

IBM Database Encryption Expert for Linux, UNIX, and
Windows



User's Guide

Version 1 Release 1

IBM Database Encryption Expert for Linux, UNIX, and
Windows



User's Guide

Version 1 Release 1

Note:

Before using this information and the product it supports, read the information in "Notices" on page 303.

This edition applies to Version 1 Release 1 of IBM Database Encryption Expert for Linux, UNIX, and Windows (product number 5724-R38) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC19-1275-02.

© Copyright Vormetric Inc. 2007, 2008. All rights reserved.

Contents

About this information

Intended audience	xv
Structure	xv
Accessibility features	xvii
If you need help	xvii
Conventions	xvii
Service updates and support information	xviii
How to send your comments	xviii

1 Introducing IBM Database Encryption Expert

The Encryption Expert Security Server	19
Logging	22
Failover Security Servers	22
The Encryption Expert File System Agent	23
The Encryption Expert Database Backup Agent	24
Context-aware access control	24

2 Management Console checklist

3 Installing, removing, and upgrading Encryption Expert

Before you begin	29
Requirements	31
Hardware requirements	31
Software requirements	31
Installing the Encryption Expert Security Server	32
Installing the primary Security Server	33
What next?	36
UNIX Security Server file structure	37
Installing and configuring the Encryption Expert failover server	38
Installing the Encryption Expert agent	43
Installing agent software on a UNIX system	43
Installing agent software on a Windows system	45
What next?	47
Database Backup Agent file structure (UNIX)	47

File System Agent file structure (UNIX)	48
Upgrading Encryption Expert	49
Upgrading Encryption Expert Security Server software	49
Upgrading Encryption Expert agent software	51
Upgrading UNIX agents	51
Removing Encryption Expert	54
Removing the Encryption Expert Security Server	54
Removing the Encryption Expert Database Backup Agent	54
Removing the Encryption Expert File System Agent	55

4 Starting the Management Console

Navigating the Management Console Interface	61
Using the Encryption Expert tray icon	62
Displaying agent status	63
Entering the host password	64
Displaying the agent log	64
Displaying the agent and Security Server fingerprints	65
Displaying the agent certificate	65
Stopping and restarting the Encryption Expert tray icon	66

5 Administering users

Configuring users	69
Changing your own user password	71
Changing another user's password	71
Logging in with an expired user account	72
Creating users	72
Changing user configurations	73
Deleting user accounts	74

6 Configuring hosts

Displaying configured hosts	75
Adding a host to the Security Server database	77
Batch host configuration	80
Setting host locks	81
FS Agent Locked	84
System Locked	85
Deleting hosts	87
Changing host certificates	88
Managing Encryption Expert agents	90
Applying Encryption Expert protection to a Database Backup Agent host	90
Applying Encryption Expert protection to a File System Agent host	91
Selecting the guard point mount type	96
Displaying File System Agent guard point status	96
Entering the File System Agent host password	97
Changing the File System Agent host password	98

Deleting guard points	99
Configuring an agent to use Security Servers	100
Configuring host settings	102
Checking host settings	106
Configuring host groups	108
Displaying host groups	108
Deleting host groups	108
Configuring agent logging	109
Applying Encryption Expert protection to a Database Backup Agent host group	109
Applying Encryption Expert protection to a File System Agent host group	112
Editing host groups	114

7 Using Encryption Keys

Viewing keys	116
Creating keys	118
Will the File System Agent require a network connection?	118
Creating a symmetric key	119
Manually creating symmetric keys	121
Creating an asymmetric key	122
Exporting a public key	123
Importing a public key	123
Deleting keys	124
Importing legacy keys	124
Configuring key groups	127
Creating key groups	127
Editing key groups	128
Deleting key groups	128
Rekeying protected data	129

8 Using policies

About Encryption Expert Database Backup Agent policies	131
Displaying policies	134
Configuring Database Backup Agent policies	134
Creating Database Backup Agent policies	134
Configuring Database Backup Agent backup rules	135
Configuring Database Backup Agent restore rules	137
Editing Database Backup Agent policies	139
Deleting Database Backup Agent policies	139
About File System Agent policies	140

9 Creating and editing File System Agent policies

About File System Agent policies	141
Navigating the Policy Composer	142
Using security rule edit functions	144
Configuring file system security rules	145

Opening the Policy Composer	146
Configuring resources	146
Adding a resource set	147
Editing a resource set	152
Deleting a resource set	154
Enabling the Allow Browsing function	154
Configuring users	155
Defining users	155
Editing user sets	158
Deleting user sets	159
Configuring processes	159
Adding process sets	160
Editing process sets	161
Removing process sets	162
Configuring time	163
Configuring access methods	167
Using access actions	168
What happens now?	169
Adding and removing effects	170
Adding a default security rule	171
Running in Warn Mode	172
Selecting a remote file, path, or object	173
Configuring encryption in file system security rules	175
Navigating the Key Selection Rules tabs	177
Key Selection Rule attributes	177
Using Key Selection Rule edit functions	177
Using variables to specify encryption keys	177
Defining a Key Selection Rule	178
Defining a Data Transformation Rule	179

10 Changing encryption keys

Using dataxform	184
Running multiple dataxform sessions	188
Preparing for dataxform	189
Create Security Rules	189
Specify the current encryption key	190
Specify the new encryption key	190
Manual dataxform	191
Automatic dataxform	192
Monitoring dataxform	193
Using the Security Server log	194
Using eetdxf_* files	195
Using dataxform_* files	195

11 Backing up and restoring DB2 databases

Backing up a DB2 database	197
How Encryption Expert Database Backup Agent encrypts backups.....	200
Restoring a DB2 database.....	201
Working with online DB2 data.....	203

12 Backing up and restoring IDS databases

Requirements	207
Configuring the IDS onconfig file	208
Configuring the IDS Database Backup Agent attributes file	208
Setting the IDS environment.....	212
Making an IDS backup	212
Making an IDS database backup using ontape	214
Using ontape without onconfig	214
Using ontape with onconfig	215
Making an IDS database backup using onbar	215
Restoring an IDS backup	216
Restoring an IDS database backup using ontape.....	216
Restoring an IDS database backup using onbar	218

13 Configuring high availability

Displaying the high availability configuration	219
Adding a failover server	220
Replicating the Encryption Expert Security Server database	221
SQL replication of the Security Server database.....	221
Configuring SQL replication.....	222
Backing up SQL replication configurations	224
Deleting a failover server.....	225
What to do in the event of a failure	225
Primary Security Server failure.....	226
Failover Security Server failure.....	226
If the Security Server databases go out of sync.....	227

14 Managing Encryption Expert logs

Displaying the Security Server log.....	229
Setting general display attributes.....	231
Configuring default log attributes.....	233
Analyzing log entries	233
Analyzing Database Backup Agent log entries.....	234
Analyzing File System Agent log entries	235
Log message levels.....	237
Using log files.....	237
Encryption Expert File System Agent log files.....	239
eetvmd_usr.log	239

messages (UNIX only)	239
secfs.log	240
SECFS_Logger_secfs_log.XXX	240
secfsd.log	240
statusfile	240
vmd_upload_ <i>hostname</i> .XXX	242
Encryption Expert Security Server log files	242
badlog.log	242
boot.log	242
eetdb2_jboss.log	242
cgss.log	243
eet_server_install_YYYY-MM-DD.log	243
db2setup.log	243
panic.log	244
server.log	244
vmdb2logup	245
voradmin.CGSSDB.ASN.CAP.log	245
voradmin.CGSSDB.CGSSDB_AQ.APP.log	246
Database Backup Agent log files	247

15 Setting preferences

Setting general preferences	249
Setting password preferences	251
Setting Security Server logging preferences	253
Setting DB2 Database Backup Agent logging preferences	255
Setting IDS Database Backup Agent logging preferences	258
Setting File System Agent logging preferences	260
Limiting the number of duplicate log entries	265
Setting syslog logging preferences	266

16 Maintaining Encryption Expert

Verifying agent to server SSL authentication	269
Determining the Encryption Expert Database Backup Agent version	270
Determining the Encryption Expert File System Agent version	271
Helpful utilities	272
Server utilities:	272
backup_config	272
cleanup_replication	272
configure_replication	272
deinstall_eet_server	273
import_v1120_config	273
logdb_controller	274
restore_config	276
re_gen_cert_auth	277
re_gen_failover_cert	278

re_gen_cert	279
re_sign_cert	279
re_sign_failover_cert	279
setup_failover_cert	280
start_replication	282
stop_replication	282
synchronize_replication	282
Agent utilities:	284
buildinfo	284
check_host	284
register_host	284
secfsd	289
vmd	292
vmsec	292
Encryption Expert database maintenance	294
Backing up the Security Server configuration	295
Restoring the Security Server configuration	296
Planning file system policies	297
Troubleshooting	298
Things to check	298
Verifying your system identity	300

17 Notices

Index

About this information

The *IBM® Database Encryption Expert for Linux®, UNIX® and Windows® User Guide* is a comprehensive reference that describes how to install, configure, and use the Encryption Expert security solution. Encryption Expert is a full-featured, multi-use data protection system that provides access control to data files, directories, and executables.

In these subsections are:

- Intended audience
- Structure
- Accessibility features
- If you need help
- Conventions

Intended audience

The *IBM Database Encryption Expert for Linux, UNIX and Windows User Guide* is intended for system, Security Server, and DB2 and IDS administrators.

The system administrator needs to know basic network and file system configuration and have root permissions for the systems on which Encryption Expert software is installed.

The Security Server administrator configures the agent and server components that comprise Encryption Expert. This process requires a familiarity with basic UNIX commands, networks, and DB2 or IDS usage. In addition, the Security Server administrator configures the SSL certificates that are used to authenticate agents to servers and servers to agents. Security Server administrators do not need to be familiar with X.509 certificate generation; however, an understanding of X.509 certificate generation will help Security Server administrators to understand the agent/server authentication configuration procedures that are described here.

The DB2 or IDS database administrator needs to know basic database configuration and maintenance procedures. The database administrator must have access permissions to the databases to be encrypted by Encryption Expert.

Structure

The following topics are available:

“Introducing IBM Database Encryption Expert”	A high-level description of the IBM Database Encryption Expert. It describes the data protection process and how Encryption Expert ensures that data is accessible only under certain conditions.
“Management Console checklist”	A high-level checklist of the steps for installing, configuring, and maintaining Encryption Expert, with cross-references to the sections that describe those steps.
“Installing, removing, and upgrading Encryption Expert”	A description of how to install, upgrade, and remove Encryption Expert server and agent software.
“Starting the Management Console”	A description of how to use the Encryption Expert Graphic User Interface (GUI) to configure Encryption Expert servers, clients, keys, and policies.
“Administering users”	A description of how to configure Encryption Expert users with specific roles and administrative controls.
“Configuring hosts”	A description of how to register and configure systems to use Encryption Expert protection.
“Using Encryption Keys”	A description of the keys used to encrypt data and how to administer them.
“Using policies”	A description of configuring DB2 and IDS Database Backup Agent policies.
“Creating and editing File System Agent policies”	A description of how to configure File System (FS) Agent policies.
“Changing encryption keys”	A description of how to rekey the encrypted files that are protected by Encryption Expert.
“Backing up and restoring DB2 databases”	A description of how to make and restore DB2 backups using Encryption Expert protection.
“Backing up and restoring IDS databases”	A description of how to make and restore IDS backups using Encryption Expert protection.
“Configuring high availability”	A description of the Security Server replication and failover process.

“Managing Encryption Expert logs”	A description of how to use the logging interface and how to read log entries.
“Setting preferences”	A description of how to set general Security Server and logging preferences.
“Maintaining Encryption Expert”	A description of miscellaneous management tasks, such as backing up and restoring the embedded Encryption Expert database and configured SSL authentication certificates.
“Notices”	Legal notices and license information.

Accessibility features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products successfully. IBM Database Encryption Expert uses standard operating system keyboard shortcuts.

If you need help

Please refer to the README file for current contact information.

Conventions

The following documentation conventions are used:

Text	Description
Bold	Bold text indicates a literal string in the Graphic User Interface (GUI) such as a button label or text-entry box label.
<i>Italic</i>	<i>Italic</i> text indicates a new term, the name of a window in the Web interface, a variable, or a document title.
[]	Bracketed items are optional and are generally used in command-line syntax.
	The pipe character indicates a choice and is generally used in command-line syntax.
...	Ellipses (...) indicate a break in information continuity. This notation is used in examples, such as log files, to display just pertinent data.

Also,

- Due to line-length restrictions, some one-line examples in the user guide are shown on multiple lines.

Service updates and support information

To find service updates and support information, including software fix packs, PTFs, Frequently Asked Question (FAQs), technical notes, troubleshooting information, and downloads, refer to the following Web page: www.ibm.com/software/data/db2imstools/support.html

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book, use either of the following options:

- Use the online reader comment form, which is located at: www.ibm.com/software/data/rcf.
 - Send your comments by e-mail to comments@us.ibm.com. Include the name of the book, the part number of the book, the version of IBM Database Encryption Expert, and, if applicable, the specific location of the text you are commenting on (for example, a page number or table number).
-

Introducing IBM Database Encryption Expert

IBM® Database Encryption Expert for Linux®, UNIX® and Windows® is a full-featured, multi-use data protection system. Encryption Expert provides access control to data files, directories, and executables. It uses strong encryption methodologies that allow you to process files but not access their contents. The same policy-based controls can be applied to ensure that DB2 and IDS database backups are secure and encrypted.




Figure 1.1: Encryption Expert architecture for DB2

The Encryption Expert Security Server

The Encryption Expert comprises a policy engine and a central key and policy manager, with complete security, performance, and scalability. Backup and restore requests are evaluated by the Encryption Expert server before the agent can perform the requested operation. The requests are evaluated by using agent-system parameters and user-defined policy constraints. Agents that run on Encryption Expert-protected hosts note every attempt made to access protected data and, based upon a set of rules, either permit or deny the access attempt. The set of rules is defined in a policy that is configured on the

Encryption Expert Security Server system and downloaded to the agent through a secure SSL network connection.

Secure Socket Layer (SSL) is the industry standard for system authentication. It relies upon signed certificates to identify systems to each other. SSL authentication is used to encrypt all communications between the Encryption Expert agent and server, as well as between server and server. The Encryption Expert employs X.509 digital certificates for agent/server and server/server communication. Tools and instructions for generating X.509 locally signed certificates are included with the Encryption Expert software.




Figure 1.2: Server and agent interaction

The Encryption Expert is a set of agent and server software packages that is administered by a Web-based user interface and command-line utilities. The Encryption Expert backup agent encrypts DB2 backups, and the Encryption Expert File

System Agent encrypts DB2 data files according to the policies defined on the Security Server.

The Encryption Expert Security Server runs on an embedded IBM DB2 Universal Database, v9.x. All database backup functions that are normally performed on the agent system are supported by the Encryption Expert server, including native database compression. Other than an additional command-line argument, DB2 backup operators are unaware of Encryption Expert intervention. The Encryption Expert administrator configures policies comprised of sets of security rules that must be satisfied in order to allow or deny access. Each security rule evaluates who, what, when, and how protected data is accessed and, if these criteria match, the Security Server either permits or denies access.

The security rules ask:

- Who is attempting to access protected data? You can configure one or more users. Users can be identified by user name, identification number, group number, instance name, and more.
- What data is being accessed? You can configure a mix of files and directories by specifying them individually or by using variables. You can include executables so, if the executables are not in the right location, they cannot access protected data.
- When is the data being accessed? You can configure a range of hours and days of the week to allow tasks, like nightly backups, to occur at specific intervals.
- How is the data being accessed? You can configure a security rule that considers how files and directories, and their attributes, are being accessed. The security rule can note attempts to read, write, delete, rename, create, and more.

When the conditions specified in a security rule match, the policy dictates whether to permit or deny access. If encryption is used, the policy can be configured to permit read access but without including the key to decrypt encrypted data. This way data can be backed up but not read.

An effective security solution also provides rich logging capabilities. The Encryption Expert agent notifies security administrators of policy violations in near real time. Encryption Expert records all context attributes of an access attempt, enabling complete traceability of host intrusion and data access events at the application and user level, and maintains an extensive log for detailed forensic analysis. The Encryption Expert lets you set the level of verbosity that you want. Levels range from just fatal errors to extremely verbose debug messages. In addition, the Encryption Expert provides audit logging to monitor all activities and transactions.

Both the Encryption Expert File System Agent and the Encryption Expert Database Backup Agent share many of the same features.

- Server/agent architecture
 - Policies with adjustable controls
 - Public and private encryption keys that are administered by an Encryption Expert Security Server
 - Centralized key management that allows you to control how and when keys are used
 - Secure communication channel between server and agent (SSL authentication)
 - Authentication certificates that are generated automatically during software installation
 - Mutual agent/server authentication that occurs during policy evaluation.
 - Configurable audit logging
-

Logging

Though transparent to the people making database backups and accessing system files, Encryption Expert agent activity is closely monitored and logged. Encryption Expert provides rich forensic capabilities through a centralized audit logging facility. All auditable events, including backups, restores, and security administration operations can be logged. Encryption Expert logs the following types of auditable events:

- Encryption Expert system events including initialization, shut down and restart
- Network connects and disconnects between different Encryption Expert components
- All backup and restore operation requests
- All configuration changes related to hosts, policies, and keys
- All creation and deletion of hosts, policies, and keys

The following Encryption Expert Database Backup Agent information is provided in the logs:

- Date and time
- Event type
- Severity
- Related subsystem
- Success or failure
- Name of related object (instance, key, policy, host, and so on)
- Description

The following Encryption Expert File System Agent information is provided in the logs:

- Date and time
- Event type
- Severity
- User identity
- Process from which the access attempt is being made
- Status: success or failure
- Name of related object (key, policy, host, and so on)
- Description

Failover Security Servers

A failover Security Server is a secondary server that is designed to assume the policy and key management load should the primary Security Server become unavailable. A failover server is a mirror of a primary server, with the same keys, policies, hosts, and so on. If the primary server become unavailable, the agent switches to a configured failover server. When the failover server is located and accessed, agent management continues as if the agent were being processed by the primary server.

Database synchronization always progresses from the primary server to the failover servers. The embedded failover server database can be synchronized manually by exporting the primary server database to each failover server, or automatically by using Encryption Expert utilities that automatically replicate the primary server database onto the failover server database. During initial failover server installation, the failover server is configured to communicate with a specific primary server. Since a failover server acts

like a backup of the primary server, during synchronization the failover server imports the primary server configuration, and the primary server updates the failover server each time a change occurs on the primary server.

Utilities are provided as part of the Security Server installation to export and import database information between servers. You also configure the failover servers that Encryption Expert agents can contact.

The Encryption Expert File System Agent

The Encryption Expert File System Agent is installed at a low-level within the host operating system. Its drivers are loaded early in the boot-up process. The Encryption Expert agent can administer all applications, processes, and files that start after it.




Figure 1.3: Encryption Expert agent within the system architecture

The Encryption Expert File System Agent does not lie dormant until it is invoked, as does the Encryption Expert Database Backup Agent. Rather, it is a process that is always running in the background. A guard point is the starting point at which to apply a policy. The agent intercepts any attempt to access anything in the guard point. The Encryption Expert File System Agent uses policies obtained from the Security Server and, based upon the applied policy, the Security Server grants or denies the access attempt.

Encryption Expert protection extends beyond simply allowing or denying access. Data in the guard point can be encrypted. The Encryption Expert encryption engine encrypts just the file contents, and leaves the file metadata intact. That is, you do not have to decrypt an encrypted file to see its name, timestamps, file type, and so on. This allows data management applications to perform their functions without exposing the file contents. By leaving the metadata intact, you have both protection and manageability. Backup managers may backup specific data, but they cannot see the contents. A file may be accessed by an unauthorized user but its contents are worthless without the appropriate Security Server approval and encryption keys. However, users with the correct policies and permissions are unaware that encryption and decryption are taking place.

Policy enforcement remains in place when there is no network connection between the Encryption Expert File System Agent and the Security Server. Key usage can be configured so that the Encryption Expert File System Agent must download the

encryption key each time encrypted data is accessed, or just the first time that encrypted data is accessed, and subsequent access uses a cached version of the key.

Policies that are configured with file encryption encrypt files as they are copied into a guard point and decrypt files when they are read or copied out of the guard point. A utility is provided to rekey files so that they are encrypted with a different key or to move the files to an unencrypted state. This utility, `dataxform`, provides a convenient way to migrate large numbers of files into or out of an encrypted state, and to change between encryption keys.

The Encryption Expert Database Backup Agent

Encryption Expert Database Backup Agent backup and restore operations can be performed in the following configurations:

- Basic backup in which data is encrypted and backed up with one server and multiple agents.
- Basic restore in which data is decrypted and restored on an agent that is configured to use the same server that was originally used to make the backup.
- Single-site, multi-EE backup and restore, in which configuration data is mirrored across multiple Security Servers in a single data center. Agents in this configuration can restore backups made by any other agent in the configuration. The public and private keys used to restore and backup data are shared between the servers using database replication. Policies and host configurations are also mirrored across all the servers.
- Multi-site restore in which backups are restored on different Encryption Expert servers in different data centers. Backups made on one Database Backup Agent can be restored on another Database Backup Agent whose server is configured with the appropriate private key. Public key information can be copied from one server to another, so the Database Backup Agents with a Security Server that is configured with the public key can just make encrypted backups. The backup can be restored on any Database Backup Agent whose server is configured with the matching private key and SHA-2 hash code. The policies and host configurations vary between the servers, so you can configure a policy on one server that allows only backups and a policy on another server that allows only restores.
- Encryption Expert backs up and restores static data-at-rest and active online data.

Context-aware access control

A policy defines a set of attributes, or conditions, that need to exist before granting backup or restore permission. When an agent system sends a backup or restore request to the server system it also sends additional information that describes the database, database server, and the agent system itself. The additional information, matched against the policy rules, are evaluated by the Security Server.

A combination of the following attributes are checked to determine whether to grant or deny backup and restore requests:

- Agent host name
- DB2 database alias
- DB2 database partition name
- DB2 database instance name
- IDS database server name
- IDS database server number
- Permissible time range, relative to the Security Server

Based upon these attributes, the policy either allows or denies permission to complete the backup or restore operation. If a Security Server permits a backup operation, the Security Server returns the public key necessary to encrypt the backup metadata which, in turn, contains an additional, randomly generated encryption key that encrypts and decrypts the actual data.

See also [“About Encryption Expert Database Backup Agent policies” on page 8-131](#).

Management Console checklist

The following checklist is presented as a high-level flow for applying Encryption Expert protection.

Step	Description
Select the system on which to run the primary Security Server	This is a dedicated server. It is not running DB2 and it is accessible across the network. See “Before You Begin” on page 2-19 .
Install the primary Security Server software	Security Server installation is described in “Installing, removing, and upgrading Encryption Expert” on page 3-29 .
Verify that you can open the Management Console	Open a Web browser, then enter <code>https://</code> followed by the host name or FQDN name of the Security Server. Include <code>:.8445</code> . For example, <code>https://vmlinux.vormetric.com:8445</code> . Specify the login name and password when prompted. See “Starting the Management Console” on page 4-57 .
Change the administrative password	Change the admin password. See “Changing another user’s password” on page 5-71 .
Create administrative user accounts	Create user accounts for the people you want to administer the Security Server. The user accounts can have complete or specialized administrative permissions. User account creation is described in “Creating users” on page 5-72 .
Log in as a user with all permissions and who can perform all Encryption Expert tasks	Log into the Security Server Web interface as a user with all the administrative permissions. You can also log in as a user with permission to perform individual tasks, but you will have to switch between users to perform all the tasks necessary to configure hosts, policies, keys, etc.
Create encryption keys	Create the keys that will be used to encrypt data on protected hosts. See “Using Encryption Keys” on page 7-115 .
Create encryption key groups	Group the keys so multiple keys can be applied as a single unit to a policy. See “Configuring key groups” on page 7-127 .
Create policies that govern the types of access and action allowed on each protected host	Specify sets of rules for host resources using a wide-range of constraints. Policies and their preparation are described in “Using policies” on page 8-131 .
Set the Security Server, File System Agent, and Database Backup Agent log levels	Set logging to the appropriate levels for the type and volume of log data you wish to generate. See “Setting preferences” on page 15-249 .

Add host systems to the Security Server database	Add the name and a description of each host system to be protected. This information is required to register the host with the Security Server. You can add the hosts to the Security Server after installing the host software, however, you will have to manually configure certificates. See “Adding a host to the Security Server database” on page 6-77.
Install the File System Agent and/or Database Backup Agent software on the host systems	You can install one or both agents on a host system. The order in which you install the agents does not matter. See “Installing, removing, and upgrading Encryption Expert” on page 3-29.
Configure hosts and groups in the Management Console	Configure the hosts to be protected by the Security Server individually or in groups. See “Configuring hosts” on page 6-75.
Configure a failover Security Server	(optional but highly recommended) Configure an alternate Security Server to use should the primary Security Server fail. See “Configuring high availability” on page 13-219.
Steps to take if there is a firewall between the host and the Security Server	Open the firewall administration utility for your network and permit access from the Security Server to the host through TCP/IP ports 7024, 8443, 8444, and 8445. You can change the number of the listening port used to exchange policy requests, policy decisions, and encryption keys between a Security Server and its File System Agents. See “Setting Security Server logging preferences” on page 15-253 for details. Windows users can open Control Panel - >Windows Firewall and make port 7024 an exception. The method for UNIX users depends upon the firewall application in use.
Test the policy and host configuration	See “Backing up and restoring DB2 databases” on page 11-197 , and “About File System Agent policies” on page 8-140.
View the Security Server log	See “Managing Encryption Expert logs” on page 14-229.

Installing, removing, and upgrading Encryption Expert

A single Encryption Expert Security Server can be configured to support multiple agents. The Security Server installation utility is a flexible tool that allows you to configure the Security Server as a standalone, primary, or failover server. Standalone and failover servers can be used if a primary Security Server fails. Also, the Security Server installation utility gives you the option to automatically create your own locally signed certificates for use in server-client authentication.

As part of the installation process, you can configure the agent to use a specific Security Server. At which time, the utility performs all of the configuration work that is necessary for server-client authentication with a minimum of user input.

Before you begin

Read this section before installing Encryption Expert software. It lists initial system setup and preparation requirements for improving the software installation process.

- Check the “[System requirements](#)” section of the README file before you install and configure any Encryption Expert agent or server software.
- Check the “[Installation notes](#)” section of the README file for directions on how to download the Security Server prerequisite package from the Vormetric™ Web site. This package is required to run Encryption Expert.
- Get the root password for the system on which you are installing the software. You must have system administrator (root) access to the system. If you are installing failover servers, you will need the root password for both the failover and primary servers.

The Encryption Expert software installer executes multiple commands, scripts, and utilities -- some of which are part of the Encryption Expert software package and some are part of the operating system. These must be run as root. To install Encryption Expert as a sudo user, the sudo user must be allowed to execute both Encryption Expert and regular system commands. If you want to use `sudo`, it is unlikely that you will have sufficient access to all the system commands and you will have to manually configure the server/client SSL certificates. Manual certificate generation is described in “[Changing host certificates](#)” on page 6-88.

- Verify that the system clock of every system that is to run Encryption Expert software, both server and agent, is correct and synchronized to the timezone in which it is located. Set the time relative to Greenwich Mean Time (GMT) or Coordinated Universal Time (UTC). Use the same time standard when you set the agents and

servers. That is, if you set one system to GMT, set all the systems to GMT. Incorrect times can affect certificate usage and make it harder to evaluate log entries. You might need to synchronize the systems with a Network Time Protocol (NTP) server.

- Set the correct time and time zone of every system in a Security Server configuration before installing any Encryption Expert software. Once set, maintain the correct time by configuring each system to access an NTP server at regular intervals. Incorrect times can make it difficult to sequence log entries, can prevent SSL communication between the Security Server and agent, and can prevent the Security Server from applying policy updates to agent systems. If Encryption Expert software is installed and you need to set the Security Server time, stop the Encryption Expert Security Server (i.e., `/etc/init.d/cgss stop`), set the clock, and then restart the Encryption Expert Security Server (i.e., `/etc/init.d/cgss start`).
- Copy installation files onto the local hard disk or use the Encryption Expert installation CD-ROM.
- Install and maintain the Encryption Expert Security Server in a “safe” environment. A safe environment has stringent user access controls, a limited number of users with root access, and system access is monitored. The Security Server can be compromised by an internal attack by someone with root access.
- Check the system and network identity of each system you want to use as an agent or Security Server. Refer to [“Verifying your system identity” on page 16-300](#) for details.
- This information in this document regarding Database Backup Agents is very DB2-centric. Almost all the references are to DB2 paths, executables, and output files. The `./agent/db2` and `./agent/ids` installation directory contain the same basic files and executables. The examples in this documentation use the `./agent/db2` path but, if you are installing Encryption Expert for IDS, replace `db2` with `ids` in the path specification and the result will be the same. Note that output file names also change accordingly. IDS-specific information is contained in [“Backing up and restoring IDS databases” on page 12-205](#).
- The Security Server Web interface is run on a Java-enabled Web browser, such as Windows Internet Explorer and Mozilla Firefox. Run Internet Explorer version 6.0 or higher. Run Firefox version 1.8 or higher.
- Microsoft Internet Options allows you to set privacy parameters for the Web browser to block cookies. If the privacy level is set too high, you will block the Security Server cookie and you will be unable to initiate a Management Console session on the local system. You are prompted to login but your login information will be ignored.
- (Windows only) Agent software for Windows platforms is installed in two folders. The Encryption Expert application is installed in `\Program Files`. The second folder in which Encryption Expert data is placed is `\Documents and Settings\All Users\Application Data`.

The `\Documents and Settings\All Users\Application Data` folder is hidden by default. Encryption Expert stores agent logs, configuration data, and certificates in the folder. If you cannot browse the `C:\Documents and Settings\All Users\Application Data\IBM` folder, enable the **Show hidden files and folders** radio button in the Folder Options menu to view the folder and its contents.

- (Windows only) If there is a firewall between the Security Server and a Windows XP or Windows 2003 agent, configure `vmd.exe` as a firewall exception on the Windows agent; otherwise, the Security Server will be unable to browse the agent. For example, if using Windows Firewall, select **Control Panel->Windows Firewall->Exceptions->Add Program...** and browse for `vmd.exe`. The default location is

```

\Program
Files\IBM\DBTOOLS\LUWEncryptionExpert\agent\vmd\bin\vmd.
exe.

```

Requirements

Hardware requirements

Minimum Security Server hardware requirements

- 4 GB of RAM
- 80 GB of disk space*
- Intel P4 2.4 GHZ equivalent, or later

* The actual amount of disk space required to install the Security Server software is considerably less. The extra space will be used to accommodate log files, which can quickly proliferate and consume large amounts of disk space. This requirement includes all the directories used by the Security Server, comprising /opt/IBM, /var/log, /home (voradmin, db2fenc1, dasusr1), and /tmp.

Agent software has been tested on the following hardware

- AIX: PowerPC 4 (64-bit)
- AIX: PowerPC 5 (64-bit)
- Linux: Intel x86_64 or AMD Opteron (64-bit)
- Solaris: Sun4c 64-bit Sparc
- Windows: Intel x86_32 (32-bit) or x86_64 (64-bit)
- Windows: AMD Opteron or Intel Xeon 64, or higher

Software requirements

Encryption Expert Security Server has been tested on the following hardware

- Linux (32-bit): Red Hat Advanced Server 4.0 Update 4

UNIX agent software has been tested on the following operating systems

- AIX 5.2 TL10 SP2, plus APAR IZ10719
 - AIX 5.2 TL10, plus APAR IZ21543
 - AIX 5.3 TL6, plus APAR IZ20619
 - AIX 5.3 TL7, plus APAR IZ07461
 - AIX 5.3 TL8, plus APAR IZ15965
 - AIX 5.3 TL6 SP3, plus APAR IZ07461
 - Linux (64-bit): Red Hat Advanced Server 4.0 Update 4
 - Solaris 9 and 10
-

Windows agent software has been tested on the following operating systems

Windows on x86 32-bit hardware running:

- Pentium 4 or later
- Windows 2000 Server, Professional, or Advanced Server, all with SP4
- Windows 2000 Advanced Server with Microsoft Cluster Service (MSCS)
- Windows Server 2003, Standard Edition, with SP1 or SP2
- Windows Server 2003, Enterprise Edition, with Microsoft Cluster Service (MSCS)
- Windows XP with SP1 or SP3
- Windows XP with SP2 and Microsoft update KB885894

Windows on x86 64-bit hardware running:

- AMD Opteron or Intel Xeon 64, or higher
- Windows Server 2003, 64-bit Enterprise Edition, with SP1
- Windows Server 2003, 64-bit Enterprise Edition, with Microsoft Cluster Service
- Windows XP, 64-bit Edition

The following DB2 versions have been tested with Encryption Expert

DB2 V8 Fixpack 14

DB2 V9.1

DB2 V9.5

The following IDS versions have been tested with Encryption Expert

Informix IDS Dynamic Server V11.10 FC2

Installing the Encryption Expert Security Server

The installer performs the following major tasks:

- Installs the Security Server database and related components. You will assign passwords at this time for some administrative users.
- Generates authentication certificates. This process is important for Encryption Expert agent/server communication. The installer first creates a signer certificate from information that you supply when prompted, and it uses the signer certificate to create a server certificate.
- Starts the JBoss® Web application server.

The Encryption Expert Security Server package includes a version 9.1 DB2 installation that is installed under `/opt/ibm/db2/V9.1`. The `voradmin` user is the default database administrative user. All configuration data is stored in this Encryption Expert embedded database.

Installation takes 5 to 10 minutes, or longer, depending upon the system load. Do not press the <Enter> key to try to speed things up. You must supply input during the installation process, and whatever you enter will be used the next time the installer requests user input.



The system that hosts the Security Server should be a dedicated server. Do not install the Security Server on a system that is hosting other applications because those applications can conflict with the Security Server. One such conflict involves DB2. The embedded

Security Server database and a regular DB2 database share common components, such as user IDs and user groups in `/etc/passwd` and `/etc/group`. Cross-edits and database corruption can occur when both applications run on the same system. The Encryption Expert Security Server installation wizard will abort if an existing DB2 installation is present. Also, do not install DB2 onto the system after installing Encryption Expert Security Server. Doing so will destroy the embedded database and make Encryption Expert Security Server unusable. These restrictions do not apply to Encryption Expert agent software installation.

