a value. Three typical calls are illustrated below:

```
set_context('publishing_application', 'role_name', 'manager');
set_context('publishing_application', 'user_name', 'jsmith');
set_context('publishing_application', 'company', 'guardium');
```

In the examples above, the second statement illustrates the format of the call we are interested in. The second parameter (the property name) is the parameter that needs to be tested, 2 is entered in the Condition1 Location box, and *user_name* in the Condition1 Value box.

If a second format of the call also sets the user name, then the Condition2 Location and Value boxes can be used. For example, assume that the following format of the procedure call is sometimes used to set a user name:

```
set_context('admin_application', 'admin_name', 'wjones');
```

To use this procedure to set the application user name, enter 2 in the Condition2 Location box, and *admin_name* in the Condition2 Value box.

Note: If two conditions are used, the user name or any other information being extracted (see below) must be in the same parameter position for both types of calls.

8. For a *Clear action*, use only the **Event Type Position / Application Username Position** fields in the Parameter Position pane. Set the Event Type Position field to 1 and Application Username Position to 0 to clear the application event, or set the Application Username Position to 1 and Event Type Position field to 0 to clear the application user.

For a *Set action*, use the Parameter Position pane to indicate which stored procedure parameters map to which SQL Guard application event attributes. The first procedure parameter is numbered 1. Use 0 (zero – the default) for all attributes *not* set by the call.

- **Application Username Position** – Enter the parameter position of the application user name you want associated with database activity from this point forward (until reset, as described previously).
- **Event Number Value Position** – Enter the parameter position of a numeric value for the event (for a transaction, this might be a dollar amount).
- **Event Date Position** – Enter the parameter position of a date/time value for the event. The format must be **yyyy-mm-dd hh:mm:ss**. The time portion (hh:mm:ss) is optional, and if omitted will be set to 00:00:00.
- **Event String Value Position** – Enter the parameter position of a string value for the event (for a login, this might be a user or account name).
- **Event Type Position** – Enter the parameter position of a name for the event type (Login, Logout, Credit Request, etc.).

9. In the Server Information pane:

- Select the database server type from the Server Type list.

- Enter the database user name in the DB Username box.
 - Optional: Enter a database name in the Database Name box. If omitted, all databases will be monitored.
 - Optional: Identify one or more servers. If no server is specified, all servers will be monitored.
 - *To select a specific server only*, enter the server IP address and network mask in the Server IP and Server Net Mask boxes
OR
 - *To select a group of servers*, select a server group from the Server IP Group list or click the Groups button to define a new group of servers.
10. When you are done, click the Add button to add the Custom Identification Procedure.

Using the Task Scheduler

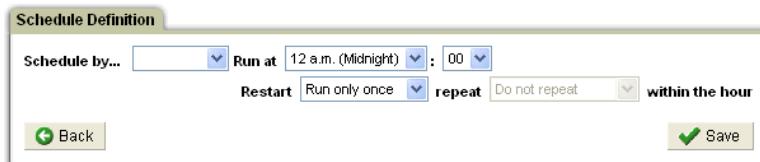
This section describes how to use the general-purpose task scheduler that is used for many SQL Guard functions (Data Export, Data Import, etc.). The specific task illustrated here is the IP-to-Hostname Aliasing. However, the Scheduling pane illustrated below and the panels used by the scheduler all contain the same controls.

Defining a Schedule



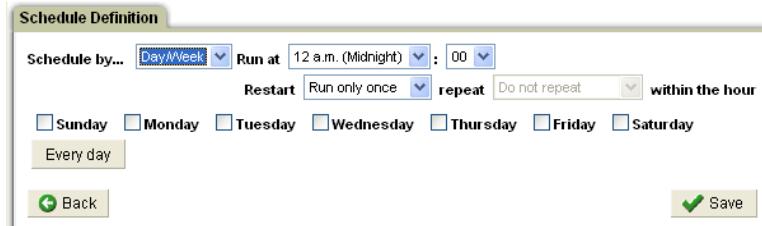
1. To define a schedule for a configured task, click the Define Schedule button:

Note: If the Define Schedule button is not active, the configuration has not yet been saved. A task cannot be scheduled until it has been saved (and in the process, validity checked).



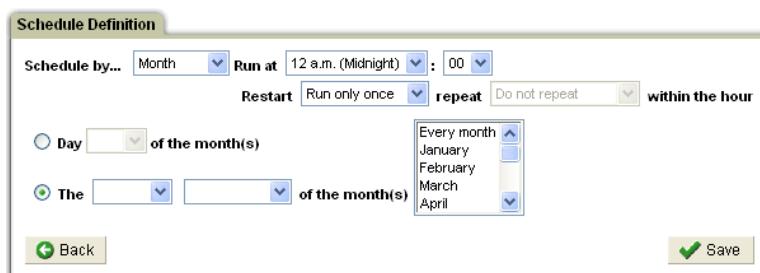
2. Select from the *Schedule by* list the Day/Week to build a schedule based on the days of the week or select Month to build a schedule based on a days of a month. Both options are described separately, below, following a description of how to use the common fields (Run at, Restart, and Repeat).
3. Select from the *Run at* list the hour of the day during which you want to run the task for the first time.
4. Select from the minutes list (following the *Run at* list) the minute within the selected *Run at* hour when you want to run the task for the first time.
5. Select from the Restart list when to restart the task: *Run only once* (do not restart the task), or from 1-12 hours from the selected *Run at* hour.

6. Select from the Repeat list the number of minutes within the hour to repeat the task: *Do not repeat*, or repeat every 1-59 minutes.
7. If you selected Day/Week from the *Schedule by* list:



- o Mark each day of the week you want the task run.
- OR**
- o Click *Every day* to select all days (or to clear all days if already selected).

If you selected Month from the *Schedule by* list:



Do one of the following:

- o **To select a numbered day (the 15th, for example):**
 - Select the *Day* button
 - Select a day: 1-31, depending on the month selected
 - Select from the list of months to the right either *Every month*, or one or more specific months: January, February, March, etc.
- o **To select a weekday occurrence (the first Monday, for example):**
 - Select the lower button (beside the word *The*).
 - From the first list, select a week relative to the start of the month: First, Second, Third, etc.
 - From the second list, select a weekday: Sunday, Monday, Tuesday, etc.

- From the third list, select either *Every month* or one or more specific months: January, February, March, etc.
8. Click the Save button to save the schedule or click Back to close the window. In the latter case, you are prompted to save your changes before closing the window.

Removing a Schedule

Once you have saved a schedule, a Remove button appears on the panel. Click the Remove button to remove the schedule. You are prompted to confirm the action. Removing the schedule clears the defined schedule.

Pausing a Schedule

Once you have saved a schedule and activated it, a Pause button may appear on the panel for some applications. Click the Pause button to pause the schedule. You are prompted to confirm the action.

Monitoring Disk Space Use

Use the Current Status Monitor graphical report (described below) on the Daily Monitor tab to monitor the amount of free disk space on the SQL Guard system.

Near Capacity Alert

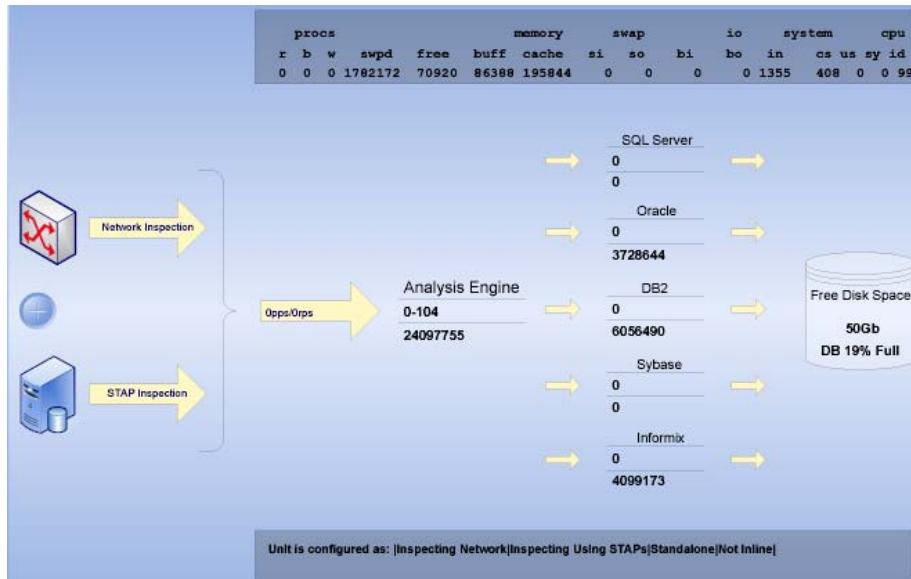
If disk storage reaches 85% of capacity, SQL Guard issues an alert so that archiving operations can be performed to free disk space. When this condition arises, and thereafter every five minutes until the capacity goes below 85%, SQL Guard:

- Writes a message to the syslog file to inform the system administrator.
- Sends an e-mail message to Guardium Technical Support.

Current Status Monitor

The Current Status Monitor graphical report on the Daily Monitor tab displays the current state of the SQL Guard unit: how many packets per second and requests per second it is processing, how much disk space, memory, etc. Each field is described below.

Note: To view the Current Status Monitor, you need the Adobe SVG Viewer. See [Software Downloads from Adobe](#).



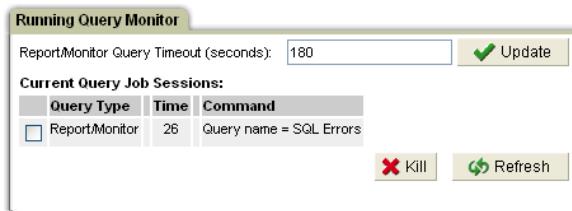
The top box displays the output of the Linux VMSTAT command. If you are familiar with that command, these statistics should be familiar to you.

Field	Description
procs	The number of processes: r : Waiting for run time. b : In uninterruptable sleep (blocked, waiting for another event). w : Swapped out but otherwise runnable.
memory	Memory use (kB): swpd : Amount of virtual memory used. free : Amount of idle memory. buff : Amount used as buffers. cache : Amount reserved for cache.
swap	Amount of memory (kB): si : Swapped in from disk. so : Swapped out to disk.

Field	Description
io	Input/Output blocks (kB/s): <ul style="list-style-type: none"> bi: Blocks received from a block device bo: Blocks sent to a block device
system	System: <ul style="list-style-type: none"> in: Interrupts per second, including the clock cs: Context switches per second
cpu	Percentage of total CPU time used by: <ul style="list-style-type: none"> us: Users sy: System id: Idle
(n)pps / (m)rps	In the arrow to the left of the Analysis Engine, two averages are calculated for the last five seconds: <i>n</i> is the average number of network packets per second, and <i>m</i> is the average number of network database requests per second.
Analysis Engine (queued – dropped) ----- (processed)	For the Analysis Engine, the upper line lists the total number of messages awaiting processing, followed by the number of messages dropped because the buffer was in danger of becoming filled. The lower line lists the total number of messages processed. The number processed will be reset to zero whenever the inspection engine is restarted.
Server Type (queued) ----- (processed)	For each server type, the number of messages awaiting processing is listed above, and the number of messages processed is listed below.
Free Disk Space	The number of bytes free.
DB n% Full	The percentage of the database space allocation that is used.

Using the Running Query Monitor

The Running Query Monitor displays the status of active user queries, and allows you to set a timeout value for all Report/Monitor queries. To open the Running Query Monitor panel, click Running Query Monitor from the Daily Monitor tab menu:



From the Running Query Monitor, you can:

- Set the query timeout for all reports and monitors *running in a portlet*. Other query processes, such as policy simulations, audit processes, baseline generations and internal processes *are not affected* by this timeout value. The default is 60 seconds. It has been modified in the example above.
- Kill any currently running user query. Some queries listed here – audit processes, for example, may exceed the query timeout specified. That is expected, because the Report/Monitor query timeout applies to reports and monitors running in a portlet only.

We do not recommend setting the Query Timeout above the default setting (60 seconds) for an extended period of time. If you set this limit upwards, it will increase the chances of overloading the system with ad hoc reporting activity.

To change the timeout setting, type a number of seconds in the Report/Monitor Query Timeout box, and click the Update button. You will be informed when the update has been completed.

To kill running query, mark it in the list and click the Kill button.

The query type will be one of the following: Report/Monitor, Audit Process, Policy Simulation, Configuration, or Definitions.

Monitoring via SNMP

There is an SNMP agent installed on Guardium systems, and read-only access is provided using the SNMP community name of **guardiumsnmp**. When querying, a value of -1 (minus one) indicates a NULL in the database. The table at the end of this section lists the available SNMP OIDs.

SNMP Examples

From a Unix session, you can display SQL Guard SNMP information using the **snmpget** or **snmpwalk** commands. (Use *snmpget -h* or *snmpwalk -h* to display command syntax.)

Under Windows (and also under Unix), various GUI-based software packages are available for displaying SNMP information. Those alternatives are not described here.

Disk space used and available:

```
> snmpget -v 1 -c guardiumsnmp g3.guardium.com UCD-SNMP-MIB::dskAvail.1  
UCD-SNMP-MIB::dskAvail.1 = INTEGER: 1043856  
> snmpget -v 1 -c guardiumsnmp g3.guardium.com UCD-SNMP-MIB::dskUsed.1  
UCD-SNMP-MIB::dskUsed.1 = INTEGER: 914856
```

To list total memory and used memory:

```
> snmpget -v 1 -c guardiumsnmp g3.guardium.com  
HOST-RESOURCES-MIB::hrStorageSize.101  
HOST-RESOURCES-MIB::hrStorageSize.101 = INTEGER: 2067352  
> snmpget -v 1 -c guardiumsnmp g3.guardium.com HOST-RESOURCES-MIB::hrStorageUsed.101  
HOST-RESOURCES-MIB::hrStorageUsed.101 = INTEGER: 1017548
```

To list the available memory:

```
> snmpwalk -v 1 -c guardiumsnmp g3.guardium.com memAvailReal  
UCD-SNMP-MIB::memAvailReal.0 = INTEGER: 1049564
```

To list values relating to cpu usage:

```
> snmpwalk -v 1 -c guardiumsnmp g3.guardium.com ssCpuRawUser  
UCD-SNMP-MIB::ssCpuRawUser.0 = Counter32: 89240  
> snmpwalk -v 1 -c guardiumsnmp g3.guardium.com ssCpuRawSystem  
UCD-SNMP-MIB::ssCpuRawSystem.0 = Counter32: 195310  
> snmpwalk -v 1 -c guardiumsnmp g3.guardium.com ssCpuRawNice  
UCD-SNMP-MIB::ssCpuRawNice.0 = Counter32: 11
```

Note: Adding the RawUser, RawSystem, and RawNice numbers provides a good approximation of total CPU usage.

```
> snmpwalk -v 1 -c guardiumsnmp g3.guardium.com ssCpuRawIdle  
UCD-SNMP-MIB::ssCpuRawIdle.0 = Counter32: 26734332
```

Guardium SNMP OID

SNMP OID	Description
.1.3.6.1.4.1.2021.9.1.7.1 UCD-SNMP-MIB::dskAvail.1	Disk space available in / directory

SNMP OID	Description
.1.3.6.1.4.1.2021.9.1.7.2 UCD-SNMP-MIB::dskAvail.2	Disk space available in /var directory
.1.3.6.1.4.1.2021.9.1.8.1 UCD-SNMP-MIB::dskUsed.1	Disk space used in / directory
.1.3.6.1.4.1.2021.9.1.8.2 UCD-SNMP-MIB::dskUsed.2	Disk space used in /var directory
.1.3.6.1.2.1.25.2.3.1.5.1 HOST-RESOURCES-MIB::hrStorageSize.1	Total memory available
.1.3.6.1.2.1.25.2.3.1.6.1 HOST-RESOURCES-MIB::hrStorageUsed.1	Memory in use
.1.3.6.1.4.1.2021.8.1.101.1 UCD-SNMP-MIB::extOutput.1	Open monitored session count
.1.3.6.1.4.1.2021.8.1.101.2 UCD-SNMP-MIB::extOutput.2	Requests logged by the current sniffer process (set to zero for each restart)
.1.3.6.1.4.1.2021.8.1.101.3 UCD-SNMP-MIB::extOutput.3	Last session timestamp
.1.3.6.1.4.1.2021.8.1.101.4 UCD-SNMP-MIB::extOutput.4	Last construct timestamp
.1.3.6.1.4.1.2021.8.1.101.5 UCD-SNMP-MIB::extOutput.5	Memory used by the sniffer process
.1.3.6.1.4.1.2021.8.1.101.7 UCD-SNMP-MIB::extOutput.7	Packets in on ETH 1 / out on ETH 2 ; usually only one number (inbound) when a SPAN port or TAP is used
.1.3.6.1.4.1.2021.8.1.101.8 UCD-SNMP-MIB::extOutput.8	Same as above, for ETH 3 / ETH 4
.1.3.6.1.4.1.2021.8.1.101.9 UCD-SNMP-MIB::extOutput.9	Same as above, for ETH 5 / ETH 6

Other MIBs accessible in the machine are: SNMPv2-MIB, IF-MIB, RFC1213-MIB, and HOST-RESOURCES-MIB. These MIB descriptions are available on the Internet.

Chapter 3: User Management

About Users

To make the SQL Guard system workable, you must define those users who will be employing it. Creating and modifying users involves deciding both who will be using the SQL Guard system and to what roles they will be assigned. A *role* is a group of users, all of whom are granted the same access privileges. For more information on roles, see [Chapter 4: Security Role Management](#).

User definitions can be imported from an LDAP server, on demand, at a specific time, or on a periodic basis. For more information, see [Importing Users from an LDAP Server](#), at the end of this chapter.

Before you connect and configure the SQL Guard system, identify which groups of users will use that system and what their function will be. For example, an information security group might use the SQL Guard system for alerting and troubleshooting purposes; a database administrator group might use the SQL Guard system for reporting and monitoring.

Once you decide which general roles/groups of users will use the SQL Guard system (and for what purpose), collect the following information for each user:

- User's first and last name
- User account name (the name they will use to log in)
- User's email address
- User's function/role with SQL Guard

Note: When deciding who will access the SQL Guard system, keep in mind that sensitive company data can be picked up by the system. Therefore, be very aware of who will be able to access that data.

User Account Security

Several SQL Guard system settings can be changed to provide additional security for user accounts. These additional features are disabled by default. You can enable or modify them using the CLI, as described in Chapter 6.

- By default password validation is enabled. For more information, see the descriptions of the [show password](#) and [store password](#) commands in Chapter 6.

- You can enable password expiration. For more information, see the descriptions of the [show password](#) and [store password](#) commands in Chapter 6.
- You can enable automatic account disabling following a specified number of failed login attempts. For more information, see the descriptions of the [show account](#) and [store account](#) commands in Chapter 6

Notes: The SQL Guard administrator can enable a disabled user account by clearing the Disabled check-box for that user in the User Maintenance panel (described later in this chapter).

If the special **admin** user account becomes locked, use the **unlock admin** CLI command to unlock it.

Sample User Accounts

During the installation process, SQL Guard is preconfigured with:

- A set of default roles.
- A set of sample user accounts, which are provided as examples of the roles for which they are named.
- The *admin* user account, an account with special privileges to be used by the SQL Guard administrator. For more information, see [About the admin User](#) in Chapter 4.

The administrator should remove any unnecessary sample user accounts. All remaining user accounts should be updated to identify their actual owners, have strong passwords assigned, and have correct email addresses and roles set. The user accounts should be configured to match the enterprise security policies for user management.

The administrator can create new roles and accounts or alter the predefined roles and accounts. The administrator can also create users associated with the predefined roles. Creating and modifying users involves deciding both who will be using the SQL Guard system and to what roles they will be assigned.

User accounts that access the Web-based GUI are granted access to functions and features of the SQL Guard GUI through the use of roles. For more information on roles in general, see [Chapter 4: Security Role Management](#).

Managing User Accounts

The User Browser allows you to add, edit, and delete users, and to assign roles to users.

Note: In Central Manager environments, all User Accounts, Roles, and Permissions are controlled by the Central Manager. To administer any of these definitions, you must be logged into the Central Manager (and not to a managed unit).

To access the User Management screens:

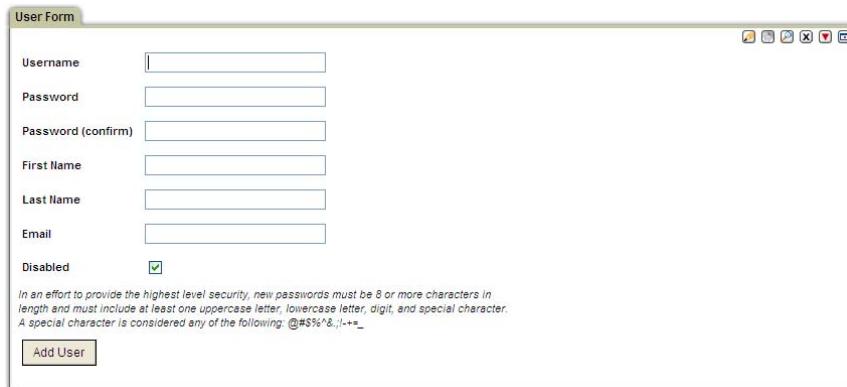
1. Click the Access Management tab to bring it to the front.
2. Click the User Browser command in the menu to the left, to open the User Browser panel (not shown).

Adding a New User

If you are logged into a unit that is controlled by a Central Manager, you cannot add new users. That function can only be performed when logged into the Central Manager.

To add a new user:

1. Click the Add User link at the bottom of the User Browser to display the User Form panel:



The screenshot shows the 'User Form' window. It contains fields for Username, Password, Password (confirm), First Name, Last Name, and Email. There is also a 'Disabled' checkbox which is checked. Below the form, there is a note about password security and an 'Add User' button.

In an effort to provide the highest level security, new passwords must be 8 or more characters in length and must include at least one uppercase letter, lowercase letter, digit, and special character. A special character is considered any of the following: @#\$%^&.,!+=_-

Add User

2. Enter a unique username in the Username box. Do not include apostrophe characters in the name. User names are *not* case sensitive.
3. Enter a password for the new user in the Password box. Passwords *are* case sensitive.

When password validation is enabled, the password must be eight or more characters in length, and must include at least one uppercase alphabetic character (A-Z), one lowercase alphabetic character (a-z), one digit (0-9), and one special character from the table below.

Table of Special Characters for SQL Guard Passwords

Special Character	Character Name
@	Commercial at
#	Number sign
\$	Dollar sign
%	Percent sign
^	Circumflex accent (carat)
&	Ampersand
.	Full stop (Period)
;	Semicolon
!	Exclamation mark
-	Hyphen (minus)
+	Plus sign
=	Equals sign
_	Low line (underscore)

If password validation is disabled, any characters are allowed.

If password expiration is enabled (see [Password Expiration](#) earlier in this chapter), the password you assign will be temporary and the user will be required to change it following the first login.

4. Enter the user's first name in the First Name box.
5. Enter the user's last name in the Last Name box.

Note: [About Investigation Users](#)

An *investigation user* must have the restore-to database name of **INV_1**, **INV_2** or **INV_3** as the Last Name. This is not enforced by the GUI, but is required for the application to function properly.

An investigation user must have the role **inv**, and no other roles, including **user**. This is the only case where the user or admin role is not required.

6. Enter the user's email address in the Email box.
7. Clear the disabled box to enable the user. The new user will not be able to log in if this box is marked.

8. Click Add User to save the new user account definition and close the panel.

Adding Roles to a User

For information on roles and their function, see [Chapter 4: Security Role Management](#). When you add a role to a user definition, you grant that user access to all components that are also assigned that role.

For a new user who has not logged into the system for the first time, the assignment of a new role will affect that user's initial portal layout, if that role also includes default tabs, menus, portlets, etc.. For more information, see [Chapter 5: Portal Management](#).

For a user who has logged into the system at least once, the assignment of a new role grants access privileges, but does not change the portal layout. If an existing user wants have all of the same tabs, menus, portlets, etc. that a new user of the role would have, the existing user will have to manually add those components the portal. Refer to Chapter 3 of the *User Guide* for detailed instructions for customizing a portal.

If you are logged into a unit that is controlled by a Central Manager, you cannot modify role assignments. That function can only be performed when logged into the Central Manager.

To associate a role with a user:

1. In the User Browser panel, click the Roles link in the row of the user whose roles you want to change.

Role Name	Assign
user	<input checked="" type="checkbox"/>
admin	<input type="checkbox"/>
guest	<input type="checkbox"/>
dba	<input type="checkbox"/>
infosec	<input type="checkbox"/>
netadm	<input type="checkbox"/>
appdev	<input checked="" type="checkbox"/>
audit	<input type="checkbox"/>
cas	<input type="checkbox"/>

2. In the User Role Form panel, for each role assigned to this user, mark the corresponding checkbox in the Assign column. *Always mark the user role.*
3. Click the Update button to save the changes and close the panel.

Editing User Information

If you are logged into a unit that is controlled by a Central Manager, you cannot edit user information. That function can only be performed when logged into the Central Manager.

To edit a user's account information:

1. In the User Browser, click the Edit link on the row of the user whose account information you want to edit to open the User Form panel:

The screenshot shows the 'User Form' window. It contains the following fields:

Username	newuser
Password	[REDACTED]
Password (confirm)	[REDACTED]
First Name	new
Last Name	user
Email	newuser@guardium.com
Disabled	<input type="checkbox"/>
Last Login	2007-04-13 09:57:41.0
Password Last Changed	2007-04-13 09:57:57.0

In an effort to provide the highest level security, new passwords must be 8 or more characters in length and must include at least one uppercase letter, lowercase letter, digit, and special character. A special character is considered any of the following: @#\$%^&.,!-=_

Update User

2. Edit the Password, First and Last Names, and Email address fields if you want.
3. Mark the Disabled box to disable this user account; or clear the Disabled box to enable a disabled account.
4. Click Update User to save all changes and close the panel.

Removing a User

If you are logged into a unit that is controlled by a Central Manager, you cannot remove users. That function can only be performed when logged into the Central Manager.

To remove a user:

1. In the User Browser, click the Remove link on the row of the user whose account you want to remove. The User Form panel opens with a Confirm Deletion button:

The screenshot shows a 'User Form' window with the following fields:

Field	Value
Username	newuser2
Password	[REDACTED]
First Name	new
Last Name	user2
Email	newiser2@guardium.com
Disabled	<input type="checkbox"/>
Last Login	2007-04-13 10:12:18.0
Password Last Changed	2007-04-13 10:20:04.0

A yellow button labeled 'Confirm Deletion' is visible at the bottom left.

2. Click the Confirm Deletion button to remove the user and close the panel.

Importing Users from an LDAP Server

You can import SQL Guard user definitions from an LDAP server by configuring an import operation to obtain the appropriate set of users. You can run the import operation on demand, or schedule it to run on a periodic basis. You can elect to have only new users imported, or you can have existing user definitions replaced. In either case, LDAP groups can be imported as SQL Guard roles.

When importing LDAP users:

- The SQL Guard **admin** user definition will not be changed in any way.
- Existing users will not be deleted (in other words, the entire set of users is not replaced by the set imported from LDAP).
- SQL Guard passwords will not be changed.
- New users being added to SQL Guard:
 - Will be marked *Inactive*
 - Will have blank passwords
 - Will have the User role marked

Configuring LDAP User Import

To configure user import from an LDAP server, follow the procedure outlined below.

1. Click the Access Management tab to bring it to the front.
2. Select LDAP Import from the menu to open the LDAP User Import panel:

3. In the LDAP Host Name box, enter the IP address or host name for the LDAP server to be accessed.
4. In the Port box, enter the port number for connecting to the LDAP server.
5. Select the LDAP server type from the Server Type list.
6. Mark the Use SSL Connection checkbox to have SQL Guard connect to your LDAP server using an SSL (secure socket layer) connection. *We recommend using an SSL connection for SQL Guard – LDAP communications.*

Note: If an SSL connection is not used, the query results returned by your LDAP server will not be encrypted, exposing potentially sensitive information to anyone with access to network traffic.

7. In the Base DN box, specify the node in the tree at which to begin the search (in the example above, DC=encore,DC=corp,DC=root).
8. For Import Mode, select either *Add on* to add (but not replace) users, or select *Override* to replace existing user definitions (except for passwords). In either case, if LDAP groups are imported as roles (see below), any new roles will be merged with existing roles.
9. Mark the Import Roles box to add LDAP user groups as SQL Guard roles.
10. In the Role Filter box, optionally enter LDAP search criteria, to limit the roles returned. For example, *syy** would return only those roles beginning with the characters *syy*. See your LDAP server documentation more information about search criteria.