Install the Encryption Expert Security Server first, followed by the Encryption Expert agents, because the agent installation utility automatically performs agent/server X.509 certificate generation. If you install the Encryption Expert agent without a working Encryption Expert Security Server available, you must skip automatic certificate generation and then manually generate the X.509 certificates later when the Security Server becomes available. Manual certificate generation is described in [“Changing host certificates” on page 6-88](#).

The Security Server installation is placed in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server`. Additional files are also installed in the following locations:

- `/tmp`: Software installation and removal logs are placed in `/tmp`.
- `/home`: Home directories are created for the default Security Server database administrator, `voradmin`, as well as `db2fenc1` and `dasusr1`. `db2fenc1` and `dasusr1` are standard DB2 administrative users. `voradmin` is the Encryption Expert database user equivalent of `db2inst1`. Do not remove the directories or `passwd` file entries for these users.
- `/etc/init.d`: This directory contains the utility that starts and stops the JBoss application server, the embedded DB2 database, and anything else that is related to running the Encryption Expert Security Server. This utility is named `cgss`.
- A more complete list of installed files is available in [“UNIX Security Server file structure” on page 3-37](#).

You can delete the log files in `/tmp`. You can also delete log files in other directories as long as they are not active logs. An active log file has a simple name, such as the agent log, `eetdb2_db2inst1.log`. Inactive log files are generally appended with some string, like the date, to imply that they have been cycled from active logs to inactive logs, such as `eetdb2_db2inst1.log.2008-03-07-16-28-57`. Do not attempt to manually modify or remove active logs. Doing so can render the Security Server inoperable. Use the Management Console interface to configure server and agent logs. Backup and delete inactive logs with some regularity to maximize available hard disk space.

If you are updating an existing Security Server installation, see [“Upgrading Encryption Expert Security Server software” on page 3-49](#).

Installing the primary Security Server

Install the Security Server on a 32-bit Linux server running Red Hat Advanced Server 4.0 Update 4.

The installer searches for, and uses utilities in, `/opt/ibm/db2/V.9.1`. It also uses `/tmp` as a software staging area and to write a log file.

To install Security Server software:

- 1: Log onto the system as the root user.
-

- 2: Insert the Encryption Expert media CD-ROM in the CD-ROM drive or copy the installation package onto the system.
- 3: Verify that there is sufficient disk space in both /tmp and in the partition on which you intend to install the software.

- 4: If you haven't done so already, download and install the Security Server prerequisite package, as described in the README file.

The Security Server prerequisite package comprises an executable named `install_eet_prereq` and a source file named `eet_prereq.zip`.

- 5: If you are installing the software remotely, set your `DISPLAY` environment variable to point to an X-Windows server.

The Security Server installation utility opens a graphic window in which you specify installation parameters.

NOTE: Skip this step if you want to install the software on a system that does not support X-Windows graphics. You can install the software in text-mode by including the “-i console” argument when you start the installation utility. You are asked the same questions regardless if you run the installation utility in graphic-mode or text-mode.

- 6: Insert the Encryption Expert media CD-ROM in the CD-ROM drive or copy the installation package onto the system.
- 7: Verify that there is sufficient disk space in both /tmp and in the partition on which you intend to install the software.
- 8: Change to the directory that contains the Security Server installation software.
- 9: Start the installation utility.

```
# ./install_eet_server.bin
```

or

```
# ./install_eet_server.bin -i console
```

The *Introduction* window opens.

- 10: Click **Next**.

The *Panel Custom Code* window opens. It displays the IBM license agreement.

- 11: Read the license agreement, then select the **I accept both the IBM and the non-IBM terms**.

You cannot proceed if you do not accept the terms of the license agreement.

(Optional) Click **Print** to send the license agreement to a printer. This button opens the printer interface in which to specify print attributes such as orientation, number of copies, and paper size.

(Optional) Click **Read non-IBM terms** to open a window to view the license agreements for third-party products such as HP-UX SDK and OpenSSL.

- 12: Click **Next**.

The *Select Server Role* window opens. It prompts you to indicate if the local Security Server installation is to be the primary server or a failover server which is used if the primary server fails or becomes inaccessible.

You should install and configure the primary Security Server before installing the failover Security Servers because the failover server needs to copy database and configuration information from the primary server at this stage of installation. If this is a failover Security Server, continue installation using the instructions in [“Installing and configuring the Encryption Expert failover server” on page 3-38](#).

13: Select the **Primary** radio button.

14: Click **Next**.

The *Administrator Password* window opens.

15: Enter a password for the Encryption Expert database administrator.

The Encryption Expert database administrator is `voradmin`. It is a database user much like `db2inst1`. The `voradmin` user has a login in `/home` and it can be used to open the Encryption Expert databases, CGSSDB, LOGDB and JMSDB, to debug database issues.

16: Click **Next**.

The installation utility checks system resources and then opens the *Pre-Installation Summary* window. This window displays the name of the product being installed, the target destination, the required amount of disk space, and the available amount of disk space.

17: Click **Install**.

The installation process begins. Files are unpacked and the keystore is generated. The installer then enters the certificate generation stage and opens the *Server Configuration* window. This window prompts you to enter the host name of the local Security Server.

18: Enter the host name of the primary Security Server in the text-entry box.

Enter the network identity that agents and the failover Security Server will use to connect to the primary Security Server. This name is also used to identify the server to agents in the server SSL certificate.

19: Click **Next**.

The *Keystore Information* window opens.

20: Enter information to identify the certificate issuer.

The information that you provide is displayed when the server signer-certificate is viewed. You are prompted to specify:

- Your organizational unit, which is frequently a department or group name
- Organization name, which is frequently the company name
- City or locality in which the organization is located
- State or province in which the organization is located
- The country in which the organization is located

21: Click **Next**.

The agent request-certificate is generated and submitted to the Security Server for signing. The signed certificate is returned to the agent. Upon the successful transfer of certificates, the *Important Message* window opens. This window displays the installation status, plus instructions for opening the Security Server Web interface.

22: Click **Next**.

The *Installation Complete* window opens.

23: Click **Done**.

The installation utility closes.

24: Check that the Security Server has “settled down” before you run Security Server applications.

After the installation utility completes, Java and DB2 processes continue to run to complete the embedded database and Security Server configuration. Run a command like `top` or `slabtop` to monitor system activity. You can open the Management Console and administer the Security Server after the Java and DB2 processes become inactive. It will take a few minutes.

- 25: Start a Management Console session by entering the URL returned by the installation utility (e.g., `https://vmlinux01:8445`) in a Web browser.

What next?

After installation completes, you can:

- Inspect the Encryption Expert installation log. The installation log is `/tmp/eet_server_install_YYYY-MM-DD.log`. The log lists the status of the DB2 database manager, the database server, the name of the user with administrative access (for example, `admin`), the server database alias, the Security Server databases created, and the commands used to create the Security Server database schema. At the end it lists the status of the DB2 and JBoss servers.
- Inspect the JBoss configuration log. The configuration log is `/tmp/cgssdb2_jboss.log`. The log lists the JBoss files installed in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA`.
- Inspect the DB2 configuration log for the Encryption Expert Security Server database. The configuration log is `/tmp/db2setup.log`. This log lists installed database components, language, service name and port number for the administrative user, license checks, etc.
- Inspect the installation hierarchy in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server`.
- Check that the `jboss` and `java` processes are running. For example,

```
# ps -aef | grep jboss
voradmin 30638      1   0 Nov14 pts/6      00:00:00 /bin/sh
/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-
4.0.4.GA/bin/run.sh
voradmin 30644 30638  0 Nov01 pts/6      00:10:02 /opt/IBM/
DB2TOOLS/LUWEncryptionExpert/server/ibm-java2-i386-50/
jre/bin/java -server -Xms128m -Xmx512m .../opt/IBM/
DB2TOOLS/LUWEncryptionExpert/server/ibm-java2-i386-50/
jre/lib/tools.jar org.jboss.Main
```

If JBoss is not running, restart it using the `cgss` command. For example,

```
/etc/init.d/cgss restart
```

- Wait a few minutes after installing the Security Server software to give the Web application server an opportunity to settle down, then open a Web browser to configure the Security Server and agents.



The default Security Server administrator is `admin` and the password is `admin123`. Log on as this administrator the first time you start a Security Server Web session. Change the

password and create other administrative users. For information about administrative roles and their configuration, see [“Configuring users” on page 5-69](#).

- Install and configure a failover Security Server. See [“Installing and configuring the Encryption Expert failover server” on page 3-38](#) and [“Replicating the Encryption Expert Security Server database” on page 13-221](#).
- Restore a backup of a Security Server database. See [“Restoring the Security Server configuration” on page 16-296](#).

UNIX Security Server file structure

The Security Server directories and files are placed in:

`/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server`

This directory contains:

- `./server/apache-2.2.4`
 - Directory that contains the Apache Web server installation.
- `./server/.bashrc`
 - An installation artifact used to ensure that the Encryption Expert software is installed in the appropriate environment.
- `./server/bin/`
 - Directory that contains the executable utilities and binaries to configure and administer the Security Server installation. You may want to add this and `./server/bin/internal` to your `$PATH` variable definition.
- `./server/db2ee0101.sys2`
 - An ASCII text file that contains the name and version of the installed software.
- `./server/deinstall_eet_prereq`
 - A utility to remove the Security Server prerequisite package.
- `./server/etc/`
 - Directory that contains Security Server configuration data, such as the access file. (The access file is created during installation and contains an encrypted master passphrase that is used to decrypt other passphrases. Do not edit or delete the access file.)
- `./server/ibm-java2-i386-50/`
 - Directory that contains the Java Runtime Environment for the Security Server.
- `./server/jboss-4.0.4.GA/`
 - This directory is created and populated by the Security Server prerequisite package that you installed earlier. This directory contains the JBoss Web application server for the Security Server. It also contains binaries to administer JBoss manually and several important log files. This directory and its contents are created by the Security Server prerequisite package, which is described in the README file. See also [“Using log files” on page 14-237](#).
- `./server/jre`
 - Directory that contains the Java Runtime Environment for the Security Server.
- `./server/license`

- Directory that contains the Encryption Expert license agreement in various languages.
- `./server/log`
- Directory that is linked to the primary Security Server logging directory `./server/jboss-4.0.4.GA/server/default/log`. This directory contains the log files that you should check first to troubleshoot the Security Server.
- `./server/nodelock`
- Text file that contains the Encryption Expert license key and the Encryption Expert license expiration date.
- `./server/pem/`
- Directory that contains Privacy Enhanced Mail (pem) certificates, certificate requests, signer certificates, and more. These files control SSL authentication between the Encryption Expert server and associated agents.
- `./server/sql/`
- Directory that contains SQL run files used to replicate the Security Server embedded databases. The files are called by the replication utilities in `./server/etc` and you should have no occasion to use these files directly.
- `/tmp`
- `./db2setup.log`: See “[db2setup.log](#)” on page 14-243.
 - `./eetdb2_jboss.log`: See “[eetdb2_jboss.log](#)” on page 14-242.
 - `./logdb_controller.log`: See “[logdb_controller](#)” on page 274.
- `/home/adminUser`
- Administrative user home directory created during Security Server installation. The default is `voradmin`.
- `/etc/init.d/`
- `./cgss`: Executable utility that starts, stops, and restarts the JBoss Web application server and the embedded DB2 database server.



Use only the files and utilities that are documented. The files and utilities that are not documented are used internally by Encryption Expert and are deliberately omitted.

Installing and configuring the Encryption Expert failover server

The installation utility allows you to designate the local Security Server installation as either a primary or a failover server. When designated as a failover server, the installer prompts you to identify the primary server and then the installer generates the SSL certificates for primary and failover server communication.

Limited Web access to the failover server Management Console is provided. You can view status and log information on the failover server, but you cannot change the failover server configuration. Security Server configuration changes are made through the Web interface on the primary server and then propagated to the failover servers. Changes that occur on a failover server are not propagated back to the primary server. This includes log data. When an agent switches to a failover server, the events logged on the failover server

remain on the failover server. To view the events logged on a failover server, open the Web interface on the failover server or configure a syslog server. Syslog configuration is described in [“Setting Security Server logging preferences” on page 15-253](#) and [“Setting syslog logging preferences” on page 15-266](#).

By default, an agent is automatically configured to use the primary Security Server when you install the agent software. However, you must manually add the failover Security Servers to the `agent.conf` file of each agent.

Agent authentication certificates do not have to be generated and exchanged with the failover servers. During failover-server software installation, the primary-server root certificate is copied to the failover servers. This gives the failover server permission to talk to the same agents as the primary server and the agents certificates remain valid, regardless which server it communicates with.

The failover server databases can be synchronized with the primary server database manually or automatically. Utilities are provided to replicate the primary server database onto the failover servers. (See [“Replicating the Encryption Expert Security Server database” on page 13-221](#) for details.)

During failover server installation, the failover server generates a certificate request and sends it through an SSH tunnel to the primary server. The primary server returns the signed certificate, and other certificates, back to the failover server through the SSH tunnel. If SSH is not properly setup, an SSH tunnel cannot be established between the primary and failover servers, the primary and failover servers will be unable to exchange certificates, and failover server software installation will not complete successfully.

Before you install failover server software:

- Enable TCP forwarding. Open the SSH daemon configuration file, `/etc/ssh/sshd_config`, on both the primary and failover servers and ensure that `AllowTcpForwarding` is set to `yes`.
- If they are not already running, start the SSH daemon, `sshd`, on both the primary and failover servers.
- Verify that a user, any user, can ssh from the failover to the primary server.

After the failover server is installed and the certificates exchanged, you can disable TCP forwarding and shutdown the `sshd` processes.

To install and configure a failover server:

- 1: Install and configure a primary Security Server.
- 2: Add the failover server to the primary Security Server database.

Add the failover Security Server to the primary Security Server database before you install the failover Security Server software. This will reduce the number of steps required to install and configure the failover server. See also [“Adding a failover server” on page 13-220](#).

a: Log into the Management Console as a user with `server` role permissions.

b: Select **High Availability** in the menu bar.

The *High Availability Servers* window opens.

c: Click **Add**.

The *Add High Availability Server* window opens.

d: Enter the name to identify the failover server in the **Server Name** text-entry box.

This is the same name that will be used to identify the failover server during the certificate generation phase of failover server software installation. It will

also be used as the Common Name (CN) in the failover server certificate. Enter the hostname or Fully Qualified Domain Name (FQDN) of the failover server.

e: Click **Ok**.

The *High Availability Servers* window opens and displays both the primary and failover servers. The **Registered** checkbox for the failover server is not enabled yet.

- 3: Log onto the system that is to act as the failover server as the `root` user.
The failover server must be able to access the primary server across the network.
 - 4: Install the Security Server prerequisite package by executing the `install_eet_prereq` utility.
See the `README` file for details.
 - 5: Install Encryption Expert Security Server software on the local system, as described in [“Installing the Encryption Expert Security Server” on page 3-32](#). Stop when you are prompted to indicate whether the installation is for a primary Security Server or a failover Security Server, and continue with the steps below.
 - 6: Select the **Failover** radio button to indicate that the local system is to be configured as a failover server.
 - 7: Click **Next**.
The *Administrator Password* window opens.
 - 8: Enter a password for the Encryption Expert database administrator.
The Encryption Expert database administrator is `voradmin`. It is a database user much like `db2inst1`. The `voradmin` user has a login in `/home` and it can be used to open the Encryption Expert databases, `CGSSDB`, `LOGDB` and `JMSDB`, to debug database issues.
 - 9: Click **Next**.
The installation utility checks system resources and then opens the *Pre-Installation Summary* window. This window displays the name of the product being installed, the target destination, the required amount of disk space, and the available amount of disk space.
 - 10: Click **Install**.
The installation process begins. Files are unpacked and the keystore is generated. The installer then enters the certificate generation stage and opens the *Server Configuration* window. This window prompts you to enter the primary and failover Security Server domain names.
 - 11: Enter the host name or FQDN of the primary Security Server in the **Primary Server Host Name** text-entry box.
This is the recognized network identity that agents and the failover server will use to communicate with the Security Server.
 - 12: Enter the host name or FQDN of the failover Security Server in the **Failover Server Host Name** text-entry box.
This is the recognized network identity that the failover server will use to communicate with the primary server.
 - 13: Click **Next**.
The *Keystore Information* window opens.
 - 14: Enter information to identify the certificate issuer.
-

The information that you provide is displayed when the failover server signer-certificate is viewed. You are prompted to specify:

- Your organizational unit, which is frequently a department or group name
- Organization name, which is frequently the company name
- City or locality in which the organization is located
- State or province in which the organization is located
- The country in which the organization is located

15: Click **Next**.

The *Setup SSH* window opens. This window opens only when installing failover servers. It is skipped during primary server software installation. This window prompts you to run the `ssh_setup` utility on the failover server and, after running the utility, you can continue the failover server installation process.

The `ssh_setup` utility creates RSA authentication credentials for the `voradmin` user so that the `voradmin` user on the failover server does not have to supply a password each time it uploads information to the primary server. This utility copies the public key to the `/home/vormetric/.ssh` directory on the primary server and creates assorted ssh configuration files in `/home/vormetric/.ssh` on the local failover server.

16: Run the `ssh_setup` utility on the failover server.

- a: Enter the full or relative path to the `./server/bin/ssh_setup` utility on the failover server. The utility takes one argument -- the host name or FQDN of the primary server. For example:

```
# ./bin/ssh_setup vmlinux01
```

Hint: If you are unsure of the primary server network identity, execute the `./server/bin/internal/check_host` utility on the primary server and enter the name exactly as shown.

The `ssh_setup` utility might question the authenticity of the server.

```
The authenticity of host 'vmlinux01 (10.2.0.37)' can't be
established. RSA key fingerprint is
84:d8:b3:3f:b9:20:c9:e5:94:e0:13:37:e5:55:1f:19.
Are you sure you want to continue connecting (yes/no)?
```

- b: Check the system name, FQDN, and/or IP address displayed, and, if you are comfortable that the primary server network information is correct, enter `yes`.

A message is displayed that indicates that the primary server has been added to the table of hosts that the failover server can connect to via ssh.

```
Warning: Permanently added 'vmlinux01' (RSA) to the list of
known hosts.
```

You are then prompted to enter the password for the `voradmin` user on the primary server.

```
voradmin@vmlinux01's password:
```

- c: Enter the primary server `voradmin` user password.
d: Return to the failover server installation currently in progress.



You should have no occasion to use the `ssh_setup` utility except during failover server installation.

17: Click **Next**.

The failover server request-certificate is generated and submitted to the primary server for signing. The signed certificate is returned to the failover server. Upon the successful transfer of certificates, the *Important Information* window opens.

This window displays the installation status, plus additional steps to perform in order to complete failover server configuration.

- 18: Click **Next**.

The *Installation Complete* window opens.

- 19: Click **Done**.

The installation utility closes and you are done installing failover server software; however, there are a few more steps that are necessary to make the failover server operational.

- 20: (optional) Verify that the failover server was able to successfully exchange certificates with the primary server.

a: Open the *High Availability Servers* window on the primary server.

b: The **Registered** checkbox for the failover server should be enabled.

- 21: Backup the failover server database.

It is recommended that you backup the failover server database before you enable replication. The backup is a pristine, “zero-state” image of the failover server database that is used to synchronize the failover server database with the primary server database.

Run the `./server/bin/backup_config` utility. The backup is placed in `/tmp`. Give the backup file an intuitive name and copy it to a safe place.

- 22: Start replication on the primary Security Server.

At this stage, the failover Security Server is fully configured and ready for replication. You must manually start replication on the primary Security Server. See [“Configuring SQL replication” on page 13-222](#) for details.

- 23: For each agent, modify the URL in the `agent.conf` to point to the new server hostname.

The `agent.conf` file is located at:

(UNIX) `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/etc/agent.conf`

(UNIX) `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2/etc/agent.conf`

(Windows) `\Documents and Settings\All Users\Application Data\IBM\DB2TOOLS\LUWEncryptionExpert\Agent\vmd\etc\agent.conf`

See step 4 in [“Configuring an agent to use Security Servers” on page 6-100](#) for an example. See [“Configuring an agent to use Security Servers” on page 6-100](#) for information about configuring multiple Encryption Expert servers.

- 24: (UNIX) Restart the `vmd` process to refresh the File System Agent configuration.

a: Kill the `vmd` process if it is still running:

```
# ps -aef | grep vmd
# kill -9 procID
```

b: Start the `vmd` process.

```
# ./agent/vmd/bin/vmd
```

If you are unable to kill the process, disable the **FS Agent Locked** checkbox for the host in the Management Console.

25: (Windows) Restart the `vmd` process to refresh the File System Agent configuration.

a: Open a CMD window.

b: Enter “`net stop vmd`”.

The `vmd` service is stopping.

The `vmd` service was stopped successfully.

c: Enter “`net start vmd`”.

The `vmd` service is starting.

The `vmd` service was started successfully.

Installing the Encryption Expert agent

The File System Agent is composed of active processes that run continuously in the background. The Database Backup Agent is a library that stays inactive until it is specified on a DB2 backup command line or when restoring an Encryption Expert-generated backup. The best way to verify Database Backup Agent installation is to execute the DB2 backup command with the Encryption Expert library, and then check the agent log file, `/var/log/eet/cgsbeetdb2_usr.log`. You can also run `commtest`. The `commtest` utility is described in [“Verifying agent to server SSL authentication” on page 16-269](#). See also [“Database Backup Agent log files” on page 14-247](#).

Always install the Encryption Expert Security Server first, followed by the Encryption Expert agents, because the agent installation utility automatically performs X.509 certificate generation and exchange. The agent must be able to access the Security Server across the network. If you install the Encryption Expert agent without a working Encryption Expert server available, you must skip automatic certificate generation and then manually generate the X.509 certificates later when the server becomes available.

If you are building a fresh system, and you have neither the DB2 software nor the Encryption Expert agent software installed, installation order does not matter. You can install and configure the Encryption Expert agent software before or after installing DB2.



It is highly recommended that you add host systems to the Security Server database before you install the agent software. This will reduce the number of steps required to install and configure agents.

Installing agent software on a UNIX system

To install Encryption Expert agent software:

- 1: Check [“Before you begin” on page 3-29](#) before proceeding.
 - 2: Get the information you need to log into the system as the `root` user or as someone with `sudo` authority.
You will need this during the latter stages of installation when the agent and server exchange certificates.
 - 3: Add the host to the Security Server database, as described in [“Adding a host to the Security Server database” on page 6-77](#).
 - 4: Log onto the agent system as the administrative or root user.
-

- 5: Either insert the installation media in a CD drive or copy the software bundle onto the agent system.
- 6: Change to the directory that contains the installation utility.
- 7: (optional) Set the IATEMPDIR environment variable to point to a disk partition or storage area that has sufficient space to unpack and store the temporary installation files.

The installation utility checks /tmp for sufficient space, and then it checks \$HOME. If neither target can provide sufficient space, the installation utility will exit.

Csh and Bourne shell examples of setting IATEMPDIR are:

```
setenv IATEMPDIR /root/buildU0084_eetSolarisAgt
export IATEMPDIR=/opt/apps/b0084_solaris
```

- 8: If you are installing the software remotely, set your DISPLAY environment variable to point to an X-Windows server.

The Encryption Expert Database Backup Agent installation utility opens a graphic window in which you specify installation parameters.

NOTE: Skip this step if you want to install the software on a system that does not support X-Windows graphics. You can install the software in text-mode by including the “-i console” argument when you start the installation utility. You are asked the same questions regardless if you run the installation utility in graphic-mode or text-mode.

- 9: Start the agent installation utility.

For X-Windows graphic mode:

```
# ./install_eet_agent.bin
```

For text-mode:

```
# ./install_eet_agent.bin -i console
```

This example assumes you are installing using a graphic interface.

The *Introduction* window opens. It displays information for using the installation utility.

- 10: Click **Next**.

The *Panel: Custom Code* window opens and displays the IBM license agreement.

- 11: Read the license agreement, then click **I accept both the IBM and non-IBM terms**.

You must accept the terms of the license agreement to continue.

(Optional) Click **Print** to send the license agreement to a printer. This button opens the printer interface in which to specify print attributes such as orientation, number of copies, and paper size.

- 12: Click **Next**.

The *Choose Install Set* window opens. It allows you select the Encryption Expert agent software to install.

Your choices are DB2 Agent, IDS Agent, and File System Agent. You may select any one of the agent packages, or, you can select File System Agent and either DB2 Agent or IDS Agent. You cannot install both DB2 Agent and IDS Agent on the same system.

- 13: Click **Next**.

The *Pre-Installation Summary* window opens. It displays the name of the products being installed, the target destination, the required amount of disk space, and the available amount of disk space.

14: Click **Install**.

The Encryption Expert agent packages are unpacked and the agent configurations are prepared. The *Server Hostname* window opens.



Solaris agent installation is done at this point. The agents are registered but the SSL certificates have not been exchanged. A reboot is required after installing agent software on a Solaris system. Click **Done**, reboot the Solaris system, and then manually register the agents using their respective utilities:

```
./agent/vmd/bin/register_host
./agent/db2/bin/setup_eet_agent_cert_db2
./agent/ids/bin/setup_eet_agent_cert_ids
```

15: Enter the host name or FQDN of the Security Server that is to administer this agent.

16: Click **Next**.

If you are installing the File System Agent, a dialog box opens that prompts you to accept the server CA certificate.

If you are installing the Database Backup Agent, a dialog box opens that prompts you to accept the server CA certificate.

The fingerprint displayed in both dialog boxes should be identical to the fingerprint displayed in the Management Console *Dashboard* window.

17: Click **OK** in each dialog box if the fingerprint displayed matches that of the Security Server.

The agents are registered with the Security Server when the certificates are exchanged. Then the *Install Complete* window opens. The window tells you the success or failure of the installation process.

18: Click **Done**.

19: Log into the Management Console as a user with *Server* role permissions.

20: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens.

21: Click the host you just upgraded in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

22: Enable the **Communication Enabled** checkboxes for the type of agents you just installed.

Your choices are **FS**, **IDS**, and **DB2**.

23: Click **Ok**.

Installing agent software on a Windows system

Windows installation is a two-stage process. The first stage comprises installing the agent software and then rebooting the Windows system. The second stage is the generation and exchange of Security Server and agent certificates.

To install the Windows agent software:

- 1: Check “Before you begin” on page 3-29 before proceeding.
- 2: Get the information you need to log into the Security Server system as the `root` user or as someone with `sudo` authority.
 You will need this during the latter stages of installation when the agent and server exchange certificates.
- 3: Add the host to the Security Server database, as described in “Adding a host to the Security Server database” on page 6-77.
- 4: Log onto the agent system as an administrative user.
- 5: If you are copying Encryption Expert agent software across the network, copy the `install_eet_agent.exe` executable onto the Windows system.
- 6: If the Encryption Expert agent software is provided on external media, such as a CD-ROM or DVD, insert the media in a local drive and locate the installation software in a Windows Explorer window.
- 7: Double-click the `install_eet_agent.exe` executable to begin Windows agent software installation.

The *Welcome* window opens.

- 8: Click **Next**.

The *Choose Installation Folder* window opens. This window prompts you for the path in which to install the Windows agent software. The default path is `C:\Program Files\IBM\DBTOOLS\LUWEncryptionExpert`.




Figure 3.1: Specifying the agent installation directory

- 9: Click the **Choose...** button to locate and select the desired path.
 You can click the **Restore Default Folder** button to discard path changes and restore the path that was displayed when the window was initially opened.
- 10: Click **Next**.
 The *Pre-Installation Summary* window opens. It displays the name of the products being installed, the target destination, the required amount of disk space, and the available amount of disk space.
- 11: Click **Install**.

If the path is incorrect, click **Previous** to return to the previous window and change the path.

After clicking **Install**, the installation utility begins the agent installation process. When finished, the *Installation Complete* window opens. The window prompts you to reboot the Windows system in order to start the Windows agent. The **Yes, restart my system** radio button is enabled by default.

12: Click Done.

The Windows system is rebooted. After the system reboots, and you log back into the Windows system, the *Register Host* window opens and you are prompted to complete the second part of the Windows agent installation process.

13: After you log back into the Windows system, register the host, as described in [“register_host” on page 16-284](#).

Skip the preliminary steps that describe how to start `register_host`, and advance to the steps that describe how to complete the *Register Host* window.

What next?

After installing the agent software, you can:

- Inspect the installation hierarchy in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent` or `\Program Files\IBM\DBTOOLS\LUWEncryptionExpert\agent` and `\Documents and Settings\All Users\Application Data\IBM\DBTOOLS\LUWEncryptionExpert\agent`.
- Read `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2/README` for usage and additional instructions.
- Start a Security Server Web session in which to configure and apply policies to the new agent. See also [“Starting the Management Console” on page 4-57](#).
- Verify that the agent system can communicate successfully with its Security Server. See [“Verifying agent to server SSL authentication” on page 16-269](#).
- Verify the Database Backup Agent installation by making a backup with the agent library file. On UNIX systems, check the agent log file, `/var/log/eet/eetdb2_usr.log` after making the backup. See also [“Backing up a DB2 database” on page 11-197](#), [“Backing up and restoring IDS databases” on page 12-205](#), and [“Database Backup Agent log files” on page 14-247](#).

Database Backup Agent file structure (UNIX)

The DB2 and IDS Database Backup Agent directories and files are placed in:

```
/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2
/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/ids
```

The `./agent/db2` and `./agent/ids` directories are almost identical. The examples in this guide use the `./agent/db2` path but, if you are installing Encryption Expert for IDS, replace `db2` with `ids` in the path specification and the result should be the same. Note that the output file names also change accordingly.

These installation directories contain:

```
./agent/db2/bin
```

- Directory that contains the executable utilities and binaries to configure and administer the agent installation.
-

- (IDS only) The `./agent/ids/bin` directory contains the IDS executable that is must be used to apply Encryption Expert protection to IDS backup and restore operations, `eetidsagt`.

`./agent/db2/etc`

- Directory that contains agent configuration data, such as the `access` file, `agent.conf` file, and error code-to-message tables. The `access` file is created during installation and contains an encrypted master passphrase that is used to decrypt other passphrases. Do not edit or delete the `access` file. The `agent.conf` file specifies the Security Servers the agent may communicate with and their listening port numbers. (Port numbers must be set to 8443.) The tables associate internal codes with meaningful log entries and error messages.

`./agent/db2/lib`

- Directory that contains the shared library object that is used to make DB2 backups.
- (IDS only) The `./agent/ids/lib` directory is empty because the library is embedded into the `eet` executable.

`./agent/db2/log`

- Directory that contains the agent log file. This directory is symbolically linked to `/var/log/eet`.

`./agent/db2/pem`

- Directory that contains Privacy Enhanced Mail (pem) certificates, certificate requests, signer certificates, and more. These files control SSL authentication between the agent and associated Encryption Expert servers.

`./agent/db2/README`

- A valuable quickstart document that describes the installed files in more detail, and provides examples on how to execute DB2 or IDS.

`/var/log/eet/`

- Directory that contains agent log files in the form of `eetdb2_usr.log.date.num`. See [“Database Backup Agent log files” on page 14-247](#).

File System Agent file structure (UNIX)

The File System Agent directories and files are placed in:

```
/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs
/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd
```

A file system installation comprises a `secfs` component and a `vmd` component. The `secfs` component manages the Encryption Expert file system on the agent host. Encryption Expert uses a proprietary file system and mounts guard points with the `secfs` file type. (You can see active guard points and their file type by running the `mount` or `df` command on the agent host.) The `vmd` component manages the host configuration at the kernel level. It manages SSL certificates, checks host configuration, registers the host with the Security Server, and it also identifies the Security Servers that the agent can access.

The `./agent/secfs` directory contains:

`./tmp:`

- This directory contains the Encryption Expert secure file system log, `secfsd.log`. This file is a record of when the guard points are mounted and unmounted. If the guard point is not mounted as a `secfs`

file system, protection is not applied. On AIX systems, this directory also contains the `secfs.log` and `secfs.log.archive` files, and a `log_archive` directory that contains older `secfsd.log` files.

`./sec:`

- This directory contains the agent configuration. The files and subdirectories contain files that can report the guard status, guard point, mount point, and more. The `./sec/current/bin` directory contains the executables used to mount the guard point, run `dataxform`, shared library objects, and more.

`./bin:`

- This directory contains the utility to remove Encryption Expert agent software from the system. The utility name is `uninstallsfcs`. Run this utility to delete the complete File System Agent installation. On AIX systems, there are also `vmmount` and `vmumount`. These are Encryption Expert mount and unmount utilities that you should not run.

The `./agent/vmd` directory contains:

`./bin:`

- This directory contains utilities to check host information, query `secfs` status, self-register the host, etc. It also contains the executable that starts the File System Agent process, `vmd`.

`./etc:`

- Directory that contains agent configuration data, such as the `access` file, `agent.conf`, and error code-to-message tables.

`./lib:`

- This directory is empty.

`./license:`

- Directory that contains the Encryption Expert licenses in various languages.

`./log:`

- Directory that contains the File System Agent log file. This directory is symbolically linked to `/var/log/eet`.

`./pem:`

- Directory that contains Privacy Enhanced Mail (pem) certificates, certificate requests, signer certificates, and more. These files control SSL authentication between the Encryption Expert agent and Security Server.

Upgrading Encryption Expert

Upgrading Encryption Expert Security Server software

The installation utility automatically initiates the Security Server upgrade feature if the utility detects an existing Encryption Expert 1.1.2 Security Server. The software upgrade process requires less user input than a full installation because it uses the current Security Server configuration. It also leaves the Security Server database intact so you do not have to recreate or restore keys, certificates, policies, host configurations, etc.

If you are upgrading a high availability Security Server configuration, there is no need to disable SQL-replication on the primary Security Server or disable the agents that access

the primary Security Server because the agents will automatically switch to a failover Security Server during the upgrade process.

The following approach is used to upgrade from Security Server version 1.1.2 to version 1.1.3. See [“import_v1120_config” on page 16-273](#) if you are upgrading the Encryption Expert Security Server from version 1.1.1.



If you upgrade the Encryption Expert Security Server from version 1.1.1 to 1.1.2, all configured users are assigned a new temporary password. The temporary password is “welcome”. Each user must enter a new password immediately after starting a Management Console session.

To upgrade the Security Server:

- 1: Log onto the system as the root user.
- 2: Back up the primary Security Server installation using the `backup_config` utility.

There is no need to backup a failover server installation.

See [“Backing up the Security Server configuration” on page 16-295](#) for details.

- 3: Start the installation utility, as described in [“Installing the primary Security Server” on page 3-33](#).

The *Introduction* window opens.

- 4: Click **Next**.