11. In the Log In As box, enter the user account to use for the connection from the SQL Guard server.
12. In the Password box, enter the password for the above user.
13. In the Search Filter box, optionally enter LDAP search criteria, as illustrated above. Typically, imports will be based on membership in an LDAP group, so the filter might use the **memberOf** keyword and look something like this:
memberOf=CN=syyTestGroup,DC=encore,DC=corp,DC=root
See your LDAP server documentation if you need help in this area.
14. For Search Filter Scope, select One-Level to apply the search to the base level only, or Sub-Tree to apply the search to levels beneath the base level.
15. In the Limit box, enter the maximum number of items to be returned. We suggest that you use this field to test new queries or modifications to existing queries, so that you do not inadvertently load a large number
16. Click the Apply or Update button to save the configuration. (After you have saved the configuration once, the Apply button becomes the Update button.)

Where to go from here...

After saving an LDAP user import configuration, you can perform the following tasks, each of which is described in a separate section. Because it is easy to miscode LDAP queries, we suggest that you test each new or modified query by using the Limit field (described in the previous section) and by running the query at least once on demand, to verify that the correct set of members is being returned.

- [Run an LDAP Import On Demand](#)
- [Schedule LDAP Import](#)

Run an LDAP Import On Demand

When you run an LDAP import on demand, you have the opportunity to accept or reject each of the users returned by the query. This is especially useful for testing purposes.

1. Click the Access Management tab to bring it to the front.
2. Select LDAP Import from the menu to open the LDAP User Import panel (illustrated previously), and click Run Once Now on the Scheduling pane.

Note: If the LDAP server is unavailable, or if SQL Guard cannot connect with it using the configuration defined, you will receive an error message: Communication with LDAP server failed. Please check your configuration and try again.

3. After the query is processed by the LDAP server, the set of users satisfying your criteria will be displayed in the LDAP Query Results panel, as illustrated to the right:
4. You can optionally change the Import Mode for this operation only. Changing the selection here will not change the configuration for subsequent operations.
5. Mark the checkbox for each user you want added, and click Import (or click Cancel to return without importing any users). If a selected user has LDAP roles defined (see guarduser20, above), those will be added as well, regardless of the Import Mode selected.



You will be notified that the users have been saved.

6. Click User Browser on the Administration Console menu to view the User Browser panel. You will notice that any users just added will be disabled – which is indicated by use of the strike-through text effect for those users, as illustrated below:

User Browser					
Filter string (case sensitive): <input type="text"/> User Name <input type="button" value="Filter"/>					
Username	First Name	Last Name	Email	Actions	
admin	admin	admin	admin@guardium.com	Edit	Roles Remove
dba	dba	dba	dba@guardium.com	Edit	Roles Remove
infosec	infosec	infosec	infosec@guardium.com	Edit	Roles Remove
netadmin	netadmin	netadmin	netadmin@guardium.com	Edit	Roles Remove
appdev	appdev	appdev	appdev@guardium.com	Edit	Roles Remove
audit	audit	audit	audit@guardium.com	Edit	Roles Remove
jvincent	Jean	Vincent	jvincent@guardium.com	Edit	Roles Remove
guarduser2-	guarduser2-	guarduser2-	guarduser2@encore.corp.root	Edit	Roles Remove
guarduser20-	guarduser20-	guarduser20-	guarduser20@encore.corp.root	Edit	Roles Remove
guarduser21-	guarduser21-	guarduser21-	guarduser21@encore.corp.root	Edit	Roles Remove

[Add User](#)

7. You can enable the new users, one at a time. To enable each new user:
 - o Click the Edit link in the Actions column, to open the User Form panel.
 - o Clear the Disabled checkbox.
 - o Add a password for the new user. (This is a common best practice).
 - o Click Update User to activate the user.

After all users have been enabled, the list of users would appear as follows:

User Browser				
Filter string (case sensitive): <input type="text"/> User Name <input type="button" value="Filter"/>				
Username	First Name	Last Name	Email	Actions
admin	admin	admin	admin@guardium.com	Edit Roles Remove
dba	dba	dba	dba@guardium.com	Edit Roles Remove
infosec	infosec	infosec	infosec@guardium.com	Edit Roles Remove
netadmin	netadmin	netadmin	netadmin@guardium.com	Edit Roles Remove
appdev	appdev	appdev	appdev@guardium.com	Edit Roles Remove
audit	audit	audit	audit@guardium.com	Edit Roles Remove
jvincent	Jean	Vincent	jvincent@guardium.com	Edit Roles Remove
guarduser2	guarduser2	guarduser2	guarduser2@encore.corp.root	Edit Roles Remove
guarduser20	guarduser20	guarduser20	guarduser20@encore.corp.root	Edit Roles Remove
guarduser21	guarduser21	guarduser21	guarduser21@encore.corp.root	Edit Roles Remove

[Add User](#)

You can verify that LDAP groups have been added as roles, either by opening the security Role Browser or by clicking the Roles link for any of the new users.

Schedule LDAP Import

To define an LDAP import schedule:

1. Click the Access Management tab to bring it to the front.
2. Select LDAP Import from the menu to open the LDAP User Import panel (illustrated previously), and click the Modify Schedule button.
3. For instructions on how to use the general-purpose task scheduler, see [Using the Task Scheduler](#) in Chapter 2.

Once a schedule has been defined, a Pause button appears on the LDAP User Import panel:



If you click that button, the schedule is paused, and the Pause button is replaced by a Resume button:



Chapter 4: Security Role Management

SQL Guard uses *security roles* to control access to SQL Guard components (reports, audit process definitions, or SQL Guard applications, for example). When a security role is assigned to any SQL Guard component, only those SQL Guard users who are also assigned that security role can access that component. A security role can thus be viewed as a group of SQL Guard users, all of whom have the same access privileges.

If no security roles are assigned to a component (a report, for example), only the user who defined that component and the special *admin* can access it.

At installation time, SQL Guard is configured with a default set of security roles and a default set of user accounts. The SQL Guard administrator can create new roles or use the predefined ones. The administrator also assigns security roles to SQL Guard users.

When user definitions are imported from an LDAP server, the groups to which they belong can optionally be defined as security roles. For more information, see [Importing Users from an LDAP Server](#) in Chapter 3.

Default Roles

The SQL Guard system is preconfigured by default to support users who fall into certain roles. The SQL Guard administrator can create new roles as well. Users of the Web-based GUI should be assigned the *user* role and one or more additional roles, depending on what access privileges you want to make available to that user.

Each default role comes with a set of preconfigured report templates and tabs on which those reports are displayed. New users receive these views by default, according to the roles defined. The administrator can change the definition for each role and users can modify their views and reports.

Role	Description
user	Provides access to the common user interface elements. Either this role or the admin role (see below) must be assigned to all users. The user role cannot be deleted.

Role	Description
admin	<p>SQL Guard administrators who can access the Administrator Console to perform maintenance on the SQL Guard system. This role cannot be deleted.</p>
	<p>Note: Do not confuse the <i>admin role</i> with the special, privileged user account named <i>admin</i>. The <i>admin role</i> can be granted (or denied) access to any component, but the <i>admin user</i> always remains privileged. See About the admin User, below.</p>
guest	<p>For use as an example role. It cannot be deleted.</p>
dba	<p>Users who have a database-centric view of security, allowing access to database-related reports and tracking of database objects.</p>
infosec	<p>Users who have an information security focus, including tracking access to the database, and handling network requests, audits, and forensics.</p>
netadm	<p>Users who have a network-centric view, including IP sources for database requests.</p>
appdev	<p>Application developers, architects, and QA personnel who have an application centric focus and want to track and report on SQL streams generated by an application.</p>
audit	<p>Auditors and others who need to view audit reports.</p>
cas	<p>Users who need access to the Change Audit System (CAS) query builders and configuration tools.</p>
inv	<p>This is a special role that must be assigned to investigation users. See the note below</p>
	<p>Note: About Investigation Users</p>
	<p>An investigation user must have the role inv, and no other roles, including user. This is the only case where the user or admin role is not required.</p>
	<p>An <i>investigation user</i> must have the restore-to database name of INV_1, INV_2 or INV_3 as the Last Name. This is not enforced by the GUI, but is required for the application to function properly.</p>

Users

SQL Guard users are granted access to SQL Guard functions and features through the use of the above-mentioned roles.

The SQL Guard system is preconfigured to support an *admin* user and a set of example users. The example user accounts are provided as illustrations of the roles for which they are named.

About the admin User

The SQL Guard admin user account (i.e., the user account with the username *admin*) is a privileged account to be used by the SQL Guard system administrator.

If automatic account lockout is enabled (a feature that locks a user account after a specified number of login failures), the special admin user account may become locked. If that happens, use the **unlock admin** CLI command to unlock it.

admin User Write-Access to All Components

The admin user account has full write-access privileges for all components (regardless of what privileges are granted to the sample *admin* role).

For example, when the admin user displays a list of audit task definitions, all audit tasks defined in the SQL Guard internal database display and the admin user can view, modify, or delete any of those definitions.

When any other user accesses that same list of audit tasks, that user will see only those audit tasks that have been assigned a security role that is also assigned to that user.

admin User To-Do List Considerations

The *to-do list* is a workflow management feature that controls the release of audit task results from one user to the next. The admin user has special privileges and responsibilities in this area.

For example, an audit task can be defined such that the results are released only to SQL Guard user **A**, who must review and electronically sign those results before they are released to SQL Guard user **B**, who also must review and sign the results before they are released to the next user or user group. If user **A** leaves the company or is otherwise unavailable, the results would be “stalled” in user A’s to-do list. To remedy the situation where a user account is deactivated, all audit process results intended for that user’s to-do list will be sent to the admin user instead.

In addition, the admin user has special authority access any user’s to-do list and to sign any task results, thus releasing those results to the next user on the receivers list.

Note: The To-Do List is described in the *SQL Guard User Guide*.

Example User Accounts

The administrator should delete any unnecessary users. The SQL Guard software is installed with the set of users described in the table below. Apart from the special admin user account (described previously), these are example accounts.

Note: Except for the admin account, the remaining example accounts in this table are configured with a password that matches the account name. For security reasons, these accounts should be removed or their passwords changed before the SQL Guard system is implemented in a production environment.

User	Description
admin	While this account serves as an example account for the default admin role, note that this account has special privileges and responsibilities and it cannot be deleted. See the above topic, About the admin User .
anon	The anon user is an example default account for the guest role.
dba	The dba user is an example default account for the dba role.
infosec	The infosec user is an example default account for the infosec role.
netadm	The netadm user is an example default account for the netadm role.
appdev	The appdev user is an example default account for the appdev role.
audit	The audit user is an example default account for the audit role.

Adding and Removing Roles

Roles provide access to components. When a user is assigned a role, that user is granted access to all components that are also assigned that role.

Note: In Central Manager environments, all User Accounts, Roles, and Permissions are controlled by the Central Manager. To administer any of these definitions, you must be logged into the Central Manager (and not to a managed unit).

To access the Security Role Management screens:

1. Click the Access Management tab to bring it to the front.
2. Click Security Role Browser in the menu on the left to display the Security Role Browser.

Note: If you are logged into a unit that is controlled by a Central Manager, you will receive an error message informing you that User and Role administration can be done on the Central Manager only.

Adding a Role

Although the SQL Guard system comes with a number of preconfigured roles, you can create additional roles with any permissions and default portlets desired. For more information on adding default portlets and viewing user-added portlet information, see [Chapter 5: Portal Management](#).

To add a role:

1. Click the Add Role button at the bottom of the Security Role Browser to display the Role Maintenance panel.
2. Type the name of the new Role in the Role Name field. Do not include apostrophe characters in the name.
3. Click the Add Role button to add the Role to the Security Role Browser page.

Removing a Role

To remove a Role:

1. Click the Remove link on the row for the role you want to remove. You are prompted to confirm the action:
2. Click the Confirm Deletion button to remove the role and close the panel.

Editing Application Role Permissions

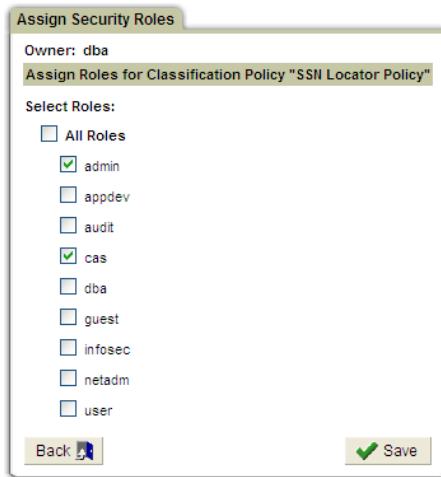
Each SQL Guard application can be assigned one or more security roles. Assigning an application to a role is the *only* way to grant a user access to that SQL Guard application. The term *application* in this context can refer to a single tool or to a collection of tools grouped on a tab or menu.

In Central Manager environments, all Users, Security Roles, and Application Role Permissions are stored on the Central Manager. The SQL Guard administrator logged on to the Central Manager *or any managed unit* can modify the Application Role Permissions stored on the Central Manager. (In contrast, the User and Security Role definitions can only be modified when logged on to the Central Manager.)

To assign security roles to applications:

1. Click the Access Management tab to bring it to the front.

2. Click Role Permissions in the menu on the left to display the Assign Roles to Applications panel, partially shown in the screenshot on the following page. The complete set of applications is described in a table at the end of this section.
3. To assign a role to an application, click that role's Roles button to open the Assign Security Roles panel.



4. Mark all roles to which you want to grant permission for this application. Click the Save button.

Application List

In the list that follows, please note the following:

- Because of space considerations, the application names that appear in the Assign Roles to Applications list do not always exactly match the application names as they appear on other tabs and menus. They are, however, similar, and you should have no problem relating the two.
- When the default *admin* and *non-admin* layouts are referred to, these references indicate the default layouts at installation time. Though it is not recommended, SQL Guard administrators can modify any default layouts following installation.
- All domains, entities, and attributes are described in an Appendix of the *SQL Guard User Guide*.

Application	Description (Assign Roles to Applications Panel)
Access Map Application	In the default <i>non-admin</i> layouts, provides access to the Access Map tab beneath the HealthGuard tab. The Access Map tab includes the Access Map Builder/Viewer, Group Builder, and Aliases Builder, all of which are described below. In the default <i>admin</i> layout, these components are all available from the Tools tab menu.
Access Map Builder/Viewer	A component of SQL HealthGuard, this application allows you to build and view access maps showing which clients access which servers. In the default <i>admin</i> layout, access this application from the Tools tab menu. In the default <i>non-admin</i> layouts, access this application from the Access Map tab beneath the HealthGuard tab.
Access Trace Tracking (Obsolete)	This query builder and its associated reporting domain are no longer supported. It is not included in the application list, but is included in this table in the event that you have old queries or reports from the access trace tracking domain. Prior to SQL Guard version 4.0, S-Tap, database server logins detected by SQL Guard could be compared to logins available from various database components (depending on the database type), by installing special S-Tap components called <i>access tracers</i> on the database sever. Beginning with version 4.0 of SQL Guard, old reports written for this domain will not work.
Access Tracking	Provides access to the Data Access domain, which allows you to create queries on entities related to database access information.
Administration Console	Provides access to the Administration Console tab on the default <i>admin</i> layout. This tab contains a menu of activities generally reserved for <i>admin</i> users.
Agg/Archive Activity Tracking	Provides access to the Aggregation/Archive domain, which allows you to create queries on entities related to the aggregation and archive process.
Alert Builder	Provides access to the Alert Builder, which you can use to create Statistical (Threshold) alerts. In the default <i>admin</i> layout, access this application from the Tools tab menu. In the default <i>non-admin</i> layouts, access this application from either the Auditing tab beneath the AuditGuard tab, the Alerts tab beneath the Reports & Alerts tab, or the Access Policies tab beneath the PolicyGuard tab.
Alert Tracking	Provides access to the Sent Alerts domain, which allows you to create queries on entities related to the aggregation, archive, backup, restore, and (Central Manager) user-table export operations.

Application	Description (Assign Roles to Applications Panel)
Alias Builder	Provides access to the Alias Builder, which allows you to create aliases for presentation purposes. In the default <i>admin</i> layout, access this application from the Tools tab menu. In the default <i>non-admin</i> layouts, access this application from the Custom Reports tab beneath the Reports & Alerts tab.
Allow Full SQL DrillDown	Provides access to drill-down reports that display full SQL. By default, all users have access to this report.
Application Tracking	Provides access to the application tracking domain.
Application User Responsibility Detection	Provides access to the application user ID detection application.
Audit Database Builder	Provides access to the audit database builder, which is a component of the value change auditing system.
Audit Process Builder	Provides access to the Policy Builder, which allows you to create, install and maintain security policies. In the default <i>admin</i> layout, access this application from the Tools tab menu. In the default <i>non-admin</i> layouts, access this application from the Access Policies tab beneath the PolicyGuard tab.
Audit Process To-Do List	Provides access to the user's own To-Do List. A user without access to this application will not be able to view or approve audit tasks. In the default <i>admin</i> layout, access this application from the Tools tab menu. In the default <i>non-admin</i> layouts, access this application from the Auditing tab beneath the AuditGuard tab.
Audit Process Tracking	Provides access to the Audit Process domain, which allows you to create queries on entities related to tracking audit processes.
Auditing Application	Provides access to the AuditGuard tab, which is included in default <i>non-admin</i> layouts. This tab provides a process flow diagram and links to many components included in the AuditGuard tab. All of the included components are described elsewhere in this table.
Autodetect Configuration	Provides access to the database auto-discovery configuration tools – usually restricted to administrators only.
Autodetect Query Builder	Provides access to database auto-discovery query builder.

Application	Description (Assign Roles to Applications Panel)
Baseline Builder	Provides access to the Baseline Builder, which allows you to create, generate and maintain baseline definitions. In the default <i>admin</i> layout, access this application from the Tools tab menu. In the default <i>non-admin</i> layouts, access this application from the Access Policies tab beneath the PolicyGuard tab.
CAS Application	Provides access to the Change Audit System application.
CAS Configuration	Provides access to the Change Audit System configuration tool.
CAS Lost Target	Provides access to the Change Audit System Lost Target tool.
CAS Query Builder	Provides access to the Change Audit System.
Catalog	Provides access to the Catalog application.
Classifier	Provides access to the Classifier application.
Comment Tracking	Provides access to the Comment Tracking domain, which can be used to create queries and reports listing user comments.
Custom Domain Builder	Provides access to the Custom Table builder.
Custom Query Builder	Provides access to the Custom Query Builder, used to report on custom tables or to build queries used to load user groups from custom tables.
Custom Reporting	Provides access to the Custom Reports tab, which is included in default <i>non-admin</i> layouts. This tab provides a process flow diagram and links to many of the components used to create custom reports. All of the included components are described elsewhere in this table.
Data Access Policy Application	Allows access to the Access Policies tab beneath the PolicyGuard tab, which is included in the default <i>non-admin</i> layouts. This tab provides a process flow diagram and links to components related to the creation and use of security policies. All components are described elsewhere in this table.
Database Analyzer	Enables the Auto Generated Calling Prox button of the Group Builder, which allows access to the Database Analyzer. This can be used to populate groups of objects or fields by analyzing stored procedures. The stored procedures can be analyzed by accessing the database where they are defined, or as they are encountered in the database traffic.

Application	Description (Assign Roles to Applications Panel)
Database Intrusion Detection	Allows access to the Alerts tab beneath the Reports & Alerts tab, which is included in the default <i>non-admin</i> layouts. This tab provides a process flow diagram and links to components related to the creation of threshold alerts. All components are described elsewhere in this table.
Database Security Health Assessment	Allows access to the Security Assessment tab beneath the HealthGuard tab, which is included in the default <i>non-admin</i> layouts. This tab provides a process flow diagram and links to the Assessment Builder and Audit Process Builder, which are described elsewhere in this table (see above).
Datasource Builder	Provides access to the generic datasource builder.
Exception Tracking	Provides access to the Exceptions domain, which allows you to create queries on entities related to exception information.
Group Builder	Provides access to the Group Builder, which allows you to create logical groups of Objects, Commands, and others items. These groups are used in queries and policies to identify certain conditions.
Group Tracking	Provides access to the Groups domain, which allows you to create queries on entities related to SQL Guard group definitions.
Installed Policy Tracking	Provides access to the installed policy tracking domain.
Investigation Data Resource	Provides access to the Investigation Data Resource.
Policy Builder	Provides access to the Policy Builder, which enables creating, modifying and installing policies.
Policy Violation Query Builder	Provides access to the Policy Violations domain, which allows you to create queries on entities related to policy violations.
Privacy Compliance	Provides access to the Privacy Sets tab beneath the AuditGuard tab, which is included in the default <i>non-admin</i> layouts. This tab provides a process flow diagram and links to the Privacy Set Builder and Audit Process Builder, which are described elsewhere in this table.
Privacy Set Builder	Provides access to the Privacy Set Builder application, which allows you to create and modify privacy sets.
Report Builder	Provides access to the Report Builder.

Application	Description (Assign Roles to Applications Panel)
Retrospective Request	Provides access to the Retrospective Request tool.
Rogue Connections Tracking	Provides access to the Rogue Connections domain, which allows you to create queries on rogue connections. Rogue connections are created when Unix based S-Taps detect database connections that have circumvented “standard” access paths monitored by S-Tap.
Security Assessment Builder	Provides access to the Security Assessment Builder, which allows you to create security assessments.
SQL Guard Activity Tracking	Provides access to the SQL Guard Activity domain, which allows you to create queries on entities related to SQL Guard user activity.
SQL Guard Login Tracking Builder	Provides access to the SQL Guard Logins domain, which allows you to create queries on entities related to SQL Guard logins. (The SQL Guard Activity domain provides more detailed information about what actually happened during SQL Guard user sessions; see SQL Guard Activity Tracking, below.)
SQL Guard User Role App Tracking	Provides access to SQL Guard User and Role tracking domain.
Trigger Builder	Provides access to the Value Change Auditing configuration tool (which builds triggers on database servers).
Value Change Tracking	Provides access to the Value Change Tracking domain for reporting.

Chapter 5: Administrator Tools

This chapter describes a number of tools that by default are restricted to SQL Guard administrators. You can make these tools more widely available by changing the roles assigned, as described in a previous chapter.

Custom Tables, Domains and Queries

A **custom table** contains one or more attributes that you want to have available on the SQL Guard server. For example, you may have an existing database table relating encoded user names to real names. In the network traffic, only the encoded names will be seen. By defining a custom table on the SQL Guard server, and uploading data for that table from the existing table, you will be able to relate the encoded and real names. You will also be able to keep the SQL Guard server table synchronized with your existing database table by running the upload on a scheduled basis.

Before defining a custom table, first verify that the data you need from the existing database is a supported data type. For each database type, the following table summarizes the supported and unsupported data types for uploading to a custom table.

Supported and Unsupported Data Types for Uploading to a Custom Table

Database	Supported Data Types			Unsupported Data Types	
Oracle	float varchar2 nvarchar2	number date	char nchar	long raw longraw rowid blob	clob nclob bfile urowid
DB2	char integer double time	varchar smallint decimal timestamp	bigint real date	blob longvarchar	clob datalink
Sybase	char nvarchar tinyint money numeric bit	nchar int datetime smallmoney float double precision	varchar smallint smalldatetime decimal real	text varbinary timestamp	binary image

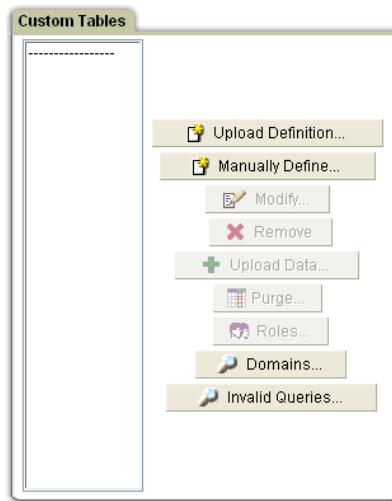
Database	Supported Data Types			Unsupported Data Types	
MSSQL	bigint datetime int numeric smalldatetime smallmoney unique identifier	bit decimal money nvarchar smallint varchar	char float nchar real tinyint	Binary Ntext timestamp	Image Text
Informix	char smallint float money	nchar decimal serial varchar	integer smallfloat date nvarchar datetime	text	
MySQL	bigint mediumint double datetime year enum	decimal smallint float timestamp char set	int tinyint date time binary	longtext tinytext text mediumtext longtext	tinyblob blob mediumblob longblob

A **custom domain** contains one or more custom tables. If it contains multiple tables, you define the relationships between tables when defining the custom domain. You cannot access the standard internal SQL Guard tables from a custom domain. Those tables are only available via the standard SQL Guard domains.

A **custom query** accesses data from a custom domain. You use the Custom Query Builder to create queries against custom domains. Custom queries can then be used like any other query to generate reports or audit tasks, populate SQL Guard groups, or to define SQL Guard aliases.

Creating Custom Tables

There are two methods for defining custom tables. Regardless of the method used, open the Custom Tables panel by selecting **Tools – Report Building – Custom Domain Builder**.



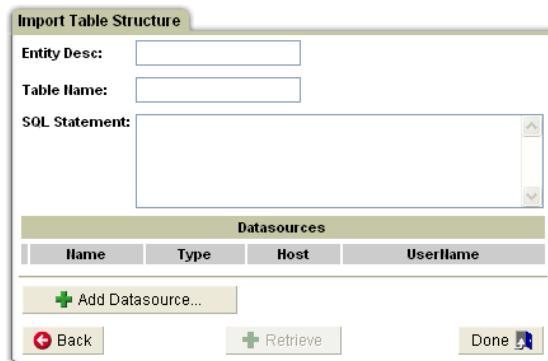
Do one of the following:

- To define a custom table by uploading metadata from a database, click Upload Definition, and see [Uploading Table Definitions](#), below.
- To define a custom table manually, click Manually Define, and see [Defining Tables Manually](#), below.

Uploading Table Definitions

To upload a table definition by accessing its metadata from the database server on which it is defined, follow the procedure outlined below.

1. If it is not already opened, open the Import Table Structure panel, as follows:
 - Select **Tools – Report Building – Custom Domain Builder** to open the Custom Tables panel (not shown).
 - Click Upload Definition to open the Import Table Structure panel:



2. In the Entity Desc box, enter a description for the table. This is the name you will use to reference the table when creating a custom query.
3. In the Table Name box, enter the database table name for the table. This is the name you will use to include the table in a custom domain.
4. In the SQL Statement box, enter a valid SQL Statement for the table (for example, `select * from <tablename>`).
5. Click Add Datasource to open the Datasource Finder in a separate window:



Use this window to identify the database from which the table definition will be uploaded.

6. Select a datasource from the list and click Add, or click New to define a new datasource. (For detailed instructions on how to define a new datasource, see *Defining Datasources* in the SQL Guard User Guide.) After the datasource has been selected, the Retrieve button on the Import Table Structure panel becomes active.
7. Click the Retrieve button to upload the table definition. Remember that only the definition is being uploaded – you will upload data later.

What to do next...

In the Custom Tables panel, select the table just uploaded and click the Modify button to verify that the expected attributes have been defined.

Defining Tables Manually

To define a custom table manually, follow the procedure outlined below.

1. If it is not already opened, open the Define Entity panel, as follows:
 - o Select Tools – Report Building – Custom Domain Builder to open the Custom Tables Panel (not shown).
 - o Click Manually Define to open the Define Entity panel:

Column Name	Display Name	Type	Size	Group Type
		Text	255	--none--

2. In the Entity Desc box, enter a description for the table. This is the name you will use to reference the table when creating a custom query.
3. In the Table Name box, enter the database table name for the table. This is the name you will use to include the table in a custom domain.
4. For each column in the table to be defined:
 - o Enter a name in the Column Name box. This will be the name of the attribute in the database table.
 - o Enter a name in the Display Name box. This is the name you will use to reference the attribute in the Custom Domain Builder.
 - o Select a data type (Text, Date, Integer, Float, or Time).
 - o For a Text attribute, enter the maximum number of characters in the Size box. (The Size box is not available for other data types.)
 - o If the attribute being defined corresponds to a group type recognized by SQL Guard, select that group type from the Group Type list.
 - o Click the Add button.
5. Click the Done button when you have added all columns of the table.