The installation utility detects the current Security Server installation and prompts you to update it.




Figure 3.2: Prompt to upgrade the current Security Server installation

- 5: Click **Ok**.

The *Panel Custom Code* window opens. It displays the IBM license agreement.

- 6: Read the license agreement, then select the **I accept both the IBM and the non-IBM terms**.

You cannot proceed if you do not accept the terms of the license agreement.

(Optional) Click **Print** to send the license agreement to a printer. This button opens the printer interface in which to specify print attributes such as orientation, number of copies, and paper size.

(Optional) Click **Read non-IBM terms** to open a window to view the license agreements for third-party products such as HP-UX SDK and OpenSSL.

- 7: Click **Install**.

The installation utility unpacks and installs the Security Server software. It starts the Security Server after the software is installed.

The *Installation Complete* window opens.

- 8: Click **Done**.

The installation utility closes.

- 9: Check that the Security Server has “settled down” before you run Security Server applications.

After the installation utility completes, Java and DB2 processes continue to run to complete the embedded database and Security Server configuration. Run a command like `top` or `slabtop` to monitor system activity. You can open the Management Console and administer the Security Server after the Java and DB2 processes become inactive. It will take a few minutes.

- 10: Start a Management Console session by entering the URL and password that were returned by the installation utility in a Web browser.



Encryption Expert 1.1.3 provides more stringent user password controls than Encryption Expert 1.1.2. The first time any user logs into the Management Console, the user password automatically enters an “expired” state and must be reset. See [“Logging in with an expired user account” on page 5-72](#) for details.

- 11: Update the failover Security Server installation in the same manner as the primary server.

The failover server upgrade process is identical to that of the primary server.

- 12: (Optional) Add additional failover servers.

Encryption Expert 1.1.2 supports only one failover server. Encryption Expert 1.1.3 supports multiple failover servers. If you want to add additional failover servers:

- a: Add the additional failover servers to the Management Console *High Availability* window of the Encryption Expert 1.1.3 primary Security Server.
- b: If not already installed, install Encryption Expert 1.1.2 Security Server software on the systems designated as failover servers.
- c: Install the 1.1.3 Security Server software on the 1.1.2 failover servers.

- 13: After the failover servers are updated, enable SQL replication on the primary server using the `./server/bin/configure_replication` utility.

See [“Configuring SQL replication” on page 13-222](#) for details.

Upgrading Encryption Expert agent software

As new software releases become available, you should upgrade the current installation to take full advantage of the latest features and new functionality. You can upgrade the Security Server installation and leave the currently configured keys, hosts, and policies intact.

Upgrading UNIX agents

To upgrade current File System Agent and Database Backup Agent installations:

- 1: If you are also upgrading the Security Server installation, upgrade it now.

It is easier to upgrade the agents after you upgrade the server. See also [“Upgrading Encryption Expert Security Server software” on page 3-49](#).

- 2: Verify that the **Registration Allowed** checkboxes are enabled for the agents to be upgraded.
-

Existing agent certificates and fingerprints are left intact during the Security Server upgrade process, so you should not have to re-enable the **Registration Allowed** checkboxes. If the **Registration Allowed** checkbox is disabled, the agent authentication credentials have been removed from the server and new credentials need to be generated during the agent upgrade process.

- 3: Log onto the agent system as the root user.
- 4: Either insert the installation media in a CD drive or copy the software bundle onto the agent system.
- 5: Change to the directory that contains the installation utility.
- 6: Start the agent installation utility.

For X-Windows graphic mode:

```
# ./install_eet_agent.bin
```

For text-mode:

```
# ./install_eet_agent.bin -i console
```

This example assumes you are installing using the graphic installation utility.

The *Introduction* window opens. It displays information for using the installation utility.

- 7: Click **Next**.

The installation utility detects existing Encryption Expert installations and prompts you to upgrade.




Figure 3.3: Prompt to upgrade the current agent installations

- 8: Click **Yes**.

If you click **No**, the installation utility aborts. If you want completely new agent software, you must delete the current agent installations and then restart the installation utility.

- 9: Click **Next**.

The *Panel: Custom Code* window opens and displays the IBM license agreement.

- 10: Read the license agreement, then click **I accept both the IBM and non-IBM terms**.

You must accept the terms of the license agreement to continue.

(Optional) Click **Print** to send the license agreement to a printer. This button opens the printer interface in which to specify print attributes such as orientation, number of copies, and paper size.

- 11: Click **Next**.

The *Choose Install Set* window opens. It allows you specify the Encryption Expert agent software that you want to upgrade.

12: Enable the checkboxes for the agents you want to upgrade.

Your choices are File System Agent and either DB2 Agent or IDS Agent.

13: Click **Next**.

The *Pre-Installation Summary* window opens. It displays the name of the products being installed, the target destination, the required amount of disk space, and the available amount of disk space.

14: Click **Install**.

The *Server Hostname* window opens.

15: Enter the host name or FQDN of the Security Server that is to administer this agent.

16: Click **Next**.

If you are upgrading the File System Agent, a dialog box opens that prompts you to accept the server CA certificate.

If you are upgrading a Database Backup Agent, a dialog box opens that prompts you to accept the server CA certificate.

The fingerprint displayed in both dialog boxes should be identical to the fingerprint displayed in the Management Console *Dashboard* window.

If you are not prompted to accept certificates, run the certificate generation utilities manually:

a: Remove the agent certificates.

The certificate removal utilities are:

```
(FS) # ./vmd/bin/register_host clean
(DB2) # ./db2/bin/setup_eet_agent_cert_db2 clean
(IDS) # ./ids/bin/setup_eet_agent_cert_ids clean
```

b: Regenerate agent certificates and reregister the agent with the Security Server.

The certificate generation and reregistration utilities are:

```
(FS) # ./vmd/bin/register_host
(DB2) # ./db2/bin/setup_eet_agent_cert_db2
(IDS) # ./ids/bin/setup_eet_agent_cert_ids
```

17: Click **OK** in each dialog box if the fingerprint displayed matches that of the Security Server.

The agents are registered with the Security Server when the certificates are exchanged. Then the *Install Complete* window opens. The window tells you the success or failure of the installation process.

18: Click **Done**.

19: Log into the Management Console as a user with `Server` role permissions.

20: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens.

21: Click the host you just upgraded in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

22: Enable the **Communication Enabled** checkboxes for **FS**, **IDS**, and **DB2**.

23: Click **Ok**.

Removing Encryption Expert

Use **Add and Remove Programs** to delete agent software from Windows systems.

Removing the Encryption Expert Security Server

You can remove the Security Server installation from a system. Removing the Security Server software also removes the Security Server database that contains the encryption keys required to restore encrypted backups. Once the encryption keys are gone you cannot decrypt backups made with those keys. Back up the Security Server database, or replicate the database on another Security Server, before removing the Security Server software, otherwise you will be unable to decrypt existing backups.

To remove the Security Server and related components:

- 1: Log onto the Encryption Expert Security Server as the `root` user.
- 2: Run the software removal utility:

```
# /opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/bin/
deinstall_eet_server
```

The utility prompts you to verify that you really want to remove the Security Server installation.

- 3: Enter yes when the utility prompts you to continue:

```
Do you wish to proceed? (yes|no) [no]: yes
```

The utility progresses without any further input. It stops the embedded database, stops the JBoss Web application server, and then removes all related components, including the Security Server prerequisite package that you installed separately, the administrative user accounts that were placed in `/etc/passwd`, and the user groups placed that were placed in `/etc/group`.

Details about the Security Server removal process are placed in `/tmp/db2_deinstall.log.num` and `/tmp/deinstall_eet_server.log`.

Removing the Encryption Expert Database Backup Agent

Authentication certificates are deleted along with the Encryption Expert agent software. Back up the Encryption Expert `/agent/pem` and `/agent/etc` directories if you plan to reinstall the Encryption Expert agent software on the same system at a later time. Otherwise, there is no need to back up the agent installation before removing it. You can simply reinstall the software and reconfigure the agent host in the Security Server. You can reuse the backup certificates if the agent and server system names and network addresses are not changed.

To delete the Encryption Expert Database Backup Agent from a host:

- 1: Log onto a host that is running the Encryption Expert Database Backup Agent as the `root` user.

(Windows) Log onto the Windows host system as an administrative user and run the Windows “Add or Remove Programs” utility to delete Encryption Expert software.

- 2: Either change to the Encryption Expert agent `bin` directory and enter the software removal utility, or enter the full path to the software removal utility.

```
# /opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2/bin/
deinstall_eet_agent_db2
Checking user permissions.
```

```

Encryption Expert Agent for DB2 will be uninstalled.
Do you wish to proceed? (yes|no) [no]:

3: Enter yes to remove the Encryption Expert agent software.
Do you wish to proceed? (yes|no) [no]: yes

Encryption Expert Agent for DB2 has been uninstalled
successfully.
#

```

Removing the Encryption Expert File System Agent

To delete the Encryption Expert File System Agent from a host:

- 1: Log into the Management Console as a Security Administrator with **Host** role permissions.
 - 2: Open the *Edit Host* window for the agent to be deleted.
 - 3: Disable **System Locked** and **FS Agent Locked** in the *General* tab.
Do not remove a host record from the Security Server database before you remove the File System Agent software from that host. If you delete the host record first, and **System Locked** and **FS Agent Locked** are enabled at the time of deletion, you must manually disable these features on the host before you can delete the agent software and regain access to guard points.
 - 4: Evaluate the current guard points in the *Guard FS* tab.
Consider the encrypted data in each guard point. Encrypted data will remain encrypted after you remove the File System Agent software from a host. You may want to copy the encrypted data out of each guard point, with an appropriate policy, so that the data is decrypted during the copy process. Or, you may want to run `dataxform`.
 - 5: Disable the **Registration Allowed** checkbox for the File System Agent in the *General* tab.
 - 6: Click **Apply**.
A dialog box will open prompting you to verify the change.
 - 7: Click **Ok**.
 - 8: Log onto the host that is running the Encryption Expert File System Agent as the root user.
(Windows) Log onto the Windows host system as an administrative user and run the Windows “Add or Remove Programs” utility to delete Encryption Expert software.
 - 9: Execute the software removal utility.
For example:

```
# cd /opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent
# ./secfs/bin/uninstallsfs
```

The utility stops associated agent processes and removes the secure file system (secfs) portion of the agent software.

```
# secfs/bin/uninstallsfs
Do you really want to uninstall Encryption Expert? (yes|no)
```
 - 10: Enter yes when prompted to continue.
Uninstalling Encryption Expert
-

```
Stopping kernel module secfs2
secfs2 is loaded, unloading...
Successfully stopped kernel module
Successfully uninstalled the kernel space component of
the Encryption Expert
Successfully removed
/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd
Successfully removed /usr/bin/vmd
Successfully removed /usr/bin/vmsec
Successfully uninstalled the user space component
of the Encryption Expert.
#
```

- 11: Verify the removal process by attempting to list (ls) the Security Server installation directory.

It should be gone entirely.

- 12: Verify that the vmd process has been stopped by listing the active processes and piping the output to the grep command.

This should return just the grep command. For example:

```
# ps -aef | grep vmd
root      1107 17932  0 19:49 pts/2 00:00:00 grep vmd
#
```

- 13: If the host is not running other Encryption Expert agents, you can delete the host record from the *Hosts* window.
-

Starting the Management Console

The Encryption Expert includes a Web-based interface that you use to create policies, configure hosts, and assign keys. This is referred to as the “Management Console”. The Web application provides a secure connection between the Encryption Expert Security Server and the host administering that Security Server. Connection security is provided through SSL authentication. Use the following Web browsers to access the Web interface:

- Mozilla Firefox, version 1.8 or higher
- Windows Internet Explorer, version 6.0 or 7.0

The system on which you run the Web browser must also run Java 1.5, or Java 1.6 with Update 2.

After opening your Web browser, ensure that the Transport Layer Security (TLS) protocol, version 1.0, is enabled. TLS 1.0 is used by the Web browser to communicate with the JBoss Web application server that drives the Management Console. Also, ensure that SSL 3.0 is enabled.

Internet Explorer 7 (IE7) users: The Policy Composer is a Java applet. You must configure your Web browser to accept the Security Server as a trusted site in order to run this applet.



Only one person should use the Security Server configuration interface at a time. Multiple persons working simultaneously on the same Security Server can potentially cross-edit configurations, overwriting the changes made by another administrator. The changes of the administrator who last saves his or her work are used.



Enable Java, JavaScript, and cookies in your Web browser before starting a Security Server Web session.

To start the Management Console:

- 1: Open a Web browser.

You can open the Web browser on the same system that is hosting the Security Server or remotely from any system that has network access to the system that is hosting the Security Server.

- 2: Enter the URL and port number for the Security Server Web interface.

`https://hostname:8445`

Replace *hostname* with the hostname or FQDN of the system hosting the Security Server.

- 3: If you cannot open the Management Console, check that the JBoss and embedded DB2 server are running.

Execute the `ps` command and `grep` for “voradmin”. You should see JBoss and DB2 processes running. In the excerpt below, processes 998 and 1004 indicate that JBoss is running. The other processes indicate that DB2 is running.

```
# ps -aef | grep voradmin
...
voradmin 32444 32443 0 Feb12 pts/1 00:00:01 db2sysc
voradmin 32449 32444 0 Feb12 pts/1 00:00:01 db2gds
voradmin 32450 32444 0 Feb12 pts/1 00:00:00 db2licc
voradmin 32582 32449 0 Feb12 pts/1 00:00:16 db2stmm (JMSDB)
voradmin 998 1 0 Feb12 pts/1 00:00:00 /bin/sh /opt/IBM/
DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/bin/
run.sh
voradmin 1004 998 1 Feb12 pts/1 00:19:54 /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/server/ibm-java2-i386-50/jre/bin/java
-server -Xms256m -Xmx512m -
Dsun.rmi.dgc.client.gcInterval=3600000 -
Dsun.rmi.dgc.server.gcInterval=3600000 -DCGSS_HOME=/opt/
IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/
server/default -Dprogram.name=run.sh -
Djava.endorsed.dirs=/opt/IBM/DB2TOOLS/LUWEncryptionExpert/
server/jboss-4.0.4.GA/lib/endorsed -classpath /opt/IBM/
DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/bin/
run.jar:/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/ibm-
java2-i386-50/jre/lib/tools.jar org.jboss.Main
voradmin 1476 32451 0 Feb12 pts/1 00:00:00 db2agent (JMSDB)
voradmin 1529 32449 0 Feb12 pts/1 00:00:01 db2dlock (LOGDB)
voradmin 32321 32449 0 Feb12 pts/1 00:00:31 db2agent (idle)
...
#
```

If you still can't open the Management Console, restart the Security Server using the following command line:

```
/etc/init.d/cgss restart
```

- 4: If you are starting a new Management Console session, the *Security Alert* window may open. Click **Yes** to place a temporary security certificate on the local system or click **View Certificate->Install Certificate...** and then complete the Certificate Import Wizard to install the certificate permanently. The *Security Alert* window is not displayed again if you install the certificate permanently.

If you do not install the certificate, you must click **Yes** to initiate a Web session with the Security Server Web server. An X509 digital certificate is issued by Encryption Expert to provide secure data transmission between the Web browser and the Security Server. The certificate is signed using the Encryption Expert private key. The security certificate is required to start the Management Console session. A temporary security certificate is only good for the current Web browser session and you will be prompted to install the certificate each time you start the Management Console.



Figure 4.1: Security Alert window

The Encryption Expert opens the *Management Console Login* window.

- 5: If this is the first time the Security Server Web interface has been started, login as the default administrative user, `admin`. If other administrative users have been added to the Security Server, login as one of those administrative users.

Enter the user name and password in the **User Name** and **Password** text-entry boxes. The default user name and password are `admin` and `admin123`.




Figure 4.2: Management Console login window

If after the configured number of attempts to login you do not successfully enter valid user credentials, the Management Console login window displays: “You have failed to login after *X* attempts. Login is disabled for *Y* minutes. Wait *Y* minutes before trying again.” The Management Console becomes inoperable and ignores further login attempts for the specified number of minutes. The default number of allowable tries is 3 and the default timeout period is 30 minutes.

- 6: If one of the following is true:
 - this is the first time the user is logging into the Management Console
 - an administrator changed the user password
 - the password has expired
 you must reset the password. See [“Logging in with an expired user account” on page 5-72](#).
- 7: If you are starting a new Management Console session in a fresh Internet Explorer window, and the signer certificate for the previous Management Console session is no longer in the system cache, you may be prompted to accept a signer certificate that is used for transmitting encrypted data between the host and Security Server.



Figure 4.3: Accepting the signer certificate

Click **Yes** to install a temporary signer certificate that is good only for the current Web browser session or enable the **Always trust content from this publisher** to install a permanent certificate.

The Encryption Expert *Management Console Login* window opens.



Figure 4.4: Management Console Dashboard window

- 8: If you are starting a remote Security Server session, check the Web browser to ensure that you are working across a secure connection.
Check the bottom of the browser for a closed lock.



Figure 4.5: Checking for a secure connection (Explorer, left - Firefox, right)
You can double-click the lock icon to display the certificate being used.

Navigating the Management Console Interface

The *Dashboard* window is the starting point for configuring the Security Server and creating policies.

You can exit the Management Console session at any time by clicking the **Log Out** text that is located in the top-right corner of the Management Console banner. For improved security, log out when you are done rather than relying upon the Management Console to timeout your login session.



Figure 4.6: Management Console banner

The Management Console interface is divided into two major parts, a menu bar and a work area. The menu bar displays the main features available to apply Security Server protection. Select a feature in the menu bar to work on that feature. The contents of the work area change relative to the selected feature in the menu bar. For example, select **Policies** in the menu bar to display configured policies in the work area, and click a specific policy to configure the rules for that policy.




Figure 4.7: The Management Console graphic interface

Select a feature in the menu bar to open a window for configuring the selected item. For example, if you select **Configuration->Preferences** in the menu bar, the *Preference* window opens. In it you can configure preferences such as the timeout interval and the number of policies to display at one time.



The options displayed in this and other windows are determined by your Security Server user login. The graphics in this document were captured by an administrative user that is configured with all administrative privileges so all the available options are displayed. If you do not see the features that you think you should be seeing, check your user login.

Using the Encryption Expert tray icon

The Encryption Expert tray icon is used to view agent configuration and status. The icon is  , and is located in the Windows taskbar. The tray icon allows you to:

- View agent logs, fingerprints, and certificates
- Enter the host password when the Security Server is not available
- Display agent installation status
- Register agents with the Security Server

Unavailable options in the Encryption Expert tray are not displayed or they are displayed in opaque form and cannot be selected.




Figure 4.8: Parts of the Encryption Expert tray

Displaying agent status

Check the agent status to display installed agent software, applied guard points, policies, host settings, and more. Status information can be displayed by using the Encryption Expert tray on Windows systems or by entering the `vmsec status` command in a UNIX terminal window.

To display the Windows agent status, right-click the  tray icon and select **Status**. An example status window is displayed in [Figure 4.8](#).

Entering the host password

A password must be provided when there is no connection between the File System Agent and the Security Server before you can decrypt/encrypt files that are encrypted with an offline key (Cached on Host). The offline password feature controls access to the encryption keys that are stored on the local host, so as to allow encryption/decryption when the Security Server is not accessible. Regardless if encryption is applied or not, the policy is still in force and if the policy does not permit access to the guard point, you will be unable to access it.

The host password is set when the host is added to the Security Server database. Passwords are set on a host-by-host basis, rather than on a host group basis. For information about setting the initial host password, see [“Adding a host to the Security Server database” on page 6-77](#).

You can also enter the host password by executing `\Program Files\IBM\DBTOOLS\LUWEncryptionExpert\agent\vm\bin\vmsec.exe passwd` in a CMD or UNIX window.

To enter the host password when there is no Security Server connection:

- 1: Select **Right-click icon->Password...**
The Encryption Expert *Offline Password* window opens.
- 2: Enter the password in the text-entry box.



Figure 4.9: Entering the password for an offline host

- 3: Click **Ok**.
You should be able to decrypt/encrypt protected data now.

Displaying the agent log

There are two ways to view the agent log:

- You can use the tray icon to open the log file, `vorvmd.log`, in an editor. The default is Notepad. This is a static display that is good for viewing past events.
Select the following to display the agent log, `vorvmd.log`:
 - Right-click tray icon->**View->File System->Log**
 - In a terminal window, you can

```
tail -f \Documents and Settings\All Users\Application
```
-


```
Data\IBM\DBTOOLS\LUWEncryptionExpert\agent\log\etvmd.log
```

This is a dynamically updated display that is good for viewing events as they occur.

An excerpt of the `etvmd.log` file is shown below:

```
2008-06-19 19:14:32.687 [COM] [DEBUG] [2604] SDK build
date: 2008-06-18 16:40:33 -0700 (PDT)
2008-06-19 19:14:32.906 [POL] [INFO ] [2604] [POL4553I]
Internal initialization successful.
2008-06-19 19:14:32.906 [COM] [DEBUG] [2604] Initialize
Agent VMD succeeded
2008-06-19 19:14:35.281 [COM] [DEBUG] [2604] Server
communication established
2008-06-19 19:14:36.281 [VMD] [INFO ] [2604] [VMD3640I]
STARTED on Windows Server 2003, Standard Edition PORT
7024 VERSION: 4.2.0.0
2008-06-20 10:45:50.046 [VMD] [INFO ] [2604] [VMD3794I]
Query security config succeeded.
2008-06-20 10:45:59.656 [VMD] [INFO ] [2604] [VMD3794I]
Query security config succeeded.
```

Displaying the agent and Security Server fingerprints

The tray icon is used to display the respective agent and Security Server authentication fingerprints. Verify that the Security Server fingerprint displayed matches the fingerprint displayed in the Management Console *Dashboard* window. Also verify that the agent fingerprint displayed matches the fingerprint in the *Edit Host* window, *General* tab.

Select the following to display the agent and Security Server fingerprints:

- Right-click tray icon->**View->File System->Fingerprints**

A set of File System Agent and Security Server fingerprints is shown below.




Figure 4.10: Viewing File System Agent and Security Server fingerprints

Displaying the agent certificate

There are two ways to view the agent certificate:

- You can use the tray icon to display the certificate.
Select the following to display the agent certificate:
 - Right-click tray icon->**View->File System->Certificate**
- In a terminal window, you can change to the `\Documents and Settings\All Users\Application`

```
Data\IBM\DBTOOLS\LUWEncryptionExpert\DataSecurityExpert\
agent\vmd\pem directory and execute:
openssl x509 -in agent-cert.pem -noout -text
```

An example File System Agent certificate is shown below:

Certificate:

```
Data:
  Version: 3 (0x2)
  Serial Number:
    0c:82:c8:de:4c:d9:f8:67:73:3f:ff:d2:c2
  Signature Algorithm: sha1WithRSAEncryption
  Issuer: CN=CG CA (S) on vmlinux01, OU=Really Fine Stuff,
  O=Vormetric, Inc., L=Santa Clara, ST=CA, C=US
  Validity
    Not Before: Jun 19 02:13:32 2008 GMT
    Not After : Jun 19 02:13:32 2012 GMT
  Subject: C=US, OU=This agent certificate was automatically
  generated, CN=CGSB FS_VMD Agent on vmwindows110
  Subject Public Key Info:
    Public Key Algorithm: rsaEncryption
    RSA Public Key: (1024 bit)
      Modulus (1024 bit):
        00:95:25:67:0a:e3:02:5f:c6:8c:4e:3a:e4:e6:a8:
        ...
        0c:0f:16:6a:bf:8f:a3:91:d6:e5:fa:2a:9c:36:93:
        da:3e:ce:79:92:fe:09:02:2f
      Exponent: 65537 (0x10001)
  X509v3 extensions:
    X509v3 Extended Key Usage:
      TLS Web Client Authentication
    X509v3 Basic Constraints: critical CA:FALSE
    X509v3 Key Usage: critical Digital Signature,
    Key Encipherment
    X509v3 Subject Key Identifier:
      F9:5F:0E:66:8E:D7:29:83:02:AA:B2:A1:23:73:6C:6D:08:8B
    ...
    1.3.9999.0.0.0.1:
      FS_VMD
  Signature Algorithm: sha1WithRSAEncryption
    4a:69:85:8f:f5:93:38:0f:6e:9b:8e:c2:f6:62:31:c7:0c:d2:
    ...
    f0:b0:a1:01:ce:94:f5:7f:bf:30:23:ad:0b:aa:a9:9a:3f:c1:
    67:c7:6b:34
```

Stopping and restarting the Encryption Expert tray icon

To stop the tray icon process and close the tray icon, right-click tray icon->**Exit**.

To restart the tray icon, perform one of the following:

- In Windows Explorer, double-click \Program Files\IBM\DBTOOLS\LUWEncryptionExpert\agent\vmd\bin\etray.exe.
 - In a UNIX window, enter the full path to the etray executable, followed by an ampersand so that the process runs in the back ground. For example:
-

```
# \Program  
Files\IBM\DBTOOLS\LUWEncryptionExpert\agent\vmd\bin\etr  
ay.exe &
```


Administering users

Plan and create users who are to perform specific tasks in the management of Encryption Expert servers and agents. Users can be configured with roles that are limited to just one aspect of the product, such as policy configuration or encryption key management. Users can also be configured with a combination of roles.

The name of the current user is displayed in the Management Console *Dashboard* window. If the current user is also configured with the `Admin` role, the user can open the *Administrators* window, in which are listed all the users on the Security Server and their roles.

Users are configured in the *Administrators* window. This section describes user and password administration. The policy for setting user passwords is configured in the *Preferences* window. See [“Setting password preferences” on page 15-251](#) for information about configuring the default password policy.

Configuring users

Several system users are created during Security Server installation. These are `voradmin`, `vorapply`, `db2fenc1`, and `dasusr1`. These are administrative accounts for the embedded DB2 database. They each have a `/home` directory and, except for `vorapply`, they each have an entry in `/etc/passwd`. Do not remove or modify these users.

Security Server users are stored in the embedded database, CGSSDB; and, the database is stored in `/home/voradmin`. If you remove the `voradmin` user you will also remove the embedded database, including all the Security Server users, and the Security Server will fail to operate.

The default Security Server administrative user that is created during server installation is `admin` and the default password is `admin123`. This user is assigned one role, `Admin`, and the `Admin` role can only be used to create other users. Log into the Management Console as this administrative user the first time you start a Management Console Web session. Change the `admin` password and then create other administrative users.

A Security Server administrative user can be configured with one or more roles. The roles are:

Role	Description
Admin	The admin role can create and modify Security Server user accounts. Only the admin user can display and configure user accounts. In addition to creating user accounts, the admin user can change other administrator passwords. It can also view log data and set Encryption Expert preferences. Though the admin user can change and access very little in the Management Console, it can create users with complete access permissions. A master administrator, which is all roles enabled in one user account, can see and do everything. No feature or command is off-limits to a master administrator.
Audit	The audit role can only view log data. See “Managing Encryption Expert logs” on page 14-229 for information about configuring server and agent logs.
Key	The key role can create, edit, and delete local key-pairs, public keys only, and key groups. It can also view log data. See “Using Encryption Keys” on page 7-115 for information about creating and managing keys.
Policy	The policy role can create, edit, and delete policies. It can also view log data. See “Creating and editing File System Agent policies” on page 9-141 for information about configuring policies.
Host	The host role can configure, modify, and delete hosts and host groups. It can also view log data. See “Configuring hosts” on page 6-75 and “Configuring host groups” on page 6-108 for information about configuring hosts and host groups.

User passwords are a vital part of a good security system. A Management Console user password, at a minimum, can contain standard ASCII alphabet characters (i.e., a-z, A-Z) and integers (i.e., 0-9). Blank spaces are also permitted. The individual elements in this combination of characters, integers, and blanks cannot occur in adjacent sequence. That is, a password cannot contain two instances of the same element if they are next to each other. For example, `mississippi` will not be accepted, but `misSiSipPi` will.

Passwords can comprise from *X* to 256 characters. *X* is set in the **Minimum Password Length** text-entry box of the *Password* tab.

Additional restraints can be applied that require all new passwords to contain at least one uppercase alphabet character, at least one special character, and the minimum number of characters that must be used.

The following are a few notes and conditions for using passwords:

- Password requirements are configurable. You can set password parameters such as the number of login attempts to permit, the type of characters required in a password,

and password expiration controls. See [“Setting password preferences” on page 15-251](#) for details.

- Users must provide their current password in order to change their password. Only users that are assigned the `admin` role can change the passwords of other users.
- If a user forgets their password, the administrator can reset the password. When a password is changed, the password enters an “expired” state. The user must enter a new password the next time s/he starts a Management Console session.
- Password changes and the users performing those changes are audited and logged.
- If you upgrade the Encryption Expert Security Server from version 1.1.1 to 1.1.2, all configured users are assigned a new temporary password. The temporary password is “welcome”. Each user must enter a new password immediately after starting a Management Console session.

Changing your own user password

To change your user password:

- 1: Log into the Management Console.
- 2: Click **Change Password** on the Management Console Dashboard.
The *Change Password* window opens.
- 3: Enter your current password in **Old Password** text-entry box.
- 4: Enter your new password in the **New Password** and **Confirm New Password** text-entry boxes.

Depending upon the settings in **Configuration->Preferences->Passwords**, a password:

- Must include at least one uppercase alphabet character.
 - Must include at least one of the following special characters:
[! @ # \$ % ^ & * () { } \ []]
 - Cannot repeat characters consecutively. A character can be used several times in a password but it cannot be adjacent to the same character.
 - Can contain non-adjacent blank spaces.
- 5: Click **Ok**.

The password is applied to the current user. The password is current and does not have to be reset the next time the user logs into the Management Console.

Changing another user’s password

Only administrative users can change the passwords of other users. The `admin` user is used to create other users. The default administrative user and password are `admin` and `admin123`. Change this password immediately after installing the Security Server software. The `admin` user cannot be deleted from the Security Server.

Note that changing a password places the new password in an “expired” state. The user must reset the password the next time s/he logs in. See also [“Logging in with an expired user account” on page 5-72](#).

To change a user password:

- 1: Log into the Management Console as a user with `admin` role permissions.
If this is a new installation, there is only the `admin` user.

- 2: Select **Administrators** in the menu bar.

The *Administrators* window opens. It displays the names of all configured users in blue regular text. The currently logged in administrative user is displayed in black italicized text.

- 3: Select a user in the **Login** column.

The *Details* window opens.

- 4: Enter the password in the **Password** and **Confirm Password** text-entry boxes.

- 5: Click **Ok**.

The password is applied to the user.

Logging in with an expired user account

If your password expires, or an administrator changes your password for you, you must reset the password when you start the next Management Console session. A warning is issued when you log in with an expired or modified password, and the Management Console prompts you to change the password.




Figure 5.1: Password expiration notice

To reset the password:

- 1: Log into the Management Console as you would normally.

The Management Console notes the expired password and prompts you to change it.

- 2: Enter the current user password in the **Old Password** text-entry box.

The current password is either the last password that you were using prior to its expiration, or the password that the administrator assigned to you.

- 3: Enter a new password in the **New Password** and **Confirm New Password** text-entry boxes.

The number of passwords that you must cycle through before reusing an old password is set by the `Password History` parameter.

- 4: Click **Ok**.

The password is applied and you logged into the Management Console.

Creating users

To create a new Security Server user:

- 1: Log into the Management Console as a user with `Admin` role permissions.

- 2: Select **Administrators** in the menu bar.

The *Administrators* window opens.

- 3: Click **Add**.

The *Add Administrator* window opens.

- 4: Enter a name for the new user in the **Login** text-entry box.
A valid user name is 5 to 12 ASCII characters long; for example, `policyAdmin1`.
- 5: Enter a password for the new user in the **Password** text-entry box.
The length of the password is administrator configurable. The minimum length is 8 characters and the maximum is whatever your system will support.
Allowed characters are uppercase and lowercase alphabet characters, numbers, and special characters. The special characters are: `!@#$$%^&*(){}[]`. There must be at least one uppercase alphabet character and at least one special character.
- 6: Re-enter the password in the **Confirm Password** text-entry box.
- 7: Select one or more roles for the new administrative user by enabling user role checkboxes.
- 8: Click **Ok**.
The new user is added to the Security Server database.



Passwords cannot be viewed. If a person forgets his or her password, edit the user configuration and assign a new user password.

Changing user configurations

You can delete users, change user passwords, and change user roles. You cannot rename users or view their passwords.

To modify an existing user account:

- 1: Log into the Management Console as a user with `Admin` role permissions.
 - 2: Select **Administrators** in the menu bar.
The *Administrators* window opens.
 - 3: Select a user in the **Login** column.
The *Details* window opens.
 - 4: If you want to change the user password, enter the new password in the **Password** and **Confirm Password** text-entry boxes.
If you leave the password text-entry boxes blank, the current password will remain in effect.
The length of the password is administrator configurable. The minimum length is 8 characters and the maximum is whatever your system will support.
Allowed characters are uppercase and lowercase alphabet characters, numbers, and special characters. The special characters are: `!@#$$%^&*(){}[]`. There must be at least one uppercase alphabet character and at least one special character.
 - 5: Assign roles by enabling checkboxes in the **Roles** column.
The check-boxes for the currently assigned roles are displayed. At least one **Roles** checkbox must be enabled for you to save the changes.
 - 6: Click **Ok**.
The changes are applied to the user. If you changed the user password,
-

Deleting user accounts

You can delete Security Server users. If you delete the same user account as the one you are currently logged in as, the account is removed immediately but you are able to continue administrative tasks until you log out. You cannot log back in as the deleted user.



The default administrative user, `admin`, cannot be deleted.

To delete existing administrative users:

- 1: Log into the Management Console as a user with `Admin` role permissions.
- 2: Select **Administrators** in the menu bar.

The *Administrators* window opens.

- 3: Enable the checkboxes in the **Select** column of those users that you want to delete.
- 4: Click the **Delete** button.

A dialog box opens that asks if you are sure you want to proceed with the operation.

- 5: Click the **Ok** button.

The user is removed from the Security Server database.


Configuring hosts

Hosts are computer systems that run the DB2/IDS Database Backup Agent and/or the File System Agent.

Add hosts to the Security Server database before you install the agent software. This makes certificate exchange and host registration quicker and easier. Once installed, define the individual hosts to be protected as well as create host groups for easier administration. Create host groups so you can manage multiple hosts simultaneously.

Displaying configured hosts

The *Hosts* window displays the hosts in the Security Server database and some details about their configuration.



The screenshot shows the 'Hosts' window with a table of configured hosts. The table has columns for 'Select', 'OS Type', 'Host Name', 'FS Agent' (Reg. Allowed, Comm. Enabled), 'DB2 Agent' (Reg. Allowed, Comm. Enabled), 'IDS Agent' (Reg. Allowed, Comm. Enabled), and 'Description'. Two hosts are listed: 'vmlinux02' and 'vmlinux05'.

Select	OS Type	Host Name	FS Agent		DB2 Agent		IDS Agent		Description
			Reg. Allowed	Comm. Enabled	Reg. Allowed	Comm. Enabled	Reg. Allowed	Comm. Enabled	
☐	Linux	vmlinux02	☒	☐	☒	☒	☒	☒	VM Linux System 02
☐	Linux	vmlinux05	☒	☒	☒	☐	☒	☒	64-bit 2.6 Linux

Figure 6.1: Sample Hosts window

The table in the *Host* window indicates the following for each host:

Column	Description
Select	This checkbox, when enabled, indicates that the host to be removed when the Delete button is clicked. Multiple checkboxes can be enabled at one time.

Column	Description
OS Type	Indicates if the Security Server was able to identify and register the host system. Part of a host identity is its operating system. Currently, a host may be <code>Unknown</code> and <code>Linux</code> . <code>Unknown</code> indicates that the host has not been registered.
Host Name	The name of the host in the Security Server database.
FS Agent	This column consists of two subcolumns, Registration Allowed and Communication Enabled . The Registration Allowed checkbox indicates that you are allowing the host to register and be configured to run File System Agent software. It only indicates that the host can be configured for File System Agent protection. The Communication Enabled checkbox indicates that a policy can be applied to a host. Registration Allowed must be enabled for Communication Enabled to be enabled.
DB2 Agent/ IDS Agent	This column consists of two subcolumns, Registration Allowed and Communication Enabled . The Registration Allowed checkbox indicates that you are allowing the host to register and be configured to run Database Backup Agent software. It only indicates that the host can be configured for Database Backup Agent protection. The Communication Enabled checkbox indicates that a policy can be applied to a host. Registration Allowed must be enabled for Communication Enabled to be enabled.
Description	A textual description of the host that facilitates host identification.

The checkboxes on the *Host* window are:

- **Select All:** Selects all the hosts in the Security Server database. Select this checkbox to delete all the hosts at one time. If you have enabled the **Select** checkbox for many individual hosts, a quick way to deselect them is to enable and then disable the **Select All** checkbox.
- **View All:** Displays all the hosts in the Security Server database on one Web page, regardless of the display number specified in the preferences.

The **Type** scroll-list indicates the hosts to display based upon several configuration parameters:

- `All` displays all the configured hosts. `All` is the default.
- `FS_ALLOWED` displays those hosts with the **FS Agent Registration Allowed** checkbox enabled.
- `FS_DISALLOWED` displays those hosts with the **FS Agent Registration Allowed** checkbox disabled.
- `DB2_ALLOWED` displays those hosts with the **DB2 Agent Registration Allowed** checkbox enabled.
- `DB2_DISALLOWED` displays those hosts with the **DB2 Agent Registration Allowed** checkbox disabled.

- **IDS_ALLOWED** displays those hosts with the **IDS Agent Registration Allowed** checkbox enabled.
- **IDS_DISALLOWED** displays those hosts with the **IDS Agent Registration Allowed** checkbox disabled.

Click the **Go** button to refresh the window so that only those hosts that match the specified type are displayed.

If **View All** is not enabled, VCR-style navigation buttons are displayed in the *Host* window. Use these buttons to advance between pages. The VCR-style buttons are:

	First. Display the first group of hosts in the <i>Host</i> window. The number of host entries in each group is determined by the display attributes defined in the <i>General</i> tab of the <i>Preferences</i> window.
	Previous. Display the previous group of hosts in the <i>Host</i> window.
	Next. Display the next group of hosts in the <i>Host</i> window.
	Last. Display the last group of hosts in the <i>Host</i> window.

Adding a host to the Security Server database

A host may be configured to run just the Database Backup Agent, just the File System Agent, or both agents. Hosts must be added to the Security Server database before they can be registered and administered by the Security Server. Hosts can be added to the Security Server database manually or in batch mode. This section describes how to add them manually. For details about adding hosts in batch mode, see [“Batch host configuration” on page 6-80](#).

To add a host to the Security Server database and apply Encryption Expert protection:

- 1: Log into the Management Console as a user with **Server** role permissions.
- 2: Select **Hosts->Hosts** in the menu bar.
The *Hosts* window opens.
- 3: Click **Add**.
The *Add Host* window opens.
- 4: Enter the name to identify the host in the **Host Name** text-entry box.
This is the same name that will be used to identify the host during the certificate generation phase of agent software installation. It will also be the CN name in the host certificate. Enter the hostname or FQDN.
- 5: Enter a password for the Database Backup Agent system in the **Host Password** and **Confirm Host Password** text-entry boxes.
Depending upon the type of encryption keys that you configure for the host, you may be prompted to supply this password in order to decrypt encrypted data when there is no network connection between the host and the Security Server. See also [“Entering the File System Agent host password” on page 6-97](#).
- 6: Enter a phrase or string in the **Description** text-entry box that helps you to identify the host.
This field is optional. The maximum number of characters is 256.

- 7: Enable the **Registration Allowed Agents** checkboxes for the agents that will run on the host system.

Your choices are **FS IDS**, and **DB2**. Enable the checkboxes for the agent(s) you want to configure. You may enable the checkboxes now or later; however, the checkbox for an agent must be enabled before you can register that agent with the Security Server. Note that only one Database Backup Agent (**DB2** or **IDS**) can be installed on a host at one time.

- 8: Click **Ok**.
- 9: Install the agent software, as described in [“Installing the Encryption Expert agent” on page 3-43](#).

Agent software installation should complete successfully. If it does not, check that there is not a mismatch between the host information configured in step 4 above and the host information configured during agent software installation.

You can manually configure agent certificates later if the certificate generation and exchange phase of agent software installation fails. See [“Changing host certificates” on page 6-88](#) for details about regenerating new File System Agent and Database Backup Agent certificates.

- 10: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens.

- 11: Click the host in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

Note that the host **Name** and **OS Type** boxes now display the host information. The host **Name** value is derived from the CN value of the agent certificate and **OS Type** is derived from the system.

Also, because you had installed and registered agents, in this case FS and DB2/IDS, additional information is displayed.

Edit Host - vmlinux100

General Guard FS Guard DB Host Settings DB2 Log IDS Log FS Log Member

Host Information

Name: vmlinux100 Description: 64-bit Linux, DB2 Se:
 OS Type: LINUX Communication Port: 7024
 FS Agent Locked: ☒ System Locked: ☒
 Password: Password Confirm Password: Password

Agent Information

Agent	Version	Certificate Fingerprint	Registration Allowed	Communication Enabled
FS	1.1.3.0	73:F0:6F:C6:F6:26:B0:9B:82:FB:EC:FE:E2:45:B3:77:DC:EC:BF:31	☒	☒
DB2	1.1.2.0	D2:F5:2D:8D:B5:9B:B3:94:B0:2B:6C:4D:5B:46:20:17:81:56:3B:A4	☒	☐
IDS			☐	☐

Ok Apply Cancel

agent software version
 agent type
 FS is registered and policies can be applied
 DB2 is registered but policies cannot be applied

Figure 6.2: Sample *General* tab display

- 12: You can change the port number used to exchange policy enforcement data between the Security Server and the File System Agent.

Generally, you change the port number when the default port number is already in use or if your firewall requires a different port number. It is recommended that you use the same port number for all your File System Agents to simplify maintenance.

The **Communication Port** text-entry box sets the port number through which the Security Server and the File System Agent exchange policy-related data. When you change the port number and click **Ok**, the configuration change is downloaded to the File System Agent host after the interval set by the `Update Host Frequency` parameter. The configuration change is downloaded to the File System Agent host using port 8443.

You must manually restart the File System Agent after the update is downloaded. Restarting the File System Agent refreshes the configuration data used by the File System Agent process, `vmd`. Execute one of the following commands on the File System Agent host to restart the File System Agent:

- On Linux:

```
# service secfs restart
```
- On AIX:

```
# /etc/rc.d/rc2.d/S99secfs restart
```

The default File System Agent communication port number is 7024. See also [“Default Host Communication Port” on page 15-255](#).

- 13: You can optionally add the host to a host group.

- 14: Create the encryption keys that the host is going to use to backup and restore DB2 data or protect file system files.
- 15: Add the encryption keys that the host is going to use to backup and restore DB2 data, or protect file system files, to a key group.
- 16: Create the policy that Encryption Expert is to use to encrypt and decrypt DB2 data or protect file system files. See [“About Encryption Expert Database Backup Agent policies” on page 8-131](#).
- 17: Enable the **Communication Allowed** checkboxes for **FS** and/or **DB2**, and/or **IDS**.
 Successfully registered agents display a hash value in the **Certificate Fingerprint** column. The **Registration Allowed** checkbox must be enabled before you can enable the **Communication Enabled** checkbox. An agent must be registered and the **Communication Enabled** checkbox enabled before you can apply policies to that agent.
 Configure keys and policies before enabling the host. You can optionally configure the host in a host group.
- 18: Enable the **Communication Allowed** checkboxes for **FS** and/or **DB2**, and/or **IDS**.
- 19: Click **Ok** to finalize the changes.

Batch host configuration

A feature is provided that reads a specially formatted text file to add multiple hosts to the Security Server database. After they are in the database, you can install and register the agents that run on those hosts.

Note that this batch process does not reduce the time it takes to add hosts to the Security Server database. It only makes it easier to add a large number of hosts by reducing key strokes and permitting unattended operation. Also note that the batch input file is not verified as it is read. If there are errors in the batch input file, such as malformed passwords or inappropriate characters, this operation can fail and hang, and not one host will be added.

The lines in a batch file are individual host definitions that follow an identical format. Each line is a comma-separated (,) list comprised of six fields. If you want to include a comma as part of the field value, enclose the whole field in double-quotes (“”). Do not enclose other special characters, such as the colon (:), in double-quotes.

The format of a batch file line is:

hostname,description,password,allow_fs_agent,allow_db2_agent,allow_ids_agent

where:

- *hostname*: An alphanumeric string that represents the hostname or FQDN of the host being added. This is the network identity of the host.
- *description*: A text string that describes the host.
- *password*: This is not a regular login or user password. This is the host password to be used to decrypt cached keys when the Security Server is not accessible, and when the host is configured with `Cached on Host` keys.
- *allow_fs_agent*: A boolean string that is either “yes” or “no” to enable or disable File System Agent registration.
- *allow_db2_agent*: A boolean string that is either “yes” or “no” to enable or disable DB2 Database Backup Agent registration.
- *allow_ids_agent*: A boolean string that is either “yes” or “no” to enable or disable IDS Database Backup Agent registration.

An example batch file is shown below:

```
host1,This is host 1,onlyMe78,yes,yes,no
host2,This is host 2,bobsNum1,yes,yes,no
host3,This is host 3,good4u678,yes,yes,no
host4,This is host 4,some1else,yes,no,no
host5,This is host 5,qwerty123,no,yes,no
host6,This is host 6,ooPB2AUoo,no,yes,no
```

To add hosts using a batch configuration file:

- 1: Create a configuration file.
- 2: Log into the Management Console as a user with `Server` role permissions.
- 3: Select **Hosts->Hosts** in the menu bar.
The *Import Hosts* window opens.
- 4: Click **Browse**.
The *Choose file* window opens.
- 5: Navigate to, and select, the configuration file.
- 6: Click **Open**.
- 7: Click **Ok** in the *Import Hosts* window.

After a while, the *Import Hosts* window closes and the *Hosts* window is redisplayed with the hosts specified in the configuration file.

Agent software can now be installed on these systems and the agents can be registered with the Security Server.

Setting host locks

FS Agent Locked and **System Locked** are two options that are used to protect the File System Agent and certain system files. File System Agent protection includes preventing some changes to the File System Agent installation directory and preventing the unauthorized termination of File System Agent processes. These options appear in *General* tab of the *Edit Host* and *Edit Host Group* windows. These options are disabled by default.

active locks

Host Information

Name: vmlinux100
OS Type: LINUX
FS Agent Locked: ☒
Password:

Description: 64-bit Linux, DB2 Se
Communication Port: 7024
System Locked: ☒
Confirm Password:

Agent Information

Agent	Version	Certificate Fingerprint	Registration Allowed	Communication Enabled
FS	1.1.3.0	73:FD:6F:C6:F6:26:80:9B:82:FB:EC:FE:E2:45:B3:77:DC:EC:BF:31	☒	☒
DB2	1.1.2.0	D2:F5:2D:8D:B5:9B:83:94:80:2B:6C:4D:5B:46:20:17:81:56:3B:A4	☒	☐
IDS			☐	☐

Ok Apply Cancel

Figure 6.3: Checking active locks

To apply locks:

- 1: Check that no one is currently in or accessing the Encryption Expert installation hierarchy; otherwise, the Security Server may be unable to lock Encryption Expert software.
- 2: Log into the Management Console as a user with `server` role permissions.
- 3: To set the locks on an individual host:
 - a: Select **Hosts->Hosts** in the menu bar.
The *Hosts* window opens.
 - b: Click a host name in the **Host Name** column.
The *Edit Host* window opens to the *General* tab.
 - c: To protect File System Agent files from modification and deletion, enable the **FS Agent Locked** checkbox.
See “FS Agent Locked” on page 6-84.
 - d: To protect a set of system files from modification and deletion, enable the **Host->System Locked** checkbox.
System Locked is automatically enabled when **FS Agent Locked** is enabled. You can enable and disable **System Locked** only when **FS Agent Locked** is enabled.
See “System Locked” on page 6-85.
 - e: Select **Ok** to finalize the changes.
- 4: To set the locks on the hosts in a host group:
 - a: Select **Hosts->Host Groups** in the menu bar.
The *Host Groups* window opens.
 - b: Click a host group in the **Name** column.
The *Edit Host Group* window opens to the *General* tab.
 - c: To protect File System Agent files from modification and deletion, enable the **FS Agent Locked** checkbox.

See “FS Agent Locked” on page 6-84.

- d: To protect a set of system files from modification and deletion, select **Host->System Locked**.

System Locked is automatically enabled when **FS Agent Locked** is enabled. You can enable and disable **System Locked** only when **FS Agent Locked** is enabled.

See “System Locked” on page 6-85.

- e: Select **Ok** to finalize the changes.

- 5: (Optional) Verify that the locks have been applied to the agent.

- a: Log onto the agent system as the root user.

- b: Execute the `secfsd` command with the `lockstat` argument.

```
# secfsd -status lockstat
FS Agent Lock: true
System Lock: true
#
```

Sometimes there is a discrepancy between what the Security Server reports as the File System Agent configuration and the actual File System Agent configuration. This can be due to several factors, such as the time delay between log uploads to the Security Server or a guard point is in use when the lock is applied.



If the locks are enabled and the Security Server cannot administer the host, such as can occur after changing authentication credentials, you must unlock the host manually. To unlock the host manually, boot the host into single-user mode and edit the `./secfs/.sec/conf/configuration/secfs_config` file.

Set both `coreguard_locked` and `system_locked` to `false`. Save the file. Boot the system into multi-user mode. You should now be able to administer the host again.

On Windows systems, boot in safe mode, rename

`C:\Windows\system32\drivers\vmgmt.sys` and `.\drivers\vmfiltr.sys` to something else, then boot in regular mode.



There must be communication between the Security Server administrator and the host system administrator. The system administrator must inform the Security Server administrator of changes to the system hierarchy. The system administrator can request to have the locks temporarily disabled to perform some administrative functions. The system administrator can remove directories and files, then, later when the lock is reapplied, the Security Server is protecting non-existent data. Another common administrative issue pertains to mounted guard points. The system administrator can remove or unmount an unlocked, non-automounted guard point. The Management Console interface is not aware of this change and does not issue a warning when you reapply the lock to the now non-existent mounted guard point. To recover an unmounted guard point:

- Disable the guard point for the file system in the Management Console.
- Mount the file system on the host.
- Enable the guard point for the file system.

FS Agent Locked

FS Agent Locked locks the contents of the Encryption Expert File System Agent directories on the host. These directories are `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs` and `./agent/vmd`.

Files in these directories cannot be modified or removed when **FS Agent Locked** is enabled; however, the Security Server can still propagate updates to the host system.

When **FS Agent Locked** is enabled:

- **System Locked** is automatically enabled.
- Certificates are exchanged and the host is bound to the Security Server.
- The File System Agent installation directory cannot be deleted or overwritten.
- The process or service for the File System Agent kernel, `vmd`, cannot be killed.
- The process or service for the File System Agent application, `secfs`, cannot be killed.
- The File System Agent guard points cannot be forcefully unmounted.

When **FS Agent Locked** is disabled:

- **System Locked** is automatically disabled.
- The File System Agent software on the host is not protected.



Do not unregister or delete the File System Agent while locks are applied. The locks stay in effect after the agent is unregistered and, without agent credentials, the Security Server cannot administer that agent and it cannot disable the locks. You will have to boot the host into single-user mode and manually modify the agent configuration in order to disable the locks.

On Linux systems, the following files and the contents of the following directories cannot be changed or moved when **FS Agent Locked** is enabled.

- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs/bin`
- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd`

On Linux systems, the contents of the following files and directories can be modified, but not removed or renamed when **FS Agent Locked** is enabled.

- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs/`
- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs/tmp`
- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs/bin`
- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd`

On AIX systems, the following files and the contents of the following directories cannot be changed or moved when **FS Agent Locked** is enabled.

- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd`

On AIX systems, the contents of the following files and directories can be modified, but not removed or renamed when **FS Agent Locked** is enabled.

- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs/`
- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs/tmp`

On Windows systems, when **FS Agent Locked** is enabled, the following folder cannot be moved and its contents cannot be modified:

- C:\Program Files\IBM\DB2TOOLS\LUWEncryptionExpert\Agent\secfs\sec

Also, the Encryption Expert entries in the registry cannot be modified or deleted when **FS Agent Locked** is enabled on a Windows system.

System Locked

System Locked applies an internal policy to the host to lock host system directories, like /var, /bin, /etc, etc. You must enable **FS Agent Locked** before you can enable **System Locked**.



To upgrade or install third-party software, add new applications, open an SSH session remotely, or modify system directories, you must disable **System Locked**.



(Windows only) Verify that the volume letter and the path for the Windows system are correct before proceeding. When Windows File System Agent software is installed, the volume letter defaults to “C:.” It is possible that the executables in the *Host Settings* tab are on a different volume or in a different folder. If the volume or path information is incorrect, the Security Server cannot sign the applications, and it will be unable to apply **FS Agent Locked** and **System Locked**.

When **System Locked** is enabled:

- Operating system directories on the host are protected.
- Microsoft Update cannot be run on Windows systems to protect the host. Microsoft update and other installation-related executables are specifically blocked. Executables like wuacult.exe and msiexec.exe cannot be run.

When **System Locked** is disabled:

- The internal policy is disabled.
- **FS Agent Locked** remains enabled.
- You can install or update system software.

The following files, directories, and subdirectories are, by default, automatically protected when **System Locked** is enabled. NB: Asterisks (*) indicate pattern matching.

On Linux systems, the following files and the contents of the following directories cannot be changed or moved when **System Locked** is enabled.

- /etc/pam.d
- /etc/rc*
- /etc/security
- /usr/lib/security

On Linux systems, the contents of the following files and directories can be modified, but not removed or renamed when **System Locked** is enabled.

- /etc
 - /etc/init.d/secfs
 - /usr
 - /usr/bin/vmd
-

- /usr/bin/vmsec
- /usr/bin/secfsd
- /usr/bin/dataxform
- /usr/lib
- /usr/lib/*pam*
- /usr/lib/security
- /var/log/eet

On Solaris systems, the following files and the contents of the following directories are protected when **System Locked** is applied.

- /usr/lib/fs*
- /usr/ker*
- /usr/pla*
- /usr/lib/securi*
- /etc/rc*
- /etc/ns*
- /etc/vfs*
- /etc/init.d/secfs
- /etc/system
- /ker*
- /pl*
- /sbin*

On AIX systems, the following files and the contents of the following directories cannot be changed or moved when **System Locked** is enabled.

- /etc/rc.d
- /etc/security
- /usr/lib/security
- /sbin/helpers/mount_secfs

On AIX systems, the contents of the following files and directories can be modified, but not removed or renamed when **System Locked** is enabled.

- /usr/bin/vmd
- /usr/bin/vmsec
- /usr/bin/secfsd
- /usr/bin/dataxform
- /etc
- /usr
- /usr/lib
- /var/log/eet

When **System Locked** is applied, a protected file or path cannot be renamed or deleted; however, if it is a directory, other files may be added to it. For example, /etc cannot be deleted nor renamed, though you can add files to it. A file that cannot be modified cannot be opened and edited in any way.

On Windows systems, files with the following extensions in the Windows OS installation folder (e.g.: \Windows, \WinNT, etc.) cannot be moved or modified when **System Locked** is enabled:

```
.exe
.dll
.sys
.cmd
.com
```

Deleting hosts

When you delete a host in the Management Console, you are deleting the host record and configuration from the Security Server database only. The agent installations on the host system remains intact, complete with applied policies. To completely remove an Encryption Expert host, run the software removal utility on the host system before you delete the host record in the Management Console.

When you delete a host you also delete that host from any host groups that it may be a member.



Do not unregister or delete the File System Agent while locks are applied. The locks stay in effect after the agent is unregistered and, without agent credentials, the Security Server is unable to administer that agent and it cannot disable the locks. You will have to boot the host into single-user mode and manually modify the agent configuration in order to disable the locks.

To remove everything associated with a host, including the agent software that runs on the host:

- 1: Run `dataxform` on the host files that you want unencrypted.
- 2: Log into the Management Console as a user with `Server` role permissions.
- 3: Disable the locks for the host.
- 4: Remove all the guard points for the host.
- 5: Select **Hosts->Hosts** in the Management Console menu bar.
The *Hosts* window opens.
- 6: Enable the **Select** checkbox of each host to be deleted.
- 7: Click **Delete**.

A dialog box opens that asks if you are sure you want to proceed with the operation.

- 8: Click **Ok**.
 - 9: Log into the host system and delete the agent software.
-

Changing host certificates

You must regenerate host certificates when you:

- move the agent to a new Security Server
- want to update agent certificates as part of regimented update process
- delete and reinstall agent software
- delete and reinstall the Security Server

To change host certificates:

- 1: Log into the Management Console as a user with `server` role permissions.
- 2: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens.

- 3: Click the host in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

- 4: Disable the **Registration Allowed** checkboxes for the **FS** and/or **DB2**, and/or **IDS** agents whose certificates you want to change.
- 5: Click **Apply**.

A dialog box opens warning you that the agent certificates will be removed and agent-server communication will be disabled. You will have to reregister the agents. Note that the agent configuration stays in place so you do not have to reconfigure policies, etc.

- 6: Click **Ok**.

The **Certificate Fingerprint** for **FS** and/or **DB2**, and/or **IDS** should be removed.

- 7: Re-enable the **Registration Allowed** and **Communication Enabled** checkboxes for **FS** and/or **DB2**, and/or **IDS**.

You can enable just the **Registration Allowed** checkboxes now and the **Communication Enabled** checkboxes later when you have policies to apply, or enable both of them now.

- 8: Click **Ok**.

- 9: If you are creating a new File System Agent certificate:

a: Log onto the File System Agent host as the root user.

b: (UNIX only) Remove the old File System Agent certificates.

```
# ./bin/register_host clean
Encryption Expert Configuration
Certificate data successfully removed.
```

The `register_host` utility works differently on Windows. The Windows `register_host` utility automatically removes current certificates.

c: Re-register the File System Agent.

```
# ./bin/register_host
```

You are prompted to continue certificate generation.

Do you wish to proceed with host registration? (Y/N) [Y]:

d: Enter `y` or `<Enter>` to continue certificate generation and agent registration.

You are prompted to enter the exact same hostname or FQDN that was used to configure the Security Server in order. The agent system will use this to contact the Security Server system.

Please enter the same host name here: vmlinux01

- e: Enter the hostname or FQDN of the Security Server that will be used to administer this File System Agent.

If you are unsure of what to enter, look at the Security Server *Dashboard* window and enter the same name that is displayed in the `Server Name` field.

The Security Server CA signing certificate fingerprint is displayed and you are prompted to indicate if it is correct and safe to use.

Is the fingerprint correct? (Y/N) [N]:

- f: Enter y to indicate that the Security Server certificate is valid.

The certificate fingerprint that is displayed should match the fingerprint displayed in the Security Server *Dashboard* window.

After a moment, you are prompted to accept the fingerprint of the CA certificate.

- g: Enter y.

After a moment, the following message is displayed, indicating that the certificates have been successfully generated and exchanged between the agent and Security Server.

```
The Encryption Expert agent has been successfully
configured to communicate with the Security Server at
'hostname/FQDN'.
```

- 10: If you are creating a new DB2 or IDS Database Backup Agent certificate:

- a: Log onto the File System Agent host as the root user.

- b: Remove the Database Backup Agent certificates.

(IDS users, replace each instance of “db2” with “ids” to achieve the same results.)

```
# ./db2/bin/setup_eet_agent_cert_db2 clean
Encryption Expert Configuration
Certificate data successfully removed.
```

- c: Re-register the Database Backup Agent.

```
# ./db2/bin/setup_eet_agent_cert_db2
```

You are prompted to continue certificate generation.

Do you wish to proceed with host registration? (Y/N) [Y]:

- d: Enter y or <Enter> to continue certificate generation and agent registration.

You are prompted to enter the exact same hostname or FQDN that was used to configure the Security Server. This will be used by the agent system to contact the Security Server system.

Please enter the same host name here: vmlinux01

- e: Enter the hostname or FQDN of the Security Server that will be used to administer this Database Backup Agent.

Look at the Security Server *Dashboard* window and enter the same name that is displayed in the `Server Name` field.

The Security Server CA signing certificate fingerprint is displayed and you are prompted to indicate if it is correct and safe to use.

Is the fingerprint correct? (Y/N) [N]:

- f: Enter y to indicate that the Security Server certificate is valid.

The certificate fingerprint that is displayed should match the fingerprint displayed in the Security Server *Dashboard* window.

After a moment, the Database Backup Agent certificate fingerprint is displayed, followed by the Security Server CA fingerprint.

g: Compare the displayed CA certificate fingerprint with the CA fingerprint on the Security Server *Dashboard* window.

h: Enter *y* if they match.

After a moment, the following message is displayed, indicating that the certificates have been successfully generated and exchanged between the agent and Security Server.

Encryption Expert Agent for DB2 successfully configured to communicate with Security Server on *hostname/FQDN*.

11: Inspect the *Edit Host* window for new fingerprints.

A new fingerprint indicates that the host successfully registered with the Security Server. If the **Certificate Fingerprint** for **FS** and/or **DB2**, and/or **IDS** is still empty, an error has occurred. Inspect the log files on the agent system to determine the cause of the failure.

Managing Encryption Expert agents

An Encryption Expert protected host can be configured to run the File System Agent and just one Database Backup Agent (DB2 or IDS).

Applying Encryption Expert protection to a Database Backup Agent host

There are two ways to apply database protection. You can apply protection on a host by host basis or you can configure multiple hosts into a group and apply the same protection to all the hosts in the group. Host groups are a convenient way to assign policies and keys simultaneously to a collection of hosts, rather than configuring each host individually. This section describes the process on a host by host basis. For group protection, see [“Configuring host groups” on page 6-108](#).

This section describes Database Backup Agent configuration for a host that is running DB2 or IDS software. You cannot apply policies to both DB2 and IDS agents simultaneously on the same host.


Before you apply Encryption Expert protection:

- Add the host to the Security Server database, as described in [“Adding a host to the Security Server database” on page 6-77](#).
- Install the Encryption Expert agent software on the host system.
- Create encryption keys.
- Configure the policies you want to apply to the host, as described in [“About Encryption Expert Database Backup Agent policies” on page 8-131](#).

To apply protection to a host running the Database Backup Agent:

- 1: Log into the Management Console as a user with `server` role permissions.
- 2: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens. The **Registration Allowed** checkbox in the **DB2 Agent** column should be enabled. The **Communication Allowed** checkbox may or may not be enabled at this time.



Hosts									
☐ Select All ☒ View All Type: ALL Go									
Add Delete Import									
Select	OS Type	Host Name	FS Agent Reg. Allowed	Comm. Enabled	DB2 Agent Reg. Allowed	Comm. Enabled	IDS Agent Reg. Allowed	Comm. Enabled	Descri
☐	Linux	vmlinux02	☒	☒	☒	☒	☒	☒	
Add Delete Import									

Figure 6.4: Displaying configured hosts

- 3: Click the host in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

- 4: Select the *Guard DB* tab.

This tab displays the applied policies, the host groups to which the policies are being applied, and their enforcement status.


- 5: Click the **Guard** button.

A dialog box opens that prompts you to select the policy to apply to the host.

- 6: Select a policy from the scroll-list.

All configured offline policies (DB2/IDS) are displayed. Select an offline policy for the type of Database Backup Agent running on the host system. Online policies (FS) are not displayed because they cannot be applied to the DB2/IDS program.

- 7: Click **Ok**.



Edit Host - vmlinux02

General Guard FS **Guard DB** Host Settings DB2 Log IDS Log FS Log Member

☒ View All Total: 1

 Page 1 of 1

Select	Policy	Host Group	Enabled
☒	allowAllOps_db2		☒

 Page 1 of 1

Figure 6.5: Applying a policy to a DB2 host

- 8: Click the *General* tab.
- 9: If it is not already enabled, enable the **Communication Allowed** checkbox for the Database Backup Agent.
- 10: Select **Ok** to finalize the changes.
- 11: Make a DB2/IDS backup to verify that the server and host can communicate, as well as to verify the policy itself.
See [“Backing up and restoring DB2 databases” on page 11-197](#) or [“Backing up and restoring IDS databases” on page 12-205](#).
- 12: Check the Security Server log to monitor the backup process.

Applying Encryption Expert protection to a File System Agent host

This section describes File System Agent host configuration.

Before you apply Encryption Expert protection:

- Add the host to the Security Server database, as described in [“Adding a host to the Security Server database” on page 6-77](#).
- Install the Encryption Expert File System Agent software on the host system.
- Create encryption keys.
- Configure the policies you want to apply to the host, as described in [“Creating and editing File System Agent policies” on page 9-141](#).
- Check that no one is using the directory before making it a guard point. If someone is in the directory when it is made into a guard point, the user can continue to use data in memory rather than use the actual data in the guard point. Tell users to save their work, to close applications that are running in the directory, and to exit the directory before applying the guard point. When they re-enter the directory they will use protected data and Encryption Expert will work appropriately.
- If you want to change a policy or rekey a guard point, be prepared to temporarily stop access to the guard point. Changing policies for a guard point requires an interruption of service because the transition process entails disabling one policy and then enabling another policy. The guard point must be inactive during the transition period to ensure guard point integrity. The same rule applies to moving a host between host groups when it includes a change in policies. Coordinate policy changes during a maintenance outage window.



You cannot apply File System Agent protection to already mounted and guarded directories, nor can you nest guard points. The `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs/.sec` directory is automatically mounted and guarded by secfs when the File System Agent process starts on the host. You cannot apply a guard point to `/opt` because it contains the existing guard point, `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs/.sec`. You can display mounted and guarded directories using the `df` command.

To apply protection to a host running the File System Agent:

- 1: Log into the Management Console as a user with `server` role permissions.
- 2: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens.

Select	OS Type	Host Name	FS Agent Reg. Allowed	Comm. Enabled	DB2 Agent Reg. Allowed	Comm. Enabled	Description
☐	Unknown	vmlinux02	☒	☐	☒	☐	VM Linux System 02
☐	Linux	vmlinux05	☒	☒	☒	☒	64-bit Linux

Figure 6.6: Displaying configured hosts

- 3: Click the host in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

- 4: Select the *Guard FS* tab.

This tab displays applied policies in a tabular format. Each policy line in the table comprises:

- a selection radio button
- the policy name
- the host group of which the current host is a member
- the path of the file or directory to which the policy is being applied
- the type of policy being applied, `DIR` or `CLUSTER_DIR`
- an indicator of the file system mount type, whether a regular mount or an automount
- policy enforcement status is either enabled or disabled
- connection status to the host

The policy table is empty if this is a new host configuration or if no policies are applied.

- 5: Select the **Guard** button.

The *Guard File System* window opens.

- 6: Select the policy to apply from the **Policy** drop-down menu.
- 7: Select the kind of object to be protected from the **Type** drop-down menu.
- Your choices are **Directory** or **Cluster directory**. Select **Directory**. Veritas **Cluster directory** is not supported at this time.
- 8: If applicable, enable the checkbox to indicate that the guard point is a UNIX automount by enabling the **Auto Mount** toggle.

See “[Selecting the guard point mount type](#)” on page 6-96 for information on the use and application of automounted file systems.



Under random circumstances, NFS file systems can be mounted before the File System Agent drivers are loaded. When this occurs, the File System Agent is unable to protect guard points on the file system. The **Auto Mount** feature prevents this from occurring. Enable the **Auto Mount** toggle in the *Edit Host* window when the guard point is in an automounted file system.

- 9: You can:
- Enter the full path to the guard point in the **Path** text-entry box and click **Ok** to apply the policy to the target guard point.
 - Enter part of the path to the guard point in the **Path** text-entry box and click **Browse** to locate and select the rest of the path.
 - Click **Browse** to locate and select the entire path. This option is used in this example.

It is recommended that you use the browser to locate the target guard point, to avoid typographical errors, and to verify host availability.

- 10: Click the **Browse** button to locate the host directory to guard.

The *Remote File Browser* window opens.

- 11: Find the target guard point.

Click the plus symbol (+) next to a folder to display the next level of the directory hierarchy. Click the minus symbol (-) to collapse the hierarchy. Click a folder or file name to select that directory or file.




Figure 6.7: Browsing for guard points

Active guard points are displayed as folders overlaid with a shield. If you suspect that the guard point status is incorrectly indicated, note that the agent status displayed in the window shows the status as it is configured in the Security Server database. It is not a real-time indication of the actual status. For actual status, log onto the agent system and run Encryption Expert utilities, like “`vmsec status`” and “`secfsd -status guard`”. Compare the two to ensure that the guard point status on the Security Server and agent match. If the two do not match, go with what you see on the agent.

To quickly traverse different directory hierarchies, you can enter part of the path to the guard point in the **Start Directory** text-entry box and click **Go**, or press <Enter>, to display and select the rest of the path.