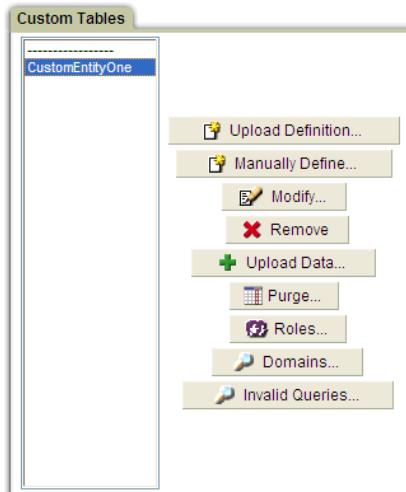
About Modifying Table Definitions

If you modify the definition of a custom table, you may invalidate existing reports based on queries using that table. For example, an existing query might reference an attribute that has been deleted, or whose data type has been changed. When applying changes to a

custom table, if any queries have been built using attributes from that table, the Queries are displayed in the Query List panel:

Query Name	Main Entity
CustomDomainQueryOne	CustomEntityOne
CustomQuerySampleTwo	CustomEntityOne
CustmSampleOne	CustomEntityOne
CustomReportTwo	CustomEntityOne

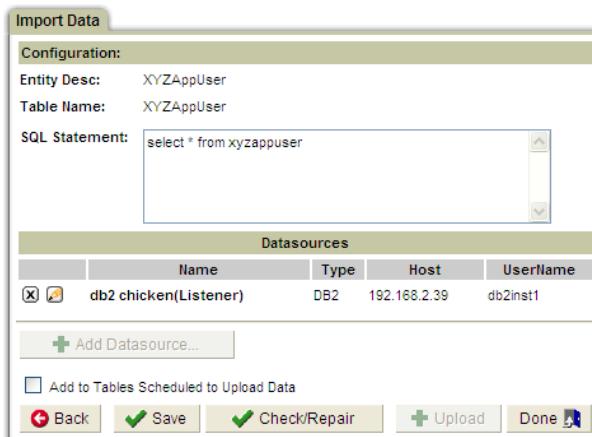
You can open and modify each query as necessary, by clicking its name in the Query Name column. If you cannot make all of the changes right away, you can come back to this panel later, by selecting the changed table in the Custom Tables panel, and clicking the Invalid Queries button.



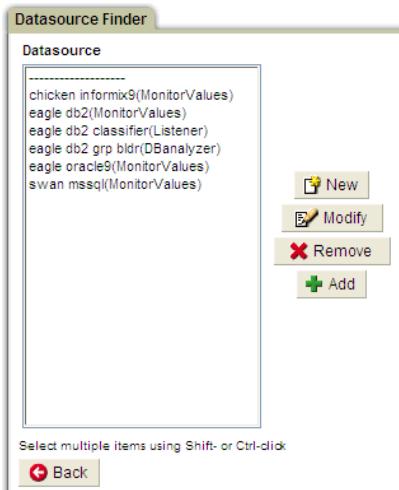
Uploading Data to Custom Tables

Once a custom table definition is in place, follow the procedure below to upload data.

1. If it is not already opened, open the Import Data panel by selecting **Tools – Report Building – Custom Domain Builder**.
2. From the list of custom tables, select the table to which you want to upload data.
3. Click the Upload Data button to open the Import Data panel:



4. In the SQL Statement box, enter a valid SQL Statement for the table (for example, `select * from <tablename>`). The result set returned by the SQL statement must have the same structure as the custom table defined.
5. Click Add Datasource to open the Datasource Finder in a separate window:



Use this window to identify one or more databases from which the table data will be uploaded.

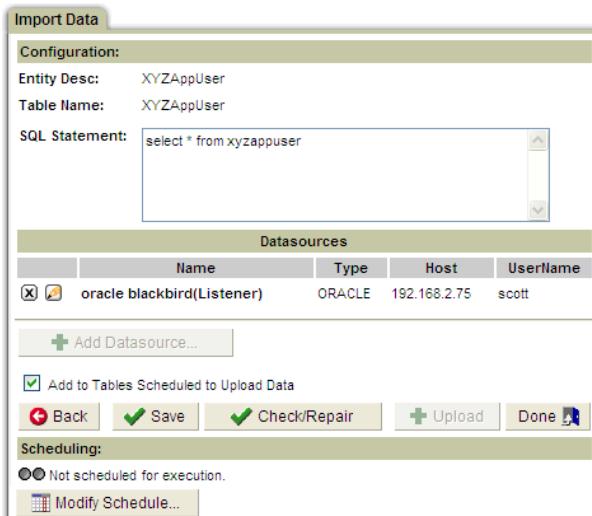
6. **For central management environments:** In a central management environment, the custom table definition resides on the central manager, and the custom table may not exist on the local (managed unit) database. Click the Check/Repair button to check if the custom table exists locally, and create one if it does not.

7. Click the Save button.
8. To upload data to this custom table, do one of the following:
 - **To upload data now:** Click Upload.
 - **To add this table to the schedule of custom table uploads:** mark the *Add to Tables Scheduled to Upload Data* checkbox, and see the Scheduling Custom Data Uploads topic, below.
9. Click the Save button.

Scheduling Custom Data Uploads

Data can be uploaded to custom tables on the SQL Guard server on a scheduled basis. Once a custom table definition is in place, follow the procedure below to schedule data uploads. Note that there is only one job that does this, and all custom tables that have been flagged for scheduled data upload will be updated by the scheduled job, one at a time. The total amount of disk space reserved on the SQL Guard server for custom tables is 4GB.

1. If it is not already opened, open the Import Data panel by selecting **Tools – Report Building – Custom Domain Builder**.
2. From the list of custom tables, select a table to which you want to upload data on a scheduled basis.
3. Click the Upload Data button to open the Import Data panel:



4. Mark the *Add to Tables Scheduled to Upload Data* checkbox. The Scheduling pane will be displayed (as illustrated above).
5. To modify the schedule, click the Modify Schedule button. This opens the standard Schedule Definition panel. See [Using the Task Scheduler](#).
6. Click Done when you are finished.

Purging Data from Custom Tables

Data can be purged from custom tables on the SQL Guard server on demand, or on a scheduled basis. To purge data from a selected table:

1. If it is not already opened, open the Custom Tables panel by selecting **Tools – Report Building – Custom Domain Builder**.
2. From the list of custom tables, select the table from which you want to purge data.
3. Click the Purge button to open the Custom Table Data Purge panel:



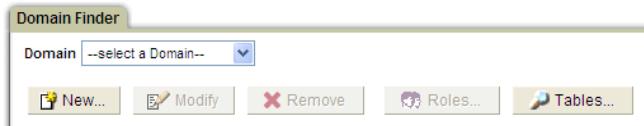
4. In the Configuration panel, enter the age of the data to be purged, as a number of days, weeks or months prior to the purge operation date.
5. To run the purge operation once, click Run Once Now.
6. To schedule a purge operation, click the Modify Schedule button to open the standard Schedule Definition panel. See [Using the Task Scheduler](#).
7. Click Done to close the panel.

Defining Custom Domains

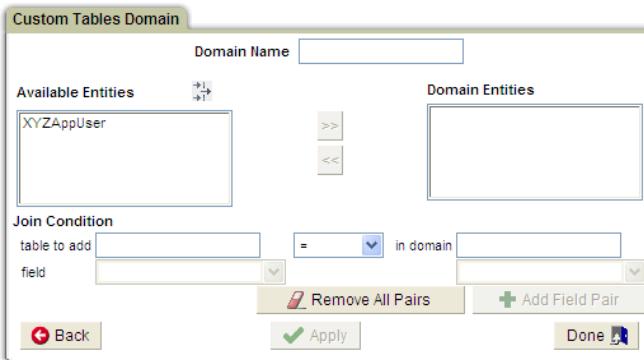
After defining one or more custom tables, define a custom domain so that you can perform query and reporting tasks using the custom data. To define a custom domain:

1. If it is not already opened, open the Custom Tables panel by selecting **Tools – Report Building – Custom Domain Builder**.

- Click the Domains button to open the Domain Finder panel:



- Click the New button to open the Custom Tables Domain panel:



- Enter a Domain Name. Typically, you will be including a single custom table in the domain, so you may want to use the same name for the domain.
- The Available Entities box lists all custom tables that have been defined (and to which you have access). Select an entity.

Optionally, click the (Filter) tool to open the Entity Filter:



Enter a Like value to select only the entities you want listed, and click Accept. This closes the filter window and returns you to the Custom Tables Domain panel, with only those entities matching the Like value listed in the Available Entities box. Select the entity you want to include.

- Click the right arrow button to move the entity selected in the Available Entities list to the Domain Entities list.

Optional: To remove one or more entities from the Domain Entities list, select them and click the left arrow button.

7. To add an entity to a domain that already has one or more tables, follow the procedure outlined below. You will need to use the Join Condition to define the relationship between the entities.

For each additional entity:

- From the Domain Entities box on the right, select an entity. All of the attributes of that entity will become available in the field drop-down list below the Domain Entities box. Select the attribute from that list that will be used in the join operation.
 - From the Available Entities list on the left, select the entity you want to add. All of the attributes of that entity will become available in the field drop-down list below the Available Entities box. Select the attribute from that list that will be used in the join operation.
 - Select = (the equality operator) if you want the join condition to be equal (e.g., domainA.attributeB = domainC.attributeD). Select *outer join* if you want the join condition to be an outer join using the selected attributes.
 - Click Add Field Pair.
 - Repeat the above steps for any additional join operations.
8. Select the Timestamp attribute for the custom domain entity.
 9. Click Apply.

What to do next...

Build a custom query to view the data or to use that data to populate a group.

Working with Custom Queries

This section describes how to open the Custom Query Builder, which is like any other SQL Guard query builder. For detailed instructions on how to build queries, see *Chapter 7: Building Queries and Reports* in the *SQL Guard User Guide*.

Use the Custom Query Builder to build queries against data from *custom domains*, which contain one or more *custom tables*.

To begin working with a custom query:

1. Select **Tools – Report Building – Custom Query Builder** to open the Domain Finder:
2. Select a custom domain from the list, and click



the Search button to open the Query Finder:

To view, modify or clone an existing query, select it from the Query Name list, or select a report using that query from the Report Title list.

To view all of the queries defined for a specific custom table, select that custom table from the Main Entity list and click the Search button (only the custom tables included in the selected custom domain will be listed).

Value Change Auditing

Overview

SQL Guard's Value Change Auditing feature monitors changes to values in database tables. You configure Value Change Auditing to monitor specific tables. For each table, you select which SQL value-change commands to monitor (insert, update, delete). Each time a value-change command is executed against a monitored table, before and after values are captured. On a scheduled basis, the change activity is uploaded to a SQL Guard server, where all of SQL Guard's reporting and alerting functions can be used.

The basic steps that you perform to use the Value Change Auditing feature are:

1. From the SQL Guard administrator console, create an audit database on the database server. This is where value-change data will be stored until it is uploaded to the SQL Guard server. See [Defining Audit Databases](#).
2. Identify the tables to be monitored, and for each table select the value-change commands (insert, delete, update) for which changes will be recorded. To record the changes, a trigger will be created for each table to be monitored, and that trigger will write the value-change data to the audit database. To allow updates to the audit database (via the trigger), all users with update privileges for the monitored table will be given appropriate privileges for the audit database. This has implications for users who may be given update privileges for that table later (see step 4, below). For detailed instructions on how to define the monitoring activities, see [Defining Monitoring Activities](#).
3. Schedule uploads to transfer value-change data from the database server to the SQL Guard server. See [Scheduling Value-Change Uploads](#).
4. Maintain audit database access privileges. After a trigger has been created, a new user may be given access to the table on which the trigger is based. If that user issues a monitored value-change command, it will fail because that user will not have appropriate privileges to update the audit database. . See [Maintaining Privileged Users Lists](#).

5. Monitor change activity from the administrator console, or use the Value Change Tracking query domain to create custom reports on the SQL Guard server. See [Value-Change Reporting](#).

Defining Audit Databases

To define an audit database and value-change monitoring activities, for each database to be monitored, you will need to have a user account with appropriate permissions to:

- Create a database.
- Create a user.
- Log in to the database.
- Create tables and triggers.

Before Defining an Audit Database Under Informix or Sybase

This topic applies for Informix (9.4 or later) and Sybase (except for Sybase IQ, which does not support triggers). Depending on the operating system for the database server, you must perform one of the following procedures *before* defining the audit database.

Locate or Create a New Informix Database Space

Under Informix, we strongly recommend that you avoid using the default root database space, *root_dbs*. You cannot drop this space or reduce its size. You can use any other database space that has been defined, or to create a new database space, perform one of the following procedures (depending on the operating system):

Create an Informix Database Space on a Windows Server

This procedure is performed outside of the SQL Guard GUI, and applies for Informix version 9.4 or later.

1. Verify that the database server is online and listening.
2. Create a zero-byte file named *guardium_dbs_dat.000* in the *C:\IFMXDATA\server-name* directory (server-name is the name of the Informix server or the service name). You can do this by saving an empty text file, and then renaming the file, replacing the *txt* suffix with *000*.
3. Make the following directory the working directory:
C:\Program Files\Informix\bin
4. Execute following command:

```
C:\Program Files\Informix\bin>onspaces -c -d guardium_dbs -p  
C:\IFMXDATA\server-name\guardium_dbs_dat.000 -o 0 -s 150000
```

If the file is created successfully, you will receive the following messages:

```
Verifying physical disk space, please wait ...
```

```
Space successfully added.
```

```
** WARNING ** A level 0 archive of Root DBSpace will need to be  
done.
```

5. Restart the Informix server, and use a suitable tool (Aqua Data Studio remote client, for example) to connect and verify that the space named *guardium_dbs* has been created. Your first connection attempt may fail with a message about the server running in *Quiescent Mode*. If this happens, attempt to re-connect at least two more times, and it should work fine. Verify that the *guardium_dbs* database space has been created (using Aqua Data Studio, look under *Storage*).

Create an Informix Database Space on a Unix Server

This procedure is performed outside of the SQL Guard GUI, and applies for Informix version 9.4 or later.

1. From a command-line window, enter the following commands:

```
su - informix  
cd demo/server  
vi guardium_dbs    => save it to create an empty file.
```
2. Without adding any text, save the empty *guardium_dbs* file.
3. Enter the following commands:

```
chmod 660 guardium_dbs  
cd ../../bin  
onspaces -c -d guardium_dbs -p  
/home/informix10/demo/server/guardium_dbs -o 0 -s 100000
```

Initialize Disks on a Sybase Server

Depending on the operating system of the database server, perform one of the following procedures to initialize disks on a Sybase server.

Initialize Disks on a Windows Sybase Server

1. Connect to the server on which you want to create *guardium_audit*.
2. Create a folder named *guardium_audit*, under the c: drive.
3. Connect to the database using *datastudio* or using *isql*.
4. Execute the following statements:

```
use master
```

```
go
disk init name="guardium_auditdev",
physname="c:/guardium_audit/guardium_auditdev", size=8192
go
disk init name="guardium_auditlog",
physname="c:/guardium_audit/guardium_auditlog", size=8192
go
```

Initialize Disks on a Unix Sybase Server

1. Connect to the database using *datastudio* or using *isql*.
2. Execute the following statements:

```
use master
go
disk init name = 'guardium_auditdev', physname
    ='home/sybase/data/guardium_auditdev' , size = 8192
go
disk init name = 'guardium_auditlog', physname
    ='home/sybase/data/guardium_auditlog' , size = 8192
go
```

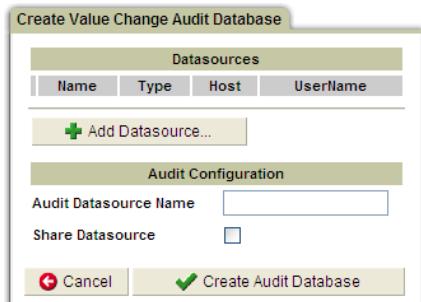
Creating the Audit Database

To create a new audit database:

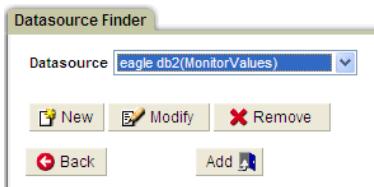
1. Do one of the following to open the Value Change Database Builder:
 - From an *admin* portal, select Tools – Config & Control – Value Change Database Builder.
 - By default, the Value Change Database Builder is only available to users having the *admin* role. If your SQL Guard administrator has made this application available to other roles, it can be placed on a custom layout by any user who also has that role. If this is the case, from the *user* portal, open the custom layout containing the Value Change Database Builder, and open that application.



2. Click Create Audit Database. This opens an empty Create Value Change Audit Database panel:



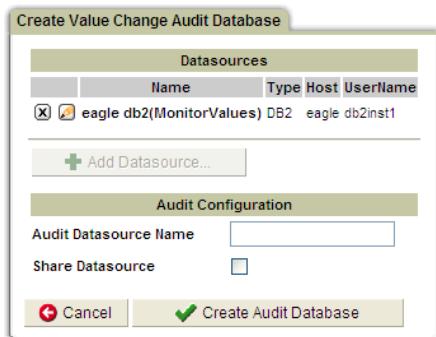
3. Click Add Datasource to open the Datasource Finder panel:



Note: Datasources that have been defined from the Value Change Auditing application are labeled *Monitor Values*, as illustrated above. Shared datasources that have been defined for other applications will have different labels (*Listener*, or *DBanalyzer*, for example), and those datasources may *not* have the appropriate set of database access permissions for Value Change Auditing application, which requires a user account having database administrator authority.

If a suitable datasource is not available, click the New button to define a new one for the database to be monitored (see [Defining New Audit Datasources](#)).

4. If a datasource that uses an administrator account is available for the database server, select it from the list and click Add, to add it to the Datasources pane on the Create Value Change Audit Database panel, as illustrated below:



The example above is for a DB2 type database. For other database types, the content of the Audit Configuration panel varies, as described later.

5. Enter an Audit Datasource Name. This is the name you will use to identify the datasource later, to define monitoring tasks and to upload data. Do not confuse this name with the name of the Datasource from the Datasources panel.
6. Mark the Share Datasource box if you want to share this datasource with other applications (Classification, for example). The default is not to share the datasource. This type of datasource requires administrator privileges, so you may not want to share this datasource with other applications.

Note: To share a datasource with other *users*, assign security roles.

7. For any database type other than DB2, there will be additional fields in the Audit Configuration pane. All fields are required. Referring to the following table, enter the appropriate values.

Additional Audit Configuration Fields Table

Database Type	Description
Informix	Database Space: Enter the name of an existing database space to use, or enter the name of the database space you created for the audit database (<i>guardium_dbs</i> in the example shown previously). If you leave this blank, the default <i>root_dbs</i> space will be used, which we do not recommend.
MS SQL Server	Audit User Name: Enter a new database user name to use when accessing the audit database. This user will be given the <i>sysadmin</i> role. Audit Password: Enter a password for the above.
Oracle	Audit Password: Enter the password for the <i>system</i> user, which will be the database account used to access the audit

Database Type	Description
	<p>database.</p> <p>Default Tablespace: Enter a name for the default tablespace (we use <i>system</i> in our examples).</p> <p>Temp Tablespace: Enter a name for the temporary tablespace (we use <i>temp</i> in our examples).</p>
Sybase	<p>Audit User Name: Enter a new database user name to use when accessing the audit database. This user will be granted the <i>sa_role</i>.</p> <p>Audit Password: Enter a password for the above.</p> <p>Data Device Name: Enter the same data device name used when initializing the disk for the audit database (<i>guardium_auditdev</i> in the disk initialization procedure described earlier).</p> <p>Log Device Name: Enter the same log device name used when initializing the disk for the audit database (<i>guardium_auditlog</i> in the disk initialization procedure described earlier).</p>

8. Select an action from the table below.

Create Value Change Audit Database Controls

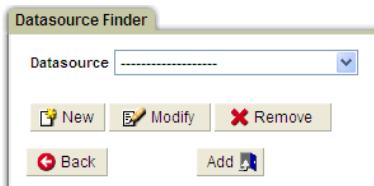
Control	Description
 (Remove)	Click to remove the datasource from the Datasources pane. Before you can do anything else, you will need to select another datasource using the Add Datasource button (see below).
Add Datasource	Only one datasource can be selected for this operation, so this button is only enabled when the Datasources pane is empty. Click the Add Datasource button to select a datasource using the Datasource Finder (see Defining New Audit Datasources).
 (Edit this Datasource)	Click to edit this datasource definition in the Datasource Definition panel (see Modifying Datasources in the SQL Guard User Guide).
Cancel	Click to cancel the operation and return to the Audit Datasource Finder panel.
Create Audit Database	Click to create the audit database.

What to do next...

After an audit database has been created on the database server, it will be available for use by the Value Change Auditing Builder (see [Chapter 3: Defining Monitoring Activities](#)).

Defining New Audit Datasources

1. Navigate to the Datasource Finder panel from the Value Change Database Builder (see Steps 1-3 of [Creating the Audit Database](#), above):



2. Click the New button to open the Datasource Definition panel:

3. Detailed instructions for completing this panel can be found elsewhere. For the audit database datasource, note the following:

Login Name must be a database user account with the authority to:

- Create a database.
- Create a user.
- Log in to the database.
- Create tables and triggers.

Share Datasource: Since this datasource requires *administrative* authority, you may not want to share it with other applications (leave this box unmarked).

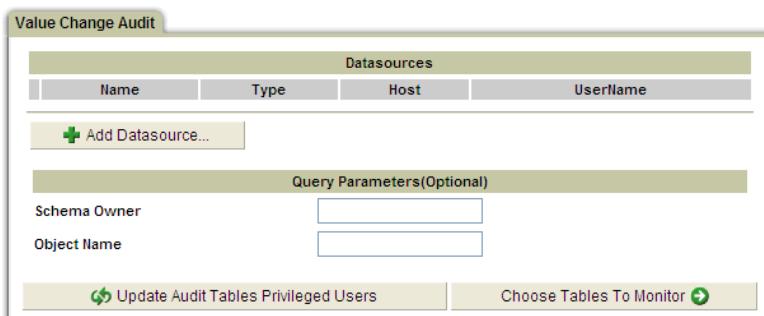
To complete the datasource definition, refer to Defining Datasources in the *SQL Guard User Guide*.

Defining Monitoring Activities

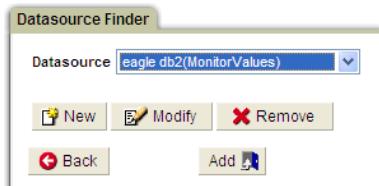
Use the Value Change Auditing Builder to add, modify or remove audit datasource definitions. Be aware that *removing* an audit datasource removes SQL Guard's definition of the audit datasource, but it does *not* remove the audit database on the database server.

After defining an audit database, use the Value Change Auditing Builder to identify the tables to be monitored, and to select the types of changes (inserts, updates, deletes) to be recorded.

1. Do one of the following to open the Value Change Auditing Builder:
 - From an *admin* portal, select Tools – Config & Control – Value Change Auditing Builder.
 - From a *user* portal, open the custom layout containing the Value Change Auditing Builder.



2. Click Add Datasource to open the Datasource Finder:



3. Select a datasource on which an audit database has been defined. If an audit database has not yet been defined, see [Creating the Audit Database](#).
4. Click Add to close the panel and add the selected datasource to the Value Change Audit panel:

Datasources				
	Name	Type	Host	Username
<input checked="" type="checkbox"/>	eagle db2(MonitorValues)	DB2	eagle	db2inst1

Add Datasource...

Query Parameters(Optional)

Schema Owner: []

Object Name: []

Update Audit Tables Privileged Users **Choose Tables To Monitor**

Note: To remove or edit the selected datasource, click the (Remove) or (Edit this Datasource) button.

5. Optionally enter a Schema Owner and/or Object Name to limit the number of tables that will be displayed when choosing the tables to be monitored. You can use the % (**percent**) wildcard character. For example, to limit the display to all tables beginning with the letter a:

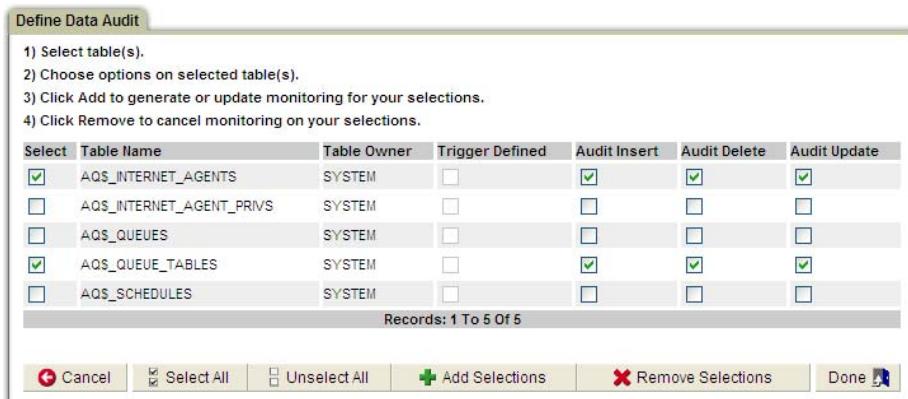
Query Parameters(Optional)

Schema Owner: []

Object Name: [a%]

Update Audit Tables Privileged Users **Choose Tables To Monitor**

6. Click Choose Tables To Monitor to open the Define Data Audit panel:



- Mark the Select box for each table you want to monitor. We have marked two in the example above.

Note: You cannot define a trigger for a table that contains one or more user-defined data types.

The **Trigger Defined** column indicates if a trigger has already been defined for the table. The **Audit Insert**, **Audit Delete**, and **Audit Update** checkboxes indicate if the trigger will record changes for that command.

If the Trigger Defined column is not marked, marking the Select checkbox for a table automatically marks all three the Audit checkboxes (Audit Insert, Audit Delete, and Audit Update). If you do not want to monitor one or two of those commands, clear the appropriate checkbox.

- Click the Add Selections button to define triggers for the selected tables. You will be informed of the action taken. For example, after clicking Add Selections on the panel illustrated previously:



- Click OK to re-display the Define Data Audit panel:

Define Data Audit

- 1) Select table(s).
- 2) Choose options on selected table(s).
- 3) Click Add to generate or update monitoring for your selections.
- 4) Click Remove to cancel monitoring on your selections.

Select	Table Name	Table Owner	Trigger Defined	Audit Insert	Audit Delete	Audit Update
<input checked="" type="checkbox"/>	AQS_INTERNET_AGENTS	SYSTEM	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input type="checkbox"/>	AQS_INTERNET_AGENT_PRIVS	SYSTEM	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	AQS_QUEUES	SYSTEM	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	AQS_QUEUE_TABLES	SYSTEM	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input type="checkbox"/>	AQS_SCHEDULES	SYSTEM	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Records: 1 To 5 Of 5

Buttons:

The selected tables remain selected, and the Trigger Defined column is now marked for those tables.

Note: The instant a trigger has been defined for a table, it is active and recording changes for the selected commands in the audit database. The configuration of triggers is done entirely on the database server, which is unlike most other SQL Guard configurations, which are defined on the SQL Guard database, and then activated or deactivated as a separate task.

10. You can define additional actions by repeating the steps above, or you can remove triggers by marking the appropriate Select checkboxes and clicking Remove Selections.
11. Click Done or Cancel after you have completed all changes. Be aware that the Cancel button does not back out any changes that you have to triggers using the Add or Remove Selections buttons.

What to do next...

If you have added value-change monitoring activities to a datasource for the first time, you should schedule uploads for this datasource, because the audit database will be emptied only after the data recorded there has been uploaded to the SQL Guard server. See [Chapter 4: Scheduling Value-Change Uploads](#).

Scheduling Value-Change Uploads

Once an audit database has been created, and value-change monitoring activities have been configured, you must define an upload schedule to transfer the collected data from the audit database to the SQL Guard server. Once data has been uploaded from the audit database to the SQL Guard server, it is removed from the audit database.

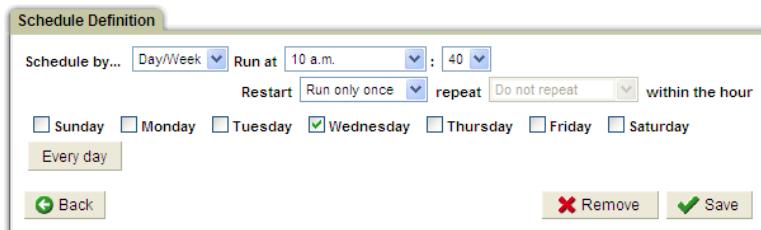
Defining an Upload Schedule

To define an upload schedule:

1. Do one of the following to open the Value Change Database Builder:
 - o From an *admin* portal, select Tools – Config & Control – Value Change Database Builder.
 - o By default, the Value Change Database Builder is only available to users having the *admin* role. If your SQL Guard administrator has made this application available to other roles, it can be placed on a custom layout by any user who also has that role. If this is the case, from the *user* portal, open the custom layout containing the Value Change Database Builder, and open that application.



2. Select the audit datasource for which you want to schedule uploads (*oracle9 eagle* in the example above), and click Schedule Upload to open the Schedule Definition panel:



3. The Schedule Definition panel is used for a number of tasks, and its use is described in detail in Chapter 2 of the Administrator Guide (see Using the Task Scheduler). Follow the instructions under that topic, and click Save when you are done.

Modifying or Removing an Upload Schedule

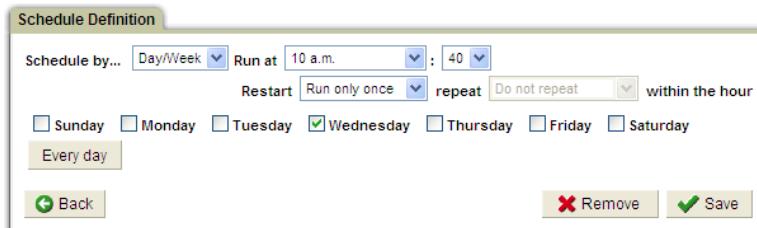
To modify or remove an upload schedule, follow the procedure outlined below.

1. Do one of the following to open the Value Change Database Builder:
 - o From an *admin* portal, select Tools – Config & Control – Value Change Database Builder.

- From a *user* portal, open the custom layout containing the Value Change Database Builder.



2. Select the audit datasource whose upload schedule you want to modify (*oracle9 eagle* in the example above), and click Schedule Upload to open the Schedule Definition panel:



3. The Schedule Definition panel is used for a number of tasks, and its use is described in detail in Chapter 2 of the Administrator Guide (see *Using the Task Scheduler*). Do one of the following:
 - **To modify the schedule:** Follow the instructions under the topic referenced above, and click Save when you are done.
 - **To remove the schedule,** click Remove. You will be prompted to confirm the action.

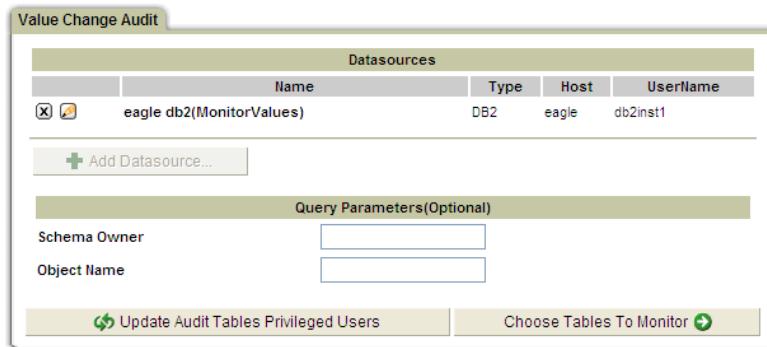
Maintaining Privileged Users Lists

When you use the value-change feature to add a trigger for a database table, all current users with permission to update that table will be granted permission to update the audit database table as well. This is required because the trigger updates the audit database with new and/or old values. If a new user is granted update permission for a monitored table, when that user attempts an update, the update will not be allowed because that user will not also have permission to update the audit database. When this happens, you can update the audit database privileged users list via the Value Change Auditing Builder.

Note: To update the audit database privileged users list, the database user ID that is used to log into the monitored database must be the *creator* of any role to which new users have been added. Otherwise, the members of that role will not be available.

For any monitored database, to update the list of users allowed to update the audit database, follow the procedure outlined below:

1. Do one of the following to open the Value Change Auditing Builder:
 - o From an *admin* portal, select Tools – Config & Control – Value Change Auditing Builder.
 - o From a *user* portal, open the custom layout containing the Value Change Auditing Builder.
2. Click Add Datasource to open the Datasource Finder panel (not shown), select the appropriate Datasource from the list, and click Add.



Value-Change Reporting

You can view value-change data from the default Values Changed report, or you can create custom reports using the Value Change Tracking domain. By default, the Value Change Tracking domain is restricted to users having the *admin* role.

The following sections describe the Value Change Tracking Domain and all default reports. The procedures for creating custom reports are described in Chapter 7 of the User Guide and are not repeated here.

Value Change Tracking Domain

The Value Change Tracking domain contains all of the data uploaded from value-change audit databases on database servers. It contains the entities described below.

Value Change Tracking Domain Entities

Entity	Description
Monitor Values	For each insert, update or delete recorded, contains the details of the change (table name, action, SQL text, etc).
Changed Columns	For each column changed, records the old and new values

Monitor Values Entity

A monitor values entity is created for each insert, update or delete recorded, contains the details of the change (table name, action, SQL text, etc).

Monitor Values Entity Attributes

Attribute	Description
Timestamp	Date and time the change was recorded <i>on the SQL Guard server</i> . This timestamp is created during the data upload operation. It is <i>not</i> the time that the change was recorded on the audit database. To obtain that time, use the Audit Timestamp entity (described below).
Last Modified Date	Date portion of the above timestamp.
Last Modified Time	Time portion of the above timestamp.
Last Modified Weekday	Day of week for above timestamp.
Last Modified Year	Year portion of above timestamp.
Server IP	IP address of the database server.
DB Type	Database type: DB2, Informix, etc.

Attribute	Description
Service Name	Oracle only. Database service name.
Database Name	DB2, Informix, Sybase, MS SQL Server only. Database name.
Audit PK	For Sybase and MS SQL Server only. A primary key used to relate old and new values (which must be logged separately for these database types).
Audit Login Name	Database user name defined in the datasource.
Audit Table Name	Name of the table that changed.
Audit Owner	Owner of the changed table.
Audit Action	Insert, Update or Delete.
Audit Old Value	A comma-separated list of old values, in the format: <i>column-name=column_value, ...</i>
Audit New Value	A comma-separated list of old values, in the format: <i>column-name=column_value, ...</i>
SQL Text	Available only with Oracle 9. The complete SQL statement causing the value change.
Triggered ID	Unique ID (on this audit database) generated for the change.
Audit Timestamp	Date and time that the trigger was executed.
Last Modified Date	Date portion of the above timestamp.
Last Modified Time	Time portion of the above timestamp.
Last Modified Weekday	Day of week for above timestamp.
Last Modified Year	Year portion of above timestamp.

Changed Columns Entity

This entity contains the name of a column, and the old and new values. It is used to supply multiple changed values for the Monitor Values entity (see above).

Changed Columns Entity Attributes

Attribute	Description
Changed Column Name	The name of the column that changed.
Old Value	The value before the change was applied.

Attribute	Description
New Value	The value after the change was applied.

Value Change Domain Default Reports

There is one default report available on the administrator portal, and one drill-down report available from that one. These are described below.

Values Changed Report

On an administrator portal, you can view the Values Changed report by selecting Values Changed from the menu of the Daily Monitor tab. The main entity for this report is the Changed Columns entity, which means that there will be a separate row of the report for every column change detected for every audit action (insert, update, delete).

By default, all changes will be listed for the reporting period. To limit report output, you can use any of the following runtime parameters.

Runtime Parameters

Entity	Attribute	Operator	Default Value
Monitor Values	DB_Type	Like	%
Monitor Values	Server IP	Like	%
Monitor Values	Audit Login Name	Like	%
Monitor Values	Audit Table Name	Like	%
Monitor Values	Audit Owner	Like	%
Monitor Values	Audit Action	Like	%

Sample Report

Values Changed														
Timestamp	Server IP	DB Type	Service Name	Database Name	Audit Login Name	Audit Timestamp	Audit Table Name	Audit Owner	Audit Action	Audit Old Value	Audit New Value	SQL Text	Triggered ID	Count of Changed Columns
2007-04-02 10:00:00	192.168.2.12	ORACLE	on9eagle		SYSTEM	2007-04-02 10:08:38	BILLTEST	SYSTEM	INSERT	F1=Bill,F2=1,F3=1		insert into billtest values ('Bill',1,1)	7	3
2007-04-02 10:00:00	192.168.2.12	ORACLE	on9eagle		SYSTEM	2007-04-02 10:08:59	BILLTEST	SYSTEM	INSERT	F1=Bill,F2=2,F3=2		insert into billtest values ('Bill',2,2)	8	3
2007-04-02 10:00:00	192.168.2.12	ORACLE	on9eagle		SYSTEM	2007-04-02 10:07:15	BILLTEST	SYSTEM	INSERT	F1=Bill,F2=3,F3=3		insert into billtest values ('Bill',3,3)	9	3
2007-04-02 10:00:00	192.168.2.12	ORACLE	on9eagle		SYSTEM	2007-04-02 10:08:19	BILLTEST	SYSTEM	UPDATE F1=Bill,F2=1,F3=1	F1=Gregg,F2=5,F3=5		update billtest set f1=Gregg,f2=5,f3=5	10	3
2007-04-02 10:00:00	192.168.2.12	ORACLE	on9eagle		SYSTEM	2007-04-02 10:08:19	BILLTEST	SYSTEM	UPDATE F1=Bill,F2=2,F3=2	F1=Gregg,F2=5,F3=5		update billtest set f1=Gregg,f2=5,f3=5	11	3
2007-04-02 10:00:00	192.168.2.12	ORACLE	on9eagle		SYSTEM	2007-04-02 10:08:19	BILLTEST	SYSTEM	UPDATE F1=Bill,F2=3,F3=3	F1=Gregg,F2=5,F3=5		update billtest set f1=Gregg,f2=5,f3=5	12	3

Records: 1 To 6 From 6

Aliases: OFF

Drill-Down Report

Report	Description
Values Changed Details	Displays the details for a single row. This report is available only as a drill-down report. It displays the column name and old and new values for each change in the row. For example, the following report is produced when drilling down on the fifth row of the report above.

Guardium®										
Server IP	DB Type	Audit Owner	Audit Table Name	Audit Action	Changed Column Name	Old Value	New Value	Count of Changed Columns		
192.168.2.12	ORACLE	SYSTEM	BILLTEST	UPDATE	F1	Bill	Gregg	1		
192.168.2.12	ORACLE	SYSTEM	BILLTEST	UPDATE	F2	1	5	1		
192.168.2.12	ORACLE	SYSTEM	BILLTEST	UPDATE	F3	1	5	1		

Records: 1 To 3 From 3

Chapter 6: Command Line Interface

Introduction

The SQL Guard command line interface (CLI) is an administrative tool that allows for configuration, troubleshooting, and management of the SQL Guard system. The command line examples listed below use certain conventions for variables and syntax.

Most SQL Guard CLI commands have changeable parameters that affect the outcome of command execution. The SQL Guard documentation represents these parameters as variables. Such variables may include names, IP addresses, subnet masks, dates, etc.

The command language is structured in a specific way so that the administrator can specify verbs that command the CLI to perform its action on product subsystems or functions and arguments that these commands may need. For example, a CLI command might be structured as follows:

```
store alerter state on
```

In this example, the administrator is enabling the alerter subsystem.

Correspondingly, the administrator can display the currently configured value for the alerter subsystem's state by:

```
show alerter state
```

In this example, there is no argument being passed, but simply a command verb with a subsystem and specific value.

Note: CLI syntax elements are *not* case sensitive.

CLI Command Arguments

Commands that handle a “state” setting accept and use the following state arguments:

on or **off**

up or **down**

enabled or **disabled**

active or **inactive**

1 or **0**

All commands can accept “?” as an argument. Using a “?” displays the command usage and option choices.

CLI Command Abbreviations

You may abbreviate commands and subcommands as long as you provide enough characters so the commands are not ambiguous.

For example:

`show`

Can be shortened to:

`sho`

Accessing the CLI

An administrator can access the CLI though:

- A physically connected PC console or serial terminal
- OR
- A network connection using an SSH client

Physical Console Access

Interactive access to the SQL Guard is through the serial port or the system console.

PC keyboard and monitor – A PC video monitor can be attached to either the front panel video connector or the video connector on the back of the appliance.

A PC keyboard with a PS/2 style connector can be attached to the PS/2 connector on the back of the appliance. Alternatively, a USB keyboard can be connected to the USB connectors located at the front or back of the appliance.

Serial port access – Using a NULL modem cable, connect a terminal or another computer to the 9-pin serial port at the back of the appliance. The terminal or a terminal emulator on the attached computer should be set to communicate as 19200-N-1 (19200 baud, no parity, 1 stop bit).

A login prompt displays once the terminal is connected to the serial port, or the keyboard and monitor are connected to the console:

```
SQL Guard  
Unauthorized access is prohibited  
guard login:
```

The administrative user for interactive command line access is `cli`. The SQL Guard administrator should set a strong password to protect this account and that password must be entered at the password prompt to complete the login. Initially, if you have just rebuilt the system from a CD, the SQL Guard `cli` user has a default password of `guardium`.

After logging in, the following CLI prompt is displayed:

```
SQL Guard  
Unauthorized access is prohibited  
guard login: cli  
Password:  
Last login: Wed Oct 22 14:26:38 on ttym1  
guard.yourcompany.com>  
guard.yourcompany.com>  
guard.yourcompany.com> _
```

The CLI is now ready to accept commands.

Network SSH Access

Remote access to the CLI is available on the management IP address or domain name using Secure Shell (SSH).

An SSH client is needed to access the CLI through the network. SSH is an encrypting protocol that can allow for secure communications to systems. It is a de-facto industry standard for interactive access to a remote system on a TCP/IP network.

SSH clients are freely or commercially available for most desktop and server platforms.

To access the SQL Guard CLI through an SSH connection, connect the SSH client to the CLI user at the IP address of the management IP.

A Unix command line example:

```
ssh -l cli 192.168.2.16
```

The SSH client may ask for the administrator to accept the cryptographic fingerprint of the remote server. Accept the fingerprint to proceed to the password prompt.

Note: If, after the first connection, you are asked again for a fingerprint, someone may be trying to induce you to log into the wrong machine.

At the password prompt, enter the password for the *cli* user account.

CLI Command Categories

You can enter the **commands** command to produce an alphabetical listing of all commands at any time. In the SQL Guard documentation, the CLI commands are divided into the following categories:

Show Commands – Used to display configuration values

Store Commands – Used to set configuration values

Operational Control Commands – Used to administer the Guardium system

Documentation Conventions

All CLI command examples are written in courier text (for example, `show system clock`).

To illustrate syntax rules, some command descriptions use dependency delimiters. Such delimiters indicate which command arguments are mandatory, and in which context. Each syntax description shows the dependencies between the command arguments by using special characters:

The < and > symbols denote a required argument.

The [and] symbols denote an optional argument.

The | (vertical bar) symbol separates alternative choices when only one can be selected.

For example:

```
store full-bypass <on | off>
```

Show Commands

You can enter the **commands** command to produce an alphabetical listing of commands.

show

Displays stored configuration values or live, operational settings.

Example: show network interface ip

Example: show system clock

show account Commands

These commands display the automatic account lockout settings.

- **show account lockout**

Indicates if automatic account lockout following a specified number of login attempt failures is enabled.

- **show account strike count**

Displays the number of failed login attempts in the specified strike interval (see below) before disabling the account.

- **show account strike interval**

Displays the number of seconds during which the failed login attempts (the strike count above) must occur in order to disable the account.

- **show account strike max**

Displays the maximum number of failed login attempts to be allowed for an account over the life of the server.

show alerter Commands

These commands display the alert transmitter settings.

- **show alerter poll**

Displays the alerter polling interval in minutes. This value represents the amount of time the alerter waits before checking its outgoing message queue to send SNMP traps or transmit email via SMTP.

- **show alerter smtp authentication type**

Displays the alerter SMTP email authentication configuration. Supported authentication types are:

NONE – Standard SMTP transfer

AUTH – Username/password authentication

- **show alerter smtp authentication username**
Displays the alerter SMTP email authentication configuration for the remote username used to authenticate SMTP transport connections.
- **show alerter smtp port**
Displays the alerter SMTP email configuration for the port number used to connect for SMTP transport connections. Default: 25, the standard SMTP port.
- **show alerter smtp relay**
Displays the alerter SMTP email configuration for what remote gateway to use for SMTP transport.
- **show alerter smtp returnaddr**
Displays the alerter SMTP email configuration for a return address for email alerts. Any bounce messages or email failures will be returned to this address.
- **show alerter snmp**
Displays the alerter SNMP configurations.
- **show alerter snmp traphost**
Displays the alerter SNMP configurations for the remote trap server to receive alert traps.
- **show alerter state operational**
Displays whether the alerter subsystem is running or stopped.
- **show alerter state startup**
Displays whether the alerter is configured to start on system start-up.

show anomaly-detection Commands

These commands display anomaly detection engine settings. The anomaly detection engine executes user-defined alerts queries and queues any alert messages for the alerter subsystem to deliver.

- **show anomaly-detection poll**
Displays the anomaly detection polling interval in minutes.
- **show anomaly-detection state**
Displays whether the anomaly detection subsystem is enabled or disabled.

show auth Command

This command displays the type of authentication used for login to SQL Guard.

- **show auth**

show buffer Command

This command displays a report of buffer use for the inspection engine process.

- **show buffer**

show build Command

This command displays build information for the installed software.

- **show build**

show defrag Command

This command displays defragmentation parameters for the system disk.

- **show defrag**

For example:

```
> show defrag
Defrag parameters:
    Packet size:      0
    Time interval: 604800
    Trigger level: 10000000
    Release level: 10000000
```

show fail-policy Command

This command indicates what the unit does with messages when operating in inline mode, when the inspection engine is not running (it may be stopped for maintenance or there may have been a failure). Open, which is the installation-time default, indicates that all messages will be passed. Closed indicates that all messages will be blocked. This setting is also displayed by the show unit type command, in a slightly different format (fail-open or fail-close).

- **show fail-policy**

Note: Inline mode is controlled using the *store unit type* command, or by enabling the database firewall (see [store firewall](#)).

show firewall Command

This command indicates whether the database firewall is turned on or off.

- **show firewall**

show gui port Command

This command displays the TCP/IP port on which the Web interface accepts connections; at installation time it defaults to 8443.

- **show gui port**

show ignored port list Command

This command displays a list of ports ignored by all inspection engines.

- **show ignored port list**

show inspection-engines Commands

These commands display settings for the network database protocol inspection engines.

- **show inspection-engines all**

Displays all inspection engines.

- **show inspection-engines configuration <id>**

Displays the inspection engine identified by the specified **id** (which is taken from the list displayed by the *show inspection-engines all* command).

- **show inspection-engines log sqlstrings**

Displays a configured value for Inspection Engines to save complete SQL strings while logging data.

- **show inspection-engines type <mssql | mssql-np | oracle | sybase | db2 | informix | cifs | ftp>**

Displays all Inspection Engines configured for the specified database protocol type.

show installed security policy Command

This command displays the name of the installed policy.

- **show installed security policy**

show license Command

This command displays the SQL Guard license.

- **show license**

show log Commands

These settings (ON or OFF) control whether the inspection engine logs entire SQL commands.

- **show log exception sql**
If ON, logs the entire SQL command when logging exceptions.
- **show log request sql string**
If ON, logs the entire SQL command for each request.

show logging granularity Command

This command displays the logging granularity in minutes.

- **show logging granularity**

show maximum query duration Command

This command displays the maximum query duration, in seconds.

- **show maximum query duration**

show network arp-table Command

This command displays the address resolution protocol (ARP) table:

- **show network arp-table**

Note: This value is an operational system value.

show network interface all Command

This command shows settings for the network interface used to connect the SQL Guard unit to the desktop LAN. The IP address, mask, state (enabled or disabled) and high availability status will be displayed.

- **show network interface all**

If IP high-availability is enabled, the system will use two interfaces (**ETH 0** and **ETH 3**). Otherwise, only **ETH 0** will be used.

show network interface inventory Command

This command shows the port names and mac addresses of all installed network interfaces. For example:

```
g3.guardium.com> show network interface inventory
eth0 00:13:72:50:CF:40
eth1 00:13:72:50:CF:41
eth2 00:04:23:CB:11:84
eth3 00:04:23:CB:11:85
eth4 00:04:23:CB:11:96
eth5 00:04:23:CB:11:97
ok
```

- **show network interface inventory**

show network interface port

Use this command to locate a physical connector on the back of the machine. After using the show network interface inventory (see above), use this command to blink the light on for that physical port.

- **show network interface port <n>**

The output will appear like the example below:

```
g3.guardium.com> sho net int port 0
The orange light on port eth0 will now blink 20 times.
```

show network macs Command

Displays a list of mac addresses (like the show network interface inventory command, see above). For example:

```
g3.guardium.com> sho net macs
eth0:    00:13:72:50:CF:40
eth1:    00:13:72:50:CF:41
eth2:    00:04:23:CB:11:84
eth3:    00:04:23:CB:11:85
eth4:    00:04:23:CB:11:96
```

```
eth5:      00:04:23:CB:11:9
ok
```

- **show network macs**

show network resolver Command

This command displays the IP addresses for a specific resolver, or for all resolvers.

- **show network resolver <1 | 2 | 3 | all >**

show network routes Commands

These commands display the IP address for the default router, or display operational information for the router.

- **show network routes defaultroute**
- **show network routes operational**

Note: This value is an operational system value.

show password Commands

These commands indicate if password disabling, expiration and validation are in use for SQL Guard portal user passwords. (These settings do not apply to the *cli* user password.)

- **show password disable**
Displays the number of days without activity after which user accounts will be disabled. A value of zero indicates that user accounts will not be disabled. The default is zero.
- **show password expiration**
Displays the number of days set for user password expiration. Zero indicates that passwords never expire. The default is 90.
- **show password validation**
Indicates if password validation is turned on or off. It is **on** by default.

When password validation is enabled, the password must be eight or more characters in length, and must include at least one uppercase alphabetic character (A-Z), one lowercase alphabetic character (a-z), one digit (0-9), and one special character from the table below.

Table of Special Characters for SQL Guard Passwords

Special Character	Character Name
@	Commercial at

#	Number sign
\$	Dollar sign
%	Percent sign
^	Circumflex accent (carat)
&	Ampersand
.	Full stop (Period)
;	Semicolon
!	Exclamation mark
-	Hyphen (minus)
+	Plus sign
=	Equals sign
_	Low line (underscore)

show product gid Command

This command displays the product global identifier(GID) value.

- **show product gid**

Displays the stored unique product global identified (GID) value.

show purge objects age Command

This command displays a table showing the age that has been set for purging various types of non-critical objects. You can use the index displayed for each object type to modify the purge age for that object (see [store purge object](#)).

- **show purge objects age**

Sample output:

```
> show purge objects age
Index Name
Age
1   Central Management Persistant Operations      7
2   S-Tap Event Log                            14
4   Assessment Tests                           7
5   Central Management Temporary Policies       7
6   S-Tap Change History                      14
7   Kerberos Authentication Info.              14
8   Comment History                           60
9   Comment Local History                     60
10  Call Graph History                        90
11  CAS Host Event History                   7
12  Unused CAS Access Names                  7
13  Unused CAS Access Name Templates        7
14  Custom Table Operations Log             7
15  table in custom db without def         7
16  Baseline entries referred to user      30
17  Internal Archive/Restore/Purge event log 30
18  Classification Process Results          7
```

show remotelog Command

If remote logging has been enabled (see [store remotelog](#)), this command displays the name of the system to which syslog messages are written. If remote logging is not enabled, the message *Not configured* displays.

- **show remotelog**

show security policies Command

This command displays the list of security policies.

- **show security policies**

show storage-system Command

Displays the storage systems configured for backup and archiving operations. Multiple systems (Network and Centera, for example) may be available for each type of backup and archiving operation.

show support state Command

This command indicates whether or not system alerts will be sent to the support email address, which can be configured using the [forward support email](#) command. By default, the support state is enabled, and the default email address is support@guardium.com.

- **show support state**

show support-email Command

This command displays the email address to which system alerts will be sent. It can be configured using the [forward support email](#) command. By default, the support email address is support@guardium.com.

- **show support-email**

show storage-system Command

Displays the types of storage systems that are enabled for archive and backup operations.

- **show storage-system**

Displays storage systems enabled for archive and backup.

Example:

```
show storage-system
NETWORK : archiving and backing-up
CENTERA : archiving and backing-up
TSM      : archiving and backing-up
```

show system Commands

These commands display system clock, NTP, domain name, and host name settings.

- **show system clock all**

Displays a summary of the settings for the system clock and time.

- **show system clock datetime**

Displays the current system clock's date and time.

Note: This value is an operational system value.

- **show system clock timezone**

Displays the configured time zone for this system.

- **show system domain**
Displays the system's configured DNS domain name.
- **show system fullname**
Displays the complete host and domain name for this system.
- **show system hostname**
Displays the system's host name.
- **show system ntp all**
Displays a summary of configurations for the NTP subsystem.
- **show system ntp server**
Displays the configured NTP time source server.
- **show system ntp state**
Displays whether the NTP subsystem is enabled or disabled.

show system public key Commands

Displays the public key for cli or tomcat. If none exists, this command creates one.

- **show system public key <cli | tomcat>**

show throttle Command

Displays throttle parameters.

- **show throttle**

For example:

```
> show throttle
Throttle parameters:
    Packet size: 228000
    Time interval: 604800
    Trigger level: 10000000
    Release level: 10000000
```

show transfer-method Command

This command displays the method to use for all *network* back-up and archive operations. Note that you can also enable Centera or TSM as an alternative to *network*.

- **show transfer-method**

show unit type Command

This command displays the SQL Guard unit type attributes. Some attributes listed are set using the [store unit type](#) command, and cleared using the [clear unit type](#) command.

Others are set by the [store fail-policy](#) command or the [store firewall](#) command, as noted below. And one attribute is set only at installation time, and cannot be modified except by re-installing the SQL Guard software (also as noted below).

- **show unit type**

Attribute	Description
aggregated	Unit collects data and sends it to an aggregator.
aggregator	This property is set only during installation of the SQL Guard software. When the unit is configured as an Aggregator, it cannot inspect network traffic, provide database firewall protection, or serve as an S-Tap host.
fail-closed	If the inspection engine is down, all messages will be blocked. This attribute is set by the store fail-policy command.
fail-open	If the inspection engine is down, all messages will be passed. This attribute is set by the store fail-policy command.
inline	Port forwarding (messages are read on one port and forwarded onto another). This attribute is added automatically when the firewall option is enabled (see store firewall).
load-balancer	Unit manages load balancing for other units.
manager	Central manager functions are enabled for this unit.
netinsp	Inspection of network traffic is enabled.
standalone	Local management (independent of a central manager).
stap	The unit can be used as an S-Tap host, i.e., the unit can receive data from and remotely manage SQL RemoteGuard or S-Tap components

Store Commands

The store commands set configuration values or live, operational settings. You can enter the **commands** command to produce an alphabetical listing of commands.

store account Commands

These commands control the automatic user account lockout settings.

- **store account lockout <on | off>**

Controls if automatic account lockout following a specified number of login attempt failures is enabled.

Example: store account lockout on

Note: If the **admin** user account becomes locked, use the **unlock admin** CLI command to unlock it.

- **store account strike count <n>**

Sets the number of failed login attempts in the specified strike interval (see below) before disabling the account.

Example: store account strike count 3

- **store account strike interval <n>**

Sets the number of seconds during which the failed login attempts (the strike count above) must occur in order to disable the account.

Example: store account strike interval 300

- **store account strike max <n>**

Sets the maximum number of failed login attempts to be allowed for an account over the life of the server. The default is 10.

Example: store account strike max 99

Note: If automatic account lockout is enabled, setting the strike **count**, **interval**, or **max** value to zero does **NOT** disable that type of check. On the contrary, it means that after just one failure the user account will be disabled!

store alerter Commands

These commands store alerter settings.

- **store alerter poll <n>**

Sets the alerter polling interval in minutes. **n** the number of minutes that the alerter waits before checking its outgoing message queue to send SNMP traps or transmit email via SMTP.

Example: store alerter poll 30

store alerter smtp authentication Commands

These commands set the alerter SMTP email configuration.

- **store alerter smtp authentication password <pw>**

Sets the alerter SMTP authentication password to **pw**.

Example: store alerter smtp authentication password zXcvb123

- **store alerter smtp authentication type <none | auth>**

Sets the SMTP authentication type used by the alerter SMTP server to the specified value. Supported authentication types are:

NONE: standard SMTP transfer

AUTH: username/password authentication

Example: store alerter smtp authentication type none

- **store alerter smtp authentication username <name>**

Sets the alerter SMTP email authentication username to the specified **name**.

- **store alerter smtp port <n>**

Sets the port on which the alerter SMTP server listens, to the value specified by **n**. Default: 25, the standard SMTP port.

Example: store alerter smtp port 25

- **store alerter smtp relay <ip address>**

Sets the gateway for the SMTP server to the IP address to the specified value.

Example: store alerter smtp relay 10.1.2.51

- **store alerter smtp returnaddr <email address>**

Sets the alerter return **email address** for email alerts. Any bounce messages or email failures will be returned to this address.

Example:

store alerter smtp returnaddr admin@sagroup.mycompany.com

store alerter snmp Commands

These commands configure the alerter for SNMP communications.

- **store alerter snmp community <name>**

Sets the alerter SNMP trap community to the **name** specified.

Example: store alerter snmp community public

- **store alerter snmp traphost <ip address>**

Sets the alerter SNMP trap server to receive alert traps, to the specified IP address or DNS host name.