12: Select the directory or file to become a guard point.

The selected path is displayed in the **Selection** text-entry box.

You should check that no file or directory below the selected guard point is currently being accessed. If something under the guard point is being used or accessed, the Security Server may not be able to take control of the directory and apply protection.



The maximum number of characters allowed in a guard point path is determined by your operating system. You can specify a guard point path up to the restriction imposed by the host operating system. However, it is recommended that you keep it below 1,000. Beyond 1,000 characters the path information for the Resource field in the Management Console log and host messages file (e.g. `/var/log/messages`) is truncated, and the Key and Effect fields that normally follow the Resource field are not displayed.



The directory (or directory path) specified in a resource set is appended to the guard point. This means if the guard point is `/mnt/remote2` and the resource set directory path is `/remoteDir`, then the policy is applied to the files and directories in `/mnt/remote2/remoteDir`.



Policies are applied beyond the local host file system, if the guard point is properly specified. Mounted directories on the local host must be on a single file system. For example, if `/mnt/remote2` is a regular mount on the local host and it contains an NFS mount to `remoteDir` on a remote host, do not make `/mnt/remote2` the guard point. Specify the full path to the NFS mount, `/mnt/remote2/remoteDir`.

13: Click **OK**.

The *Guard File System* window opens to display the selected values.




Figure 6.8: Completed guard point selection

14: Click **OK**.

The *Edit Host* window is updated to display the new guard point.




Figure 6.9: New guard point without policy enforcement

Note the guard point status.

- A green circle indicates an active connection to the agent system.
- A red rectangle indicates that the policy is not yet applied or the host is unavailable.
- A yellow triangle indicates that an attempt to delete a guard point is still pending.

15: Wait a few moments then click the **Refresh** button to update the display.

The red rectangle should change to a green circle.

It may be easier to execute the `df` command repeatedly on the host system until you notice a `secfs` mount for the new guard point, or execute `tail -f /var/log/eet/eetvmd_root` and wait until a message like the following is displayed:

```
Successfully received and implemented a new security configuration.
```

16: Redisplay the *Guard FS* tab.




Figure 6.10: New guard point with policy enforcement

Selecting the guard point mount type

- Under random circumstances, NFS file systems can be mounted before the File System Agent drivers are loaded. When this occurs, the File System Agent is unable to protect guard points on the file system. The **Auto Mount** feature prevents this from occurring. Enable the **Auto Mount** toggle in the *Edit Host* window when the guard point is in an automounted file system.
- Under random circumstances, NFS file systems can be mounted before the EE drivers are loaded. When this occurs, EE is unable to protect guard points on the file system. The **Auto Mount** feature prevents this from occurring. Enable the **Auto Mount** checkbox in the *Edit Host* window when the guard point is in an automounted file system.
- When applying File System protection to an automounted file system, do not apply the guard point to the link-target directory. Rather, apply the guard point to the full path to the directory underneath it. For example, if the automounted directory `"/Auto"` mounts a link-target directory named `"/documents"`, do not set the guard point to `"/documents"`. Instead, set the guard point to `"/Auto/documents"`.
- Do not configure Linux 64-bit hosts to automount directories with the `"/net"` option. The automounter uses the automount map associated with each mount point to locate each file system as it is accessed. The PEM cannot resolve file system selection for guard points, including any directory below a guard point, that are configured with the `"/net"` option.

Displaying File System Agent guard point status

You can display File System Agent guard point status on the Security Server and on the host that is running the File System Agent. The agent status displayed in the Management Console shows the status as it is configured in the Security Server database. It is not a real-time indication of the actual status. For actual status, log onto the agent system and run Encryption Expert utilities. Compare the two to ensure that the guard point status on the Security Server and agent match. If the two do not match, go with what you see on the agent.

You can view agent status information several ways, a few of the ways is by logging into the agent system and:

- Executing “vmsec status”
- Executing “secfsd -status” and viewing the status file `./agent/secfs/tmp/statusfile`
- Executing “secfsd -status guard” to display just guard points, the applied policies, and the guard point status
- (UNIX) Running `df` to display Encryption Expert mounted file systems
- (Windows) Pop-up the Encryption Expert and select **Status**

To view File System Agent guard point status on the Security Server:

- 1: Log into the Management Console as a user with `server` role permissions.
- 2: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens.

- 3: Click the host in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

- 4: Click the status indicator of a guard point.

The status indicator is a green circle, a yellow triangle, or a red rectangle in the **Status** column.

The *Guard Point Status* window opens.




Figure 6.11: Guard Point Status window

Do not click a guard point with a red rectangle status indicator. The *Guard Point Status* window will not display any configuration or status data when the red rectangle is displayed.

The window is not automatically updated. You must close and reopen the window after the Security Server and File System Agent synchronize and the status indicator turns green.

- 5: Click the “X” on the *Guard Point Status* window to close it.

Entering the File System Agent host password

A host password is required to access encrypted data when there is no network connection to the Security Server and the host system is configured with offline keys (Cached on Host). If a guard point is accessed before entering the password, the accessing application will appear to hang.

The password must be entered when a host system with offline keys is booted without a network connection to the Security Server. Enter the password before accessing protected data. Data access is allowed after the password is entered and if the policy permits. After

the password is entered, other users can access protected files without having to re-enter the password because they are using the cached keys; however, access is still controlled by the applied policy.

To specify the File System Agent host password:

- 1: Open a UNIX terminal on the host system.
- 2: Enter `vmsec passwd` on the command line.

You are prompted to enter the password. The password was set by the Security Server administrator who created the initial host configuration.

- 3: Enter the password.

The terminal window can appear to hang, and interrupts like `<Ctrl>-c` and `<Ctrl>-z` will have no effect. If you experience this problem, open a new terminal window and enter the File System Agent password. The hung application in the other window will resume operation. Note that the same behavior may occur in the event that the File System Agent is configured with online (*Store on Server*) keys and the host cannot access the Security Server. In this case entering the password will have no effect. The agent will get the keys when the network connection to the Security Server is restored, and the hung application will resume. See also “[Creating keys](#)” on page 7-118.



Virus-scan software can cause the File System Agent password-entry window to pop-up when it accesses protected guard points. This is because the File System Agent has encountered an unknown and untrusted application. Enter the File System Agent login credentials to continue with the virus scan.

Changing the File System Agent host password

The offline password feature is designed to protect a laptop or similar portable system from being stolen and then used to access encrypted data. A password must be provided when there is no connection between the File System Agent and the Security Server before being allowed to decrypt/encrypt files that are encrypted with an offline key (*Cached on Host*). The offline password feature controls access to the encryption keys that are stored locally on a particular machine, so as to allow encryption/decryption when the Security Server is not accessible.

The host password is set when the host is added to the Security Server database. Passwords are set on a host-by-host basis, rather than on a host group basis. For information about setting the initial host password, see “[Adding a host to the Security Server database](#)” on page 6-77.

To change the host password:

- 1: Log into the Management Console as a user with `server` role permissions.
- 2: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens.

- 3: Click the host in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

- 4: Enter the new password in the **Password** and **Confirm Password** text-entry boxes.
- 5: Click **Apply** or **Ok**.

- 6: Tell the user(s) of the host the new password or they will be unable to access encrypted data when there is no network connection between the File System Agent system and the Security Server.

Deleting guard points

The process of deleting a guard point is relatively simple, and involves just a few key clicks. However, there are some preliminary steps to take before deleting a guard point:

- Encrypted data in a guard point will still be encrypted when the guard point is removed. If you are not going to reuse the guard point for any reason, such as uninstalling the File System Agent software from a host, either, copy the encrypted files out of the guard point so that they are saved as unencrypted files, or, rekey the encrypted files while the guard point is still applied.
- Take the guard point out of service so that no user or application is accessing the directories and files in the guard point. A guard point is a mounted file system. Removing a guard point involves unmounting the file system. File systems cannot be unmounted when in use.
- Delete all the guard points and disable the locks for a host before deleting the host record from the Security Server database. This ensures that there are no residual guard points in effect on the host.

To delete a guard point:

- 1: Log into the Management Console as a user with `server` role permissions.
- 2: Select **Hosts->Hosts** in the menu bar.
The *Hosts* window opens.
- 3: Select a host in the **Host Name** column of the *Hosts* window.
The *Edit Hosts* window opens.
- 4: Select the Guard FS tab.
The guard points are displayed.
- 5: Select the radio button in the **Select** column for the guard point to be deleted.
Only one guard point at a time can be selected at a time.
- 6: Click **Unguard**.
- 7: Note the guard point status.
 - A green circle indicates an active connection to the agent system.
 - A red rectangle indicates that the policy is not yet applied or the host is unavailable.
 - A yellow triangle indicates that an attempt to delete a guard point is still pending.




Figure 6.12: Guard point awaiting deletion

- 8: Click **Refresh** to update the tab.

After the File System Agent acknowledges that the guard point has been removed from the host, it is removed from the Management Console *Guard FS* tab.

- 9: Check the mount points on the File System Agent host to ensure that the guard point has been removed.

There are several ways to check, two of them are shown below:

```
# df
#/usr/bin/secfs -status guard
```

Configuring an agent to use Security Servers

The primary Security Server is configured on the agent system during agent software installation. This section describes how you can change the agent configuration so that it uses a different Security Server, uses different SSL certificates, or to add a failover Security Server for the agent to use in case the primary Security Server is unresponsive.

The Security Servers that the agents access are configured in `./agent/db2/etc/agent.conf` and `./agent/vmd/etc/agent.conf`. The path of the first `agent.conf` file is for Database Backup Agents. The second is for File System Agents.

Before you start, you must have a working Database Backup Agent and/or File System Agent installation and a working primary Security Server installation. The following example assumes that both the File System Agent and Database Backup Agent are installed on one host system. If only one agent is installed on the system, just perform the steps for that agent.

To configure the servers for the File System Agent and Database Backup Agent to use:

- 1: Log onto the agent host as the root user.
- 2: If you want to change the primary server:
 - a: Remove the existing agent/server certificates and configuration from the agent system.


```
# ./db2/bin/setup_eet_agent_cert_db2 clean
# ./vmd/bin/register_host clean
```

This step assumes that you want to configure the File System Agent and Database Backup Agent that are running on the host to use a Security Server other than the currently configured Security Server. These utilities remove all the agent certificates from the `./db2/pem` and `./vmd/pem` directories, and they remove all the server URLs from the `./db2/etc/agent.conf` and `./vmd/etc/agent.conf` files.

- b: Log into the Management Console as a user with `server` role permissions.
- c: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens.

- d: Click the host in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

- e: Disable the **Registration Allowed** checkboxes for **FS** and/or **DB2** and/or **IDS**. The **Communication Enabled** checkboxes are automatically disabled.

This step removes the existing agent certificates from the Security Server database. The host record and the policies that it uses are left intact.

- 3: Click **Ok**.

A dialog box opens warning you that the agent certificates will be removed and there will be no further agent-server communication. You will have to

reregister the host. The host configuration stays in place so you do not have to reconfigure policies, guard points, etc.

a: Click **Ok**.

The fingerprints displayed in the **Certificate Fingerprint** fields for **FS** and/or **DB2** and/or **IDS** should be removed.

b: Re-enable the **Registration Allowed** and **Communication Enabled** checkboxes for **FS** and/or **DB2**, and/or **IDS**.

c: Click **Ok**.

d: Run the agent certificate generation utilities without any arguments on the agent system.

```
# ./db2/bin/setup_eet_agent_cert_db2# ./vmd/bin/
register_host
```

e: For each certificate generation utility you run, enter the hostname of the Security Server when prompted. For example:

```
Enter the same host name here: vmlinux01
```

Enter the primary Security Server in this step. Add the failover Security Servers later.

The certificate generation utility submits a certificate request to the Security Server, the Security Server signs it, and returns it and a Security Server certificate to the host. The agent certificate generation utility should complete with a message like:

```
Encryption Expert Agent for DB2 successfully configured to
communicate with Security Server on vmlinux01.
```

and

```
The Encryption Expert agent has been successfully
configured to communicate with the Security Server at
'vmlinux01'.
```

4: Add the failover servers to the Database Backup Agent and File System Agent configurations.

At this point, the agent is configured to interact with the primary Security Server. Add the failover server URLs to a configuration file that already contains the primary server URL. The first URL in the list is the primary server and the second and remaining URLs are the failover servers. If the primary server is unavailable, the agent tries the failover servers in the order in which they are entered in the file.

a: Open `./agent/db2/etc/agent.conf` in a text editor.

One non-commented line that begins with “URL” should be present. This line contains the URL for the primary server. For example,

```
URL = https://vmlinux01:8443/communication-server/
SSControllerDB2
```

b: Add failover server URLs to the end of the primary server line. Separate the URLs with a comma (.). For example,

```
URL = https://vmlinux01:8443/communication-server/
SSControllerDB2,https://vmlinux02:8443/communication-
server/SSControllerDB2,https://vmlinux22:8443/
communication-server/SSControllerDB2
```

c: Save and exit the `agent.conf` file.

d: Open `./agent/vmd/etc/agent.conf` in a text editor.

There is an uncommented line at the bottom of the file that begins with “URL”. This line contains the URL for the primary server. For example,

```
URL = https://vmlinux01:8443
```

e: Add the failover server URLs to the end of the primary server line. Separate the URLs with a comma (.). For example,

```
URL = https://vmlinux01:8443,https://vmlinux22:8443
```

f: Save and exit the `agent.conf` file.

5: Restart the File System Agent process.

The File System Agent is an active process that runs in the host system cache. It must be restarted to refresh the configuration information. Skip this step if you modified only the Database Backup Agent `access.conf` file. The Database Backup Agent does not need to be restarted because it is not loaded until it is invoked by the `db2 backup` or `restore` command. Examples of restarting the File System Agent process are:

- On Linux:

```
# service secfs restart
```

- On AIX:

```
# /etc/rc.d/rc2.d/S99secfs restart
```

Both the Database Backup Agent and the File System Agent should be configured to use the primary and failover servers after the File System Agent restarts.

Configuring host settings

PAM (Pluggable Authentication Module) and LAM (Loadable Authentication Module) systems are integral parts of the Security Server authentication process. Without PAM/LAM on UNIX systems, the user `uid` can be displayed in the Security Server logs with the message `User Not Authenticated`. Thus, no matter if a valid user tries to execute an application, if the **User** or **Resource** security rule attribute requires specific user credentials, permission to execute is denied.

You can just remove the **User** and **Resource** security rule attributes from the policy, but that would grant access to anyone. You have two choices: You can either trust a specific application or give the application valid user credentials.

Applications like `su`, `ssh`, and `login` have a sense of identity and they require user credentials before they run. An application that attempts access to guarded data and applications, but lacks valid user credentials, can be trusted by the File System Agent by adding the application to the *Host Settings* tab. The objects listed in the *Host Settings* tab are signed applications that require user authentication before they are allowed to run.

These keywords are used to accept the user credentials only. The application must still be granted permission by the other security rules in the policy.

Since the applications in the *Host Settings* tab are signed, apply host settings after:

- installing File System Agent software
- upgrading File System Agent software
- configuring PAM or LAM
- changing the `pam.conf` file
- changing any of the files listed in the *Host Settings* tab

To configure application authentication credentials:

- 1: Log into the Management Console as a user with `Server` role permissions.
- 2: Select **Hosts->Hosts** in the menu bar.

The *Hosts* window opens.

- 3: Click the host in the **Host Name** column.

The *Edit Host* window opens to the *General* tab.

- 4: Select the *Host Settings* tab.

This tab displays a default set of system applications that require user authentication. The default applications and settings for a Linux system are:

```
/bin/su
/usr/sbin/sshd
/usr/sbin/in.rlogind
/bin/login
/usr/bin/gdm-binary
/usr/bin/kdm
```

The default applications and settings for a Windows system are:

```
C:\WINDOWS\system32\winlogon.exe
|lock|C:\WINDOWS\system32\msiexec.exe
|lock|C:\WINDOWS\system32\wuauclt.exe
|lock|C:\WINDOWS\system32\wupdmgr.exe
|lock|C:\Program
Files\IBM\DBTOOLS\LUWEncryptionExpert\agent\secfs\sec\
bin\vminstall.exe
|exempt|C:\WINDOWS\explorer.exe
|exempt|C:\WINDOWS\regedit.exe
|exempt|C:\WINDOWS\system32\regedt32.exe
|exempt|C:\WINDOWS\system32\svchost.exe
|exempt|C:\WINDOWS\system32\services.exe
|exempt|C:\WINDOWS\system32\smss.exe
```



(Windows only) Verify that the volume letter and the path for the Windows system are correct before proceeding. When Windows File System Agent software is installed, the volume letter defaults to “C:.” It is possible that the executables in the *Host Settings* tab are on a different volume or in a different folder. If the volume or path information is incorrect, the Security Server cannot sign the applications, and it will be unable to apply **FS Agent Locked** and **System Locked**.

- 5: Add one of the following keywords to the beginning of an application path whose authentication method you want to change.

The following table lists the keywords that you can enter in the *Host Settings* tab that override different authentication requirements:

Keyword	Description
authenticator	(UNIX only) The Security Server uses PAM/LAM authentication on UNIX File System Agent hosts to validate user identity. If PAM/LAM is not properly configured, or you do not wish to configure PAM/LAM, a valid user may be denied access with a "User Not Authenticated" or "FAKED USER" log entry. Apply the authenticator keyword to those signed executables you wish to trust without PAM/LAM validation.
exempt	(Windows only) When processes or applications are started, the internal policy and regular policies are checked locally or by the Security Server. When a policy check is performed and <code>exempt</code> is applied to the process, a 6 second timeout is imposed on the check. Without <code>exempt</code>, an application can wait indefinitely for a policy access check to complete, as in the case when the Security Server is required but is not accessible. If the check times-out because the Security Server is unavailable for any reason, access is denied. Exempt host processes are also "exempt" from pop-up messages that describe the occurrence of access violations. An example of what causes such pop-ups is an application that tries to memory map a file for which it does not have encryption permission (i.e., memory map with no <code>viewability</code> key on Windows). The only reasons to include <code>exempt</code> in the configuration are shorter wait periods and blocked pop-ups.
lock	(Windows only) Specifies an application that cannot be executed on the host. An application defined with <code>lock</code> does not go through an internal policy check. It is simply not allowed to run at all. A default set of applications are locked on the Windows host to prevent their execution and causing potential failure during bootup. The same effect can be achieved by configuring the Resource and Process security rule attributes in a policy; however, certain default applications are automatically locked in the <i>Host Settings</i> tab as a precautionary measure in the event you fail to include these applications in the policy. Sometimes problems occur when installing software on a locked host, such as installation failure or application lockup. Specific processes can be identified where, when they are locked, they cannot be started and the failure goes away. For example: <pre> lock c:\winnt\system32\msiexec.exe</pre>

Keyword	Description
trust	(UNIX only) Apply this feature to authenticate the credentials of applications and processes that do not have a user login. Add an application or process to the <i>Host Settings</i> tab with trust in order for the File System Agent to accept the application or process credentials. Security Server trust is only applied to the specified application or process. Any subprocesses that it may spawn are not trusted. Use trustfrom if you want to trust the subprocesses. Use trust for applications and processes that do not use a login, but are started by the system, such as cron and db2sysc.
trustfrom	(UNIX only) trustfrom is identical to trust except that trustfrom also authenticates the credentials of the subprocesses that are spawned by the process or application. Use restraint in applying this keyword because, though you may trust the initial process, you have to consider the subprocesses it can spawn.



Create entries in the *Host Settings* tab for Windows applications running under Wow64 that require some form of user authentication. The syswow64 paths are created by default during Windows file agent installation. \WINNT is for Windows 2000 and \Windows is for Windows XP and Windows Itanium operating systems.

In Wow64, all file-access to C:\Windows\System32 is redirected to C:\Windows\syswow64, and is implemented using the File System. Redirected syswow64 paths are effective only for 64-bit Windows file agents. This is the path where programs compiled for 32-bits are stored in order to run on a 64-bit system.

- 6: You can remove or modify entries to control their respective access permissions. To modify an entry:
 - a: Click and drag the mouse cursor across the text to be changed.
You can also double- or triple-click text to select words or the entire string.
 - b: Type or paste the new text.
You can also press the <Delete> key to remove the string.
- 7: Add additional keywords and application paths, as desired.




Figure 6.13: Host Setting window with authentication keywords

- 8: Click **Apply** or **Ok** to apply the changes. The applications are resigned.
- 9: If there are any users currently logged into the system, have them log out and then log back in to refresh their user authentication credentials.
- 10: Verify the change by logging onto the host and accessing a guard point, then check the user information in the Management Console log.



You should not continue to get User Not Authenticated errors after updating the host settings. If you update host settings for ssh or sshd on AIX or Solaris systems (for example, by adding |authenticator| to /usr/lib/ssh/sshd in the *Host Settings* tab), you must restart secfs. The users can log back onto the host system after secfs restarts. Linux and Windows hosts do not require a secfs restart after changing the host settings, but still they require the user to log off and then log back on.

Checking host settings


Check log entries for failed user authentication entries if a user is unable to access the guard point. Failed authentication log entries consist of the user name, actual user ID (if applicable), and the message, “User Not Authenticated”. This occurs when the entries in the *Host Settings* tab and the PAM configuration are not synchronized, or the `pam.conf` file is not configured.

Compare the two log entries below. The first log entry is the result of executing the `xcalc` command in a guard point without PAM configured. The policy requires a user named “bc”. The second log entry is the result of doing the same thing with the same policy and guard point, but with the PAM properly configured:

```
EET2602E: [SecFS, 0]
allowAdminsOnly_fsPolicy[bc,uid=1001 (User Not
Authenticated)] User[/home/apps/app1/bin/xcalc]
Process[read_file] Action[/home/apps/app1/bin] Res[/
xcalc None] Key[DENIED] Effect[Code (1U,2M) {9}]
```

```
EET2603I: [SecFS, 0]
allowAdminsOnly_fsPolicy[bc,uid=1001,gid=10\staff\]
User[/bin/csh] Process[read_dir] Action[/home/apps/app1/
bin] Res[/] PERMIT Effect[Code (1M) {9}]
```

If User Not Authenticated messages are generated, and the policy and `pam.conf` file are properly configured, open the *Host Settings* tab and click **Apply**. This forces a refresh of the signatures for the applications listed in the *Host Settings* tab. You must refresh the signatures after making any changes to the `pam.conf` file.



The UNIX File System Agent captures user credentials when the user provides authentication with a valid user name and password during login with applications like `ssh` and `ftp`. This is achieved by using the PAM, which is configured in the `/etc/pam.conf` file. The problem is that applications such as DB2 and Oracle run as daemons and do not authenticate to the operating system via the PAM. Instead, these applications are denied access to guarded directories because their credentials are not trusted. Since the File System Agent does not trust the credentials supplied by any application that does not pass PAM authentication, adding an application to the host setting with `|trust|` directs the File System Agent to trust the user credentials presented by the application and, if the policy permits, grants access to the protected data. `|trustfrom|` directs the File System Agent to trust the credentials presented by the application, as well as its subprocesses. For example, if the Sybase database requires protection, add the full path to the `dataserver` application by entering something like the following entry into the *Host Settings* tab and click **Apply**:

```
|trust|/opt/SyBase/ASE-12_5/dataserver
```

The keywords `|trust|` and `|trustfrom|` simply accept whatever user credentials are provided. `|authenticator|` is different in that it sets the credentials to those of the user or application that is attempting to run the program. For example, a person attempting to login to a host via the `ssh` command may need to execute `/usr/local/sbin/sshd` on the host. Without PAM configured, or improperly configured, the user `uid` will include `User Not Authenticated`, indicating an unauthenticated login. If you add the line `|authenticator|/usr/local/sbin/sshd` to the *Host Settings* tab, then the user will log onto the host as the user executing the `ssh` command, regardless of the PAM/LAM configuration. With `|authenticator|` configured, if a user named `bob` `ssh`'es to the host, then the user for that `ssh` session will have `bob`'s credentials. This means the **User** and **Resource** security rule attributes can still be used to determine who can and cannot access protected data.

Host Settings examples

The following example shows the application of the `|trust|` and `|trustfrom|` keywords in the *Host Settings* tab. Leave the `pam.conf` file unconfigured on a UNIX host. Create and apply a policy that permits a user, such as `bcroot`, access to the guard point.

SSH onto the UNIX host and perform some action, such as `ls` inside the guard point. The log shows the user name and the message, “User Not Authenticated”, and the File System Agent denies all access to the unidentified user. Logically, you may want to make `sshd` a trusted application but that does not work. If you add `|trust|/usr/local/sbin/sshd` and/or `|trustfrom|/usr/local/sbin/sshd`, click the **Update Setting** icon, log off the UNIX host, `ssh` back on, and `ls` the guard point again, you get the same error.

You need to make the application that is running the command a trusted application. In this case, the `ls` command is being executed from a `csch` shell, so you need to make the `csch` shell a trusted application. As a test, add `|trust|/usr/bin/csch`, click **Ok**, log off the UNIX host, `ssh` back on, and `ls` the guard point again. The File System Agent notes the user change, labels the user as a fake, and denies access. (HOST[sys-techpub]:ALARM: Policy(unixPolicy002):

```
User[root,uid=0,gid=1/
other,root,bin,sys,adm,uucp,mail,tty,lp,nuucp,daemon/ (faked
as(bcroot,1001)] Process[/usr/bin/ls] Access[read_attr]
Res[/file2] [FAKED USER!! DENIED] Code (1U,2M)). |trust| is not
the correct choice in this case because |trust| passes the user credentials to just the
specified process or application (e.g., csch), not the subprocesses it may spawn (e.g., ls).
```

Change `|trust|` to `|trustfrom|`, as in `|trustfrom|/usr/bin/csch`, click **Ok**, log off the UNIX host, `SSH` back on, and `ls` the guard point again. This time access is permitted because the correct user credentials have been passed from a specified process (e.g., `csch`) to its subprocess (e.g., `ls`). (HOST[sys-techpub]:AUDIT:

```
Policy(unixPolicy002): User[bcroot,uid=1001,gid=10/staff/]
Process[/usr/bin/ls] Access[read_attr] Res[/file2] [PERMIT]
Code (1M))
```

Configuring host groups

Displaying host groups

To display configured host groups:

- 1: Log into the Management Console as a user with `server` role permissions.
- 2: Select **Hosts->Host Groups** in the menu bar.

The *Host Groups* window opens. All configured host groups are displayed.

Deleting host groups

As part of Security Server maintenance, you will occasionally need to remove host groups from the Security Server database. Deleting a host group removes only the group; the individual hosts that are members of that group remain intact. You cannot delete host groups that are configured in a policy. You must remove the host group from the policy before you can delete the host group itself.

To remove a host group:

- 1: Log into the Management Console as a user with `server` role permissions.
- 2: Select **Hosts->Host Groups** in the menu bar.
The *Host Groups* window opens. All configured host groups are displayed.
- 3: Enable the selection checkboxes of those host groups that you want to delete.
The selection checkboxes are located in the **Select** column of the *Host Groups* window.
- 4: Click **Delete**.
You are prompted to verify the deletion.
- 5: Click **Ok**.

Configuring agent logging

You can configure the agent logging behavior for an existing host or for future hosts. To configure agent logging for an existing host, configure the respective *DB2 Log*, *IDS Log*, and *FS Log* tabs in the *Edit Host* window for that host. To configure the default agent logging behavior for all new hosts, configure the respective *DB2 Log*, *IDS Log*, and *FS Log* tabs in the *Preferences* window. Configuring the default agent logging behavior does not affect the current logging behavior of existing hosts.

The tabs in the *Edit Host* window and the *Preferences* window are identical and are configured identically. For information about:

Configuring DB2 Database Backup Agent logging behavior, see [“Setting DB2 Database Backup Agent logging preferences” on page 15-255](#).

Configuring IDS Database Backup Agent logging behavior, see [“Setting IDS Database Backup Agent logging preferences” on page 15-258](#).

Configuring File System Agent logging behavior, see [“Setting File System Agent logging preferences” on page 15-260](#).

Applying Encryption Expert protection to a Database Backup Agent host group

There are two ways to apply host protection. You can apply protection on a host by host basis or you can configure multiple hosts into a group and apply the same protection to all the hosts in the group. Host groups are a convenient way to assign policies and keys simultaneously to a collection of hosts, rather than configuring each host individually. This section describes the process on a host group basis. For individual host protection, see [“Applying Encryption Expert protection to a Database Backup Agent host” on page 6-90](#).

You can configure hosts either before or after configuring host groups; however, creating hosts before creating host groups is quicker and requires fewer steps.

Encryption Expert protection can be applied to a host that is running just a DB2 or IDS agent, just a File System Agent, or both. This section only describes Database Backup Agent host group configuration.

Before you apply Encryption Expert protection:

- Create the initial host configuration in the Management Console for each host to be added the host group. See [“Adding a host to the Security Server database” on page 6-77](#).
- Install the Encryption Expert agent software on each host system.
- Create encryption keys.
- Configure the policies you want to apply to the hosts in the host group, as described in [“About Encryption Expert Database Backup Agent policies” on page 8-131](#).

To create and apply protection to a group of hosts running Database Backup Agents:

- 1: Log into the Management Console as a user with `server` role permissions.
- 2: Select **Hosts->Host Groups** in the menu bar.

The *Host Group* window opens.

- 3: Click **Add**.

The *Add Host Group* window opens.

- 4: Enter the name for the new host group in the **Host Group Name** text-entry box.


This field is mandatory. The maximum number of characters is 64.

- 5: Enter a phrase or string in the **Description** text-entry box that helps you to identify this host group.

This field is optional. The maximum number of characters is 256.

- 6: Click **Ok**.

The *Host Groups* window is redisplayed and it includes the new host group. For example:



Host Groups		
☐ Select All ☒ View All		Total Host Groups: 1
Add Delete		
Select	Name	Description
☐	hostGroup02	DB2-only hosts
Add Delete		

Figure 6.14: Host groups list

- 7: Click the host group in the **Name** column.

The *Edit Host Group* window opens. It has four tabs: *General*, *Guard FS*, *Guard DB*, and *Member*.

- 8: Add hosts to the host group.

a: Select the *Member* tab.

b: Click **Add**.

The *Add Host* window opens. It displays all configured hosts.

- c: Enable the **Select** checkboxes of all the hosts that you want to add to the host group.

Select the hosts to add to the host group based upon the policies to be applied. For example, if you want to apply file system protection and DB2 database protection policies, then the hosts you select should be running file system and DB2 Database Backup Agents.

d: Indicate whether you want the host(s) to retain their current settings or inherit the group settings.

Select either the **Yes** or the **No** radio button for the question, “**Do you want to apply the selected host(s) settings to host group settings?**” Select **Yes** for the host(s) to inherit the host group system lock, FS agent lock, communication enabling, etc. That is, select **Yes** to apply the host group configuration settings to the host(s) being added. Select **No** to add the hosts as they are and retain their individual configurations.

The default is **Yes**.

e: Click **Ok**.

The *Member* tab is redisplayed with the new host group members.

9: Select the *Guard DB* tab.

This tab displays the applied policies, the host groups to which the policies are being applied, and their enforcement status. Nothing is displayed since this is a new installation and no policies are applied.

10: Click **Guard**.

The *Guard Host Group DB* window opens to display all DB2 and IDS database policies.

11: Select the radio button for the policy you want to apply to all database operations.

Only one database policy may be applied at a time. The selected policy must be appropriate for the type of hosts in the host group. That is, if the hosts in the host group are running DB2 Database Backup Agents, select a DB2 policy. If they are running IDS agents, select an IDS policy.

12: Click **Ok**.

13: You can configure file system protection at this time by selecting the *Guard FS* tab, and selecting a guard point and file system policy.

14: Select the *General* tab.

The *General* tab displays the host group name and its description. It is also used to enable the Security Server to begin administering the host group members.

15: Enable the **Enable DB2 Agent** or **Enable IDS Agent** checkbox.

The member hosts are administered as a group when you enable the **Enable DB2 Agent**, or **Enable IDS Agent**, and **Enable FS Agent** checkboxes.

16: If the hosts in the host group are also running a File System Agent, enable the **Enable FS Agent** checkbox.

17: (Optional) Enable the **FS Agent Locked** and **System Locked** checkboxes to apply protection to system files and Encryption Expert files on the host.

18: Click **Ok** to finalize the changes and close the window.

19: Make a database backup on one of the hosts in the host group to verify that the server and a host in the host group can communicate, as well as to verify the policy itself.

See [“Backing up and restoring DB2 databases” on page 11-197](#) or [“Backing up and restoring IDS databases” on page 12-205](#).

20: Display the Security Server log to monitor the backup process.

Applying Encryption Expert protection to a File System Agent host group

There are two ways to apply host protection. You can apply protection on a host by host basis or you can configure multiple hosts into a group and apply the same protection to all the hosts in the group. Host groups are a convenient way to assign policies and keys simultaneously to a collection of hosts, rather than configuring each host individually. This section describes the process on a host group basis. For individual host protection, see [“Applying Encryption Expert protection to a File System Agent host” on page 6-91](#).

You can configure hosts either before or after host groups; however, creating hosts before creating host groups requires fewer steps.

Encryption Expert protection can be applied to a system with just DB2, just file system, or both. This section only describes File System Agent host configuration.

Before you apply Encryption Expert protection:

- Create the initial host configuration in the Management Console for each host to be added the host group. See [“Adding a host to the Security Server database” on page 6-77](#).
- Install the Encryption Expert agent software on each host system.
- Create encryption keys.
- Configure the policies you want to apply to the hosts in the host group, as described in [“About File System Agent policies” on page 9-141](#).

To create and apply protection to a group of hosts that are running File System Agents:

- 1: Log into the Management Console as a user with `Server` role permissions.
- 2: Select **Hosts->Host Groups** in the menu bar.

The *Add Host Group* window opens.

- 3: Enter the name for the new host group in the **Host Group Name** text-entry box. This field is mandatory. The maximum number of characters is 64.

- 4: Enter a phrase or string in the **Description** text-entry box that helps you to identify this host group.

This field is optional. The maximum number of characters is 256.

- 5: If an Encryption Expert File System Agent is configured on one or more of the hosts in the host group, you can enable the **FS Agent Locked** and **System Locked** checkboxes to prevent the deletion or modification of File System Agent installation files and some regular system files.

- 6: Click **Ok**.

The *Host Groups* window is redisplayed and it includes the new host group.

- 7: Click the host group in the **Name** column.

The Host Group applet opens.

- 8: Add hosts to the host group.

a: Select the *Member* tab.

b: Click **Add**.