Example: store alerter snmp traphost monitorbox.mycompany.com

Example: store alerter snmp traphost 10.1.52.199

store alerter state Commands

These commands change state settings for the alerter subsystem.

- **store alerter state operational <on | off>**

Starts (on) or stops (off) the alerter subsystem.

Example: store alerter state operational off

Note: You can also use the restart or stop commands to start or stop the alerter subsystem. These commands are described in the *Operation Control Commands* section, below.

- **store alerter state startup <on | off>**

Sets whether the alerter subsystem starts on system start-up.

Example: store alerter state startup off

store anomaly-detection Commands

These commands utilize the anomaly detection engine, allowing you to execute user-defined alerts queries and queue any alert messages for the alerter subsystem to deliver.

- **store anomaly-detection poll <n>**

Sets the anomaly detection polling interval to the number of minutes specified by **n**.

Example: store anomaly-detection poll 60

- **store anomaly-detection state <on | off>**

Sets whether the anomaly detection subsystem is enabled or disabled.

Example: store anomaly-detection state on

store auth SQL_GUARD Command

Use this command to reset the type of authentication used for login to the SQL Guard server, to SQL_GUARD. LDAP or Radius authentication can be configured and enabled from the administrator portal, but not from the CLI.

- **store auth SQL_GUARD**

store certificate Command

This command stores a server certificate on the SQL Guard unit. See the [Certificate Commands](#) topic later in this chapter for a description of how to use this command.

store defrag Commands

Use the first format of this command to restore defaults, or use the second to store defragmentation settings. In either case, after entering one of these commands, you will need to issue the **restart inspection-engine core** command for the changes to take effect.

- **store defrag default**
- **store defrag size < s > interval < i > trigger < t > release < r >**
Specify the packet size **s** in bytes, up to a maximum of 2^{17} (131072).
Specify the time interval **i**, the trigger level **t**, and the release level **r** as a number of seconds, up to a maximum of 2^{31} (2147483648).

store fail-policy Command

This command controls what the unit does with messages in inline mode, when the inspection engine is not running (it may be stopped for maintenance or there may have been a failure). *open*, which is the installation-time default, indicates that all messages will be passed. This is like having a policy with a single *Allow all* rule. *close* indicates that all messages will be blocked, which is like disconnecting the SQL Guard Server from the network.

- **store fail-policy <open | close>**

Example: store fail-policy close

Note: *Inline* mode is enabled using the [store unit type](#) command and disabled using the [clear unit type](#) command. *Inline* mode is also enabled when the database firewall is enabled (see [store firewall](#)).

store firewall Command

This command turns the database firewall on or off. When the database firewall is enabled, special access blocking rule actions (*Drop* and *Terminate*) are available in security policies. See the description of the policy rules in the *SQL Guard User Guide* for more information about using these actions.

- **store firewall <on | off>**

When **on**, the unit provides database firewall security.

Example: store firewall on

Notes: Setting the firewall option **on** automatically adds the *inline* attribute to the system's unit type. See [store unit type](#).

When changing the database firewall setting, be sure to check the fail-policy setting (see [show fail-policy](#)) and change it if necessary (see [store fail-policy](#)).

store full-bypass Command

This command is intended for emergency use only, when traffic is being unexpectedly blocked by the SQL Guard server. When on, all network traffic passes directly through the system, and is not “seen” by SQL Guard.

When using this command, you will be prompted for the admin user password.

- **store full-bypass <on | off>**

When **on**, all network traffic passes through the system.

Example: store full-bypass on

store gui port Command

The store gui port command sets the port used by the SQL Guard management interface.

- **store gui port <n>**

Sets the TCP/IP port number on which the management interface accepts connections. The default is 8443. **n** must be a value in the range of 1024 and 65535.

Example: store gui port 8082

store ignored port list Command

This command sets the port numbers to be ignored by the inspection engines.

- **store ignored port list <n>[,<n2>]**

Sets one or more port numbers to be ignored by the inspection engine. The list you specify completely replaces the existing list. Each number is separated from the next by a comma, and no blanks or other white-space characters are allowed in the list. Use a hyphen to specify an inclusive range of numbers.

Example: store ignored port list 33,60-70

store inspection-engine log sqlstrings Command

This command turns the Inspection Engines log on or off.

- **store inspection-engine log sqlstrings <on | off>**

When **on**, the Inspection Engines log complete SQL strings while logging data.

Example: store inspection-engines log sqlstrings off

store installed security policy Command

This command sets the installed security policy.

- **store installed security policy <name>**

Sets the specified security policy as the installed security policy.

Example: store installed security policy SecPolAugust04

store license Command

This command sets the media from which the license is accessed each time the system starts.

- **store license <console | usb | fd>**

Sets the media type from which the system will access the license: console, USB device, or FD (floppy disk). This command works *only* from the console. When the *console* setting is used, the system will prompt for the license, and it will be stored in memory during operation (and perhaps vulnerable).

Example: store license console

store local-stap Command

This command controls whether or not the SQL Guard server will use Kerberos traffic to decode Kerberos-encrypted database user names. This feature applies to MS SQL traffic only. If you enable this feature, be sure that the SQL Guard unit *sees* the Kerberos traffic, which will be from (and to) the Windows Domain Controller.

- **store local-stap <on | off>**

When ON, Kerberos-encrypted database user names are decoded. When OFF, Kerberos-encrypted database user names display as a string of hexadecimal characters.

After you enter the *store local-stap on* command, a message displays asking you to check that the unit type includes *stap*. To use this feature, the unit type for the system *must* include the *stap* option. See [store unit type](#) for more information.

If you change the unit type to include *stap*, be sure to issue the [restart inspection-core](#) command and then the [restart inspection-engines](#) command.

store log Commands

These settings (ON or OFF) control whether or not the inspection engine logs entire SQL commands.

- **store log exception sql <on | off>**

When set to ON, logs the entire SQL command when logging exceptions.

- **store log request sql string <on | off>**

When set to ON, logs the entire SQL command for each request.

store logging granularity Command

This command sets the logging granularity, to a specified number of minutes. You must use one of the minute values shown in the syntax, below.

- **store logging granularity <1, 2, 5, 10, 15, 30 or 60>**

store maximum query duration Command

This command sets the maximum number of seconds for a query. The default is 60. We do not recommend that you set this value above the default, because doing so increases the chances of overloading the system with query processing. This value can also be set from the Running Status Monitor panel.

- **store maximum query duration <n>**

Example: store maximum query duration 300

store network interface Commands

These commands set the IP address and mask for the network interface that users and S-TapS use to connect to the SQL Guard Server. One or two cards can be used. The first, which is always ETH 0, is always required. A second port (always ETH 3) can be made available using the high-availability option, as described below. (See the illustration on the top/back of your system for a diagram of network connector numbering. All standard configurations are illustrated in Chapter 1: [Network Interfaces and Connectors](#).)

- **store network interface high-availability <on | off>**

Enables a second interface card, ETH 3, to be used for the network interface. The second port (always ETH 3) must first be connected to the network, just as ETH 0 is connected. There is a slight delay, caused by the switch re-learning the port configuration.

Example: store network interface high-availability on

- **store network interface inventory**

Resets the network interface MAC addresses stored in the Guardium internal tables. This command is intended for use after replacing or moving any network card.

- **store network interface ip <ip address>**

Sets the IP address used by SQL Guard.

Example: store network interface ip 10.10.10.192

- **store network interface mask <ip mask>**

Sets the subnet mask for the SQL Guard IP address.

Example: store network interface mask 255.255.255.0

store network resolver Command

This command sets the IP address for the primary (1), secondary (2), or tertiary (3) DNS server to be used by the SQL Guard network interface card. Each resolver address must be unique. To remove a DNS server, enter **null** instead of an IP address.

- **store network resolver <1 | 2 | 3> <ip address | null>**

Example: store network resolver 1 192.168.1.25

Example: store network resolver 3 null

store network routes Command

This command sets the IP address for the default router to the specified value.

- **store network routes defaultroute <ip address>**

Example: store network routes 1 192.168.4.35

store password Commands

These commands control password disabling, expiration and validation:

- **store password disable <days>**

Sets the number of days of inactivity, after which user accounts will be disabled. When set to 0 (zero), no accounts will be disabled for this purpose. For any other value, the account. At installation, the default value is zero. You must restart the GUI after changing this setting.

Example: store password disable 30
 restart gui

- **store password expiration <age>**

Sets the age (in terms of days) for user password expiration. When set to 0 (zero), the password never expires. For any other value, the account user must reset the password the first time they log in after the current password has expired. The default value is 90. You must restart the GUI after changing this setting.

Example: store password expiration 30
 restart gui

- **store password validation <on | off>**

Turns password validation on or off. The default value is **on**. You must restart the GUI after changing this setting.

Example: store password validation off
 restart gui

When password validation is enabled, the password must be eight or more characters in length, and must include at least one uppercase alphabetic character (A-Z), one lowercase alphabetic character (a-z), one digit (0-9), and one special character from the table below.

Table of Special Characters for SQL Guard Passwords

Special Character	Character Name
@	Commercial at
#	Number sign
\$	Dollar sign
%	Percent sign
^	Circumflex accent (carat)
&	Ampersand
.	Full stop (Period)
;	Semicolon
!	Exclamation mark
-	Hyphen (minus)
+	Plus sign
=	Equals sign
_	Low line (underscore)

store product gid Command

- **store product gid <n>**
Sets the stored unique product GID value.

Example: store product gid 122244

store purge object Command

This command sets the age (in days) at which non-essential objects will be purged. Use the **show purge objects age** command to display a table showing the index, object name, and age for each object type for which a purge age is maintained. Then use the appropriate index from that table in the command below to set the purge age.

- **store purge object <index> age <days>**

Example: First use the **show** command to display the indexes and current values:

```
>show purge objects age
1  Central Management Persistant Operations    7
2  S-Tap Event Log                            14
4  Assessment Tests                           7
5  Central Management Temporary Policies     7
6  S-Tap Change History                      14
7  Kerberos Authentication Info.            1
8  Comment History                          60
9  Comment Local History                   60
10 Call Graph History                      90
11 CAS Host Event History                  7
12 Unused CAS Access Names                7
13 Unused CAS Access Name Templates      7
14 Custom Table Operations Log           7
15 table in custom db without def       7
16 Baseline entries referred to user    30
17 Internal Archive/Restore/Purge event log 30
18 Classification Process Results        7
```

Then use the *store* command to set a new purge age for a specific object:

```
>store purge object 2 age 21
```

Note: The age zero (0) means that objects of this type have a different purge criteria, and cannot be reset by this means. For example, Audit Process Results (illustrated above) are purged as specified in each audit process definition.

store remotelog Command

This command controls the use of remote logging. When turned on, it identifies the system to which all *Unix syslog* messages will be written. In addition to the system messages, statistical alerts and policy rule violation messages can be written to syslog. See *Statistical Alerts* and *Building Policies* in the *SQL Guard User Guide* for more information.

If you turn on remote logging, be sure that the receiving host has enabled this capability (see below).

- **store remotelog [off | on <hostname>]**

Example: store remotelog on g116.mycompany.com

Notes: To configure the receiving machine to accept remote logging, edit */etc/sysconfig/syslog* on that machine to include the ‘-r’ option. For example:

```
SYSLOGD_OPTIONS="-r -m 0"
```

Then restart the syslog daemon:

```
/etc/init.d/syslog restart
```

The standard syslog file in Linux is named:

```
/var/log/messages
```

store stap certificate Command

Use this command to store a certificate from the S-Tap host (usually a database server) on the Guardium server. This command functions exactly like the store certificate console command, described later. See

- **store stap certificate**

You will be prompted as follows:

```
Please paste your new server certificate, in PEM format.  
Include the BEGIN and END lines, then press CTRL-D.
```

If you have not done so already, copy the server certificate to your clipboard. Paste the PEM-format certificate to the command line, then press CRTL-D. You will be informed of the success or failure of the store operation.

When you are done, use the *restart gui* command (described above) to restart the SQL Guard GUI.

store storage-system Command

This command adds or deletes a storage system type for archiving or system backup.

- **store storage-system <Centera | TSM> <backup | archive> <on | off>**

Add or delete the specified type of storage system as a destination option for backup or archive.

Example: store-system Centera archive on

store syslog-trap

This command enables or disables the syslog trap on the Guardium server.

- **store syslog-trap <on | off>**

When enabled, all syslog messages will be sent to the SNMP server configured for the unit.

store system Commands

These commands store APC parameters, system clock, NTP, domain name, host name, and shared system key settings.

- **store system apc battery-level <percent>**

Sets the minimum charge percent (0-100) before powering down.

- **store system apc timeout <seconds>**

Sets the maximum number of seconds to run on battery power before powering down.

- **store system clock datetime <YYYY-mm-dd hh:mm:ss>**

Sets the system clock's date and time to the specified value, where **YYYY** is the year, **mm** is the month, **dd** is the day, **hh** is the hour (in 24-hour format), **mm** is the minutes, and **ss** is the seconds. The seconds portion is required, but will always be set to 00.

Example: store system clock datetime 2003-10-03 12:24:00

- **store system clock timezone <list | timezone>**

Lists the allowable time zone value (list option), or sets the time zone for this system to the specified **timezone**.

Example: store system clock timezone list

Example: store system clock timezone america/new_york

- **store system domain <domain name>**

Sets the system's DNS domain name to the specified value.

Example: store system domain mycompany.com

- **store system hostname <hostname>**

Sets the system's host name to the specified value.

Example: store system hostname guardbox

store system ntp Commands

These commands identify an NTP server (optionally) used by the system and enable or disable use of that server.

- **store system ntp server <host_name>**

Sets the host name of the NTP time source server to be used to the specified value (you cannot specify an IP address here).

Example: store system ntp server ntp123.guardium.com

- **store system ntp state <on | off>**

Enables or disables use of the NTP server.

Example: store system ntp state on

- **store system patch install <cd | sys | ftp | scp >**

Installs one or more patches from a CD, a file location, or a network location.

cd – For one or more patches to be installed from a CD, insert the CD in the CD ROM drive before executing this command. Then respond to the following prompt:

Please choose one patch to apply (1-n,q to quit):

ftp or scp – For one or more patches located somewhere on the network, and to be accessed by ftp or scp, you must respond to the prompts shown below.

Be sure to supply the full path name for the patch, including the filename:

Host to import patch from:

User on hostname:

Full path to the patch, including name:

Password:

Please choose one patch to apply (1-n,q to quit):

sys – For one or more patches to be installed from the /tmp directory(?), you must respond to the following prompt:

Please choose one patch to apply (1-n,q to quit):

- **store system patch show <log | history>**

Displays the patch history or log.

Example: store system patch show log

- **store system shared key <key value>**

Sets the system's shared key value to the specified value. This key must be the same for a manager unit and all of the units it will manage. Once a managed unit has registered for management by the manager, the key is no longer used. (You cannot "unregister" a unit by changing this value.)

Example: store system shared key 123abc

store support state Command

This command enables or disables the sending of email alerts to the support email address, which is support@guardium.com by default, but can be configured using the [forward support email](#) command. By default it is enabled.

- **store support state <on | off>**

Example: store support state off

store throttle Command

Use the first format of this command to restore defaults, or use the second to store throttle settings. In either case, after entering one of these commands, you will need to issue the **restart inspection-engine core** command for the changes to take effect.

- **store throttle default**
- **store throttle size < s > interval < i > trigger < t > release < r >**
 - Specify the packet size **s** in bytes, up to a maximum of 2^{17} (131072).
 - Specify the time interval **i**, the trigger level **t**, and the release level **r** as a number of seconds, up to a maximum of 2^{31} (2147483648).

store transfer-method Command

This command sets the method used to transfer files from the SQL Guard unit, except for files sent to an aggregator, which are always sent using SCP.

- **store transfer-method <ftp | scp>**

Example: store transfer-method ftp

store trusted certificate Command

This command stores a CA or trusted path certificate on the SQL Guard unit. See the [Certificate Commands](#) topic later in this chapter for a description of how to use this command.

store unit type and clear unit type Commands

These commands store or clear SQL Guard unit type attributes (see the table below). The syntax options for setting or clearing the attributes is identical.

- **store unit type <manager | standalone > [aggregated] [inline] [load-balancer] [netinsp] [stap]**
- **clear unit type <manager | standalone > [aggregated] [inline] [load-balancer] [netinsp] [stap]**

Attribute	Description
aggregated	Enables aggregation functions for a collector. Note that the aggregator attribute is not set by the store unit type command – a unit must be built as an aggregator during the initial installation procedure.

Attribute	Description
inline	Indicates port forwarding is in use, which means that the unit should be installed inline between all clients on the one side, and the database servers on the other. Up to three pairs of network interface cards will be used (ETH 1 & 2, ETH 3 & 4, ETH 5 & 6). Messages from all clients are read on one port of the pair, and forwarded to the server the other port of the pair. This attribute is added automatically when the database firewall option is enabled (see store firewall). When inline mode is enabled, the fail-policy setting becomes important (see store fail-policy).
load-balancer	Enables load balancing functions. A Central Manager cannot also function as a load balancer. The CLI will allow this attribute but the functionality will not be allowed.
manager	Enables central manager functions
netinsp	Allows inspection of network traffic
standalone	Enables local management (independent of a central manager)
sTap	Can be used as an S-Tap server, meaning that the unit can receive data from, and manage S-Tap components

store user password Command

Use this command to reset the *cli* user password.

- **store user password**

To simplify the support process, we suggest that you keep the *cli* password assigned initially by Guardium. To change the *cli* password, use the **store user password** command. You will be prompted to enter the current password, and then the new password twice, as illustrated below. None of the password values you enter on the keyboard will display on the screen. The *cli* user password must:

- Be at least six characters in length.
- Contain at least one digit character (0-9).
- Contain at least one lowercase alphabetic character (a-z).
- Contain at least one uppercase alphabetic character (A-Z).

The **store user password** dialog should look like this, with the *guard.yourcompany.com* prompt replaced by the host and domain names configured for the SQL Guard server:

```
guard.yourcompany.com> store user password
Changing password for 'cli'.
```

```
Enter current password:  
Enter new password:  
Re-enter new password:  
Ok  
guard.yourcompany.com>
```

Note: There is no way to retrieve the CLI user password once it is set. If you lose this password, contact Guardium Technical Support to have it reset.

Inspection Engine Control Commands

This set of commands can be used to list, add, start, and remove inspection engines.

Note: See also the [show inspection-engines Commands](#) subtopic under the *Show Commands* topic earlier in this chapter.

- **list inspection-engines**
Displays a list of all inspection engines on the unit. Each unit is identified by an index number. Use the index number displayed in this list to identify an inspection engine in the *start*, *stop*, or *reorder inspection-engine* commands (but not in the *remove inspection-engine* command).
- **start inspection-core**
Starts the inspection engine core.
- **stop inspection-core**
Stops the inspection engine core.
- **start inspection-engine <first index>[,nth index]**
Starts one or more inspection engines identified using **index** values from the list produced by the *list inspection-engines* command (see above).
- **stop inspection-engine <first index>[,nth index]**
Stops one or more inspection engines identified using **index** values from the list produced by the *list inspection-engines* command (see above).
- **remove inspection-engine <name>**
Removes the single inspection engine identified by its name. The name can include only letters, numbers and blanks. If the inspection engine name contains any special characters, use the GUI to remove it.
- **reorder inspection-engine <first index>[,nth index]**
Specifies a new order for the inspection engines, using **index** values from the list produced by the *list inspection-engines* command (see above). In the

example below, if the displayed indices are 1, 2, 3, and 4, the current order of the engines will be reversed:

Example: reorder inspection-engine 4,3,2,1

- **add inspection-engine <name> <protocol>**

<fromIP/mask> [,fromIP/mask]

<port> [-port]

<toIP/mask> [,toIP/mask]

<exclude client list> <active on startup>

Adds an inspection engine configuration to the end of the inspection engine list. The arguments are described in the table below. You can re-order your list of inspection engines after adding a new one by using the *reorder inspection-engine* command. Adding an inspection engine *does not* start it running; to start it running, use the *start inspection-engine* command.

add inspection engine Command Arguments

Argument	Description
<name>	The new inspection engine name; must be unique on the unit.
<protocol>	The protocol monitored, which must be one of the following: DB2, Sybase, MSSQL, MSSQL-NP, Informix, CIFS, FTP or Oracle.
<fromIP/mask> [,fromIP/mask]	A list of clients, identified by IP addresses and subnet masks. Separate each IP address from its mask with a slash, and multiple entries by commas. An address and mask of all zeroes is a wildcard. If the <i>exclude client list</i> option (see below) is Y, the inspection engine monitors traffic from all clients <i>except</i> for those in this list. If the <i>exclude client list</i> option is N, the inspection engine monitors traffic from <i>only</i> the clients in this list.
<port> [-port]	The port or range of ports over which traffic between the specified clients and database servers will be monitored.
<toIP/mask> [,toIP/mask]	The list of database servers, identified by IP addresses and subnet masks, whose traffic will be monitored. Separate each IP address from its mask with a slash, and multiple entries by commas. An address and mask of all zeroes is a wildcard.
<exclude client list>	Y/N value; defaults to N. If Y, the inspection engine monitors traffic from all clients <i>except</i> for those identified in the client list (see above). If N, the inspection engine monitors traffic from <i>only</i> the clients listed in the client list.
<active on startup>	Y/N value; defaults to N. If Y, the inspection engine is activated

Argument	Description
	on system startup.

Operational Control Commands

You can enter the **commands** command to produce an alphabetical listing of commands.

? Command

Usage command which lists available commands or arguments. This command can be used as an argument to any command to produce a list of arguments available to that command.

Example: ?

Example: show ?

Example: show network ?

aggregator backup keys file Command

Use this command to back up the shared secret keys file to the specified location. See [About the System Shared Secret](#), in Chapter 2.

- **aggregator backup keys file <user@host:/path/filename>**

aggregator clean shared-secret Command

If no shared secret has been defined via the System Configuration panel, the shared secret value is null. All files archived or exported from a unit with a null shared secret can be restored or imported only on systems where the shared secret is null.

- **aggregator clean shared-secret**

Sets the shared secret value to null.

Note: Once a shared secret has been defined, you cannot change its value to null by blanking out the System Shared Secret field in the System Configuration panel. You can only set it to a null value using the *aggregator clean shared-secret* command, above. For more information about the shared secret, see [About the System Shared Secret](#), in Chapter 2.

aggregator debug start Command

Starts writing debugging information relating to aggregation activities. Use this command only when directed to do so by Guardium Support. Be sure to turn debugging off after the requested time interval – use the **aggregator debug stop** command.

- **aggregator debug start**

aggregator debug stop Command

Stops writing debuggin information relating to aggregation activities.

- **aggregator debug stop**

aggregator list failed imports Command

When a restore or import operation fails because of a shared secret mismatch, the offending file is moved from the /var/importdir directory to the /var/dump directory, and it is renamed using the original file name plus the suffix *.decrypt_failed*. Use the following command to list all such files:

- **aggregator list failed imports**

Lists all files with the *.decrypt_failed* suffix in the /var/dump directory.

aggregator recover failed import Command

Use this command to move and rename failed import files, prior to re-attempting an import or restore operation. Use the *all* option to move all files ending with the suffix *.decrypt_failed*, or use the *filename* option to identify a single file to be moved.

- **aggregator recover failed import <all | filename>**

Move all files ending with the suffix *.decrypt_failed*, or the specified file, from the /var/dump directory to the /var/importdir directory, renaming all moved files by removing the *.decrypt_failed* suffix.

Note: After moving the failed files, but before a restore or import operation runs, be sure that the system shared secret matches the shared secret used to encrypt the exported or archived file.

aggregator recover failed restore Command

Use this command to move and rename failed restore files, prior to re-attempting an restore operation. Use the *all* option to move all files ending with the suffix *.decrypt_failed*, or use the *filename* option to identify a single file to be moved.

- **aggregator recover failed restore <all | filename>**

Move all files ending with the suffix *.decrypt_failed*, or the specified file, from the /var/dump directory to the /var/importdir directory, renaming all moved files by removing the *.decrypt_failed* suffix.

Note: After moving the failed files, but before a restore or import operation runs, be sure that the system shared secret matches the shared secret used to encrypt the exported or archived file.

aggregator restore keys file Command

Use this command to restore the shared secret keys file from the specified location. See [About the System Shared Secret](#), in Chapter 2.

- **aggregator restore keys file <user@host:/path/filename>**

backup and restore system Commands

These commands back up and restore the SQL Guard internal database. You can back up or restore either configuration information only, or the entire system (data plus configuration information, except for the shared secret key files, which are backed up and restored separately, see the *aggregator backup keys file* and *aggregator restore keys file* commands, above). These commands stop all inspection engines and web services and restart them after the operation completes.

- **backup system**
- **import file**
- **restore system**

For all backup, import (see import file, below) and restore operations, you will receive a series of prompts to supply some combination of the following items, depending on which storage systems are configured and the type of restore operation. Respond to each prompt as appropriate for your backup or restore operation. The following table describes the prompts that you may receive:

Item	Description
1. SCP 2. FTP 3. TSM 4. CENTERA	Select the method to use to transfer the file. TSM and Centera will be displayed only if those storage methods that have been enabled (see the store storage-method command).
1. Data 2. Configuration	Select Configuration to back up definitions and configuration information only, or select Data to back up data in addition to configuration information.
1. restore from archive 2. restore from backup	Select <i>restore from archive</i> to restore archived data, or select <i>restore from backup</i> to restore configuration information.
1. normal 2. upgrade	If restoring from the same software version of SQL Guard, select <i>normal</i> . If restoring configuration information following software upgrade of the SQL Guard server, select

Item	Description
	<i>upgrade.</i>
host	The remote host for the backup file.
remote directory	The directory for the backup file. For FTP, the directory is relative to the FTP root directory for the FTP user account used. For SSH, the directory path is a full directory path. For Windows SSH servers, use Unix-style path names with forward slashes, rather than Windows-style backslashes.
username	The user account name to use for the operation (for backup operations, this user must have write/execute permission for the directory specified above)
password	The password for the above username.
file name	The file name for the archive or backup file.
Centera server	Enter the Centera server name. If using PEA files, use the following format: <Host name/IP>?<full PEA file name>, for example: 128.221.200.56?/var/centera/us_profile_rwqe.pea.txt
Centera clipID	For a Centera restore operation, the Content Address returned from the backup operation. For example: 6M4B15U4JM4LBeDGKCPF9VQO3UA

After you have supplied all of the information required for the backup or restore operation, a series of messages will be displayed informing you of the results of the operation. For example, for a restore system operation the messages should look something like this (depending on the type of restore and storage method used):

```
gpg: Signature made Thu Feb 22 11:38:01 2007 EST using DSA key ID 2348FF9E
gpg: Good signature from "Backup Signer <support@guardium.com>"
```

Proceeding to shutdown services
 Proceeding to startup services
 Safekeeping admin.xreg
 Safekeeping client.xreg
 Safekeeping controllers.xreg
 Safekeeping controls.xreg
 Safekeeping guardium-portlets.xreg
 Safekeeping local-portlets.xreg
 Safekeeping local-security.xreg
 Safekeeping local-skins.xreg
 Safekeeping media.xreg
 Safekeeping portlets.xreg
 Safekeeping security.xreg
 Safekeeping skins.xreg
 guard_sniffer.pl -reorder

```
Recovery procedure was successful.  
ok
```

commands Command

Use the **commands** command to produce an alphabetical listing of commands.

debug Command

Enable/disable debug mode. Without an argument, it toggles the debug state. Optionally, a state argument can be passed.

Example: debug on

diag Command

See the separate section for this command: [Diagnostics Command](#).

dsaa mail sender state Command

This command enables or disables the sending of a special e-mail notification to a specific user. **DSAA** is an acronym for Database Security Audit and Analysis. This notification includes two numeric amounts: a count of requests and a count of sessions. These counts begin when the inspection engine core was last started.

- **dsaa mail sender state [off | on <customer name> <email address>]**
 - off** – Disables the sending of mail.
 - customer name** – The name of the customer to receive email; should contain no spaces.
 - email address** – The email address to which the message is to be sent.

eject Command

This command dismounts and ejects the CD ROM. This is useful after upgrading or re-installing the system.

- **eject**

export audit-data Command

Use this command only under the direction of Guardium Support. This command is used to export data for a single date to a compressed archive file in the /var/dump directory.

- **export audit-data <yyyy-mm-dd>**

Exports audit data from various internal SQL Guard tables to a compressed archive file. The file created will be identified in the messages produced by the system. See the example below.

If you enter the audit-data command for the date 2005-09-16, a set of messages similar to the following will be created:

```
supp2.guardium.com> export audit-data 2005-09-16
2005-09-16
Extracting GDM_ACCESS Data ...
Extracting GDM_CONSTRUCT Data ...
Extracting GDM_SENTENCE Data ...
Extracting GDM_OBJECT Data ...
Extracting GDM_FIELD Data ...
Extracting GDM_CONSTRUCT_TEXT Data ...
Extracting GDM_SESSION Data ...
Extracting GDM_EXCEPTION Data ...
Extracting GDM_POLICY_VIOLATIONS_LOG Data ...
Extracting GDM_CONSTRUCT_INSTANCE Data ...
Generating tar file ...
/var/csvGenerationTmp ~
GDM_ACCESS.txt
GDM_CONSTRUCT.txt
GDM_CONSTRUCT_INSTANCE.txt
GDM_CONSTRUCT_TEXT.txt
GDM_EXCEPTION.txt
GDM_FIELD.txt
GDM_OBJECT.txt
GDM_POLICY_VIOLATIONS_LOG.txt
GDM_SENTENCE.txt
GDM_SESSION.txt
~
Generation completed, CSV Files saved to /var/dump/732570-
supp2.guardium.com-w20050919110317-d2005-09-16.exp.tgz
ok
```

The data from each of the named internal database tables is written to a text file, in CSV format. The name of the archive file ends with *exp.tgz* and the remainder of the name is formed as described under the topic [About Archived Data File Names](#), below.

You can use the *export file* command (see below) to transfer this file to another system.

export file Command

Use this command only under the direction of Guardium Support. This command is used to export a single file from the /var/dump, /var/log, or /var/importdir directory.

Note: To export SQL Guard data to an aggregator or to archive data, use the appropriate menu commands on the Administration Console panel.

- **export file </local_path/filename> <user@host:/path/filename>**
Exports the specified file to the specified directory.
local_path must be one of the following: **/var/log**, **/var/dump**, or
/var/importdir.

forward support email Command

Use this command to configure the email address for all messages that by default are sent to Guardium Support (support@guardium.com). To display the current setting, use the [show support-email](#) command.

- **forward support email to <email_address>**
Sets the email address for support messages.

help Command

Displays useful information about commands and command line interface usage. To display help on a topic, use “HELP command arguments”.

Example: help store network interface
Example: help restart inspection-engine

Additional help is available on using the command line interface with “help introduction”.

import file Command

This command is used to import archives to be restored on the system. Once the data has been imported, use the Data Restore command in the Data Management section of the Administration Console menu.

- **import file**
Imports an archive. For a description of the prompts that you may receive, see [backup and restore Commands](#), above.

About Archived Data File Names

When SQL Guard data is archived (or exported to an aggregator), there is a separate file for each day of data. Depending on how your export/purge or archive/purge operation is configured, you may have multiple copies of data exported for the same day. Archive and export data file names have the same format:

<daysequence>-<hostname.domain>-w<run_datestamp>-d<data_date>.dbdump.enc

daysequence is a number representing the date of the archived data, expressed as the number of days since year 0. The same date appears in *yyyy-mm-dd* format in the *data_date* portion of the name (see below).

hostname.domain is the host name of the SQL Guard unit on which the archive was created, followed by a dot character and the domain name.

run_datestamp is the date that the data was archived or exported, in *yyyymmdd.hhmmss* format.

data_date is the date of the archived data, in *yyyy-mm-dd* format.

For example:

732423-g1.guardium.com-w20050425.040042-d2005-04-22.dbdump.enc

import tsm config Command

Use this command to upload a TSM client configuration file to the SQL Guard server. You must do this before performing any archiving or backup operations using TSM. You will always need to upload a *dsm.sys* file, and if that file includes multiple *servername* sections, you will also need to upload a *dsm.opt* file. For information about how to create these files, check with your company's TSM administrator.

- **import tsm config <user@host:/path/[dsm.sys | dsm.opt]>**

Imports a TSM configuration file. *filename* must be either *dsm.sys* or *dsm.opt*. You will be prompted for a password for the user account on the specified host.

iptraf Command

IPTraf is a network statistics utility distributed with the underlying operating system. It gathers a variety of information such as TCP connection packet and byte counts, interface statistics and activity indicators, TCP/UDP traffic breakdowns, and LAN station packet and byte counts. The IPtraf User's Manual is available on the internet at the following location (it may be available at other locations if this link does not work):

<http://iptraf.seul.org/2.7/manual.html>

license check Command

This command to verifies that the SQL Guard license is valid.

- **license check**

list audit-data Command

Use this command to display any files ending with the suffix *.tgz* in the */var/dump* directory. For more information about audit-data files, see the description of the [export audit-data](#) command.

- **list audit-data**

load-balancer Commands

The following commands control load balancing operations. Load balancing cannot be performed by a Central Manager. Load balancing is enabled or disabled using the [store](#) or [clear load-balancer](#) commands.

- **load-balancer map <IP><MAC><ETH>**
Maps the IP and MAC addresses of a collector to the Ethernet port that it is connected to on the balancer.
- **load-balancer list mapping**
Lists all load mappings.
- **load-balancer unmap mapping_id**
Removes the mapping identified by a *mapping_id*, which is taken from the *load-balancer list mapping* command (above).
- **load-balancer link inspection_engine_name ethernet_port (ethX) backup_ethernet_port (ethX)**
Links a load balancer inspection engine to a primary Ethernet port and a backup Ethernet port.
- **load-balancer list links**
Lists all load balancer links.
- **load-balancer unlink link_id**
Removes the link identified by a *link_id*, which is taken from the *load-balancer list links* command (above).
- **load-balancer ping <interval>**
Sets the ping interval when determining if a collector is online.
- **load-balancer forwarding [on | off]**
Used on a collector to properly forward traffic. Must be off when the machine is a load balancing collector, and on when the machine is not participating in load balancing.

ping Command

Sends ICMP ping packets to a remote host. This command is useful for checking network connectivity.

Example: ping somehost.mycompany.com

Example: ping 10.2.15.29

quit Command

Exits the command line interface.

remove audit-data Command

Use this command only under the direction of Guardium Support. This command is used to remove exported audit data files. You will be prompted to enter an index number to identify the file to be removed. Only those files in the /var/dump directory ending with the *.tgz* suffix will be listed. (See [About Archived Data File Names](#) above for information about how archived data file names are formed.)

- **remove audit-data**

Removes a compressed audit data file.

You will be prompted to identify the file to be removed. For example:

```
supp2.guardium.com> remove audit-data  
1. 732570-supp2.guardium.com-w20050919110317-d2005-09-16.exp.tgz  
2. 732573-supp2.guardium.com-w20050920104600-d2005-09-19.exp.tgz  
Which file to delete ? (q to quit)...
```

Type the number of the file to remove, and press Enter.

register / unregister for Central Management Commands

These commands are used on managed units to register or unregister for central management. Please note the important caution notice below.

Caution: The **unregister** command restores the machine configuration that was saved when the machine was registered for central management. If that unit was registered for central management under a previous release of the Guardium software, restoring that configuration without first applying a patch to bring it to the current software release level will disable the unit, potentially causing the loss of all data stored there. Accordingly, **do not unregister a unit** until you have verified that the pre-registration configuration is at the current software release level. If you are unsure about how to verify this, contact Guardium Support **before unregistering the unit**.

- **register management <manager ip> <port>**

Registers a unit for management by the Central Manager identified by *manager ip* and *port*. The unit's pre-registration configuration will be saved.

- **unregister management**

Unregisters the unit from management by a Central Manager. The unit's pre-registration configuration will be restored, replacing any changes made

from the Central Manager. This command is intended for emergency use only, when the Central Manager is not available. See the note below.

Note : After unregistering using this command, you should also unregister *from* the Central Manager, since that is the only way the count of managed units will be reduced. For more information, see [About Central Manager Licenses](#) in Chapter 2.

restart Commands

These commands stop and restart a subsystem or process.

- **restart alerter**

Restarts the email/SNMP alert sender/transmitter.

Note: You can perform the same function using the store alerter state command to stop and then start the alerter:

```
store alerter state operational off  
store alerter state operational on
```

- **restart gui**

Restarts the Web interface.

- **restart inspection-core**

Restarts the database inspection engine core, but not the inspection engines. The collection of database traffic stops when this command is issued. To restart the collection of traffic for one or more specific inspection engines, follow this command with one or more `start inspection engine` commands (described earlier, under the Inspection Engine Control Commands topic). Alternately, to restart the collection of traffic for all inspection engines, use the `restart inspection engines` command (below).

- **restart inspection-engines**

Restarts the database inspection engine core and all inspection engines. The collection of database traffic stops temporarily while this occurs and restarts only when database connections reinitiate.

- **restart system**

Reboots the SQL Guard appliance. The system will completely shutdown and restart.

stop Commands

These commands stop a subsystem or process.

- **stop alerter**

Stops the alert message transmitter.

Note: You can perform the same function using the following command:

```
store alerter state operational off
```

- **stop gui**
Stops the Web user interface.
- **stop inspection-engines**
Stops all inspection engines. To start only a specific inspection engine, use *start inspection-engine* (described in the previous section).
- **stop system**
Stops and powers down the system.

unlock admin Command

This command unlocks the special admin user account, which may be locked out by the automatic account locking feature if the admin user has more failed login attempts than are allowed by that feature. (See the description of the [store account Commands](#).)

- **unlock admin**

Certificate Commands

Use these commands to create a certificate signing request (CSR), and to install server, CA, or trusted path certificates on the SQL Guard unit.

csr Command

This command generates a CSR for the SQL Guard unit. Do not perform this action until *after* the SQL Guard network configuration parameters have been set. Within the generated CSR, the common name (CN) will be created automatically from the host and domain names assigned. There are no arguments for this command:

- **csr**

Enter the command exactly as shown. You will receive a number of prompts, for the organizational unit (OU), country code (C), and so forth. Be sure to enter this information correctly.

After you respond to the last prompt, the system displays a description of the request, followed by the request itself, and followed finally by additional instructions. For example:

```
This is the generated CSR:  
Certificate Request:  
Data:  
Version: 0 (0x0)  
Subject: C=US, ST=MA, L=Waltham, O=XYZCorp, OU=Accounting, CN=g2.xyz.com  
-----BEGIN NEW CERTIFICATE REQUEST-----
```

```
MIICWjCCAhcCAQAwDELMAkGA1UEBhMCVVMxEDAOBgNVBAgTB1dhbHRoYW0xETAPBgNVBAoTCEd1
YXJkaXVtMRUwEwYDVQQLEwxndWFyZG11bS5jb20xCTAHBgNVBAMTADCCAbgwggEsBgchhkj00AQB
MIIIBHwKBgQD9f1OBHxUSKVLfSpwuOTn9hG3UjzvRADDHj+A1tEmaUVdQCJR+1k9jVj6v8X1ujD2
y5tVbNeBO4adNG/y2mC3a51QpaSfn+gEexAiwk+7qdf+t8Yb+DX58aophUPBpuD9tPFHsMCNvQT
WhaRMvZ1864rYdcq//IiAxmd0UgBxwIVAJdqU18VlwMspK5sqLrhAvwWBz1AoGBAPfhoIXWmz3e
y7yrXDa4V7151K+7+jrqgv1XTAs9B4JnUV1XjrrWU/mcQcQgYCOSRzxi+hMKBYTt88JMozIpuE8
FnqlVHyNKOCCjrh4rs6Z1kW6jfww6ITV18ftiegeko8yk8b6ouZCjgIPf4VrlnwaSi2ZegHtVJWQB
TDv+z0kqA4GFAAKBqCONSeB4g4/limbhku25Yln9CGM3a2evEnqjZxts4itxeTywPQvdkjdsMo
kaQ1BxmNUsZOJzrq5nC5Cg3X9spa+bzFr+PgR/Szk17nHcxKvjlk451L67K11Xv61TUfv/bU
PKmiaGKDttsp2ktG4dbFXQd1CjEGo0aNFCYn6qAAMasGByGSM44BAMFAAMwADatAhUAhHTY5z9X
NiBAuyAC9PS4GzleYakCFZ2kcfxfjXLBfY51228XWMAUOn95
-----END NEW CERTIFICATE REQUEST-----
```

Please copy and paste this output to a file, starting at the BEGIN and END lines, and use that file to work with your Certificate Authority in obtaining a certificate. I will be expecting the incoming certificate to be in PKCS#7 PEM format. Your CA will help you in receiving that format. Once you have it, please use the "store certificate" command to complete this operation.

Before continuing, check the Subject line to verify that you have entered your company information correctly. From this point forward, use whatever procedure you would normally use to obtain a server certificate from your CA. When you have obtained the server certificate, use the *store certificate* command (described below) to store the certificate on the unit.

store certificate Command

This command stores a server certificate on the SQL Guard unit. Before executing the command, obtain a server certificate (in PEM format) from your CA and copy the certificate, including the *Begin* and *End* lines, to your clipboard.

- **store certificate console**

Enter the command exactly as shown. You receive the following information and prompt:

```
Please paste your new server certificate, in PEM format.
Include the BEGIN and END lines, then press CTRL-D.
```

If you have not done so already, copy the server certificate to your clipboard. Paste the PEM-format certificate to the command line, then press CRTL-D. You will be informed of the success or failure of the store operation.

When you are done, use the *restart gui* command (described above) to restart the SQL Guard GUI.

store trusted certificate Command

This command stores a CA or intermediate trusted path certificate on the SQL Guard unit. When storing a CA and one or more intermediate certificates, you must store them in hierarchical order, beginning with the CA certificate. Before executing the command, obtain the appropriate certificate (in PEM format) from your CA, and copy the certificate, including the *Begin* and *End* lines, to your clipboard.

- **store trusted certificate**

Enter the command exactly as shown. The following prompt will be displayed:

```
What is a one-word alias we can use to uniquely identify this  
certificate?
```

Enter a one-word name for the certificate and press Enter. The following instructions will be displayed:

```
Please paste your CA certificate, in PEM format.  
Include the BEGIN and END lines, then press CTRL-D.
```

If you have not done so already, copy the certificate to your clipboard. Paste the PEM-format certificate to the command line, then press CRTL-D. You are informed of the success or failure of the store operation.

When you are done storing all certificates on the trusted path, use the *restart gui* command (described above) to restart the SQL Guard GUI.

Diagnostics Command

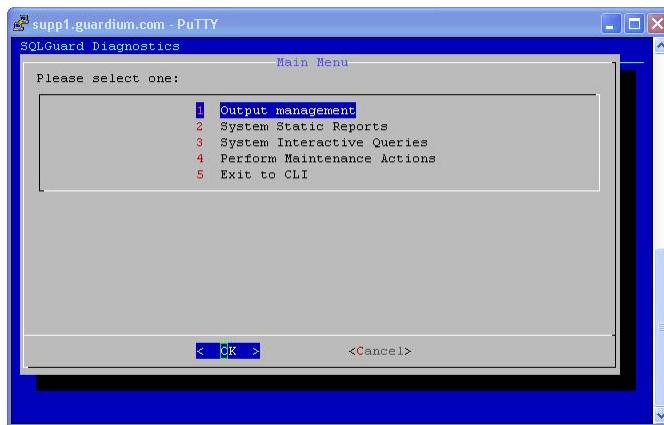
This section describes how to use the CLI **diag** command. For general instructions on how to use the CLI and detailed descriptions of all other CLI commands, see the previous sections of this chapter.

You should only use the **diag** command under the direction of Guardium Technical Support. There are no functions that you would perform with this command on a regular basis.

Opening the Diagnostics Main Menu

To use the **diag** command, follow the procedure outlined below:

1. Log into the SQL Guard unit as the *cli* user.
2. Enter the **diag** command (with no arguments) at the command line prompt.
3. Enter the same password used by the SQL Guard *admin* user when you are prompted for a password.
4. You are presented with the main command menu, illustrated below.



5. Do one of the following to move the option selection cursor (which is selecting the first item in the example above):
 - Type the desired entry's number (the selection cursor moves to the selected entry)
OR
 - Use the Up or Down arrow key to select the desired entry.
6. Press the Spacebar, the Left arrow key, or the Right arrow key to move the command selection cursor at the bottom of the display (which is selecting the OK command in the example above).
7. Perform an action by selecting the appropriate option in the display area and then doing one of the following:
 - Select the appropriate command with the command selection cursor, then press the Enter key
OR
 - Click on the appropriate action command.

About the Output

The diag command creates output in two directories:

- /var/log/guard/diag/current
- /var/log/guard/diag/depot

Each directory is described in the following subsections.

/var/log/guard/diag/current Directory

Most output from the diag commands is written in text format to the *current* directory. For most commands, this directory contains a separate output file. Each time you run the same command, output is appended to the single file for that command. For a smaller number of commands, a separate file is created for each execution, usually incorporating a date and time stamp in the filename.

Note: You should “clean up” after each session, so in subsequent sessions you are not looking at old information. When you pack files to a single compressed file for exporting (see the following topic), all files in the *current* directory are deleted. Alternatively, you can use the *Delete recordings* command of the Output Management menu to delete individual files.

The files in the *current* directory are easy to identify since the names are created from menu and command names. For example, after you use the File Summary command from the System Interactive Queries menu, a file named *interactive_filesummary.txt* is created in the *current* directory.

Note: If you look at the *current* directory while in the process of using a command, you may see a hidden temporary file with the same name as the one that will contain the output for that command. The temporary file will be removed when the output is appended to the command output file.

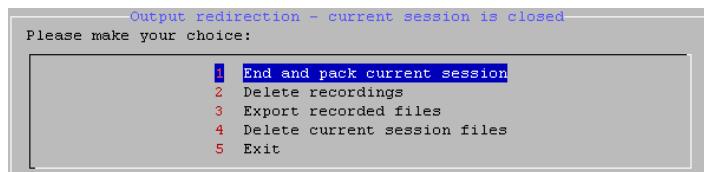
/var/log/guard/diag/depot Directory

When you pack the diag output files in the *current* directory to a compressed file (to send to Guardium Technical Support, for example), it is stored in the *depot* directory. The filename is in the format **diag_session_<dd_mm_hhmm>.tgz**, where the variable portion of the name indicates when the file was created. For example, a file created at 12:15 PM on May 20th would be named as follows: *diag_session_20_5_1215.tgz*.

After exporting files (see the *Export recorded files* topic, below), you can remove them from the *depot* directory using the *Delete recordings* command of the Output Management menu.

1: Output Management

The Output Management commands control what is done with the output produced by the diag command. Each Output Management command is described separately below.



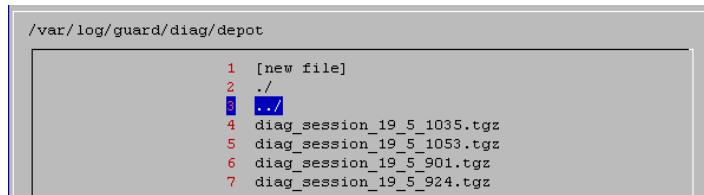
1.1: End and pack current session

Use this command to pack all diagnostic files in the *current* directory into a single compressed file, and remove those files from the *current* directory. When you enter this command, there is no feedback to indicate that the command has completed. You can verify that the command has finished by displaying the directory of the *depot* directory. When the command completes, there is a file named in the following format:

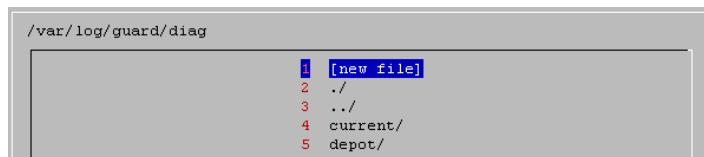
diag_session_<mm_dd_hhmm>.tgz, where the variable portion of the name is a date and time stamp, as described previously. Use the *Export recorded files* command of the Output Management menu to send the file to another system.

1.2: Delete recordings

Use this command to delete files in the *depot* or *current* directory. (To delete only the current session files, use the *Delete current session files* command.) When you enter this command, the *depot* directory structure displays:



You can navigate the directories using the Up and Down arrow keys and pressing Enter. For example, selecting **../** (as illustrated above) and pressing Enter moves the selection up one level in the directory structure:



You could then select the *current* directory and press enter, to navigate down to that folder and delete individual command output files:

```
/var/log/guard/diag/current/  
1 [new file]  
2 ./  
3 ../  
4 .interactive_filesummary.txt  
5 .tcpdump.tmp  
6 interactive_filesummary.txt  
7 interactive_list_directory.txt
```

Note: You can navigate to other directories, but you cannot delete files except from the current and depot directories.

When you have selected the file you want to delete, press Enter.

Caution: You will not be prompted to confirm the delete action.

1.3: Export recorded files

Use this command to send a file from the *depot* directory to another site. To export a file:

1. Select *Export recorded files* from the Output Management menu. The depot directory displays:

```
/var/log/guard/diag/depot  
1 [new file]  
2 ./  
3 ../  
4 diag_session_18_5_1412.tgz  
5 diag_session_19_5_822.tgz
```

2. Select the file to be sent or use the **./** and **../** entries to navigate up or down in the directory structure. (However, keep in mind that you can only export files from the *depot* directory.)
3. With the file to be transmitted selected, press Enter.
4. You are prompted to select FTP or exit:

```
Export /var/log/guard/diag/depot/diag_session_19_5_822.tgz  
Please make your choice:  
1 FTP  
2 Exit
```

Select FTP and press Enter.

5. You are prompted to supply a host name:



Enter the host name of the receiving system (or its IP address), and press Enter.

6. You are prompted for a user name:



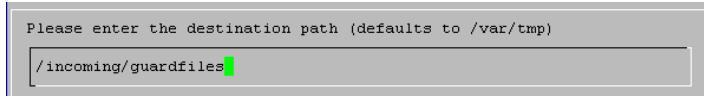
Enter a user account name for the receiving system, and press Enter.

7. You are prompted for a password:



Enter the password for the user on the receiving system.

8. You are prompted to identify a directory to receive the sent file on the receiving system:



Enter the path relative to the ftp root of the directory to contain the file on the receiving system and press Enter.

9. You are prompted to confirm the details of the transfer (the file to be sent and its destination). Press Enter to perform the transfer, or select Cancel and press Enter to start over.
10. You are informed of the success (or failure) of the operation; for example:



1.4: Delete current session files

Use this command to delete files created during the current session.

1.5: Exit

Use the Exit command to return to the main menu.

2: System Static Reports

Use the System Static Reports command of the Main Menu to produce an extensive set of reports (described below).

1. Select System Static Reports from the Main Menu. You are informed that the process is running:

```
Building report, please wait.
```

2. After the report has been created, it displays in the viewing area:

```
^ (+)
Build version: 34e1eb12eb68ba76cb49028251c9a0d6 /usr/local/guardium/etc

Patches:
2005/05/12 15:31:30: START Installation of 'Update 3.6.Op1'
2005/05/12 15:31:56: Installation Done - Successfully Installed

2005/05/12 17:24:45: START Installation of 'Update 3.6.Op2'
2005/05/12 17:24:56: Installation Done - Successfully Installed

2005/05/12 18:58:03: START Installation of 'Update 3.6.Op2'
2005/05/12 18:58:13: Installation Done - Successfully Installed

Current uptime:
10:13:17 up 7 days, 18:44, 2 users, load average: 0.32, 0.32, 0.12
System nameservers:
192.168.3.20
L (+) ( 0% )
```

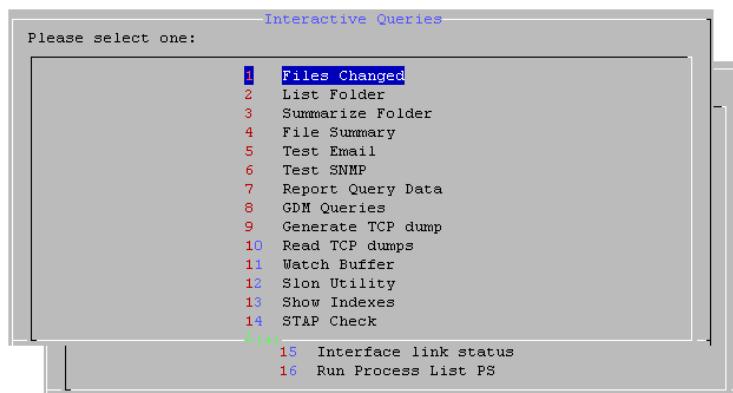
Note: This report is lengthy and may be easier to view using a text editor, after exporting it to a desktop computer (see the *Export recorded files* topic, above).

Use the Up and Down arrow keys to scroll up or down in the report. When you are done viewing the report, press Enter to return to the Main Menu.

For an outline of the information contained in this report, see the [System Static Reports Overview](#) topic at the end of this section.

3: System Interactive Queries

Select System Interactive Queries from the main menu to open the Interactive Queries menu, which is illustrated below. (Use the Down arrow key to scroll past the tenth item to see all items on this menu.)



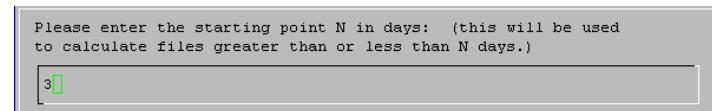
In addition to displaying the requested information, each interactive query command creates output in a separate text file in the *current* directory. See the Overview topic above for more information about the files created.

Each command is described in the following sections.

3.1: Files Changed

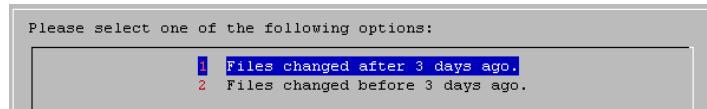
Use the Files Changed command to display a list of files changed either before or after a specified number of days.

1. Select Files Changed from the Interactive Queries menu. You are prompted to enter a number days:



Type a number and press Enter.

2. You are asked if you are interested in the files changed before or after that number of days:



Select 1 or 2 and press Enter.

3. The full directory path for each changed file is displayed. For example:

```
^ (4)
/var/log/lastlog
/var/log/messages
/var/log/secure
/var/log/wtmp
/var/log/sa/sa17
/var/log/sa/sa17
/var/log/sa/sa18
/var/log/guard/tomcat_log/localhost_log.2005-05-18.txt
/var/log/guard/diag/current/interactive_filechanged_-3.txt
/var/log/guard/diag/current/interactive_filechanged_+1.txt
/var/log/guard/diag/current/interactive_filechanged_-1.txt
/var/log/guard/diag/depot/diag_session_18_5_1412.tgz
/var/log/guard/snif_output.txt
/var/log/guard/snif_err.txt
/var/log/guard/snif_stderr.txt
/var/log/guard/snif_buf_usage.txt
L (+) < EXIT > ( 52%)
```

Note: If not all data fits in the display area, use the Up and Down arrow keys to scroll through the data. The current position in the file is indicated by the number in the lower right part of the display (52% in the example above). The white bars above and below the display area indicate the presence of more data with a plus sign.

3.2: List Folder

Use this command to list the contents of various directories.

1. Select List Folder from the Interactive Queries menu.
2. You are prompted to select a directory:



Select a directory and press Enter. The selected directory is displayed. Remember that if multiple commands of the same type are issued, the data for each execution of the command is appended to the single text file maintained for that command.

In the example below, the command has been issued for the /usr/local/guardium directory:

```
DIRECTORY: /usr/local/guardium
total 24
drwxr-xr-x    4 root      root        4096 May 20 18:38 upgrade
drwxrwxr-x    4 root      tomcat     4096 May 13 15:42 etc
drwxrwxr-x    2 root      tomcat     4096 May 13 15:21 scripts
drwxrwxr-x    2 root      tomcat     4096 Apr 28 16:36 bin
drwxrwxr-x    2 root      tomcat     4096 Apr 28 16:36 cli
drwxrwxr-x    2 root      tomcat     4096 Apr 28 16:36 diag
```

3. Press Enter or click Exit when you are done.

3.3: Summarize Folder

Use the Summarize Folder command to display the output of the **du** (Disk Usage) command, as illustrated below:

1. Select Summarize Folder from the Interactive Queries menu. There are no prompts. You are presented with a display of disk use for various directories:

```
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/hdc3           2064200    1127540   831804  58% /
/dev/hdc1            99043       9369    84560  10% /boot
none                515232        0    515232  0% /dev/shm
/dev/hdc2           74215848   1030144   69415684  2% /var

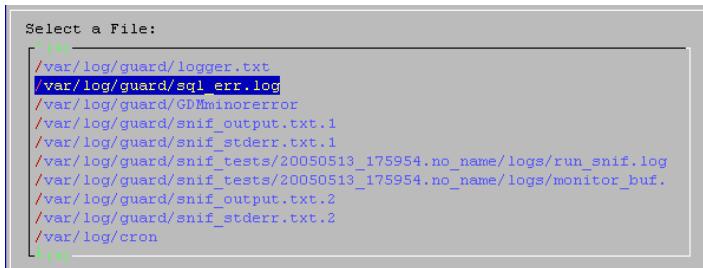
DIRECTORY SUMMARY IN MBYTES:
--- dir: /var/log/guard/jetspeed_log/* ---
0      /var/log/guard/jetspeed_log/README
0      /var/log/guard/jetspeed_log/access.log
1      /var/log/guard/jetspeed_log/jetspeed.log
--- dir: /var/log/guard/tomcat/* ---
--- dir: /usr/local/guardium/* ---
1      /usr/local/guardium/cli
1      /usr/local/guardium/diag
1      /usr/local/guardium/scripts
```

2. Use the Up and Down arrow keys to scroll through the directories.
3. Press Enter or click Exit when you are done.

3.4: File Summary

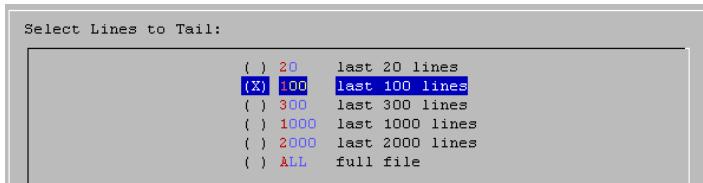
Use this command to list all or some portion of a log file.