The *Add Host* window opens that prompts you to select one or more hosts to add to the host group. The dialog box displays all hosts in the Security Server database.

c: Enable the **Select** checkboxes of one or more hosts in the table.

d: Click **Ok**.

The *Member* tab is redisplayed with all the new hosts.

9: Select the *Guard FS* tab.

This tab displays the applied policies, paths to the guard points, type of each guard point, and their enforcement status. Nothing is displayed since this is a new installation and no policies are applied.

10: Click **Guard**.

The *Guard File System* window opens.

11: Select the policy to apply from the **Policy** drop-down menu.

12: Select the kind of object to be protected from the **Type** drop-down menu.

Leave it set to **Directory**. Veritas cluster directories are not supported.

13: If applicable, enable the **Auto Mount** checkbox if the guard point is a UNIX automount mapped file system.

14: Click **Browse** to open the *Remote File Browser* window.

15: Use the browser to display the directories on the host system.

16: Select the folder to become a guard point.

17: Click **OK**.

The *Edit Host Group* window displays the new host group.



Figure 6.15: Configuring a host group

18: Select the *General* tab.

19: Enable the **Enable FS Agent** checkbox.

The member hosts are administered as a group when you enable the **Enable DB2 Agent**, **Enable IDS Agent**, and/or **Enable FS Agent** checkboxes.

20: You can also enable the **FS Agent Locked** and **System Locked** checkboxes at this time to protect the File System Agent software installation on the host system, as well as to protect system files.

System Locked is automatically enabled when **FS Agent Locked** is enabled. You can enable and disable **FS Agent Locked** only when **System Locked** is enabled.

See also [“Setting host locks” on page 6-81](#)

21: Click **Ok**.

Editing host groups

You can change the configuration of existing host groups. You can:

- Change a group description
- Enable policy enforcement
- Edit policies
- Apply policies
- Define guard points
- Add hosts to a host group
- Lock the Encryption Expert files on the host

To edit a Database Backup Agent host group:

- 1: Log into the Management Console as a user with `Server` role permissions.
- 2: Select **Hosts->Host Groups** in the menu bar.

The *Host Groups* window opens. All configured host groups are displayed.

- 3: Select the host group to edit in the **Name** column.

The *Edit Host Group* window opens. It displays four tabs: *General*, *Guard PS*, *Guard DB*, and *Member*.

The *General* tab is displayed by default. You can only change the **Description** text-entry box and the **Enable DB2 Agent**, **Enable IDS Agent**, and **Enable FS Agent** checkboxes. You cannot rename a host group.

- 4: Select the tab for the category of change you want to perform.
- 5: Make the desired changes.
- 6: Click **Ok** to finalize the changes and close the window.

Using Encryption Keys

There are two types of keys: authentication and encryption. Authentication keys verify the identity of the Security Server to the host and visa versa. The Security Server and host cannot communicate without valid authentication keys. These we refer to generically as authentication certificates. Encryption keys encrypt and decrypt data, whether the data resides in a directory or is being copied to a backup tape. (Authentication adheres to the SSL X.509 standard using asymmetrical RSA keys.)

RSA symmetric (public and private) keys are used to encrypt and decrypt the randomly generated AES or 3DES key that is used to encrypt and decrypt the backup data. The public RSA key is used to encrypt the AES or 3DES key. The private RSA key is used to decrypt the AES or 3DES key.

RSA keys are used in two ways:

- Keys that are locally generated on the current Security Server are used in backup and restore operations. These are referred to as local keys. These are your standard public and private keys, in which both halves of the key-pair reside on the same system. This is described in [“Creating an asymmetric key” on page 7-122](#).
- Public keys and their hash codes are copied from a remote Security Server to the local Security Server. The backups that are made on a Security Server configured with just the public key can be restored on a different Security Server that is configured with the matching private key. This is described in [“Exporting a public key” on page 7-123](#).

You can add one or more keys to a key group and apply that key group to a policy. A key group can be a mix of RSA key pairs and public keys.

With key management, you can:

- View key information, such as the SHA-2 hash code or the public key
- Create local keys
- Duplicate keys on remote Security Servers

Once encryption is applied, you must keep track of the encryption keys you are using. Encrypted data is unusable without the proper keys.

The key attributes and the policies you apply to a host determine if a consistent connection is required between the Security Server and File System Agent. The key attributes and policies determine whether the host can work online or offline.

Once registered and activated, the host can be protected by the Security Server. Access to protected data may require Security Server login credentials, depending upon how the host

is configured in the Management Console. The Security Server login credentials are used to decrypt locally-stored host encryption keys.

Hosts with their keys stored on the Security Server require a consistent connection to the Security Server and do not require the Security Server password to access protected data. The policies stay in effect as long as there is a connection between the Security Server and the host. When the network connection is interrupted, users cannot access encrypted data. Data can be accessed after the network connection is re-established. See [“Will the File System Agent require a network connection?” on page 7-118.](#)

Hosts with the keys locally cached are a different matter. In order to protect the keys from attack, the keys are encrypted on the host system. If the host can connect across the network to the Security Server, the required keys are downloaded to the host without any human intervention. If the File System Agent is unable to connect to the Security Server use the `vmsec password-entry` command to unlock the locally-stored encryption keys. All users can access and decrypt/encrypt data, in accordance with applied policies, after entering the Security Server password.

Viewing keys

You can view the following key information of RSA key-pairs and RSA public keys:

- Key name
- Key length
- SHA-2 hash code
- Public key

To view key information:

- 1: Log into the Management Console as a user with `Key` role permissions.
- 2: Select **Keys->Keys** in the menu bar.

The *Keys* window opens and displays all the configured keys.

- 3: Click the key to be viewed in the **Name** column.

The *Key Details* window opens. The content of this window changes based upon the type of key being viewed. The content can consist of:

Field	Description
Name	Name assigned to the key when it was first created.
Algorithm	Algorithm you had selected when you created the key. The symmetric key algorithms are 3DES, AES128, and AES256 The asymmetric key algorithms are RSA1024, RSA2048, and RSA4096.

Field	Description
Key Type	If a symmetric key algorithm is configured, <code>Stored on Server</code> or <code>Cached on Host</code> can be displayed. <code>Stored on Server</code> keys are downloaded to non-persistent memory on the host. Each time the key is needed, the host retrieves the key from the Security Server. <code>Cached on Host</code> downloads and stores (in an encrypted form) the key in persistent memory on the host. If an asymmetric key algorithm is configured, <code>Key Pair</code> or <code>Public Key</code> can be displayed. <code>Key Pair</code> is a standard RSA key in two parts: a public key and a private key. The Security Server with this type of key can allow an agent to backup and restore data. <code>Public Key</code> indicates that key contains only the public key portion of a public:private key pair. The Security Server with this key can allow an agent to backup data only.
Description	Optional text description of the key.

- 4: If you are viewing a symmetric key type, click **Back** to return to the *Keys* window.
- 5: If you are viewing an asymmetric key type, click **Click to Export**.
The *File Download* dialog box opens.
- 6: Click **Open** to display the public key portion of the asymmetric key in a Web browser.

How the public key is displayed depends upon your Web browser and what it does with files of type `.xml`. The key data can be displayed as raw XML code in a Web browser page, or, if your system is configured with an XML editor, the editor can be opened and the formatted XML file displayed.
- 7: Click **Save** to save the public key portion of the asymmetric key.
The *Save As* dialog box opens.
 - a: Specify a path and name for the file.
 - b: Click **Save**.
The *Download Complete* dialog box opens.
 - c: Click **Open** to display the public key portion of the asymmetric key in a Web browser or click **Close**.
- 8: Click **Back** to return to the *Keys* window.

Creating keys

Will the File System Agent require a network connection?

Encryption keys are required to ensure data integrity and privacy. You can create keys that provide varying levels of encryption protection.

Encryption keys can be stored exclusively on the Security Server, downloaded to the host, or downloaded to the host and stored in non-persistent memory. The keys can also be downloaded and stored (in an encrypted form) in persistent memory on the host for use when there is no network connection between the host and Security Server. If a network connection to the Security Server is unavailable, and the File System Agent is configured with persistent keys, enter the host password using the `vmsec password` utility. You can then resume accessing encrypted data.

There are three types of encryption keys: **Stored on Server**, **Cached on Host**, and **Cached on Host plus Unique to Host**. The effects of these types are described below:

For keys that are stored only on the Security Server (**Stored on Server**):

- The keys are downloaded to non-persistent memory on the host.
- Keys remain in effect if they were used before losing the Security Server connection.
- A Security Server connection is required to download keys after a system reboot in order to access encrypted data.
- Do not apply **Stored on Server** keys to offline host files because, even though you enter the offline password, when there is no network connection and an attempt is made to access the files, the window making the attempt may wait indefinitely for the online keys.

For keys that are stored on the host for offline use (**Cached on Host**):

- A host encryption key, encrypted using the host password, is downloaded to the host. All hosts using the same encryption key can access encrypted data on other hosts that use the same key.
 - When the key is needed, it is decrypted and cached. When the host goes offline from the Security Server, the current policy remains in effect because the encryption key is locally available.
 - If the File System Agent cannot connect to the Security Server after a reboot, any attempt to access the contents of an encrypted file on an unconnected host will not complete. The application will hang until the host password is provided to unlock the encryption keys. The application resumes accessing once the password is entered.
 - The key can also be reestablished after a reboot without access to the Security Server through the use of the host password. You can specify the host password using the `"vmsec passwd"` utility.
-

For unique keys that are stored on the host for offline use (**Cached on Host plus Unique to Host**):

- Using a token stored with the host record on the Security Server, the encryption key is made unique for each host.
- This unique host encryption key is downloaded to the host and stored in an encrypted manner using the host password.
- When the key is needed, it is decrypted and cached. When the host goes offline from the Security Server, the current policy remains in effect because the encryption key is locally available in the system cache.
- If the File System Agent cannot connect to the Security Server after a reboot, any attempt to access the contents of an encrypted file on an unconnected host will not complete. The application will hang until the Encryption Expert host password is entered in another terminal window. The application resumes execution once the password is entered.
- The key can also be reestablished after a reboot without access to the Security Server through the use of the host password. You can specify the host password using the “vmsec passwd” utility.
- These keys are used for locally attached devices, as files encrypted by them can only be read by one machine. Therefore, do not use **Cached on Host plus Unique to Host** keys for cloned systems, RAID configurations, clustered environments, or any environment that utilizes host mirroring.

Creating a symmetric key

The two tabs in the *Keys* window are **Symmetric** and **Asymmetric**. Symmetric keys are based upon 3DES and AES algorithms and are used to encrypt data. Symmetric keys are configurable for the File System Agent only. The symmetric keys for the Database Backup Agent are automatically generated. Asymmetric keys are public and private key-pairs that are based upon RSA algorithms and are used to encrypt symmetric keys. Asymmetric keys are used by File System Agents and Database Backup Agents.

To create a symmetric key:

- 1: Log into the Management Console as a user with `KEY` role permissions.
- 2: Select **Keys->Keys** in the menu bar.

The *Keys* window opens.

- 3: Select the *Symmetric* tab.
- 4: Complete the field in this window by using the information in the following table:

Field	Description
Name	Enter the name for the key in the Name text-entry box. This field is mandatory. The maximum number of characters is 64.

Field	Description
Description	Enter a phrase or string in the Description text-entry box that helps you to identify the key. This field is optional. The maximum number of characters is 256.
Algorithm	Select an encryption algorithm from the Algorithm scroll-list. Your choices are 3DES, AES128, and AES256. The default is AES128.
Key Type	Select the location for the generated key from the Key Type scroll-list. Your choices are <code>Stored on Server</code> and <code>Cached on Host</code>. <code>Stored on Server</code> keys are downloaded to non-persistent memory on the host. Each time the key is needed, the host retrieves the key from the Security Server. <code>Cached on Host</code> downloads and stores (in an encrypted form) the key in persistent memory on the host. The cached keys are used when there is no network connection between the host and Security Server. All hosts using the same encryption key can access encrypted data on other hosts that use the same key. The Unique to Host checkbox is displayed when <code>Cached on Host</code> is selected. The default is <code>Stored on Server</code>.
Unique to Host	This checkbox is displayed when the Key Type is set to <code>Cached on Host</code>. When enabled, this checkbox uses a token that is stored in the host record on the Security Server to make the encryption key unique. The unique host encryption key is downloaded to the host and stored in an encrypted manner using the host password. These keys are used for locally attached devices, as files encrypted by them can only be read by one machine. Therefore, do not enable this checkbox for cloned systems, RAID configurations, clustered environments, or any environment that utilizes host mirroring. The Unique to Host checkbox can be enabled only when the Key Type is set to <code>Cached on Host</code> and the Key Creation Method is set to <code>Generate</code>.

Field	Description
Key Creation Method	Select if the key is to generated automatically using a random seed or if it is to be generated by importing a file. Your choices are <i>Generate</i> and <i>Manual Input</i> . <i>Generate</i> is the default. The Unique to Host checkbox is disabled when Key Creation Method is set to <i>Manual Input</i> .
Key String.	This scroll-list is displayed when Key Creation Method is set to <i>Manual Input</i> . See “Manually creating symmetric keys” below for details.

5: Click **Ok**.

Manually creating symmetric keys

Encryption Expert allows you to import symmetric keys and to manually create symmetric keys. Importing symmetric keys is a CoreGuard 3.x legacy issue. This is described in [“Importing legacy keys” on page 7-124](#). Manual key generation is described in this section.

You can randomly create your own keys or copy third-party keys into the Security Server database. If you have access to the hex values of AES and 3DES keys, you can enter those values in the Management Console, and then use those keys to encrypt and decrypt files. This is particularly useful if you want to migrate encrypted files from other applications into Encryption Expert.

To manually create keys:

- 1: Log into the Management Console as a user with **Key** role permissions.
- 2: Select **Keys->Keys** in the menu bar.
The *Keys* window opens.
- 3: Select the *Symmetric* tab.
- 4: Enter the name of the key in the **Name** text-entry box.
The maximum number of characters is 64.
- 5: Enter a descriptive string in the **Description** text-entry box.
This field is optional. The maximum number of characters is 256.
- 6: Select an algorithm in the **Algorithm** pop-up menu.
Your choice is 3DES, AES128, or AES256.
- 7: Select where the key is to be stored from the **Key Type** pop-up menu.
You choice is *Cached on Host* or *Stored on Server*. *Stored on Server* requires a constant network connection to the Security Server. *Cached on Host* can operate without a Security Server connection.
- 8: Set **Key Creation Method** to *Manual Input*.
- 9: Enter a hex string [0-9,a-f,A-F] for the key in the **Key String** text-entry box.
 - Enter 32 hex characters (128 bits) if the selected algorithm is AES128.
 - Enter 48 hex characters (192 bits) if the selected algorithm is 3DES.

- Enter 64 hex characters (256 bits) if the selected algorithm is AES256.

10: Click **Ok**.

Creating an asymmetric key

The two tabs in the *Keys* window are **Symmetric** and **Asymmetric**. Symmetric keys are based upon 3DES and AES algorithms and are used to encrypt data. Symmetric keys are configurable for the File System Agent only. The symmetric keys for the Database Backup Agent are automatically generated. Asymmetric keys are public and private key-pairs that are based upon RSA algorithms and are used to encrypt symmetric keys. Asymmetric keys are used by File System Agents and Database Backup Agents. Also, the public half of an RSA key-pair can be imported into other Security Servers so that these other Security Servers can encrypt data but not decrypt it. A Security Server with the private half of the RSA key-pair is required to decrypt data.

To create an asymmetric key:

- 1: Log into the Management Console as a user with **Key** role permissions.
- 2: Select **Keys->Keys** in the menu bar.
The *Keys* window opens.
- 3: Click **Add**.
- 4: Select the *Asymmetric* tab.
- 5: Complete the field in this windows by using the information in the following table:

Field	Description
Name	Enter the name for the key in the Name text-entry box. This field is mandatory. The maximum number of characters is 64.
Description	Enter a phrase or string in the Description text-entry box that helps you to identify the key. This field is optional. The maximum number of characters is 256.
Key Type	Select the type of RSA key to generate. The choices are Key Pair and Public Key . Key Pair creates a standard RSA key in two parts: a public key and a private key. The Public Key File text-entry box is displayed when you select Public Key . The default is Key Pair .
Algorithm	Select an encryption algorithm from the Algorithm scroll-list. Your choices are RSA1024 , RSA2048 , and RSA4096 . The default is RSA1024 .

Field	Description
Public Key File	This text-entry box is displayed when Key Type is set to <code>Public Key</code> . Click Browse... to select the X.509 certificate file that contains the public key.

- 6: Click **Ok**.

Exporting a public key

A public key is the public key portion of a public:private RSA key-pair. The public key of an RSA key-pair is used to make backups only. The private key of an RSA key-pair is used to restore backups. The public key can be imported into other Security Servers to enable them as backup-only servers.

To export the public key of an RSA key pair:

- 1: Log into the Management Console as a user with `Key` role permissions.
- 2: Select **Keys->Keys** in the Management Console menu bar.
The *Keys* window opens.
- 3: Click the RSA key-pair or RSA public key in the **Name** column that you want to export.
The *Key Details* window opens.
- 4: Click **Click to Export**.
The *File Download* dialog box opens, prompting you to save the public key.
- 5: (Optional) Click **Open** to display the public key.
How the public key is displayed depends upon your Web browser and what it does with files of type `.xml`. The key data can be displayed as raw XML code in a Web browser page, or, if your system is configured with an XML editor, the editor can be opened and the formatted XML file displayed.
- 6: Click **Save**.
The file locator opens.
- 7: Enter the path and name for the file.
The default file name is `PublicKey.xml`.
- 8: Click **Save**.

Importing a public key

To import the public key of an RSA key pair:

- 1: Log into the Management Console as a user with `Key` role permissions.
- 2: Select **Keys->Keys** in the Management Console menu bar.
The *Keys* window opens.
- 3: Click **Add**.
- 4: Select the *Asymmetric* tab.
- 5: Select `Public Key` from the **Key Type** scroll-list.
The **Algorithm** scroll-list is replaced with the **Public Key File** text-entry box.

- 6: Click **Browse...** to open the file locator.
- 7: Locate and select the public key file.
- 8: Click **Open**.

The *Keys* window opens and displays the imported key. The key type is *Public Key*.

Deleting keys

Do not delete keys without first backing them up. Once the keys are gone, and there is no backup, data encrypted with those keys cannot be restored or accessed.

To delete keys:

- 1: Log into the Management Console as a user with *Key* role permissions.
- 2: Select **Keys->Keys** in the menu bar.
The *Keys* window opens.
- 3: Enable the **Selected** checkbox for those keys you want to delete.
- 4: Click **Delete**.

Importing legacy keys

A migration path is available for current Vormetric CoreGuard 3.x users to upgrade to Vormetric Encryption Expert 1.1.3. You can import AES and 3DES symmetric keys that were created on a CoreGuard 3.x appliance into the Encryption Expert 1.1.3 database. This allows you to upgrade to Encryption Expert 1.1.3 without having to rekey the data in File System Agent guard points.

There are three types of symmetric keys in 3.x: *Unique to Host*, *Requires-Online*, and *Allow-Offline*. You cannot import *Unique to Host* keys into 4.x. You can import *Requires-Online* or *Allow-Offline* keys. In 1.1.3, these translate to *Stored on Server* and *Cached on Host*, respectively.

To import CoreGuard 3.x symmetric keys into Encryption Expert 1.1.3:

- 1: Log into the Management Console as a user with *Key* role permissions.
- 2: Select **Configuration->Expert Server Public Key**.
- 3: The *Export Server Public Key* window opens.
- 4: Click **Click to Export**.

The *File Download* window opens. It prompts you to save or view the public key.

- 5: Click **Save**.

The *Save As* window opens.

- 6: Enter the path and name for the public key file.



The saved key file must be accessible by the legacy Security Server because it is required to encrypt the legacy key file.

The default name is *ServerKey.pem*.

- 7: Click **Save**.

The *File Download* window opens. It prompts you to view the public key or open the folder that contains the saved public key.

- 8: Click **Close**.
- 9: Start a Management Console session on the legacy Security Server.
- 10: Select **Import/Export** → **Configuration** in the navigation tree.
The *Save Configuration* window opens.
- 11: Select the **Export** radio button.
- 12: Select the **Config Type** radio button.
- 13: Select **Key(s)** in the **Config Type** drop-down menu.

Save Configuration [Save Configuration] [Load Configuration] ?

Save To
☐ Archive ☒ Export

Config Type: ☒ Key(s)
 allowOffline
 auto01_key
 bernard_key1
 key
 key001
 key002
 key050427a
 Key1
 laptop_key
 manual_key

Select Key(s):

☐ Archive
 — Select from List —

Encryption Key (public key)
 [Browse...]

[Export] [Clear]

Archive File-List

Order	Size (kB)	Date	File-Name
1	427	Nov 22 13:03:17 2006	bootactive
2	409	Sep 12 15:06:09 2006	SSA_exportNETWORK_09122006.conf
3	409	Sep 12 15:03:38 2006	NetworkUserArchive
4	1	Sep 12 14:58:30 2006	SSA_exportPOLICY_09122006.conf
5	409	Sep 12 14:45:44 2006	SSA_export_09122006.conf
6	1	Aug 29 12:37:28 2006	SSA_export_08292006.conf
7	1	Aug 29 12:32:51 2006	genericEncryption4
8	1	Aug 29 12:29:38 2006	genericEncryption3
9	1	Aug 29 12:19:08 2006	genericEncryption2
10	1	Aug 29 11:49:45 2006	arcPolRefGenEnc2
11	1	Aug 28 16:21:14 2006	08282006genericEncryption2.conf
12	1	Aug 28 16:01:59 2006	SSA_export_08282006.conf
13	1	Aug 10 15:25:44 2006	policyExportTest001.conf
14	1	Aug 01 17:33:21 2006	SSA_export_08012006.conf
15	247	Apr 10 01:13:45 2004	SSA_export_05112006.conf

Figure 7.1: Selecting symmetrical keys for export

The **Config Type** drop-down menu displays all the keys in the Security Server Appliance.

- 14: Select one or more keys for export or archive.

Alternately, you can select the **Archive** radio button in the **Save To** category and select **Key(s)**.

The maximum number of keys that can be imported from one file is 500. If you want to import more than 500 keys in Encryption Expert 1.1.3, export the keys into multiple files.

15: Use the browser for the **Encryption Key** text-entry box to locate and select the Encryption Expert public key that you had exported earlier.

16: Click **Export**.

The *Download* window opens.

17: To save the keys to a remote file, right-click the string that begins with **Right button click ...**

18: Select **Save Target As...**

19: Use the browser that opens to specify a location and name for the file.



The saved file must be accessible by the Encryption Expert Security Server.

20: Click **Continue**.

The *Save Configuration* window reopens.

21: Return to Management Console session on the Encryption Expert Security Server.

22: Select **Configuration->Import Symmetric Keys**.

The *Import Symmetric Keys* window opens.

23: Use the browser to select the legacy key file.

24: Click **Ok**.

If a key with the same name already exists in the Encryption Expert Security Server database, the imported key name is appended with `_importX`, where `X` is an incremented number starting at 1.

Keys				
There were key name conflicts and conflicted keys get renamed. Please go to log to see the details.				
☐ Select All ☒ View All Type: All Go Total Keys: 0				
Add Delete				
Selected	Name	Algorithm	Key Type	Description
☐	aes128_yy	AES128	Cached on Host	This is an imported key from CoreGuard
☐	aes128_yy_import1	AES128	Cached on Host	This is an imported key from CoreGuard
☐	aes128_yy_import2	AES128	Cached on Host	This is an imported key from CoreGuard
☐	aes256_yy	AES256	Stored on Server	This is an imported key from CoreGuard
☐	aes256_yy_import1	AES256	Stored on Server	This is an imported key from CoreGuard
☐	aes256_yy_import2	AES256	Stored on Server	This is an imported key from CoreGuard
☐	clear_key	CLEAR	Stored on Server	Clear Key for online protection. Do not delete!
☐	des3_yy	3DES	Cached on Host	This is an imported key from CoreGuard
☐	des3_yy_import1	3DES	Cached on Host	This is an imported key from CoreGuard
☐	des3_yy_import2	3DES	Cached on Host	This is an imported key from CoreGuard
Add Delete				

Figure 7.2: Importing same-name keys

Invalid keys will be skipped and noted in the 4.x log, with messages like:

Unique to host key and clear key importing are not supported. The key `unixEncKey001` has been ignored.

Configuring key groups

Individual keys are configured as members of a key group. The key group is selected in the backup rule of a policy and is added to the backup metadata. Key groups are specified in policies and each member of the key group is used to encrypt the AES key that encrypts the data.

Key groups are comprised of RSA public keys and, when the policy is applied, each member of the key group is used to encrypt the AES key that encrypts the actual data.

You must configure keys and key groups if you configure encryption in your policies.



The keys in a key group are placed in the header of each block in a backup, thus increasing the size of the backup. The larger the key group, the larger the resulting backup file. To minimize file bloating and improve processing performance, keep the number of keys in a key group small.

Creating key groups

A key group is a collection of asymmetric keys that are applied to policies as a unit.

To create a key group and add keys to the key group:

- 1: Log into the Management Console as a user with `Key` role permissions.
- 2: Select **Keys->Key Groups** in the menu bar.

The *Key Groups* window opens and displays all currently configured key groups.

- 3: Click **Add**.

The *Add Key Group* window opens.

- 4: Enter the name for the new key group in the **Key Group Name** text-entry box. This field is mandatory. The maximum number of characters is 64.

- 5: Enter a phrase or string in the **Description** text-entry box that helps you to identify this key group.

This field is optional. The maximum number of characters is 256.

- 6: Click **Add Keys**.

The *Add Keys to Key Group* window opens.

- 7: Enable the **Selected** checkbox of each key to add to the key group.

You can select local keys, public keys, or a mix of the two. Symmetrical keys (AES/3DES) are not displayed because only the keys used to encrypt the AES/3DES encryption keys (i.e., RSA asymmetric keys) are allowed in key groups.

- 8: Click **Add Selected Keys to Key Group**.

The *Add Key Group* window opens and displays the new key group.

- 9: Click **Ok**.

The *Key Groups* window is redisplayed and the new key group is displayed in the key group table.

Editing key groups

You can rename key groups, change their descriptions, add keys to key groups, and remove keys from key groups.

Removing a key from a key group does not affect the key itself. The key remains intact.

To edit a key group:

- 1: Log into the Management Console as a user with `Key` role permissions.
- 2: Select **Keys->Keys** in the Management Console menu bar.
The *Keys* window opens.
- 3: Click the key group in the **Name** column that you want to modify.
The *Key Group* window opens.
- 4: To change the key group name, replace the name in the **Key Group Name** text-entry box with the desired new name.
- 5: To change the key group description, replace the text in the **Description** text-entry box with a new description.
- 6: To add more keys to the key group:
 - a: Click **Add Keys**.
The *Add Keys to Key Group* window opens.
 - b: Enable the **Selected** checkbox of each key to add to the key group.
You can select local keys, public keys, or a mix of the two. Symmetrical keys (AES/3DES) are not displayed because only the keys used to encrypt the AES/3DES encryption keys (i.e., RSA asymmetric keys) are allowed in key groups.
 - c: Click **Add Selected Keys to Key Group**.
The *Edit Key Group* window opens and displays the new key in the group.
- 7: To remove keys from the key group:
 - a: Enable the **Selected** checkbox of each key to be deleted from the key group.
 - b: Click **Delete**.
A dialog box opens that asks if you are sure you want to proceed with the operation.
 - c: Click **Ok**.

Deleting key groups

To delete one or more key groups:

- 1: Log into the Management Console as a user with `Key` role permissions.
 - 2: Select **Keys->Key Groups** in the menu bar.
The *Key Groups* window opens and displays all currently configured key groups.
 - 3: Enable the **Selected** checkbox of each key group to be deleted.
 - 4: Click **Delete Selected**.
-

Rekeying protected data

Rekeying is the act of changing the keys used to encrypt resources. Rekey at some administrator-defined interval as a heightened security precaution or to copy data out of the Security Server protection system for use by other applications or for archival.

See [“Changing encryption keys” on page 10-183](#).

Using policies

This section introduces policy concepts and describes the parts of the Management Console that are used to administer policies. This section also describes how to configure database policies, which are referred to as “offline” policies. File system policies, which are referred to as “online” policies, are more complex to configure and are described in a separate section. For information about configuring file system policies, see [“Creating and editing File System Agent policies” on page 9-141](#).

About Encryption Expert Database Backup Agent policies

When the Encryption Expert Database Backup Agent library is used during a database backup or restore operation, the agent contacts the Encryption Expert Security Server to determine if the operation is allowed. The server makes this determination by matching agent-supplied data against rules in the policy. This process is described below:

- 1: If the host is not registered with the server, the policy is not evaluated and the operation is denied.

An audit event will be generated on the server: *“EET0142I: The Security Server received a request from an unrecognized host”*

- 2: If the host is registered with the server, but a policy is not applied to the host, the operation is denied.

An audit event will be generated on the server: *“EET0014I: There were no applicable policy or rule found”*

- 3: If there are one or more rules associated with a host, the server evaluates the agent-supplied data against the security rule for “matches” in database instance, alias, and partition number. If these are not defined in the policy, they are treated as wildcards and will match any value. The first rule that “matches” is used.

When a rule matches, an audit message is generated on the server: *“EET0144I: The Security Server matched policy <name>”*

- 4: When a rule matches, it can evaluate to either PERMIT or DENY.

The parameters that are evaluated in a rule include the time (matched against the “Valid from”/”To” time range) and “Additional Conditions”, which specify key word pairs that can be passed along on the DB2 backup command line with the COMPROPTS key word.

When a rule evaluates to PERMIT, an audit message is generated on the server: *“EET0006I: The Security Server permitted a backup*

request; The policy used is <name> the rule used is <rule name>

When a rule evaluates to DENY, an audit message is generated on the server:
“EET0147I: The Security Server matched rule <name> to the BACKUP request for <host> But < message>. The rule will not be applied”

A policy is a set of rules that must be satisfied before granting a backup or restore operation. Rules comprise a conditional expression, an action, and a policy attribute. The conditional expression uses one or more of the following parameters and evaluates to either true or false based on the values of the attributes at the time of evaluation:

- Permissible time range, relative to the Security Server
- Agent host name
- DB2 alias
- DB2-generated timestamp on the agent for backup operations
- DB2 partition number (DB2 9.x only)
- DB2 instance name
- IDS server name
- IDS server number

Based upon these attributes, the policy either allows or denies permission to complete the backup or restore operation. If a Security Server permits a backup operation, the Security Server returns a randomly generated AES encryption key to the agent with which to encrypt the data payload. This key is placed in the backup metadata. The Security Server also returns RSA public keys with which to encrypt the backup metadata. If the Security Server permits a restore operation, the Security Server returns the private RSA key necessary to decrypt the backup metadata. The agent then extracts the key in the backup metadata and proceeds to decrypt the data payload.




Figure 8.1: Backup operation

Inversely, if the policy allows you to restore a backup, the Security Server checks the backup metadata for public key information (public key and SHA-2 hash code), then returns the matching private key to decrypt the metadata. The randomly generated key in the metadata is decrypted with the private key and is used to decrypt the actual data.



Policies and policy evaluation are independent of SSL authentication. SSL authentication is the exchange of authentication certificates between agent and server, and are configured separately from policies and encryption keys.



Some functions and features are case-insensitive and some are case-sensitive. In general, system-level strings like user names and host names are case-sensitive, whereas the Encryption Expert-level strings you specify in the Management Console, such as for search engines and policy fields, are not. Also, SSL authentication is case-sensitive. Change the case of a host name and the Security Server will deny access. As a rule-of-thumb, just assume whatever you are entering is case-sensitive and you will have no surprises.

Displaying policies

Policies are displayed in the *Policy* window. You may choose to display all the policies on one Web page or display a set number of policies per Web page. You can then select policies for modification or deletion.

To display Encryption Expert policies:

- 1: Log into the Management Console as a user with `Policy` role permissions.
- 2: Select **Policies** in the menu bar.

The *Policies* window opens. All configured policies are displayed.

To view and select policies in the *Policy List* window:

- **Select All:** Selects all the policies in the Security Server database. Select this checkbox to delete all the policies at one time. If you have enabled the **Select** checkbox for many individual policies, a quick way to deselect them is to enable and then disable the **Select All** checkbox.
- **View All:** Displays all the hosts in the Security Server database on one Web page, regardless of the display number specified in the preferences. If **View All** is not enabled, VCR-style navigation buttons are displayed in the *Policies* window. Use these buttons to advance between pages. The VCR-style buttons are:

	First. Display the first group of policies in the <i>Policy List</i> window.
	Previous. Display the previous group of policies in the <i>Policy List</i> window.
	Next. Display the next group of policies in the <i>Policy List</i> window.
	Last. Display the last group of policies in the <i>Policy List</i> window.

Configuring Database Backup Agent policies

The Security Server uses policies to evaluate backup and restore operations. If the conditions in the policy are satisfied, the operation is allowed, and the Security Server provides the data that is needed to complete the operation. This section describes how to create new database backup policies and administer existing ones.



Multiple database policies can be applied to a single host; however, only one policy can be enabled at a time.

DB2 and IDS policy configuration are almost identical. The information presented in this chapter is applicable to both DB2 and IDS policy configuration. Configuration exceptions are noted.

Creating Database Backup Agent policies

The Security Server uses policies to evaluate backup and restore operations. If the conditions in the policy are satisfied, the operation is allowed, and the Security Server provides the data that is needed to complete the operation. Backup policies are divided in backup rules and restore rules.

If you are configuring a policy that allows backup operations only or restore operations only, you only need to configure just the backup rules or just the restore rules. By default, backup and restore operations are denied.

To create a new policy:

- 1: Log into the Management Console as a user with `Policy` role permissions.
- 2: Select **Policies** in the menu bar.

The *Policies* window opens.

- 3: Click **Add Offline Policy**.

Add Offline Policy is used exclusively for DB2 and IDS. **Add Online Policy** is used exclusively for file system.

The *Add Offline Policy* window opens.

- 4: Enter the name for the policy in the **Name** text-entry box.

This field is mandatory. The maximum number of characters is 64.

- 5: Enter a phrase or string in the **Description** text-entry box that helps you to identify the policy.

This field is optional. The maximum number of characters is 256.

- 6: Create the backup and restore rules now or later.

You must configure either, or both, a backup rule and a restore rule. If the policy is configured with just a backup rule, the host using that policy can only backup databases. If the policy is configured with just a restore rule, the host using that policy can only restore backups.