1. Select File Summary from the Interactive Queries menu.
2. You are prompted to select a file:



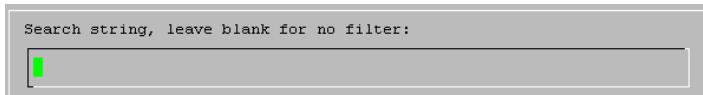
Use the Up and Down arrow keys to scroll the selection cursor to the file you want to view.

3. Press Enter or click OK.
4. You are prompted to select the number of lines to display:



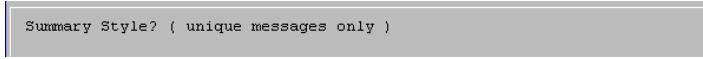
Make your selection and press Enter.

5. You are prompted to enter an optional search string:



Use this box if you are searching for a particular log message (you can enter a regular expression). Otherwise leave the box empty and press Enter.

6. The following prompt is displayed:



Press Enter to answer yes, meaning that only unique messages will be displayed. For example, if only one type of SQL error were contained in the sql_err.log file, the display might look like this:

```
^ (+)
ERROR FROM execSql # Column 'APP_EVENT_TYPE' cannot be null
****INSERT INTO GDM_APP_EVENT  (APP_USER_NAME, APP_EVENT_TYPE, EVENT_VA
```

Otherwise select No and press Enter (all messages will be displayed, as illustrated below):

```
^ (+)

****INSERT INTO GDM_APP_EVENT  (APP_USER_NAME, APP_EVENT_TYPE, EVENT_VA
ERROR FROM execSql 1 Column 'APP_EVENT_TYPE' cannot be null

****INSERT INTO GDM_APP_EVENT  (APP_USER_NAME, APP_EVENT_TYPE, EVENT_VA
ERROR FROM execSql 1 Column 'APP_EVENT_TYPE' cannot be null

****INSERT INTO GDM_APP_EVENT  (APP_USER_NAME, APP_EVENT_TYPE, EVENT_VA
ERROR FROM execSql 1 Column 'APP_EVENT_TYPE' cannot be null

****INSERT INTO GDM_APP_EVENT  (APP_USER_NAME, APP_EVENT_TYPE, EVENT_VA
ERROR FROM execSql 1 Column 'APP_EVENT_TYPE' cannot be null
L (+)  ( 80%)
```

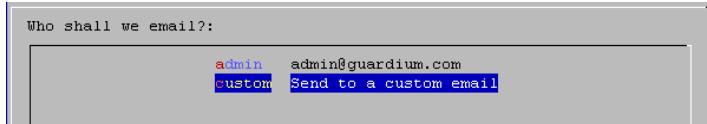
Be aware that when the Summary Style is used, variables are replaced by the pound sign character (#), as illustrated in the first example above. For some log data containing variables such as IP addresses or dates, the replacements can be extensive. For example:

```
^ (+)
May ## ##:#:#:# localhost.localdomain #####-##-## ##:#:#:#,###[QuartzWo
May ## ##:#:#:# localhost.localdomain Purge Object - MANAGEMENT_OPERATI
May ## ##:#:#:# localhost.localdomain Purge Object - SOFTWARE_TAP_EVENT
May ## ##:#:#:# localhost.localdomain Purge Object - com.guardium.audit
May ## ##:#:#:# localhost.localdomain Purge Object - com.guardium.datam
May ## ##:#:#:# localhost.localdomain Purge Object - com.guardium.datam
May ## ##:#:#:# supp# kernel: device eth# entered promiscuous mode
May ## ##:#:#:# supp# kernel: device eth# left promiscuous mode
```

3.5: Test Email

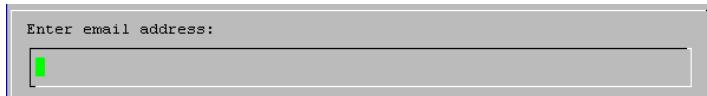
Use this command to send a test email using the configured SMTP server.

1. Select Test Email from the Interactive Queries menu.
2. You are prompted to select a recipient:



Select Custom and press Enter.

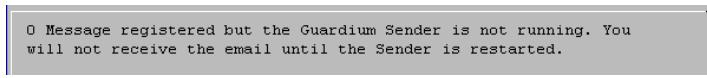
3. You are prompted to supply an email address:



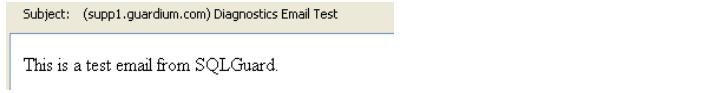
Type an email address and press Enter. If the message is sent successfully, you receive the following message:



If the mail server has not been started, you receive the following message:



Mail messages you receive from this command have the subject and message body illustrated below:



Note: On the Administration Console, the **Test Connection** link in the SMTP pane of the Alerter configuration panel only tests that an SMTP port is configured, not that mail can actually be delivered via that server. You can use this command to test email delivery without having to configure and trigger a statistical or real-time alert, or an audit process notification.

3.6: Test SNMP

Use this command to send a test SNMP trap to the configured SNMP server.

1. Select Test SNMP from the Interactive Queries menu.
2. You are informed of the activity:

```
Sending snmp trap to: 192.168.3.1, community: guardium
```

And the results:

```
Could not send snmp trap to 192.168.3.1
```

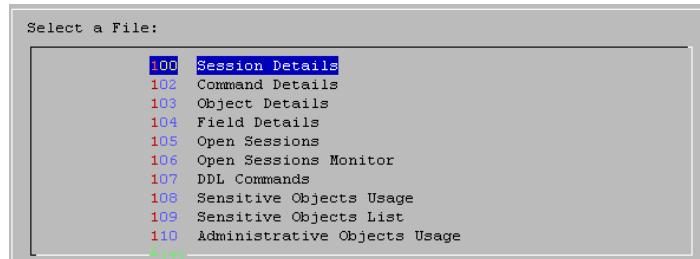
In the example above, no *guardium* trap community was defined on the specified SNMP host.

Note: On the Alerter Configuration panel, the **Test Connection** link in the SNMP pane only tests that an SNMP port is configured, not that a trap can actually be delivered via that server. You can use this command to test trap delivery without having to configure (and trigger) a statistical or real-time alert, or an audit process notification.

3.7: Report Query Data

Use this command to display the actual SELECT statement used for a report query. This might be useful if a user-written report is producing unexpected output.

1. Select Report Query Data from the Interactive Queries menu.
2. You are prompted to make a selection from a list of report titles:



Use the Up and Down arrow keys to select an entry and press the Enter key. Each entry in this list is a *Report* entity. All pre-defined reports are listed first. These are numbered in the range 100-225 (for version 3.6.1 – the numbers will most likely grow incrementally with each release, as more pre-defined reports are created).

User written reports are listed following the pre-defined reports, beginning with number 20001 (for version 3.6.1).

The selected report SELECT statement will be displayed. The select statement for *105 Open Sessions* is illustrated below:

```
^ (+)

*** SQL for report id: 105 - Open Sessions ***
Select GDM_ACCESS.CLIENT_IP, GDM_ACCESS.DB_USER, GDM_ACCESS.NET_PROTOCOL
```

3.8: GDM Queries

Use this command to display a count of observed SQL calls during a 100 second interval.

1. Select GDM Queries from the Interactive Queries menu.
2. A message displays requesting your patience:

```
This report takes 100 seconds to collect data. Press Enter on
Yes to continue. There will be no progress bar, please be
patient and wait 100 seconds for the results to be displayed in
the next screen.
```

Select yes to continue, which after a delay of about 100 seconds results in a display like the one illustrated below:

```
^ (+)
04/13/06 10:10:31
*****
CLIENT_IP      SERVER_IP      CMD_CT  SERVER_TYPE
192.168.1.18   192.168.2.100  221     DB2
192.168.1.18   192.168.2.20   28      MS SQL SERVER
192.168.1.18   192.168.2.22   11      SYBASE
192.168.1.18   192.168.2.247  23      INFORMIX
192.168.1.18   192.168.2.25   23      ORACLE
192.168.1.18   192.168.2.30   158    DB2
192.168.1.18   192.168.2.33   236    DB2
192.168.1.18   192.168.2.39   21      INFORMIX
192.168.1.18   192.168.2.40   213    DB2
192.168.1.18   192.168.2.45   10      ORACLE
192.168.1.18   192.168.2.48   230    MS SQL SERVER
192.168.1.18   192.168.2.70   6       MS SQL SERVER
192.168.1.18   192.168.2.72   185    DB2
L (+)          < EXIT >           ( 77%)
```

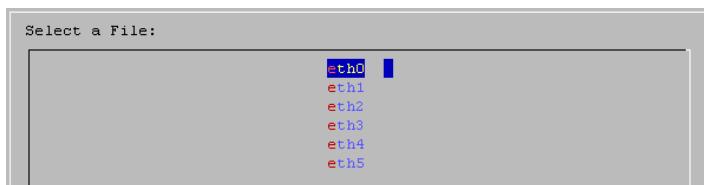
The CMD_CT column lists the number of observed SQL calls from the specified clients to the specified servers.

Press Enter when you are done viewing the report.

3.9: Generate TCP dump

Use this command to create a TCP dump. For this command, output is written to a command file only and not to the screen. Unlike most other commands, a separate file is created in the *current* directory for each execution of this command. The file name is in the format: **tcpdump_<mmyyyy-hhmss>**, where the variable portion is a date and time stamp: **mmyyyy** is the month and year, and **hhmmss** is the hours, minutes, and seconds.

1. Select *Generate TCP dump* from the Interactive Queries menu.
2. You are prompted to select an interface:



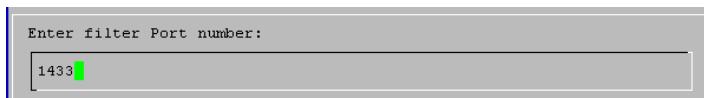
Select a port and press Enter.

3. You are prompted for an optional filter IP address:



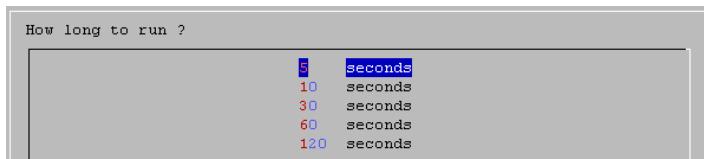
If you are interested in traffic from only a specific address, enter that IP address and press Enter. Otherwise, just press Enter.

4. You are prompted for an optional port number:



If you are interested in traffic from only a specific port, enter that port number and press Enter. Otherwise, just press Enter.

5. You are prompted to select how many seconds of traffic to capture:



Select a number of seconds and press Enter.

6. You are prompted to press Enter to start collecting data:

```
Press ENTER and wait 5 seconds for program to complete gathering  
TCP dumps.
```

Press Enter. You are returned to the menu after (approximately) the specified number of seconds.

7. To view the TCP dump data, select the Read TCP dumps command (see below) or export the file (see Export Reported Files on the Output Management menu, described previously).

3.10: Read TCP dumps

Use this command to display a TCP dump file created previously.

1. Select *Read TCP dumps* from the Interactive Queries menu.
2. You are prompted to select file. The TCP dump files are listed from oldest to newest. The file name is in the format: **tcpdump_<mmddyy-hhmmss>**, where the variable portion is a date and time stamp: **mmddyy** is the month, day, and year; and **hhmmss** is the hours, minutes, and seconds.

```
Select a TCP dump:  
/var/log/guard/diag/current/tcpdump_052005-073516  
/var/log/guard/diag/current/tcpdump_052005-074015  
/var/log/guard/diag/current/tcpdump_052005-074442  
/var/log/guard/diag/current/tcpdump_052005-074807  
/var/log/guard/diag/current/tcpdump_052005-075906
```

Select the file you want to view and press Enter.

3. The selected file displays:

```
^ (+)  
07:59:09.828518 supp1.guardium.com.ssh > 192.168.1.223.1114: P 18497745  
07:59:09.829672 192.168.1.223.1114 > supp1.guardium.com.ssh: P 1:53(52)  
07:59:09.861143 supp1.guardium.com.ssh > 192.168.1.223.1114: . ack 53 w  
07:59:10.177148 802.1d config 8003.00:0e:d7:98:07:40.80c8 root 8000.00:  
07:59:11.011529 supp1.guardium.com.32778 > 192.168.3.160.8075: udp 34 (07:59:11.012064 192.168.3.80.1083 > supp1.guardium.com.9500: P 25536622  
07:59:11.012080 supp1.guardium.com.9500 > 192.168.3.80.1083: . ack 213  
07:59:12.177051 802.1d config 8003.00:0e:d7:98:07:40.80c8 root 8000.00:  
07:59:14.176961 802.1d config 8003.00:0e:d7:98:07:40.80c8 root 8000.00:
```

Use the Up and Down arrow keys to scroll through the display and press Enter when you are done.

3.11: Watch Buffer

Use this command to watch activity in the SQL Guard buffers:

1. Select Watch Buffer from the Interactive Queries menu.

The display is updated every second:

```

Every 1s: head -13 /var/log/guard/snif_buf_usage.txt      Fri May 20 08:46:00 2005

100
204800 buffers out of 204800
89 connection used, 32679 unused, 0 dropped (sniffer), 0 ignored (analyzer)
127513478722 bytes lost, 23996 connections ended, 197708779576 bytes sent, 10338
3077 request sent
Dropped Packets: 2494650 buffer full, 0 too short , 0 ignored
time now is 1116593159
Analyzer/Parser buffers size: 0 (2156557) 0 (2213594) 0 (0) 0 (0) 0 (0) 0 (0)

ms-tsql-logger 0 (4090327)
syb-tsql-logger 0 (0)
ora-tsql-logger 0 (0)
db2-sql-logger 0 (0)
Client 192.168.1.18(44031) Server 192.168.2.21(1433) dlvr=1268952,2473719 vrdct=
1268952,2473719 goodttl=0 lastact=1116593102 block=0 dirty=0 term=0 termrep=0 ig
nr=0 pscan=0

```

2. Press Ctrl-C to close the display.

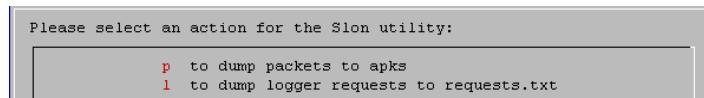
3.12: Slon Utility

Use this command to run the *slon* utility, which tracks packets. For this command, output is not written to the screen. Output is written to one of two command files in the *current* directory, for each execution of the command:

- apks.txt.<day_dd-mmm-yyyy_hh.mm.ss.ttt>
- OR
- requests.txt.<day_dd-mmm-yyyy_hh.mm.ss.ttt>

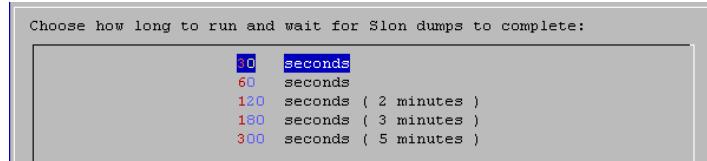
The variable portions or the file names are date and time stamps. For example, *apks.txt.Fri_20-May-2005_08.52.00.789*.

1. Select Slon Utility from the Interactive Queries menu. You will receive the following prompt:

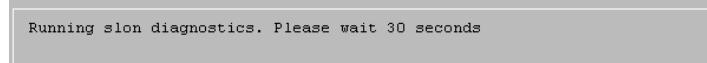


2. Select either option and click OK.

3. Regardless of your selection, you will be prompted to select the time period for the activity:

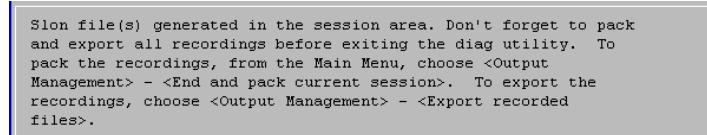


4. Select a time period and press Enter. You are notified that the program will run for one minute and prompted to press Enter:

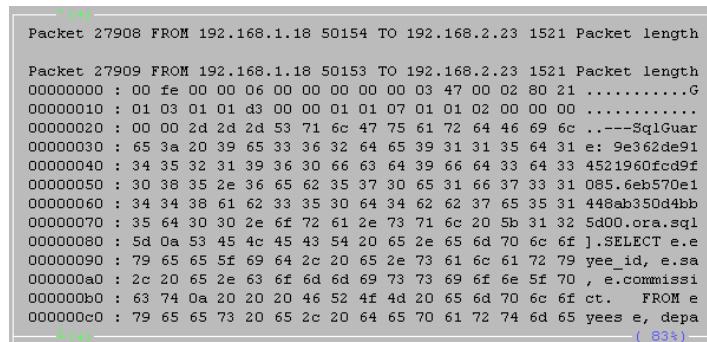


Press Enter and wait.

5. When processing completes, a message will be displayed:



You can use the File Summary command (described previously) to display the output of this command. The *apks* file output is shown below:

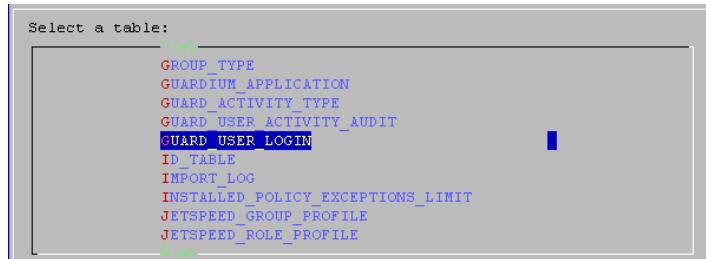


But because this command can produce a large amount of data, you will probably want to export the file to another system, where you can view the contents using a text editor. (Pack the current session data, and export the recordings as described earlier in this section.)

3.13: Show Indexes

Use this command to show indexes for various internal tables:

1. Select Show Indexes from the Interactive Queries menu.
2. You are prompted to select a table:



Select a table and press Enter to display the indexes for that table:

```
*****+
Indexes for table: GUARD_USER_LOGIN
*****
Index Name: PRIMARY
Seq in Index: 1
Column Name: LOGIN_ID
Cardinality: 25

Index Name: GUARD_USER_LOGIN_IDX2
Seq in Index: 1
Column Name: USER_NAME
Cardinality: 1

Index Name: GUARD_USER_LOGIN_IDX2
Seq in Index: 2
Column Name: LOGIN_DATETIME
L (+) < EXIT > ( 75%)
```

3. Use the Up and Down arrow keys to scroll through the display. Press Enter when you are done.

3.14: STAP Check

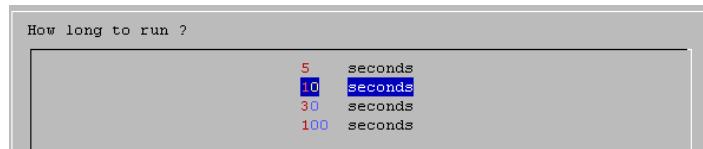
Use this command to display S-Tap definitions and traffic information:

1. Select STAP Check from the Interactive Queries menu.
2. The system's unit type displays in numeric format:

```
This system's unit type is 548
```

Press Enter.

3. You are prompted to select the number of seconds to monitor the S-Tap traffic:



Use the Up and Down arrow keys to make a selection and press Enter.

4. You are informed of approximately how long to wait for output, and prompted to press Enter:

Press ENTER and wait 10 seconds for program to complete gathering STAP information.

Press Enter.

5. The S-Tap Definitions and Server Traffic reports display:

^ (+)
STAP DEFINITIONS:
pkts bytes target prot opt in out source des
1345K 54M ACCEPT tcp -- * * 0.0.0.0/0 0.0
1726K 1335M ACCEPT tcp -- * * 0.0.0.0/0 0.0
42M 1868M ACCEPT tcp -- * * 0.0.0.0/0 0.0
57M 22G ACCEPT tcp -- * * 0.0.0.0/0 0.0

STAP SERVER TRAFFIC:
Server Name/IP Packets by Port

192.168.3.160: 9500=>0 16016=>0

Press Enter when you are done viewing the report.

3.15: Interface link status

Use this command to display interface link status.

1. Select *Interface link status* from the Interactive Queries menu.
2. The status of all interfaces displays:

^ (+)
Interface Status:
eth0 Speed: 100Mb/s Link detected: yes
eth1 Speed: 100Mb/s Link detected: yes
eth2 Speed: Unknown! (65535) Link detected: no
eth3 Speed: Unknown! (65535) Link detected: no
eth4 Speed: Unknown! (65535) Link detected: no

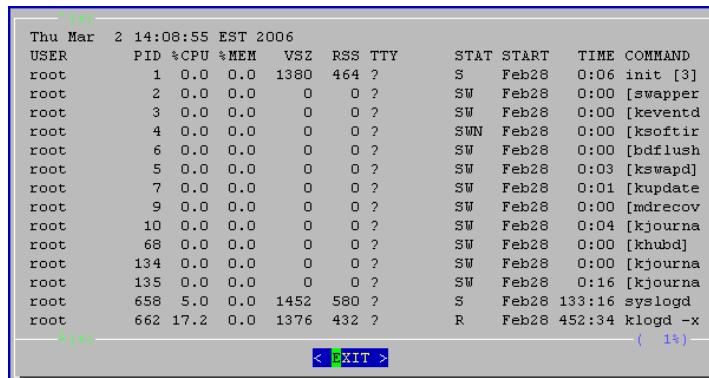
3. Use the Up and Down arrows to scroll through the display. Press Enter when you are done.

Note: This command displays the link status only. To display interface configuration information, use the *show network interface all* CLI command.

3.16: Run Process List PS

Use this command to display the results of a ps command.

1. Select Run Process List PS from the Interactive Queries menu.
2. The complete list of processes displays:



```

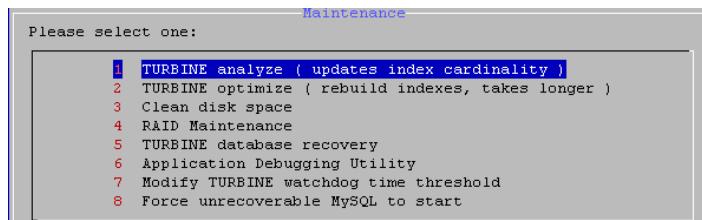
^ (4)
Thu Mar  2 14:08:55 EST 2006
USER      PID %CPU %MEM    VSZ RSS TTY      STAT START   TIME COMMAND
root        1  0.0  0.0  1380  464 ?          S   Feb28  0:06 init [3]
root        2  0.0  0.0     0   0 ?          SW  Feb28  0:00 [swapper]
root        3  0.0  0.0     0   0 ?          SW  Feb28  0:00 [keventd]
root        4  0.0  0.0     0   0 ?          SWN Feb28  0:00 [ksoftir]
root        6  0.0  0.0     0   0 ?          SW  Feb28  0:00 [bdflush]
root        5  0.0  0.0     0   0 ?          SW  Feb28  0:03 [kswapd]
root        7  0.0  0.0     0   0 ?          SW  Feb28  0:01 [kupdatate]
root        9  0.0  0.0     0   0 ?          SW  Feb28  0:00 [mdrecov]
root       10  0.0  0.0     0   0 ?          SW  Feb28  0:04 [kjournal]
root       68  0.0  0.0     0   0 ?          SW  Feb28  0:00 [khubd]
root      134  0.0  0.0     0   0 ?          SW  Feb28  0:00 [kjournal]
root      135  0.0  0.0     0   0 ?          SW  Feb28  0:16 [kjournal]
root      658  5.0  0.0  1452  580 ?          S   Feb28 133:16 syslogd
root     662 17.2  0.0  1376  432 ?          R   Feb28 452:34 klogd -x
L (4)                                     (  1%)
< EXIT >

```

3. Use the Up and Down arrows to scroll through the display. Press Enter when you are done.

4: Perform Maintenance Actions

Select the Perform Maintenance Actions option from the Main Menu to open the Maintenance menu (illustrated below). Use these commands only under the direction of Guardium Support. These *do not* need to be run on a regular basis.



Each Maintenance menu command is described below.

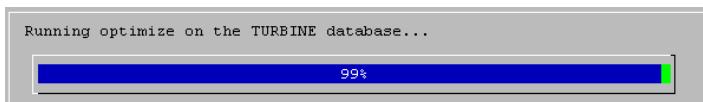
4.1: TURBINE analyze (updates index cardinality)

Use this command to optimize index cardinality on SQL Guard's internal database. A progress bar displays while the operation is running. When the operation completes, you are returned to the Maintenance menu.

4.2: TURBINE optimize (rebuild indexes, takes longer)

Use this command to analyze and re-index SQL Guard's internal database.

1. Select *TURBINE optimize (index cardinality)* from the Maintenance menu. A progress bar displays while the operation is running:

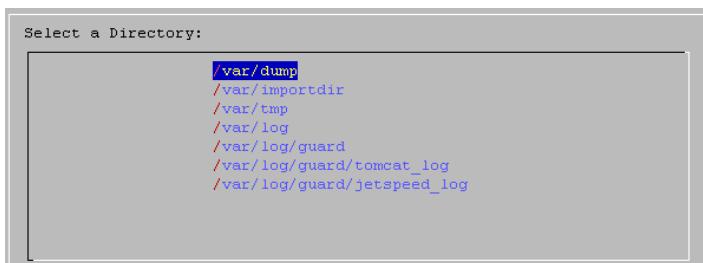


When the operation completes, you are returned to the Maintenance menu.

4.3: Clean disk space

Use this command to clean unused disk space. You are returned to the Maintenance menu when the procedure completes.

1. Select Clean disk space from the Maintenance menu. You will be prompted to select a directory:



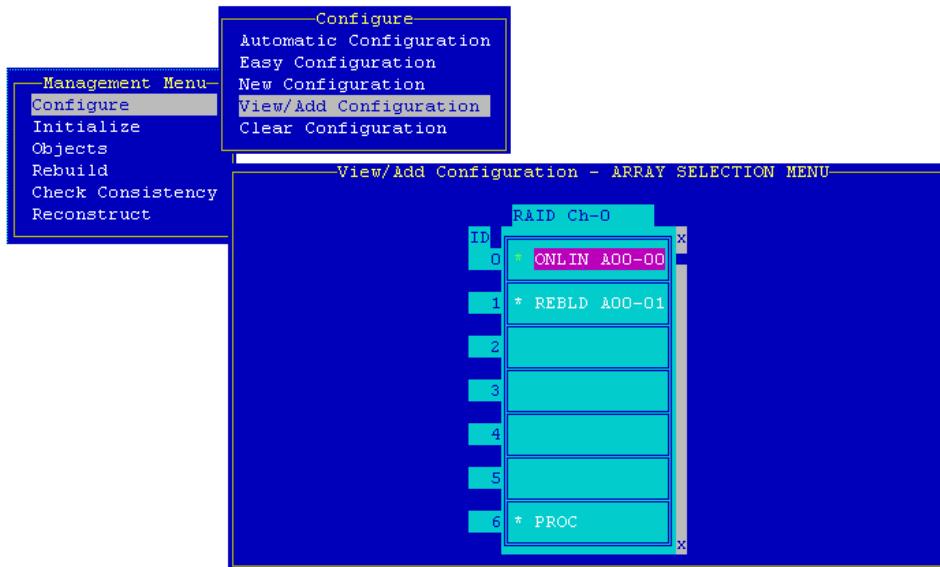
2. Select the directory from which you want to remove files. The contents of the directory will be listed, and you will be prompted to confirm that you want to remove all files.

When the operation completes, you are returned to the Maintenance menu.

4.4: RAID Maintenance

Use this command only under the direction of Guardium Support. This command provides access to the Management Menu of the RAID controller utility program, which can be used to display the status of the RAID drives, as illustrated below. If your system

does *not* have a RAID controller, an error message displays if you select this command. You must be extremely careful when using the RAID controller utility program, since several of the functions provided will erase all information on the disk.



4.5: Transfer Backups & System Recovery

Use this command to restore a backed up version of the internal database. You will be prompted to confirm the operation:

```
This procedure will shutdown all services and replace the  
database from the backup file. Are you sure you wish to perform  
this restore operation?
```

4.6: Application Debugging Utility

Use this command to turn debugging on or off. You are prompted to choose one of the following:

```
Please select an action for the Logging utility:  
true to enable  
false to disable  
reset disables logging, resets to defaults
```

4.7: Modify TURBINE watchdog threshold

```
Enter the new timeout threshold in seconds for the TURBINE sql  
session watchdog:
```

[]

Use this option to change the timeout limit for long queries.

4.8: Force unrecoverable MySQL to start

```
Would you like to start the GUI and sniffers in this highly  
degraded mode ?
```

`false` `default`
`true`

Use this option only when directed to do so by Guardium Support.

5: Exit to CLI

Select Exit to CLI on the Main Menu. Press Enter to close the `diag` command and return to the command line interface.

System Static Reports Overview

The following subtopics provide an outline of the major components of the System Static Reports output. The fragments of output shown are intended to illustrate the type and level of information contained in the report, rather than provide a detailed description of the actual contents (that is beyond the scope of this document).

System Configuration Information

The top portion of the System Static Reports output describes the build version, the patches applied, the current system uptime, and name server information:

```
Build version: 34e1eb12eb68ba76cb49028251c9a0d6 /usr/local/guardium/etc/cvstag  
  
Patches:  
2006/02/22 16:16:50: START Installation of 'Update 5.0'  
2006/02/22 16:18:04: Installation Done - Successfully Installed  
  
< lines deleted... >  
  
Current uptime:  
09:03:43 up 6 days, 17:34, 1 user, load average: 0.44, 0.50, 0.41  
  
System nameservers:  
192.168.3.20  
  
DB nameservers:  
192.168.3.20  
  
Gateway: 192.168.3.1 (system) 192.168.3.1 (def)
```

Next, the file system information displays (shown partially below):

```
Filesystem      Size  Used Avail Use% Mounted on
/dev/hdc3       2.0G  1.1G  813M  58% /
/dev/hdc1       97M   9.2M   83M  10% /boot
none           504M    0   504M  0% /dev/shm
/dev/hdc2       71G   1.2G   66G  2% /var

          total:    used:   free: shared: buffers:  cached:
Mem:  1055199232 1041711104 13488128          0 63275008 186220544
Swap: 536698880 295432192 241266688
MemTotal: 1030468 kB
MemFree: 13172 kB

< lines deleted... >
```

This is followed by information about the mail and SNMP servers configured:

```
SMTP server: 192.168.1.7 on port 25 : REACHABLE
SMTP user: undef
SMTP password: undef
SMTP auth: NONE

SNMP trapsink: undef UNREACHABLE
SNMP trap community: undef
SNMP read community: undef
```

The final section of the system configuration section describes the network configuration for the unit: IP address, host and domain names, etc:

```
eth0:          192.168.3.101 (system) 192.168.3.101 (def)
hostname:      (system) g1 (def)
domain:        (system) guardium.com (def)
mac address: 00:04:23:A7:77:F2 (MAC1) 00:04:23:A7:77:F2 (MAC2)
unit type: 548 Standalone STap
```

Internal Database Information

The next major section of the System Static Reports output contains information about the internal database status and threads (only the first few threads are shown):

```
uptime 77097  seconds.
27  threads.
78545028  queries.

+-----+-----+-----+-----+
| id | user   | host      | db    | command | time | state |
+-----+-----+-----+-----+
| 1137 | enchantedg | localhost | TURBINE | Sleep   | 26   |
| 1257 | enchantedg | localhost.localdomain:33587 | TURBINE | Sleep   | 0    |
| 1258 | enchantedg | localhost.localdomain:60409 | TURBINE | Sleep   | 7716 |
| 1259 | enchantedg | localhost.localdomain:48233 | TURBINE | Sleep   | 322  |

< lines deleted... >
```

Note: Many major sections of the System Static Reports output are separated by a line of equal sign characters (=), as illustrated above.

The list of threads is followed by an analysis of table status.

Web Servlet Container Information

The next several sections of the System Static Reports output contain information about the Web servlet container environment (Tomcat):

```
=====
Currently defined Tomcat port is 8443.

The TOMCAT daemon is running and listening on port(s): 8005 8443.

Currently OPEN ports
java run by tomcat on port *:8443

< lines deleted... >

=====

These are the nanny latest actions:
May 19 14:13:09 guard nanny:[5528]: Also checking tomcat.
May 19 14:13:09 guard nanny:[5528]: Going for my initial nap.

< lines deleted... >

This is the TOMCAT command line:
  463 sh -c ps -o pid,cmd -e | grep Dcatalina.base
21917 grep Dcatalina.base.
```

Inspection Engine Information

The next major section of the System Static Reports output contains information about the inspection engine:

```
=====
This is the SNIF (pid: 13036) command line: 13036 /usr/local/guardium/bin/snif.
This is the SNIF status:
Name:           snif
State:          R (running)
Tgid:           13036

< lines deleted... >

=====
Current timestamp is 2005-05-20 11:56:41
This is the last timestamp at GDM_CONSTRUCT_INSTANCE: 2005-05-20 11:56:41
This is the last timestamp at GDM_EXCEPTION: 2005-05-20 11:56:41
This is the last timestamp at GDM_POLICY_VIOLATIONS_LOG: 2005-05-20 11:56:41

=====
Snif buf usage at Fri May 20 11:56:44 2005:
100 204800 buffers out of 204800
126 connection used, 32642 unused, 0 dropped (sniffer), 9 ignored (analyzer)
0 bytes lost, 60 connections ended, 601752099 bytes sent, 579063 request sent
Dropped Packets: 0 buffer full, 0 too short , 451 ignored
time now is 1116604603
Analyzer/Parser buffers size: 6 (66533) 0 (62902)

ms-tsql-logger 0 (11331)
syb-tsql-logger 0 (70)
ora-tsql-logger 79 (67803)
db2-sql-logger 0 (20544)

< lines deleted... >
```

IP Tables Information

The next major section contains information about the IP tables:

```
=====
IPTABLES:
-----
      tcp  --  192.168.2.0/24      192.168.1.0/24      tcp spts:1521:60000  set 0x23
      tcp  --  192.168.1.0/24      192.168.2.0/24      tcp dpts:1521:60000  set 0x22

< lines deleted... >
```

S-Tap Information

The next major section contains S-Tap information:

```
=====
STAP:
-----
      0      0 ACCEPT    tcp  --  *      *  0.0.0.0/0      0.0.0.0/0      tcp spt:9500
      0      0 ACCEPT    tcp  --  *      *  0.0.0.0/0      0.0.0.0/0      tcp dpt:9500
  2696  148K ACCEPT    tcp  --  *      *  0.0.0.0/0      0.0.0.0/0      tcp spt:16016
  2835  175K ACCEPT    tcp  --  *      *  0.0.0.0/0      0.0.0.0/0      tcp dpt:16016

< lines deleted... >
```

IP Traffic Information

The next major section contains IP traffic information:

```
IP traffic statistics.
OUTPUT OF ETH0
Fri May 20 11:57:04 2005; ***** Detailed interface statistics started *****

*** Detailed statistics for interface eth0, generated Fri May 20 11:58:04 2005

< lines deleted... >

OUTPUT OF ETH1
Fri May 20 11:57:04 2005; ***** Detailed interface statistics started *****

*** Detailed statistics for interface eth1, generated Fri May 20 11:58:04 2005

Total:          82440 packets, 53892382 bytes
            (incoming: 82440 packets, 53892382 bytes; outgoing: 0 packets, 0 bytes)
IP:  82440 packets, 52632747 bytes
            (incoming: 82440 packets, 52632747 bytes; outgoing: 0 packets, 0 bytes)

< lines deleted... >
```

Inspection Engine STDERR and STDOUT Information

The next section contains the last messages output by the sniffer:

```
Snif  STDERR:  
< lines deleted... >  
  
Snif  STDOUT:  
Fri_20-May-2005_04:04:35 : Guardium Engine Monitor starting  
Fri_20-May-2005_04:14:37 : Guardium Engine Monitor starting  
Fri_20-May-2005_04:24:38 : Guardium Engine Monitor starting  
  
< lines deleted... >
```

Import Directory Information

The next section lists the import directory contents:

```
These are the contents of the importdir directory:  
total 0
```

Aggregator Activity Information

The last section lists aggregator activities (there are none in the example):

```
=====  
This is the aggregator last activities:
```

Appendix A: Time Zone List

The table below contains accepted time zone values for the CLI.

Timezone	Description (optional)
Africa/Abidjan	
Africa/Accra	
Africa/Addis_Ababa	
Africa/Algiers	
Africa/Asmera	
Africa/Bamako	Southwest Mali
Africa/Bangui	
Africa/Banjul	
Africa/Bissau	
Africa/Blantyre	
Africa/Brazzaville	
Africa/Bujumbura	
Africa/Cairo	
Africa/Casablanca	
Africa/Ceuta	Ceuta and Melilla
Africa/Conakry	
Africa/Dakar	
Africa/Dar_es_Salaam	
Africa/Djibouti	
Africa/Douala	
Africa/EI_Aaiun	
Africa/Freetown	
Africa/Gaborone	
Africa/Harare	
Africa/Johannesburg	
Africa/Kampala	
Africa/Khartoum	
Africa/Kigali	

Timezone	Description (optional)
Africa/Kinshasa	Western Democratic Republic of Congo
Africa/Lagos	
Africa/Libreville	
Africa/Lome	
Africa/Luanda	
Africa/Lubumbashi	Eastern Democratic Republic of Congo
Africa/Lusaka	
Africa/Malabo	
Africa/Maputo	
Africa/Maseru	
Africa/Mbabane	
Africa/Mogadishu	
Africa/Monrovia	
Africa/Nairobi	
Africa/Ndjamena	
Africa/Niamey	
Africa/Nouakchott	
Africa/Ouagadougou	
Africa/Porto-Novo	
Africa/Sao_Tome	
Africa/Timbuktu	Northeast Mali
Africa/Tripoli	
Africa/Tunis	
Africa/Windhoek	

Timezone	Description (optional)
America/Adak	Aleutian Islands
America/Anchorage	Alaska Time
America/Anguilla	
America/Antigua	
America/Araguaina	Tocantins
America/Aruba	
America/Asuncion	
America/Barbados	
America/Belem	Amapa, E Para
America/Belize	
America/Boa_Vista	Roraima
America/Bogota	
America/Boise	Mountain Time - South Idaho & east Oregon
America/Buenos_Aires	East Argentina (BA, DF, SC, TF)
America/Cambridge_Bay	Central Time – west Nunavut
America/Cancun	Central Time – Quintana Roo
America/Caracas	
America/Catamarca	Catamarca (CT)
America/Cayenne	
America/Cayman	
America/Chicago	Central Time
America/Chihuahua	Mountain Time – Chihuahua
America/Cordoba	Most locations (CB, CC, CH, CN, ER, FM, LP, LR, MN, NQ, RN, SA, SE, SF, SJ, SL, TM)
America/Costa_Rica	

Timezone	Description (optional)
America/Cuiaba	Mato Grosso, Mato Grosso do Sul
America/Curacao	
America/Danmarkshavn	east coast, north of Scoresbysund
America/Dawson	Pacific Time – north Yukon
America/Dawson_Creek	Mountain standard time – Dawson Creek and Fort Saint John, British Columbia
America/Denver	Mountain Time
America/Detroit	Eastern time – Michigan – most locations
America/Dominica	
America/Edmonton	Mountain time - Alberta, east British Columbia & west Saskatchewan
America/Eirunepe	West Amazonas
America/El_Salvador	
America/Fortaleza	NE Brazil (MA, PI, CE, RN, PR)
America/Glace_Bay	Atlantic Time – Nova Scotia – places that did not observe DST 1966-1971
America/Godthab	Most locations
America/Goose_Bay	Atlantic Time – East Labrador
America/Grand_Turk	
America/Grenada	
America/Guadeloupe	
America/Guatemala	

Timezone	Description (optional)
America/Guayaquil	Mainland
America/Guyana	
America/Halifax	Atlantic Time – Nova Scotia (most places), NB, W Labrador, East Quebec and PEI
America/Havana	
America/Hermosillo	Mountain standard time – Sonora
America/Indiana/Knox	Eastern standard time – Indiana – Starke County
America/Indiana/Marengo	Eastern standard time – Indiana – Crawford County
America/Indiana/Vevay	Eastern standard time – Indiana – Switzerland County
America/Indianapolis	Eastern standard time – Indiana – most locations
America/Inuvik	Mountain time – west Northwest Territories
America/Iqaluit	Eastern standard time – east Nunavut
America/Jamaica	
America/Jujuy	Jujuy (JY)
America/Juneau	Alaska time – Alaska panhandle
America/Kentucky/Monticello	Eastern time – Kentucky – Wayne County
America/La_Paz	
America/Lima	

Timezone	Description (optional)
America/Los_Angeles	Pacific Time
America/Louisville	Eastern time – Kentucky – Louisville area
America/Maceio	Alagoas, Sergipe
America/Managua	
America/Manaus	East Amazonas
America/Martinique	
America/Mazatlan	Mountain time – S Baja, Nayarit, Sinaloa
America/Mendoza	Mendoza (MZ)
America/Menominee	Central time – Michigan – Wisconsin border
America/Merida	Central Time - Campeche, Yucatan
America/Mexico_City	Central time - most locations
America/Miquelon	
America/Monterrey	Central time – Coahuila, Durango, Nuevo Leon, Tamaulipas
America/Montevideo	
America/Montreal	Eastern time – Ontario and Quebec – most locations
America/Montserrat	
America/Nassau	
America/New_York	Eastern time
America/Nipigon	Eastern time – Ontario and Quebec – places that did not observe DST 1967-1973

Timezone	Description (optional)
America/Nome	Alaska time – west Alaska
America/Noronha	Atlantic islands
America/North_Dakota/Center	Central time – North Dakota - Oliver County
America/Panama	
America/Pangnirtung	Eastern standard time – Pangnirtung, Nunavut
America/Paramaribo	
America/Phoenix	Mountain standard time – Arizona
America/Port-au-Prince	
America/Port_of_Spain	
America/Porto_Velho	West Para, Rondonia
America/Puerto_Rico	
America/Rainy_River	Central time – Rainy River and Fort Frances, Ontario
America/Rankin_Inlet	Eastern standard time – central Nunavut
America/Recife	Pernambuco
America/Regina	Central standard time – Saskatchewan – most locations
America/Rio_Branco	Acre
America/Santiago	Most locations
America/Santo_Domingo	
America/Sao_Paulo	South and Southeast Brazil (BA, GO, DF, MG, ES, RJ, SP, PR, SC, RS)

Timezone	Description (optional)
America/Scoresbysund	Scoresbysund / Ittoqqortoormiit
America/Shiprock	Mountain time – Navajo
America/St_Johns	Newfoundland Island
America/St_Kitts	
America/St_Lucia	
America/St_Thomas	
America/St_Vincent	
America/Swift_Current	Central standard time – Saskatchewan – midwest
America/Tegucigalpa	
America/Thule	Thule / Pituffik
America/Thunder_Bay	Eastern time – Thunder Bay, Ontario
America/Tijuana	Pacific time
America/Tortola	
America/Vancouver	Pacific time – west British Columbia
America/Whitehorse	Pacific time – south Yukon
America/Winnipeg	Central time – Manitoba & west Ontario
America/Yakutat	Alaska time – Alaska panhandle neck
America/Yellowknife	Mountain time – central Northwest Territories
Antarctica/Casey	Casey Station, Bailey Peninsula
Antarctica/Davis	Davis Station, Vestfold Hills

Timezone	Description (optional)
Antarctica/DumontDUrville	Dumont-d'Urville Base, Terre Adelie
Antarctica/Mawson	Mawson Station, Holme Bay
Antarctica/Mcmurdo	McMurdo Station, Ross Island
Antarctica/Palmer	Palmer Station, Anvers Island
Antarctica/South_Pole	Amundsen-Scott Station, South Pole
Antarctica/Syowa	Syowa Station, East Ongul I
Antarctica/Vostok	Vostok Station, South Magnetic Pole
Arctic/Longyearbyen	Svalbard
Asia/Aden	
Asia/Almaty	Most locations
Asia/Amman	
Asia/Anadyr	Moscow+10 - Bering Sea
Asia/Aqttau	Atyrau (Atirau, Gur'yev), Mangghystau (Mankistau)
Asia/Aqtobe	Aqtobe (Aktobe)
Asia/Ashgabat	
Asia/Baghdad	
Asia/Bahrain	
Asia/Baku	
Asia/Bangkok	
Asia/Beirut	
Asia/Bishkek	
Asia/Brunei	
Asia/Calcutta	

Timezone	Description (optional)
Asia/Choibalsan	Dornod, Sukhbaatar
Asia/Chongqing	Central China – Gansu, Guizhou, Sichuan, Yunnan, etc.
Asia/Colombo	
Asia/Damascus	
Asia/Dhaka	
Asia/Dili	
Asia/Dubai	
Asia/Dushanbe	
Asia/Gaza	
Asia/Harbin	Heilongjiang
Asia/Hong_Kong	
Asia/Hovd	Bayan-Olgii, Govi-Altai, Hovd, Uvs, Zavkhan
Asia/Irkutsk	Moscow+05 - Lake Baikal
Asia/Jakarta	Java and Sumatra
Asia/Jayapura	Irian Jaya and the Moluccas
Asia/Jerusalem	
Asia/Kabul	
Asia/Kamchatka	Moscow+09 – Kamchatka
Asia/Karachi	
Asia/Kashgar	southwest Xinjiang Uyghur
Asia/Katmandu	
Asia/Krasnoyarsk	Moscow+04 – Yenisei River
Asia/Kuala_Lumpur	peninsular Malaysia
Asia/Kuching	Sabah and Sarawak

Timezone	Description (optional)
Asia/Kuwait	
Asia/Macau	
Asia/Magadan	Moscow+08 – Magadan
Asia/Makassar	East and south Borneo, Celebes, Bali, Nusa Tengarra, west Timor
Asia/Manila	
Asia/Muscat	
Asia/Nicosia	
Asia/Novosibirsk	Moscow+03 – Novosibirsk
Asia/Omsk	Moscow+03 – west Siberia
Asia/Oral	West Kazakhstan
Asia/Phnom_Penh	
Asia/Pontianak	West and central Borneo
Asia/Pyongyang	
Asia/Qatar	
Asia/Qyzylorda	Qyzylorda (Kyzylorda, Kzyl-Orda)
Asia/Rangoon	
Asia/Riyadh	
Asia/Saigon	
Asia/Sakhalin	Moscow+07 – Sakhalin Island
Asia/Samarkand	West Uzbekistan
Asia/Seoul	
Asia/Shanghai	East China – Beijing, Guangdong, Shanghai, etc.
Asia/Singapore	

Timezone	Description (optional)
Asia/Taipei	
Asia/Tashkent	East Uzbekistan
Asia/Tbilisi	
Asia/Tehran	
Asia/Thimphu	
Asia/Tokyo	
Asia/Ulaanbaatar	Most locations
Asia/Urumqi	Tibet and most of Xinjiang Uyghur
Asia/Vientiane	
Asia/Vladivostok	Moscow+07 – Amur River
Asia/Yakutsk	Moscow+06 – Lena River
Asia/Yekaterinburg	Moscow+02 – Urals
Asia/Yerevan	
Atlantic/Azores	Azores
Atlantic/Bermuda	
Atlantic/Canary	Canary Islands
Atlantic/Cape_Verde	
Atlantic/Faeroe	
Atlantic/Jan_Mayen	Jan Mayen
Atlantic/Madeira	Madeira Islands
Atlantic/Reykjavik	
Atlantic/South_Georgia	
Atlantic/St_Helena	
Atlantic/Stanley	
Australia/Adelaide	South Australia
Australia/Brisbane	Queensland – most locations
Australia/Broken_Hill	New South Wales – Yancowinna

Timezone	Description (optional)
Australia/Darwin	Northern Territory
Australia/Hobart	Tasmania
Australia/Lindeman	Queensland - Holiday Islands
Australia/Lord_Howe	Lord Howe Island
Australia/Melbourne	Victoria
Australia/Perth	Western Australia
Australia/Sydney	New South Wales – most locations
Europe/Amsterdam	
Europe/Andorra	
Europe/Athens	
Europe/Belfast	Northern Ireland
Europe/Belgrade	
Europe/Berlin	
Europe/Bratislava	
Europe/Brussels	
Europe/Bucharest	
Europe/Budapest	
Europe/Chisinau	
Europe/Copenhagen	
Europe/Dublin	
Europe/Gibraltar	
Europe/Helsinki	
Europe/Istanbul	
Europe/Kaliningrad	Moscow-01 – Kaliningrad
Europe/Kiev	Most locations
Europe/Lisbon	Mainland
Europe/Ljubljana	
Europe/London	Great Britain

Timezone	Description (optional)
Europe/Luxembourg	
Europe/Madrid	Mainland
Europe/Malta	
Europe/Minsk	
Europe/Monaco	
Europe/Moscow	Moscow+00 – west Russia
Europe/Oslo	
Europe/Paris	
Europe/Prague	
Europe/Riga	
Europe/Rome	
Europe/Samara	Moscow+01 – Caspian Sea
Europe/San_Marino	
Europe/Sarajevo	
Europe/Simferopol	Central Crimea
Europe/Skopje	
Europe/Sofia	
Europe/Stockholm	
Europe/Tallinn	
Europe/Tirane	
Europe/Uzhgorod	Ruthenia
Europe/Vaduz	
Europe/Vatican	
Europe/Vienna	
Europe/Vilnius	
Europe/Warsaw	
Europe/Zagreb	
Europe/Zaporozhye	Zaporozh'ye, E Lugansk
Europe/Zurich	
Indian/Antananarivo	
Indian/Chagos	

Timezone	Description (optional)
Indian/Christmas	
Indian/Cocos	
Indian/Comoro	
Indian/Kerguelen	
Indian/Mahe	
Indian/Maldives	
Indian/Mauritius	
Indian/Mayotte	
Indian/Reunion	
Pacific/Apia	
Pacific/Auckland	Most locations
Pacific/Chatham	Chatham Islands
Pacific/Easter	Easter Island and Sala y Gomez
Pacific/Efate	
Pacific/Enderbury	Phoenix Islands
Pacific/Fakaofa	
Pacific/Fiji	
Pacific/Funafuti	
Pacific/Galapagos	Galapagos Islands
Pacific/Gambier	Gambier Islands
Pacific/Guadalcanal	
Pacific/Guam	
Pacific/Honolulu	Hawaii
Pacific/Johnston	Johnston Atoll
Pacific/Kiritimati	Line Islands

Timezone	Description (optional)
Pacific/Kosrae	Kosrae
Pacific/Kwajalein	Kwajalein
Pacific/Majuro	Most locations
Pacific/Marquesas	Marquesas Islands
Pacific/Midway	Midway Islands
Pacific/Nauru	
Pacific/Niue	
Pacific/Norfolk	
Pacific/Noumea	
Pacific/Pago_Pago	
Pacific/Palau	
Pacific/Pitcairn	
Pacific/Ponape	Ponape (Pohnpei)
Pacific/Port_Moresby	
Pacific/Rarotonga	
Pacific/Saipan	
Pacific/Tahiti	Society Islands
Pacific/Tarawa	Gilbert Islands
Pacific/Tongatapu	
Pacific/Truk	Truk (Chuuk)
Pacific/Wake	Wake Island
Pacific/Wallis	
Pacific/Yap	Yap

Appendix B: Reinstalling SQL Guard

The following section provides you with a step-by-step procedure for reinstalling the SQL Guard software. A new SQL Guard system is pre-loaded with the software necessary for its functioning. An administrator needs only to configure settings and physically install the SQL Guard system to enable its use. The procedure below should only be performed if circumstances require SQL Guard software to be reinstalled on the SQL Guard system; this procedure should *not* be used as part of the standard initial installation operation.

Note: Reinstalling the SQL Guard software reformats the unit's disk storage. Any data on that unit will be lost.

The summary steps below describe the process to reload a system from its installation CD. Be sure you have the correct CD before starting this procedure.

The CD-based installation steps do not require that a PC display and keyboard be attached to the SQL Guard unit. However, to monitor the process you should attach a display. Installation steps to connect a PC display and keyboard are described in *Chapter 1: Installation*.

Reinstalling SQL Guard Software

For security purposes, disconnect the SQL Guard system from the network while performing this procedure by disconnecting all Ethernet cables.

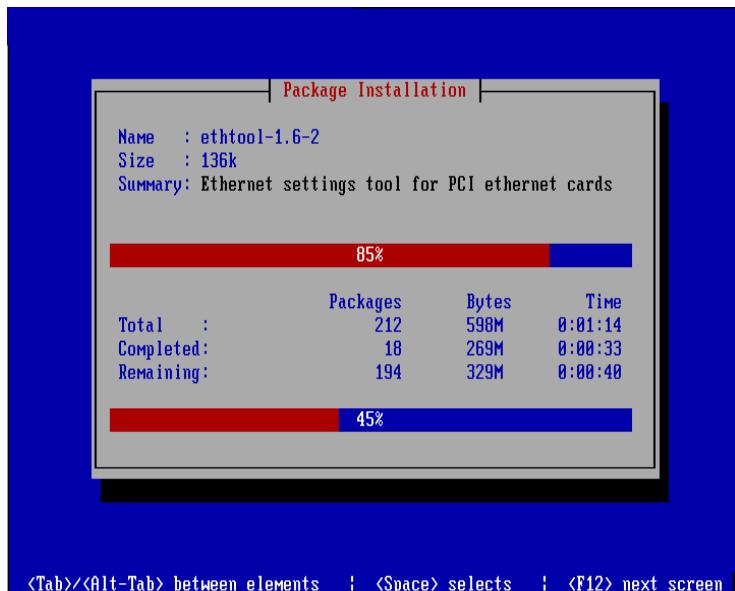
1. Insert the SQL Guard installation CD media into the CD-ROM drive located in the front of the SQL Guard machine.
2. If the system is powered on, turn off the power with the power button on the front panel.
3. Power the SQL Guard system on using the power button on the front panel of the system.

At this point, the power LEDs are lit on the front panel and the CD-ROM drive LEDs are active.

Once the system is ready to begin, a splash screen displays if a PC monitor is connected. After a few seconds the installation continues.

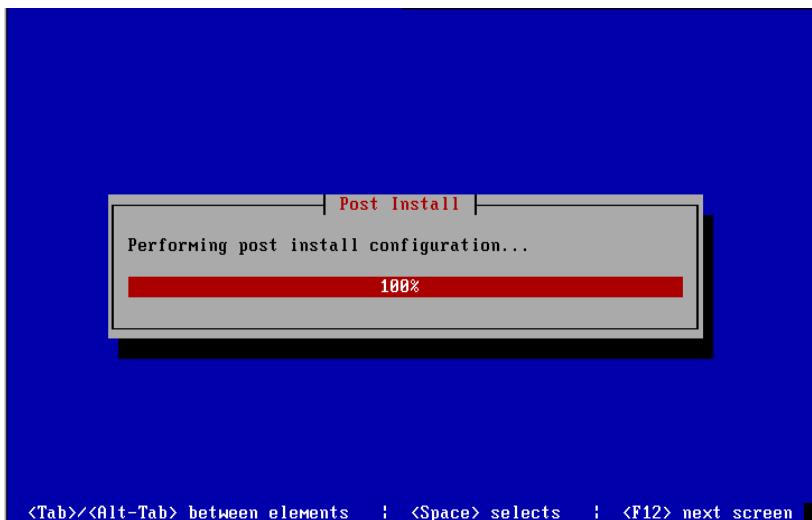


After displaying a few text messages, the monitor clears to a blue screen and the installation continues, displaying red-colored progress bars:



- The top progress bar displays progress on each individual package.
- The bottom progress bar displays overall progress for the installation.

Once the software is loaded, the monitor clears to a blue screen and displays a post-installation configuration progress bar. This may take a few minutes while the software is being configured.



A set of three tones are emitted from the appliance when the process is complete. In addition, the system ejects the CD-ROM drive and reboots.

4. Remove the CD media from the tray and close the CD-ROM drive.
5. When the system reboots, you are prompted for several items. Reply as indicated below (remember to note your admin password):

Admin password: *enter your admin password*

Master Key Y/N: *enter N*

The SQL Guard system is now reinitialized to a default configuration. For the default configuration, the *cli* password will be *guardium*. Be sure to reset this to your *cli* password after logging onto the *cli* for the first time.

The administrator should now perform the configuration and installation steps described beginning in [Step 3: Initial System Configuration](#) of *Chapter 1: Installation*.