See [“Configuring Database Backup Agent backup rules” on page 8-135](#) and [“Configuring Database Backup Agent restore rules” on page 8-137](#).

- 7: Click **Ok** to save the policy.

The *Policies* window displays the new policy.

Note the **Type** column in the *Policies* window. The icon displayed in this column indicates if the policy is configured as an offline policy or an online policy. The icon for an offline, or DB2/IDS, policy is a sheet of paper overlaid with a green database symbol (). The icon for an offline, or file system, policy is a sheet of paper overlaid with a yellow folder symbol (.

Configuring Database Backup Agent backup rules

Backup rules define the conditions that must exist, and the actions to take, if a backup request matches the conditions specified in the backup rule.

You can configure backup rules that contain only the public part of a public:private key-pair. The public-only key can be used to encrypt data but it cannot be used to decrypt the data. Since public-only keys are created by copying the public key information to the local Security Server, it means that the private key is not on the local Security Server. If you configure a key group that contains just public-only keys, the local Security Server will not be able to restore encrypted backups.

To create a Database Backup Agent backup rule:

- 1: Log into the Management Console as a user with `Policy` role permissions.
 - 2: Open an offline policy in the Management Console to display either the *Add Offline Policy* window or the *Edit Offline Policy* window.
-

The *Add Offline Policy* window is displayed when you create a new offline (DB2/IDS) policy. This is described in [“Creating Database Backup Agent policies” on page 8-134](#). The *Edit Offline Policy* window is displayed when you edit an existing offline (file system) policy. This is described in [“Editing Database Backup Agent policies” on page 8-139](#).

- 3: If you are creating a new policy, enter the policy name in the **Name** text-entry box.

This field is mandatory. The maximum number of characters is 64.

- 4: Enter a phrase or string in the **Description** text-entry box that helps you to identify the DB2 or IDS policy.

This field is optional. The maximum number of characters is 256.

- 5: If you are creating a new policy, select the radio button for the type of Database Backup Agent that is to use this policy next to the **Database Type** label.

Your choice is **DB2** or **IDS**.

- 6: If you are creating a new policy, click the **Add** button under **Backup Rules**.

The *Add Backup Rule* window opens.

- 7: If you are editing an existing policy, click the rule in the **Name** column of the **Backup Rules** table.

The *Edit Backup Rule* window opens.

- 8: Enter the name for the backup rule in the **Name** text-entry box.

This field is mandatory. The maximum number of characters is 64.

- 9: Enter a phrase or string in the **Description** text-entry box that helps you to identify the backup rule.

This field is optional. The maximum number of characters is 256.

- 10: Select **Allow** from the **Action** scroll-list to permit backups when the rules match the specified criteria.

By default, **Deny** is selected and backups are not allowed. You can leave it set to **Deny** and continue configuring the policy and hosts, but you must select **Allow** to permit database backups.

- 11: Enable the **Compression** checkbox to compress data while it is being backed up.

This step is optional. Compression is based upon the `zlib` algorithm.

- 12: Specify the time range during which the backup rule for this policy is applicable.

By default, there is no time constraint. Other rules permitting, data can be accessed at any time, 24/7. Enter a start and end time to restrict backups to a specific portion of the day. Time is specified using a 24 hour clock. The format of a time specification is `hh:mm:ss`. For example, 3 : 30 pm is specified as 15 : 30 : 00. Enter the start time in the **Valid from** text-entry box and the end time in the **To** text-entry box. Days of the week cannot be specified.

Time is relative to the Security Server system clock, not the Database Backup Agent system clock. If you configure a policy that includes a time constraint, and the Database Backup Agent system is in a different timezone than the Security Server system, you must compensate for the timezone difference. For example, if the Security Server is in the PST timezone, let's say San Jose, and the Database Backup Agent is in the EST timezone, let's say New York, and you want to permit access between 9 am to 5 pm, configure the policy to permit access between 6 am to 2 pm.

- 13: (DB2 only) Enter the DB2 instance user name in the **Instance** text-entry box.
This field is optional and is used to restrict policy application to specific database instance users. For example, if a database instance user named `rocko` performs database backups and restores, and you want only this database instance user to be recognized by the policy, enter `rocko` in the text-entry box. The maximum number of characters is 8.
- 14: (DB2 only) Enter the number of a DB2 partition in the **Partition** text-entry box.
This field is optional and is used to restrict the policy to just the specified database partition. The value must be a non-negative integer.
Note that DB2 v8.x always returns 0 to the agent, which in turn returns 0 to the Security Server. Since the value is always 0, this parameter should not be used in policies that are used to back up or restore DB2 v8.x databases. This parameter works as expected for DB2, v9.x, database backups and restores.
- 15: (DB2 only) Enter the DB2 alias in the **Alias** text-entry box.
This field is optional and is used to restrict the policy to just the databases identified by the alias. The maximum number of characters is 8.
- 16: (IDS only) Enter the IDS server name in the **Server Name** text-entry box.
- 17: (IDS only) Enter the IDS server number in the **Server Number** text-entry box.
- 18: Select the encryption algorithm to use in the **Encryption Algorithm** scroll-list.
Your choices are `None`, `AES128`, or `AES256`. `None` is the default. The **Key Group** text-entry box is displayed when you select an algorithm.
- 19: Click the **Select** button next to the **Key Group** text-entry box.
The *Select Key Groups* window opens. All configured RSA key groups are displayed. Enable the **Select** radio button of the key group to apply when making backups. The selected key group will be used to encrypt the AES key that is included in the backup metadata.
- 20: Click **Select Key Group**.
The *Add Backup Rule* window opens and the selected key group is displayed in the **Key Group** text-entry box.
- 21: Enter additional backup parameters and their values in the **Additional Conditions** text-entry boxes.
Additional conditions are optional. These are parameters and values that are placed in the DB2 backup or restore command like `COMPROPTS` options. These parameters and values are passed to the Encryption Expert Security Server and used for policy evaluation. Refer to `./agent/db2/README` or `/agent/ids/README` for additional details.
The maximum number of characters allowed in the **Name** or **Values** text-entry boxes is 64.
- 22: Click **Ok**.
The *Add Offline Policy/Edit Offline Policy* window opens and the *Backup Rules* table displays the new or edited rule.
- 23: Click **Ok** to save the policy.

Configuring Database Backup Agent restore rules

Note that you cannot use public keys to decrypt an encrypted backup.

To create a Database Backup Agent restore rule:

- 1: Log into the Management Console as a user with `Policy` role permissions.
- 2: Open an offline policy in the Management Console to display either the *Add Offline Policy* window or the *Edit Offline Policy* window.

The *Add Offline Policy* window is displayed when you create a new offline (DB2) policy. This is described in [“Creating Database Backup Agent policies” on page 8-134](#). The *Edit Offline Policy* window is displayed when you edit an existing offline (file system) policy. This is described in [“Editing Database Backup Agent policies” on page 8-139](#).

- 3: If you are creating a new policy, enter the policy name in the **Name** text-entry box.

This field is mandatory. The maximum number of characters is 64.

- 4: Enter a phrase or string in the **Description** text-entry box that helps you to identify the DB2 or IDS policy.

This field is optional. The maximum number of characters is 256.

- 5: If you are creating a new policy, select the radio button for the type of Database Backup Agent that is to use this policy next to the **Database Type** label.

You choice is **DB2** or **IDS**.

- 6: If you are creating a new policy, click the **Add** button under **Restore Rules**.

The *Add Restore Rule* window opens.

- 7: If you are editing an existing policy, click the rule number in the **Name** column of the **Restore Rule** table.

The *Edit Restore Rule* window opens.

- 8: Enter the name for the backup rule in the **Name** text-entry box.

This field is mandatory. The maximum number of characters is 64.

- 9: Enter a phrase or string in the **Description** text-entry box that helps you to identify the backup rule.

This field is optional. The maximum number of characters is 256.

- 10: Select **Allow** from the **Action** scroll-list to permit backups when the rules match the specified criteria.

By default, **Deny** is selected and backups are not allowed. You can leave it set to **Deny** and continue configuring the policy and hosts, but you must select **Allow** to perform database backups.

- 11: Specify the time range during which the backup rule for this policy is applicable.

By default, there is no time constraint. Other rules permitting, data can be accessed at any time, 24/7. Enter a start and end time to restrict backups to a specific portion of the day. Time is specified using a 24 hour clock. The format of a time specification is *hh:mm:ss*. For example, 3 : 30 pm is specified as 15 : 30 : 00. Enter the start time in the **Valid from** text-entry box and the end time in the **To** text-entry box. Days of the week cannot be specified.

Time is relative to the Security Server clock, not the agent host clock. If you configure a policy that includes a time constraint, and the File System Agent system is in a different timezone than the Security Server system, you must compensate for the timezone difference. For example, if the Security Server is in the PST timezone, let's say San Jose, and the File System Agent is in the EST timezone, let's say New York, and you want to permit access between 9 am to 5 pm, configure the policy to permit access between 6 am to 2 pm.

- 12: (DB2 only) Enter the DB2 instance user name in the **Instance** text-entry box.

This field is optional and is used to restrict policy application to specific database instance users. For example, if a database instance user named `rocko` performs database backups and restores, and you want only this database instance user to be recognized by the policy, enter `rocko` in the text-entry box. The maximum number of characters is 8.

- 13: (DB2 only) Enter the number of a DB2 partition in the **Partition** text-entry box.

This field is optional and is used to restrict the policy to just the specified database partition. The value must be a non-negative integer.

Note that DB2 v8.x always returns 0 to the agent, which in turn returns 0 to the Security Server. Since the value is always 0, this parameter should not be used in policies that are used to back up or restore DB2 v8.x databases. This parameter works as expected for DB2, v9.x, database backups and restores.

- 14: (DB2 only) Enter the DB2 alias in the **Alias** text-entry box.

This field is optional and is used to restrict the policy to just the databases identified by the alias. The maximum number of characters is 8.

- 15: (IDS only) Enter the IDS server name in the **Server Name** text-entry box.

- 16: (IDS only) Enter the IDS server number in the **Server Number** text-entry box.

- 17: Click **Ok**.

The *Add Offline Policy/Edit Offline Policy* window opens and the *Restores Rules* table displays the new or edited rule.

- 18: Click **Ok** to save the policy.

Editing Database Backup Agent policies

You can change policy attributes. You will frequently change a policy while it is being developed and tested. You will occasionally need to enable/disable policies, allow/deny the operation, change keys, change hosts and host groups, and so on.

To edit an offline (DB2/IDS) policy:

- 1: Log into the Management Console Web interface as a user with `policy` role permissions.
- 2: Select **Policies** in the menu bar.
The *Policies* window opens.
- 3: Click an offline policy in the **Name** column.

The *Edit Offline Policy* window opens.

See also [“Configuring Database Backup Agent backup rules” on page 8-135](#) and [“Configuring Database Backup Agent restore rules” on page 8-137](#).

Deleting Database Backup Agent policies

To delete policies:

- 1: Log into the Management Console Web interface as a user with `policy` role permissions.
 - 2: Select **Policies** in the menu bar.
The *Policies* window opens.
 - 3: Enable the **Selected** checkbox for each policy to delete.
-

- 4: Click **Delete**.

You are prompted to verify the deletion.

- 5: Click **Ok** to delete the selected policies.

If the policy is in use with an active guard point, the operation will fail and the *Operation Aborted* window is displayed.

If the policy is not in use, it is removed from the *Policies* window.

About File System Agent policies

Because of their complexity and flexibility, File System Agent policies are described in their own chapter. See [“Creating and editing File System Agent policies” on page 9-141](#).

Creating and editing File System Agent policies

The Policy Composer is used to configure file system policies. The Policy Composer is a Java applet that runs separately from the main Management Console interface. This allows you to edit policies and do something else in the Management Console, like view logs, simultaneously. This section describes the parts of the Policy Composer and how to configure file system policies. Database policies are configured in the Management Console. For information about configuring database policies, see [“Using policies” on page 8-131](#).

About File System Agent policies

The Policy Composer is a separate Java applet that is started within the Management Console. This section describes opening the Policy Composer to edit File System Agent policies.

Encryption Expert File System Agent policies are used to protect:

- System files
- Data files and folders
- Applications




Figure 9.1: Security rule components

Navigating the Policy Composer

To navigate and use the Policy Composer:

- 1: Log into the Management Console Web interface as a user with `Policy` role permissions.
- 2: Select **Policies** in the menu bar.
The *Policies* window opens.
- 3: Click **Add Online Policy** to create a new policy or click an existing file system policy in the **Name** column to modify the policy.

File system policies are indicated by the icon, ()

The *Policy Composer* window for file systems opens to the *Security Rules* tab.




Figure 9.2: Policy Composer window

The icons along the top of the *Policy Composer* window are described in the following table.

Icon	Tooltip	Description
	Exit	Exits the Policy Composer. Appears on every tab.
	New Policy	Opens an empty <i>Policy Composer</i> window in which to configure a new policy. If you had modified an existing policy, or you had started to configure security rules, the Encryption Expert opens a dialog box that asks if you want to save your work before continuing.
	Save Policy	Save the changes made to the policy. If this is a new policy, the Encryption Expert opens a dialog box that asks you to name the policy. If this is an existing policy, the policy is saved without any prompts.
	Validate Policy	Verifies the policy syntax. Note also that the syntax is automatically checked when you add or replace a key rule.
	Reset Policy	Discards all the changes made since the last time the policy was saved. This is the equivalent to Revert To Saved in some Windows applications.

The **Allow Editing** checkbox on the top of the *Policy Composer* window is used to allow or prevent changes to the current policy. This checkbox is enabled by default. Disable the checkbox when you are satisfied with the policy. You may re-enable it at any time.

The row of tabs across the top of the window are:

- **Security Rules** -- specifies the conditions upon which to apply the rule and whether or not to allow or deny access
- **Key Selection Rules** -- specifies the key to use to encrypt resources in the guard point, and an optional subset of resources to which to apply the key
- **Data Transformation Rules** -- specifies an alternate encryption key that is used to cycle through keys for security maintenance or to restore data to its non-encrypted state

The group of buttons, text-entry boxes, and toggles on the top of the window are used to configure all the attributes of the current security rule.

The toggle beneath the current security rule attributes group, **Warn Mode**, is used to test policies by tracking how rules are evaluated, without enforcing the policy. You can monitor the log to determine how data is being accessed and then modify the policy accordingly.

The row of buttons in the middle of the *Policy Composer* window are security rule edit functions that allow you to add and delete security rules, edit existing rules, move rules up and down the scroll-list to change rule precedence, and so on. These are described in [“Using security rule edit functions” on page 9-144](#).

The security rules scroll-list on the bottom half of the window, displays all the security rules in the policy. Select a security rule by double-clicking anywhere in the row of the desired rule. The buttons, toggles, and text-entry boxes on the top of the window change to display the attributes of the selected rule. Once selected, the rule may be edited or moved up or down the scroll-list to change its order of precedence.

Using security rule edit functions

The row of buttons in the middle of the *Policy Composer* window are edit functions used to create new security rules and manipulate existing ones.

The functions are:

Function	Description
Add	Used to create a new security rule or copy an existing security rule, Add turns the current security rule attributes displayed on the top of the <i>Policy Composer</i> window into a security rule, placing it in the scroll-list on the bottom of the window. You cannot add a security rule with the same Label . You can copy a security rule by editing the security rule and changing the label, then click Add .

Function	Description
Replace	Replaces the configuration of the selected security rule in the scroll-list with the current security rule attributes that are displayed on the top of the <i>Policy Composer</i> window. Define the attributes for a new security rule or edit an existing security rule. Change the attributes as needed. Select a security rule in the scroll-list. Click Replace . The selected security rule is reconfigured with the current security rule attributes.
Edit	Opens the selected security rule for editing. The attributes of the security rule are displayed on the top of the <i>Policy Composer</i> window. You can change the configuration of the individual attributes. Select a security rule in the scroll-list. Click Edit . The attributes of the security rule are displayed and they are ready for editing. You can perform the same function by double-clicking a security rule in the scroll-list.
Reset	Clears the security rule attributes displayed on the top of the <i>Policy Composer</i> window and stops editing the current security rule. Changes to visible attributes, such as the Label, Action, and Effect that you can see on the <i>Policy Composer</i> window, are discarded; however, lower-level changes, such as creating a new resource set or time set, remain intact.
Remove	Deletes the selected security rule from the policy. You can restore the deleted security rule by exiting the <i>Policy Composer</i> window without saving the changes. If you save the policy after deleting a security rule, the security rule is gone.
Up	By default, security rules are evaluated in sequential order, beginning with the first security rule at the top of the scroll-list, down towards the last on the bottom of the scroll-list. When a match is found, the security rule may be used and the remaining security rules ignored. Up moves the selected security up the scroll-list, giving it greater precedence.
Down	Down moves the selected security down the scroll-list, giving it less precedence.

It always good practice to save the policy after reaching major editing milestones.

Configuring file system security rules

A policy is comprised of security rules and sometimes key rules. A security rule defines the users or user groups authorized to have specified access to specific files, directory paths for a designated period of time. In short, it defines *who* is accessing data, *what* are they doing with the data, *where* the data is located, *when* the Security Rule is applicable, and *how* the data can be accessed.

See also “[Planning file system policies](#)” on page 16-297.

Opening the Policy Composer

Policies are configured with the Policy Composer.

To open the Policy Composer:

- 1: Log into the Management Console Web interface as a user with `Policy` role permissions.
- 2: Select **Policies** in the menu bar.
The *Policies* window opens.
- 3: Click **Add Online Policy** to create a new policy or click an existing file system policy in the **Name** column to modify the policy.

File system policies are indicated by the icon, ().

The *Policy Composer* window for file systems opens to the *Security Rules* tab.




Figure 9.3: A configured file system policy

The major security rule components are:

- the files and directories that require protection (optional) (“[Configuring resources](#)” on page 9-146 and “[Enabling the Allow Browsing function](#)” on page 9-154)
- the users to which to apply the rule (optional) (“[Configuring users](#)” on page 9-155)
- the programs, applications, and libraries to which to apply the rule (optional) (“[Configuring processes](#)” on page 9-159)
- the time range when this rule is active (optional) (“[Configuring time](#)” on page 9-163)
- the manner in which the resource is being accessed (mandatory) (“[Configuring access methods](#)” on page 9-167)
- the effect for the action (mandatory) (“[What happens now?](#)” on page 9-169)

Additional components may be necessary, such as including key rules for encrypting and decrypting protected data. These additional components are noted, as they are needed, with a description of their use or a cross-reference to the description.

Configuring resources

A resource specifies the files and directories to which to apply the security rule. A resource set is a collection of directories, files, or both, to be protected. If you do not configure a resource set for the security rule then all files and directories in the guard point are assumed.



If you are configuring a rekey policy, add a security rule to protect “dataxform_*” files, as described in [“Create Security Rules” on page 10-189](#).

Resource sets are reusable. A resource set can be used in other security rules.

Adding a resource set

To create a resource set:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: Click the **Resource** button.
The *Resource Set List* window opens.
- 3: Click **Add Resource Set** or select **Resource->Add a new set** to create a new resource list.
The *Add New Resource Set* window opens.
- 4: Enter the name to assign the new resource list in the **Resource Set** text-entry box.



The following ASCII characters may be used to name resource sets: a-z, A-Z, 0-9, and the underscore (_). Spaces and other non-alphanumeric characters are not allowed.

- 5: Click **Add Resource**.
The *Resource Object* window opens.
- 6: Enter a full or relative path of the start point in the **Directory** text-entry box, or, click the **Browse** button to open a browser in which to select the path on the agent system.

A resource may be a combination of a directory, a file, and patterns or special variables. This step describes how to specify resources using the browser to select specific directories and files. You can optionally include variables and patterns to specify a set of resources.



The directory (or directory path) specified in a resource set is appended to the guard point. This means if the guard point is `/mnt/remote2` and the resource set directory path is `/remoteDir`, then the policy is applied to the files and directories in `/mnt/remote2/remoteDir`.

- 7: Click the **Browse** button.
The *Remote Object Chooser* window opens.
 - 8: Select a path in this window, as described in [“Selecting a remote file, path, or object” on page 9-173](#).
The directory of the selection is displayed in the **Directory** text-entry box and the file in the **File** text-entry box.
-




Figure 9.4: Adding a new resource

- 9: You can modify the **Directory** and **File** text-entry boxes to include variables or patterns to extend resource selection.

A resource can include a directory path fragment to the data objects to be protected. The directory path you specify determines the extent to which directories match the resource specification. The following table shows several different ways to specify the same directory and how it affects the resource path.

Directory Pattern	Result
sales	There is not enough information for the Security Server to determine the file system of the host. A dialog box will open asking you to indicate the file system type. File systems types are UNIX and Windows . Once the file system type is specified, the Policy Composer automatically adds a backslash (\) for Windows or a slash (/) for UNIX to the directory name. The result is / sales for UNIX or \sales for Windows, and access is limited to only the sales directory directly under the guard point. Windows agents are not supported at this time so always indicate the UNIX file type.
sales/	Access is limited to only the sales directory directly under the guard point.
*sales/	Access is limited to the directories that end in “ sales” directly under the guard point.
*/sales/	Access is limited to only the sales directory directly under the guard point.
*/sales	Access is granted to any directories named "sales" anywhere under the guard point.

Directory Pattern	Result
*/sales	Access is granted to any directories that end in "sales" anywhere under the guard point.

The following table lists the optional arguments you can add to your resource specification.

	Usage
pattern	Two meta-characters are supported in path definitions for policy resources and processes. They are the asterisk and question mark (* and ?). The asterisk indicates one (1) to many characters, and the question mark indicates exactly one (1) character. The meta-characters are used in simple string comparisons between an expression you specify and the physical system files and directories. If there are zero instances of characters for files or directories, the string comparison will mismatch. For example, if the pattern is "/a/*/c", and the directory "/a/c" exists, there is no match because the asterisk (*) represent 0 characters in the comparison string, resulting in "/a//c". /a//c is not the same string as /a/c, so there is no match. The asterisk meta-characters is also part of a "floating directory" feature. The asterisk can stand for one or more directory levels. For example, /a/*/z/ will match /a/b/z/ and /a/b/c/d/z/. All instances of the /z directory anywhere under /a will result in a match. The question mark indicates exactly one character in a string comparison, so it cannot find multiple directories. If the pattern is /a??/z, then only the two-character directories immediately below /a that contain /z will match.
variables	The two variables are based upon the user name or the group name of the user. The variables are uname and gname , respectively. On UNIX systems, uid and gid may also be used. When the security rule is applied, the variable is replaced by the actual user name or user group name. gname and gid work like uname , except they use the user group affiliation to determine access. For example, set Directory to /opt/local/ gname . Later, when you make /opt/local the guard point, only the members of the group specified in Users are allowed access to /opt/local/engineering.





Figure 9.5: Defining the resource directory with |uname|

- 10: The path specified in this window is relative to the guard point. If you use the browser, you may need to modify the directory path in the **Directory** text-entry box because the browser returns a full path and appends it to the guard point path.
- 11: Enable **Include subfolders** to recursively protect all similarly named files defined in the resource pattern.



Include subfolders finds all occurrences of the resource pattern under the guard point and applies policy protection to them. This allows you to define one resource set and apply it to similarly named files anywhere under the guard point.



- 12: Click **Add & New** if you want to add more resources to the current resource set.

Repeat the process for each additional resource set.

If you do not provide sufficient path information for the Policy Composer to determine the file system of the host, a dialog box opens asking you to indicate the file system. File systems are UNIX or Windows. The dialog prompts with “Type of operating system file path convention.” Click **UNIX**.




Figure 9.6: Verifying host operating system.

- 13: Click **Add & Close**.

If you do not provide sufficient path information for the Policy Composer to determine the file system of the host, the dialog box in [Figure 9.6](#) opens and prompts you to indicate the file system. Click **UNIX**.

The *Add New Resource Set* window opens and displays the resources in the resource set.

- 14: Click **OK**.

The *Resource Set List* window opens and displays the new resource set and the individual resources that are in the set.




Figure 9.7: Resource with paths added

- 15: Click **Select**.

The *Resource Set List* window closes and you are returned to the main Policy Composer window. The new resource name is displayed in the **Resource** text-entry box.

Resource	resourceSet01	☐ Exclude	☒ Allow Browsing
User		☐ Exclude	
Process		☐ Exclude	
When		☐ Exclude	
Action			
Effect			

Figure 9.8: Security Rule entry with a resource

- 16: Enable or disable the **Exclude** toggle to the right of the **Resource** text-entry box.

Disable this toggle to include just the resources in the resource set. Enable this toggle to include all host resources except those resources in the resource set. That is, use the inverse of the defined resource set. The toggle is disabled by default.

Editing a resource set

To edit an existing resource set:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: Click the **Resource** button.

The *Resource Set List* window opens.

- 3: Open a resource set for editing. You may:
 - Double-click a resource set in the scroll-list.
 - Select a resource set in the scroll-list and click **Properties**.
 - Select a resource set in the scroll-list and select **Resource->Edit the set**.

The *Edit Resource Set* window opens. It displays all configured resources in the selected set.

value	attribute	subfolders
/bcobain/dir2/file1	all	yes
/(uname)/*	all	yes

Figure 9.9: Editing resource sets

- 4: To change the attributes of an existing resource:
 - a: Open the resource to be modified. You may:
 - Double-click a resource in the scroll-list.
 - Select a resource in the scroll-list and click **Properties**.

- Select a resource set in the scroll-list and select **Set->Edit the selected member**.

The *Resource Object* window opens.

b: You can :

Change the path to the resource by manually entering a new path or using the browse button to select the path in the **Directory** text-entry box.

Change the resource itself by entering a file name, a combination of text and special characters, or just special characters, in the **File** text-entry box.

Change the extent to which the rule is applied in the file hierarchy by enabling or disabling the **Include subfolders** checkbox.

c: Click **OK**.

The window closes and you are returned to the *Resource Set List* window.

5: To add a resource to the current resource set:

a: Open the resource set to be modified. You may:

- Double-click a resource set in the scroll-list.
- Select a resource set in the scroll-list and click **Properties**.
- Select a resource set in the scroll-list and select **Set->Edit the set**.

The *Edit Resource Set* window opens.

b: Open the resource to be modified. You may:

- Double-click a resource in the scroll-list.
- Select a resource in the scroll-list and click **Properties**.
- Select a resource set in the scroll-list and select **Set->Edit the selected member**.

The *Resource Object* window opens.

c: Enter the resource information.

d: Click **OK**.

The *Resource Set List* window opens.

6: To delete individual resources from the current resource set:

a: Open the resource set to be modified. You may:

- Double-click a resource set in the scroll-list.
- Select a resource set in the scroll-list and click **Properties**.
- Select a resource set in the scroll-list and select **Resource->Edit the set**.

The *Edit Resource Set* window opens.

b: Select one or more resources in the resource set.

Single-click a resource in the scroll-list to select an individual resource. Hold the <Ctrl> key down to select multiple resources. Hold the <Ctrl> and <Shift> keys down to select a range of resources.

c: Click **Remove**.

d: Click **OK**.

7: Select the resource set in the scroll-list if it is not currently selected.

8: Click **Select**.

The window closes and you are returned to the main Policy Composer window.

- 9: Enable or disable the **Exclude** toggle to the right of the **Resource** text-entry box.
Disable this toggle to include just the resources in the resource set. Enable this toggle to include all host resources except those resources in the resource set. That is, use the inverse of the defined resource set. The toggle is disabled by default.

Deleting a resource set

This section describes how to delete a resource set. Deleting a resource set removes it entirely, and other policies configured with the resource set will be affected. Do not delete a resource set if it is being used by other policies.

To delete a resource set:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: Click the **Resource** button.
The *Resource Set List* window opens.
- 3: Single-click a resource set in the scroll-list.
- 4: Delete the resource. You may:
 - Click **Remove**.
 - Select **Resource->Remove the set**.

A dialog box opens, warning you that the resource set may be part of other active policies, and that removing the resource set of an active policy can cause other policies to fail.

- 5: Click **OK** to remove the resource set.
- 6: Click **Replace** in the main *Policy Composer* window if you are editing an existing security rule, otherwise click **Add**.

The changes you made are applied to the security rule.

- 7: Click the **Save** icon or select **Policy->Save** to keep your changes.

Enabling the Allow Browsing function

The browse function grants access to directories between a guard point and resource files, allowing you to traverse the directory hierarchy to the resources in the resource set. Enable the **Allow Browsing** attribute of a security rule to access and traverse the directories that are below the guard point, leading down to the resources in the rule. Users that match the criteria set by the security rule can access the directories between the guard point and the resource. The files in these directories can be listed like any file in a browser, but they cannot be modified, copied, or deleted.

The browse function is enabled by default. To enable or disable the browse function:

- 1: Open a policy in the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: Double-click a security rule in the scroll-list on the bottom of the *Policy Composer* window, or select the rule in the scroll-list and click **Edit**.
The components of the security rule are displayed in the text-entry boxes. The security rule is now ready for editing.
- 3: Enable or disable the **Allow Browsing** checkbox next to the **Resource** text-entry box to permit or deny access to the directories between the guard point and the resources.




Figure 9.10: Enabling the Allow Browsing toggle

The effect of the **Allow Browsing** checkbox is local to the current security rule resource set.

- 4: Click **Replace** in the main *Policy Composer* window if you are editing an existing security rule, otherwise click **Add**.

The changes you made are applied to the security rule.

- 5: Click the **Save** icon or select **Policy->Save** to keep your changes.

Configuring users

A user defines the security context of **Who** is accessing resource files and directories. A user can be an individual user, set of users, or even an application. For example, when users authenticate to an DB2 database, the Encryption Expert solution relies upon database access controls and trusts that the DB2 application authenticated the user.

A user set is a collection of individual users. User sets are reusable. Any user set in the current Security Server configuration can be used by any other policy.

Note that on UNIX systems, user credentials are typically set on File System Agent hosts via PAM or LAM, or are overridden by entering special keywords in the Management Console *Host Settings* tab. Without proper PAM or LAM authentication, the user credentials may not be accepted by the Security Server, which triggers an entry in the log that states, “User Not Authenticated”, and an otherwise valid user is denied access.

Compare the two log entries below. The first log entry is the result of executing the `cat` command in a guard point without PAM configured. The policy requires a user named “bc”. The second log entry is the result of doing the same thing with the same policy and guard point, but with the PAM properly configured:

```
Oct 30 18:18:26 vmlinux02 kernel: [AUD] [INFO ] [EET2602I]
AUDIT: Policy[permitAll] User[bc,uid=1001 (User Not
Authenticated)] Process[/bin/ls] Access[read_dir] Res[/]
Effect[PERMIT Code (1M)]
Oct 30 18:19:23 vmlinux02 kernel: [AUD] [INFO ] [EET2602I]
AUDIT: Policy[permitAll] User[bc,uid=1001,gid=10\wheel\]
Process[/bin/ls] Access[read_dir] Res[/] Effect[PERMIT
Code (1M)]
```

If you configure **User** security rule attributes, and you experience this problem, see [“Configuring host settings” on page 6-102](#).

Defining users

A user can be identified using a combination of one or more attributes. A specific user can be identified using the user name, user identification number, and user group number, or,

multiple users can be identified by group affiliation alone. There are other attributes you can use to identify users. The user attributes you specify determine the individual, or group of individuals, affected by the security rule.

User identification is optional. If you do not define users in the security rule, then user identification is not used to determine access permissions.

To create a new user definition:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).

- 2: Click the **User** button.

The *User Set List* window opens.

- 3: Click **Add User Set** or select **User->Add a new set** to create a new user set.

The *Add New User Set* window opens.

- 4: Enter the name to assign the new resource list in the **User Set** text-entry box.



The following ASCII characters may be used to name resource sets: a-z, A-Z, 0-9, and the underscore (_). Spaces and other non-alphanumeric characters are not allowed.

- 5: Click **Add User**.

The *User Object* window opens. You must configure at least one field in order to use this feature.

Figure 9.11: Specifying user objects

- 6: Enter the user name in the **Uname** text-entry box.

This text-entry box is used to identify a user based upon their login name.

- (UNIX) This attribute is case-sensitive. Enter the user name exactly as it exists in the `passwd` file.
- (Windows) Enter the user login name. This attribute is not case-sensitive.

- 7: (UNIX) Enter the user identification number in the **Uid** text-entry box.

The **Uid** attribute is ignored on Windows hosts, so you can define this attribute and apply the user definition to security rules for both UNIX and Windows hosts.

- 8: (UNIX) Enter the user group number in the **Gid** text-entry box.

If you use this text-entry box, enter only the primary group ID number of the user. Do not use secondary, tertiary, etc. group ID numbers.

The **Gid** attribute is ignored on Windows hosts, so you can define this attribute and apply the user definition to security rules for both UNIX and Windows hosts.

- 9: Enter the names of the groups to be granted/denied access permissions in the **Group Names** text-entry box.

Multiple group names may be entered. Separate group names with a comma.

Enter the group name exactly as it exists in the group file or group alias file. You are not restricted to entering the primary group. You can enter any group in which the user is a member.

- 10: (Windows) Enter the network domain of the user in the **OS Domains** text-entry box.

Multiple domain names may be entered. Separate domain names with a comma.

- 11: Click **Add & New** if you want to add more users to the set.

If you click **Add & New**, the user information is added to the **Add New User Set** scroll-list. All the data on the *User Object* window is removed and the window is ready for another user definition.

- 12: Click **Add & Close** if you are done adding users to the set.

The *Add New User Set* window opens. The new user set is displayed, along with its attributes, in the scroll-list.

- 13: Click **OK**.

The *User Set List* window opens.




Figure 9.12: Viewing user sets

- 14: If it is not already selected, select the user set you just created from the scroll-list on the left of the window.

- 15: Click **Select**.

The *Policy Composer* window opens and displays the selected user set in the **User** text-entry box.




Figure 9.13: Security Rule entry with a user set

- 16: Enable or disable the **Exclude** toggle to the right of the **User** text-entry box.
Disable this toggle to include just the users in the user set. Enable this toggle to include all host users except those users in the user set. That is, use the inverse of the defined user set. The toggle is disabled by default.
- 17: Click **Replace** in the main *Policy Composer* window if you are editing an existing security rule, otherwise click **Add**. You cannot click **Add** unless there is an **Effect** defined, too.

Editing user sets

The same sets can be used in multiple policies. Changes to a set are applied to all policies that use that set. Do not modify a set unless you know how it is used in other policies and how your changes will affect them.

To modify user sets:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: Click the **User** button.
The *User Set List* window opens.
- 3: Open the user set to be modified. You may:
 - Double-click a user set in the scroll-list.
 - Select the user definition and select **User->Edit the selected member**.
 - Select a user set in the scroll-list and click **Properties**.
The *Edit User Set* window opens.
- 4: If you want to change an existing user definition:
 - a: Open the *User Object* window for a user definition. You may:
 - Double-click the user definition.
 - Select the user definition and select **User->Edit the selected member**.
 - Select the user definition and click **Properties**.
 - b: Change the fields as described in steps 6 through 13 starting on [page 156](#).
- 5: If you want to delete an existing user from the user set. You may:
 - Select one or more user definitions, right-click over the selected user definitions, and select **Remove** from the pop-up menu.
 - Select one or more user definitions and select **Set->Remove the selected members**.
 - Select the user definition and click **Remove**.

- 6: If you want to add a user to the user set. You may:
 - a: Click **Add User** or select **Set->Add new members**.
The *User Object* window opens.
 - b: Complete the fields as described in steps 6 through 13 starting on [page 156](#).

Deleting user sets

The same sets can be used in multiple policies. Do not delete a set if it impacts other policies.

To delete a user set entirely from the Security Server:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: Click the **User** button.
The *User Set List* window opens.
- 3: Select the user set to be deleted.
- 4: Click **Remove** or select **User->Remove the set**.
A dialog box opens, asking if you really want to delete the set. If the set is configured in another policy, the dialog box warns you of that possibility and recommends that you check the other policy before continuing with deletion.
- 5: Click **OK**.
- 6: Click **Close** to close the *User Set List* window without using the currently selected set.
- 7: If you want to use another user set to replace the deleted user set:
 - a: Select another set in the scroll-list.
 - b: Click **Select**.
The *Policy Composer* window opens.
 - c: Click **Replace** to add the new set to the current security rule.

Configuring processes


A process set is a collection of executables. These are any programs, utilities, or scripts that can be executed on the host system. Some examples are a mail program, backup utility, and any UNIX command. Process sets are reusable. Any process set in the current Security Server configuration can be used by any other policy.

When an attempt is made to start the executable, it is verified against the defined process. If the executable is in a specified location in the directory hierarchy, it is permitted to run and can access protected resources. If the executable does not match the criteria and does not reside under the guard point, then it is still started but is not granted access to protected resources.

A simple scenario is when you edit a text file using `vi`. The Security Server checks the resource criteria. If the text file is a defined resource and the user has access permissions to the file, the process criteria is checked. The process criteria is the directory path to the executable or just the executable name.

If the process definition includes a directory or full-path, the location of the executable being called is verified. If it resides in the directory, or the full path is correct, the executable is started and the file is opened. If the executable being called does not reside in the directory, or the path is incorrect, the executable is started but cannot open the file.

Even if you have write permission to the file, if the executable called to open the file is not in the directory, or is not in the full-path, the executable is started but cannot open the file.



Adding process sets

To create a new process set:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: Click the **Process** button.
The *Process Set List* window opens.
- 3: Click **Add Process Set** or select **Process->Add a new set** to create a new resource list.
The *Add New Process Set* window opens.
- 4: Enter the name to assign the new resource list in the **Process Set** text-entry box.



The following ASCII characters may be used to name resource sets: a-z, A-Z, 0-9, and the underscore (_). Spaces and other non-alphanumeric characters are not allowed.

- 5: Click **Add Process**.
The *Process Object* window opens.
- 6: Click **Browse** to select executables on the agent system, or enter the **Directory** and **Basename** values manually.
The *Remote Object Chooser* window opens when you click **Browse**. See [“Selecting a remote file, path, or object” on page 9-173](#) for details about completing this window.

After closing the *Remote Object Chooser* window, the *Process Set* window displays the path and executable information.

Directory is the directory path information. It may be a full path, a relative path, or nothing. **Basename** is the name of the executable. These text-entry boxes are automatically populated when you use the *Remote Object Chooser* window.

- 7: Click **Add & New** to add the process to the *Add New Process Set* window and leave the *Process Object* window open for additional processes, or click **Add & Close** to add the process and return to the *Add New Process Set* window.
- 8: Click **Ok** to close the *Add New Process Set* window and reopen the *Process Set List* window.




Figure 9.14: Completed Process Set List window

- 9: To use the new process set, select the process set if it is not already selected, and click **Select**.

The *Policy Composer* window opens. The new process set name is displayed in the **Process** text-entry box.

- 10: If you do not want to use the process set at this time, click **Close**.

The *Policy Composer* window opens.

- 11: Enable or disable the **Exclude** toggle to the right of the **Process** text-entry box.

Disable this toggle to include just the processes in the process set. Enable this toggle to include all host processes except those processes in the process set. That is, use the inverse of the defined process set. The toggle is disabled by default.

- 12: Click **Replace** in the *Policy Composer* window if you are editing an existing security rule.

The changes you made are applied to the security rule.

Editing process sets

To modify process sets:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: If the policy is using the process set to be modified, select a security rule that uses the process set from the scroll-list in the *Policy Composer* window and click **Edit**.
- 3: Click **Process**.

The *Process Set List* window opens.

- 4: Open the process set to be modified. You may:
 - Double-click a process set in the scroll-list.
 - Select a process set in the scroll-list and click **Properties**.
 - Select a process set in the scroll-list and select **Process->Edit the set**.

The *Edit Process Set* window opens.

- 5: If you want to change an existing process definition:
 - a: Open the *Process Object* window for a process definition. You may:
 - Double-click the process definition.
 - Select the process definition, right-click the process definition, and select **Edit** from the pop-up menu.
 - Select the process definition and select **Set->Edit the selected member**.
 - Select the process definition and click **Properties**.
 - b: Change the fields in the *Process Object* window as described in steps 6 through 8, starting on [page 160](#).
 - 6: If you want to delete processes from the process set, you may:
 - Select one or more processes in the *Edit Process Set* window, right-click over the selected processes, and select **Remove** from the pop-up menu.
 - Select one or more processes and select **Set->Remove the selected members**.
 - Select one or more processes and click **Remove**.
 - 7: If you want to add a process to the process set:
 - a: Click **Add Process** in the *Edit Process Set* window or select **Set->Add new members**.

The *Process Object* window opens.

 - b: Complete the fields in the *Process Object* window as described in steps 6 through 8, starting on [page 160](#).
- 8: To select the process set to apply to the current policy, select the policy set in the *Process Set List* window and click **Select**.

The *Policy Composer* window opens and displays the selected process in the **Process** text-entry box.

- 9: Click **Replace** in the *Policy Composer* window if you are editing an existing security rule.

The changes you made are applied to the security rule.

Removing process sets

To remove an existing process set:

- 1: Open the Policy Composer, as described in “[Opening the Policy Composer](#)” on [page 9-146](#).
 - 2: Click the **Process** button.
- The *Process Set List* window opens.
- 3: Select the process set to be deleted.
 - 4: Click **Remove** or select **Process->Remove the set**.

A dialog box opens, asking if you really want to delete the set. If the set is configured in another policy, the dialog box warns you of that possibility and recommends that you check the other policy before continuing with deletion.

- 5: Click **OK**.
- 6: Click **Close** to close the *Process Set List* window without using the currently selected set.
- 7: If you want to use another process set to replace the deleted process set:
 - a: Select another set in the scroll-list.
 - b: Click **Select**.
The *Policy Composer* window opens.
 - c: Click **Replace** to update the current security rule.

Configuring time

Time constraints satisfy the “when” criteria of the Security Server paradigm. A security rule can include time constraints to limit access based upon time, day of the week, calendar date, or a combination of these. You can restrict access based upon time for many reasons. A few being:

- business hours
- system or network load (e.g., limit backups to off-hours)
- unveil new applications on the enterprise network at a specific time and day
- one-day access for a limited-time promotion

By default there is no time constraint. Other rules permitting, data can be accessed at any time, 24/7.



Time is relative to the agent system clock, not the Security Server system clock. Time compensation is not required if the File System Agent system and the Security Server system are in different timezones. For example, if the Security Server is in the PST timezone, let's say San Jose, and the File System Agent is in the EST timezone, let's say New York, and you want to permit access between 9 am to 5 pm, configure the policy to permit access between 9 am to 5 pm.

Time sets are reusable. Any time set in the current Security Server configuration can be used in other security rules in other file system policies.

To create and apply a time set to the security rule:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
 - 2: If the policy is using the process set to be modified, select a security rule that uses the process set from the scroll-list in the *Policy Composer* window and click **Edit**.
 - 3: Click **When**.
The *Time Set List* window opens.
 - 4: Click **Add Time Set**.
The *Add New Time Set* window opens.
 - 5: Enter a name for the new time set in the **Time Set** text-entry box.
 - 6: Click **Add Time Object**.
-

The *Time Object* window opens.




Figure 9.15: Blank Time Object window with scroll-lists and drop-down menus

7: If you want to define a time constraint that includes the days of the week:

a: Select a start day in the left **Week Day** drop-down menu.

You can select a day of the week, like **Tue** (Tuesday) or **Fri** (Friday), or you can select **Only**. **Only** indicates that access is permitted only on the day selected in the right **Week Day** drop-down menu.

b: Select the end day in the right **Week Day** drop-down menu.

You can select a day of the week, like **Tue** (Tuesday) or **Fri** (Friday). If **Only** is selected in the left **Week Day** drop-down menu, access is permitted only on the day selected the right **Week Day** drop-down menu.



If you specify both days of the week and calendar dates, only the days of the week that occur between the specified dates are used.

- 8: If you want to define a time constraint that includes the hours of the day:
 Enter a specific start time in the left **Time** text-entry box using the time notation, *hh:mm:ss*, or click the ellipsis (...) next to the left **Time** text-entry box to open a window in which to set the desired start time. If you use the ellipsis to set the time, note that time is displayed in the window using standard A.M./P.M. notation where the maximum hour is 12, but the value entered in the **Time** text-entry box uses military notation where the maximum hour is 24.
 Instead of entering a specific time, you can select the **From** or **To** keyword from the left **Time** drop-down menu. **From** indicates that access is permitted beginning at the time entered in the right **Time** text-entry box up to 24:00:00. **To** indicates that access is permitted starting at 00:00:01 up to the time entered in the right **Time** text-entry box.
- 9: If you want to define a time constraint that includes calendar dates:
 Enter a specific start date in the left **Date** text-entry box using the date notation, *mm/dd/yy* or *mm/dd/yy/hh:mm:ss*. The first example indicates just a calendar date, such as *Jul/4/07* (leading zeros in the day can be dropped). The second example is a calendar date appended with the time. It is your choice which format to use. If you decide to specify just the date in one **Date** text-entry field, you must also specify just the date in the other **Date** text-entry field.
 Instead of entering a specific time, you can select the **From** or **To** keyword from the left **Date** drop-down menu. **From** indicates that access is permitted beginning at the date entered in the right **Date** text-entry box. **To** indicates that access is permitted up to the date entered in the right **Date** text-entry box.



If you configure **Time** and **Date**, and you also include time in the **Date** text-entry boxes, only the overlapping time is used. The same applies to **Week Days** and **Date**. Only the days of the week that fall between the dates are used.




Figure 9.16: Determining actual time in time sets

You can select **Only** from the left **Date** drop-down menu to limit access to just the date specified in the right **Date** drop-down menu. When **Only** is selected, the format of the date in the right **Date** drop-down menu is *mm/dd/yy*. You cannot

include a time with the date. Either the time range specified in the **Time** drop-down menus is used or, if **Time** is not set, the entire day is implied.




Figure 9.17: Completed Time Object window.

The example in [Figure 9.17](#) configures the security rule so that access is only allowed Monday through Friday, from 9 AM to 5 PM, until December 31, 2008.

- 10: Click **OK**.

The *Add New Time Set* window opens.

- 11: Click **OK**.

The *Add New Time Set* window closes and you are returned to the *Time Set List* window.

- 12: If it not already selected, select the desired time set from the scroll-list.

- 13: Click **Select**.

The *Policy Composer* window opens. The new time set name is displayed in the **When** text-entry box.




Figure 9.18: Viewing the specified time set

- 14: Enable or disable the **Exclude** toggle to the right of the **When** text-entry box.

Disable this toggle to use the times specified in the time set. Enable this toggle to use times opposite to those specified in the time set. That is, use the inverse of the defined time set. The toggle is disabled by default.

- 15: Click **Replace** in the *Policy Composer* window if you are editing an existing security rule.

Configuring access methods

An “action” is an attempt to access protected data in some way. It can be an attempt to read, delete, rename, etc. The Policy Manager provides a wide range of access methods to specify precisely the action required to match the security rule. The action is identified, evaluated, and used as a factor in determining whether or not to grant access. Specifying an access method is optional. If you do not specify one, the Policy Composer defaults to **all_ops**, or “all operations”.

Actions are described below:

Action	Description
d_chg_att	Change directory attributes (e.g., <code>chown usr dir1</code>).
d_chg_sec	Pertains to attempts to change any security property of a Windows folder, such as you would on the <i>Security</i> tab of the Properties window.
d_mkdir	Make a new directory (e.g., <code>mkdir dir1 dir2</code>).
d_rd	View directory contents (e.g., <code>ls dir</code> ; <code>cd dir</code>).
d_rd_att	Read the attributes of a directory (e.g., <code>ls -la dir1</code>).
d_rd_sec	Pertains to attempts to view the security properties of a Windows folder, such as you would on the <i>Security</i> tab of the Properties window.
d_ren	Rename a directory (e.g., <code>mv dir1 dir2</code>).
d_rmdir	Delete a directory (e.g., <code>rm -r dir1</code>).
f_chg_att	Change file attributes (e.g., <code>chmod</code>).
f_chg_sec	Pertains to attempts to change any security property of a Windows file, such as you would on the <i>Security</i> tab of the Properties window.
f_cre	Create a new file.
f_link	Link to a file (e.g., <code>ln file_name link_name</code>).
f_rd	Read a file.
f_rd_att	Read the attributes of a file (e.g., <code>ls -l file</code>).

Action	Description
f_rd_sec	Pertains to attempts to view the security properties of a Windows file, such as you would on the <i>Security</i> tab of the Properties window.
f_ren	Rename a file.
f_rm	Delete a file.
f_wr	Write to an existing file.
f_wr_app	Append data to a file.
read	This method is a collection of the preceding methods that are related to reading files, directories, and their attributes. This method comprises f_rd , f_rd_attr , f_rd_sec , d_rd , d_rd_attr , and d_rd_sec .
write	This method is a collection of the preceding methods that are related to writing to files, directories, and their attributes. This method comprises f_wr , f_cre , f_ren , f_rm , f_link , f_chg_attr , f_chg_sec , d_ren , d_chg_attr , d_chg_sec , d_mkdir , and d_rmdir .
key_op	Key operations. This requires 2 keys. One key is specified in the <i>Key Selection Rules</i> tab and the other in the <i>Data Transformation Rules</i> tab. Key Selection keys are the encryption keys for the current security rule. Data Transformation keys are used to migrate data from an encrypted form to a non-encrypted form or to change the encryption keys used to access the data. See “Defining a Data Transformation Rule” on page 9-179 and “Changing encryption keys” on page 10-183 .
all_ops	All operations. That is, any attempt to access the data in any way. This is the default if no action is specified.

Using access actions

To define a set of access methods that deal with how protected data is being accessed:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: Click the **Action** button.
The *Select Actions* window opens. This window displays the access methods in the scroll-list on the left side.
- 3: Select one or more access methods.
You can select multiple access methods.
- 4: Click **Add >>**.

- 5: The selected access methods are moved to the scroll-list on the right of the window.
- 6: Click **OK**.

The window closes and you are returned to the main *Policy Composer* window.

- 7: Click **Replace** in the main *Policy Composer* window if you are editing an existing security rule.

The changes you made are applied to the security rule.

What happens now?

The last stop in security rule definition is to determine what to do once all the conditions set by the other security rules in the policy have been satisfied. Do you grant access to the resource? Do you grant access and track the ways in which the resource is being accessed? Do you jump to an alternate set of security rules for additional access evaluation? Determining the access to allow or the alternate steps to take is called an “effect”. You must include an effect in the security rule. You must configure an effect before you can add a security rule to the policy.

The following table describes the available effects:

Effect	Description
deny	Deny the access attempt to the resource. For example, you can specify an effect that denies any attempt to access a resource. permit and deny cannot be used in the same effect.
permit	Grant the access attempt to the resource. For example, you can specify an effect that allows writing to a directory. permit and deny cannot be added to the same effect.
apply_key	Applies an encryption key to data in a guard point. When applied, data copied into the guard point is encrypted with the key specified in the <i>Key Selection Rules</i> tab and data that is accessed in the guard point is decrypted using the same key. When used with the key_op Action and the <i>Data Transformation Rules</i> tab, the data in the guard point can be rekeyed to a different encryption key or to no encryption at all. apply_key operates as a toggle on a per resource set basis. If a resource set contains a mixture of encrypted and non-encrypted resources, when apply_key is applied, encrypted resources are decrypted with the key when they are accessed. Non-encrypted resources are encrypted with the key when they are accessed, and the data becomes unusable. If you select this effect, you must also specify the key rules to apply for encrypting and decrypting the resources. For information about specifying key rules, see “Defining a Key Selection Rule” on page 9-178.

Effect	Description
audit	Used in conjunction with permit or deny , audit creates a entry in the Management Console log that describes what is being accessed, when it is being accessed, the security rule being applied, and other statistical information that can help you to evaluate the performance and efficacy of the security configuration. See “Running in Warn Mode” on page 9-172 and “Analyzing log entries” on page 14-233 .

Adding and removing effects

An effect is applied when the conditions set in a security rule are matched. If the policy contains more than one security rule, the security rules are executed in sequential order. The first security rule that matches completely is used.

To configure the effect:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).

- 2: If you want to add or change the actions of an existing security rule, select the security rule and click the **Edit** button.

The components of the security rule are displayed in the text-entry boxes. The security rule is now ready for editing.

- 3: Click **Effect** on the *Policy Composer* window.

The *Select Effects* window opens. This window displays available effects in the left scroll-list and the current assigned effects in the right scroll-list.

- 4: To add effects to the current security rule:

- a: Select one or more effects in the left scroll-list.

Refer to the table above to determine effect compatibility.

- b: Click **Add >>**.

The selected effects are moved to right-scroll list.




Figure 9.19: Completed Select Effects window

- 5: To remove effects from the current security rule:
 - a: Select one or more effects in the right scroll-list.

b: Click **Remove** <<.

The selected effects are moved to the right-scroll list.

6: Click **OK**.

The *Select Effects* window is closed and you are returned to the *Policy Composer* window.

7: Click **Replace** in the main *Policy Composer* window if you are editing an existing security rule.

The changes you made are applied to the security rule.

8: Click the **Save** icon or select **Policy->Save** to keep your changes.



A rootkit is a system tool that can be exploited to hide malicious code, such as viruses, worms, backdoors, and spyware. Third-party tools, like F-Secure Blacklight™ and SysInternals RootkitReveal™, are available to check the system for the presence of a rootkit. If the anti-rootkit software detects files that it cannot access, such as Encryption Expert resources, the anti-rootkit software will flag these files as Trojan Horses and will remove them the next time the system reboots.

If you are running anti-rootkit software, include the anti-rootkit application in the policies you create and assign the anti-rootkit application **permit** or **apply_key** access permission.




Figure 9.20: Example showing defined security rules

Adding a default security rule

You should include a default security rule to use in the event no other security rule in the policy matches the access being performed. This ensures that all host data is protected, even though you create specific security rules to cover a small sampling of host data.

Leave the default security rule entirely blank, with the exception of **Action** and **Effect** on the *Security Rules* tab:

- Set **Action** to **all_ops**.
- Set **Effect** to **deny audit**.

This security will match any attempt to access any data on the host. When this security rule is encountered, the Security Server will deny access and send an ALERT message to the Security Server log. This way you protect the host while tracking all access attempts not defined by the other security rules in the policy.




Figure 9.21: Example showing default security rule

Since the default security rule is a guaranteed match, it must be evaluated last. Use the global options and the location in the scroll-list to make the default security rule the last security rule to be evaluated whenever an access attempt is made. Figure 9.21 shows a completed policy. The default security rule is last in the scroll-list. If no other security rule matches the access attempt being performed, the host data is protected by the default security rule.

For example, in the case where there are two security rules in a policy, the following example shows that there was a mismatch in the first security rule, but the second rule, the default rule, was guaranteed to match any access attempt:

- Rule-1M: All components matched in the first security rule
- Rule-1U: Mismatch in User in the first security rule
- Rule-2M: All components matched in the second security rule

```
EET2602I: [SecFS, 0] AUDIT: Policy[dailyMaint_policy]
User[bcobain,uid=1001,gid=10\wheel\] Process[/bin/ls]
Access[read_dir] Res[/] Effect[PERMIT Code (1M)]
```

```
EET2602I: [SecFS, 0] AUDIT: Policy[dailyMaint_policy]
User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\]
Process[/bin/ls] Access[read_dir] Res[/] Effect[DENIED Code
(1U,2M)]
```

For details about log entries and their interpretation, see [“Analyzing log entries” on page 14-233](#).

Running in Warn Mode

A debug option is provided where you can override the assigned Effect of all the security rules in a policy. When enabled, the Effect is set to **permit**. That is, regardless if the Effect


is set to **deny**, applications are granted complete access to protected data. When an access attempt is made that matches a security rule in the policy, access is granted and a log message is issued. The log data is useful for policy debugging or policy violation auditing.

When working with encrypted data, a **deny** statement must include **apply_key** when **Warn Mode** is enabled; otherwise, though normally denied read or write access, a denied user running in **Warn Mode** can overwrite or edit data and, if not granted **apply_key** (i.e., decrypt read/encrypt write), any changes to the data will corrupt that data.

Enable the **Warn Mode** toggle to disable the policy but track each access attempt that matches any security rule in the policy. When you are done testing the policy, disable the **Warn Mode** toggle to apply the policy and use the **Effect** defined by the security rules. Also, disable the **Warn Mode** toggle to stop issuing log messages and stop permitting access when there is no matching rule.



This option generates a warning each time an access attempt is made that matches any security rule in the policy. This warning is sent as a log message and it can be viewed in the Management Console. You must also configure the Management Console to accept warnings, otherwise the warnings are not tracked in the Security Server log. See [“Displaying the Security Server log” on page 14-229](#).



Selecting a remote file, path, or object

The *Remote Object Chooser* window allows you to browse the files and directories of a remote host. Some Policy Composer windows provide this option to browse a remote host to select files and directories, rather than entering the path information manually. It is easier to use the *Remote Object Chooser* and minimizes typographical errors.

The browser communicates across a link between the Security Server and an Encryption Expert-enabled host. Since the *Remote Object Chooser* window browses hosts on the network, you must configure the hosts before using this window. Otherwise, the hosts are not displayed. See [“Configuring hosts” on page 6-75](#) for information on configuring hosts.

To use the *Remote Object Chooser* window:

- 1: Open the Policy Composer, as described in “Opening the Policy Composer” on page 9-146.
- 2: Open the *Remote Object Chooser* window.

You can:

- Click **Resource**, **Add Resource Set**, **Add Resource**, **Resource Object**, and **Browse**.
- Click **Process**, **Add Process Set**, **Add Process**, and **Browse**.

The *Remote Object Chooser* window opens.




Figure 9.22: Selecting a remote object

- 3: To enter the path and file information manually:
 - a: Enter a path and/or file in the **Selection** text-entry box.

You can enter a full path or relative path. A string without any backslashes (/) is a file or process name. If the string contains backslashes, everything up to the last backslash is a directory path. Everything after the last backslash is a file or process name. If you want to specify a path but not a file or process name, place a backslash-asterisk (/*) at the end of the path.
 - b: Click **Ok**.

The path information you entered is displayed in the **Directory** and **File** text-entry fields of the *Resource Object* window or the **Directory** and **Baseline** text-entry fields of the *Process Object* window.
- 4: To browse for a path or file:
 - a: Select the host system with the paths and files you want to use in the **Host** drop-down menu.

All Security Server configured hosts are displayed in the drop-down menu.
 - b: Select the type of host information to display from the **Dir Only** drop-down menu.

Your choices are **Dir Only** and **Dir & File**. **Dir Only** displays just host directories. **Dir & File** displays both host directories and files.
 - c: Enter the directory path in which to begin browsing in the **Start Directory** text-entry box.

This optional step can be used to specify some portion of the path to the object you are browsing. For example, if you know the object is somewhere under /

`opt/local`, you can enter `/opt/local` in this text-entry box. When the browser starts, the search will begin at the specified directory.

d: Click **Go**.

- If you selected **Dir Only**, the directory tree for the host is displayed; starting at the top-level directory.
- If you selected **Dir & File**, the directory tree and the files within the directories are displayed; starting at the top-level directory.
- If you entered a starting path in the **Start Directory** text-entry box, the browser starts at that directory level.

e: Traverse the system hierarchy in the browser until the desired object is displayed.

f: Select the object.




Figure 9.23: Selecting a remote object

g: Click **OK**.

The path information you entered is displayed in the **Directory** and **File** text-entry fields of the *Resource Object* window or the **Directory** and **Basename** text-entry fields of the *Process Object* window.

The path specified in this window is relative to the guard point. If you use the browser, you may need to modify the directory path in the **Directory** text-entry box because the browser returns a full path and appends it to the guard point path.

Configuring encryption in file system security rules

Security rules are a collection of conditions that must be satisfied before allowing some level of access. Security rules are described starting with [“Configuring file system security rules” on page 9-145](#). The *Key Selection Rules* tab defines the encryption keys to apply to resources in a security rule. The *Data Transformation Rules* tab is used to rekey resources from one set of encryption keys to another, or to rekey resources into clear data, effectively removing encryption.

Applying key rules is optional. By default, resources are not encrypted. Once you encrypt resources, they remain encrypted. If you encrypt resources and remove the encryption key,

or remove File System Agent software from the host, the data becomes unusable until the encryption key is reapplied. The resources must be rekeyed to a non-encrypted state.



Changing or deleting keys or key rules can make host data unusable. Use caution in the application of encryption keys. Keep track of the encryption keys being used, taking care not to accidentally remove the keys. Also, do not simply change the key in the *Key Selection Rules* tab. Encrypted data in a guard point must be rekeyed in order to change to a different key or it must be rekeyed in order to change to a non-encrypted state. See “[Changing encryption keys](#)” on page 10-183 for details.




Figure 9.24: Key rules within a security rule



If you want to change a policy or rekey a guard point, be prepared to temporarily stop access to the guard point. Changing policies for a guard point requires an interruption of service because the transition process entails disabling one policy and then enabling another policy. The guard point must be inactive during the transition period to ensure guard point integrity. The same rule applies to moving a host between host groups when it includes a change in policies. Coordinate policy changes during a maintenance outage window.



Limited application of NTFS compression is permitted on Windows hosts. You can apply NTFS compression only if you do not apply a policy with data encryption. If you are going to configure the **Key Selection Rules** and **Key Selection Rules** tabs, check that the file systems to be protected by the policy do not have NTFS compression enabled. Applying a policy that includes encryption to NTFS-compressed file systems can corrupt the data.

To check if NTFS compression is enabled, right-click a folder to be protected by the policy and select **Properties**. Select the **Advanced ...** button in the **General** tab of the *Properties* window. Note if the **Compress contents to save disk space** checkbox is enabled in the *Advanced Attributes* window. You can leave the checkbox enabled and not include file encryption in your policy or you can disable the checkbox and configure the *Key Selection Rules* and *Key Selection Rules* tabs.

Navigating the Key Selection Rules tabs

The interface for editing key rules is similar to that for editing security rules. The two key rule tabs, *Key Selection Rules* and *Data Transformation Rules*, on the *Policy Composer* window, have the same basic interface as the *Security Rules* tab.




Figure 9.25: Key Selection Rules window

Key Selection Rule attributes

The two attributes of a key selection rule are:

- **Resource:** Opens the *Resource Set List* window from which you can select or create the resource set whose members are to be encrypted. Specifying a resource set is optional. If you do not specify a resource set here in the *Key Selection Rules* tab, encryption is applied to the resources specified in the *Security Rules* tab. The *Resource Set List* window is described in [“Configuring resources” on page 9-146](#).
- **Key:** Provides a drop-down menu from which to select the “clear” and configured AES encryption keys. Other than `clear_key`, default keys are not provided. Key generation is described in [“Creating a symmetric key” on page 7-119](#).

Using Key Selection Rule edit functions

The edit functions on the *Key Selection Rule* tab are the same as those on the *Security Rules* tab. See [“Using security rule edit functions” on page 9-144](#).

Using variables to specify encryption keys

Encryption keys can also be specified using two variables, `|uname|` and `|gname|`. The two variables correspond to the user name and user group name, respectively. When the key rule is applied, variables are replaced by the name or group name of the user making the access attempt. In this way you can apply one key rule to multiple users, where each user has his or her own encryption key.


An example of applying variables to encryption key names, is to define the resource of a security rule with `/home/|uname|`. In the *Key Selection Rules* tab, define the key name as `|uname|_key`. Then, if a user named `bubba` tries to access his home directory, `/home/bubba`, and matches the other components of the security rule, the encryption key

is also expanded to `bubba_key`. Having already made a key named `bubba_key`, the user can view the decrypted contents of his home directory. If the key does not exist, the user can still access the directory contents but they will be encrypted.

The `|gname|` variable works like `|uname|`, except `|gname|` uses the user group affiliation to determine access. For example, In the *Key Selection Rules* tab, define the key name as `|gname|_key`. Then, when a user attempts to access an encrypted resource, the group name of the user is determined and the encryption key name is expanded to include the group name. If the user belongs to the `engineering` group, the encryption key name is expanded to `engineering_key`.

Defining a Key Selection Rule

A key selection rule defines the encryption key to apply to a specific resource set or the encryption key to use as a default in the event no other key rule matches.



Do not change an encryption key if any files that use that key are open or being accessed in cache. A key change takes effect immediately; however, open files and cached data continue to use the old key. Changing keys while protected files and cached data are in use is unpredictable and can corrupt data.

To add encryption to a policy:

- 1: Open the Policy Composer, as described in [“Opening the Policy Composer” on page 9-146](#).
- 2: Create a policy with one or more security rules, or open an existing policy.
- 3: Verify that the guard points to which the policy is/will be applied are empty.

Encryption is applied as a simple toggle. If a guard point uses encryption, data is encrypted as it is copied into the guard point and it is decrypted as it is being read from the guard point. Therefore, you have to apply encryption to an empty guard point and then copy the data into that guard point. This encrypts the data as it is placed in storage. If the data is already present in the guard point and you apply encryption, the data is “decrypted” as you copy the contents of the guard point. If you decrypt something that isn’t encrypted, you end up with garbage.

You can also use the `dataxform` utility to convert between encrypted and unencrypted data. See [“Changing encryption keys” on page 10-183](#) for details on `dataxform`.

- 4: Save the policy.
- 5: Click the *Key Selection Rules* tab.
- 6: Click **Resource**.

A resource specification is optional. If you do not specify a resource, the key rule will always match. It is recommended that every *Key Selection Rules* set include one key rule with a blank resource. The key rule with the blank resource is placed at the bottom of the key rules scroll-list and is used as a default. In the event no other key rule matches, a default encryption key will be applied. In effect, a key rule with a blank resource indicates all the resources that do not match any of the preceding key rules.

If you click **Resource**, the *Resource Set List* window opens. Complete this window, as described in [“Adding a resource set” on page 9-147](#).

- 7: Click **Key** to reload the Security Server key list.

This step is optional. It refreshes the key list information on the Security Server.

- 8: Select an encryption key from the **Key** drop-down menu.

Encryption keys must be generated for them to be displayed in the drop-down menu. Only symmetric keys (AES, 3DES) are displayed and available for selection. For information about generating encryption keys, see [“Creating a symmetric key” on page 7-119](#).

A default key, `clear_key`, is provided with the Security Server. This key indicates that clear text (no encryption) is applied to the resource. `clear_key` is typically used for debug purposes or to rekey encrypted data back into clear text and restoring the resource to an unencrypted state.

- 9: Click **Add**.

The key rule is added to the scroll-list on the *Key Selection Rules* tab.

- 10: Add as many additional key rules that you like.

- 11: (Optional) Create a key with a blank resource that uses the encryption keys you want to use as a default in case there are no matches.

a: Click **Add**.

b: Use the **Up** and **Down** buttons to move the key rule with the blank resource to the bottom of the scroll-list, if it is not there already.

- 12: Select **Policy->Save** or click the **Save Policy** icon.

- 13: Apply the policy to a guard point.

- 14: Copy the data into the guard point.


The data is encrypted as it is copied.

Defining a Data Transformation Rule

The *Data Transformation Rules* tab defines the encryption key to apply for:

- Rekeying resources from one encryption key to another encryption key
 - Rekeying resources from one encryption key into a clear, unencrypted state
 - Encrypting clear, non-encrypted data
-

The user interface for the *Key Selection Rules* tab and the *Data Transformation Rules* tab are almost identical. For details about interface layout and usage, see, “[Navigating the Key Selection Rules tabs](#)” on page 9-177.



To define Key Selection Rules:

- 1: Open the Policy Composer, as described in “[Opening the Policy Composer](#)” on page 9-146.

By default, the *Policy Composer* window opens to the *Security Rules* tab.

- 2: Configure the security rule **Action** set to include to `key_op`.

See “[Configuring access methods](#)” on page 9-167 for details. Remove `key_op` from the policy after you run `dataxform`.

- 3: Select the *Key Selection Rules* tab.

The three tabs in the *Policy Composer* window each allow you to configure resources. Each resource configuration is a subset of the other. Without a resource configuration, the policy is applied to all the resources under the guard point. If you configure resources in the *Security Rules* tab, the policy is applied to just those resources under the guard point. When you configure an encryption key in the *Key Selection Rules* tab, it can be applied to all or some of the resources configured in the *Security Rules* tab. When you configure an encryption key in the *Data Transformation Rules* tab, it can be applied to all or some of the resources configured in the *Key Selection Rules* tab.

See “[Configuring resources](#)” on page 9-146 for details about configuring resources.

- 4: Do not change the currently configured key.

A key must be configured in order to open the *Data Transformation Rules* tab.

- 5: Select the *Data Transformation Rules* tab.

Changes made to the *Data Transformation Rules* tab are ignored if the *Key Selection Rules* tab is not configured.

- 6: Configure the resources to be rekeyed.

You can configure all or some of the resources that are specified in the *Key Selection Rules* tab.

- 7: Click **Key** to reload the Security Server key list.

This step is optional. It refreshes the key list information on the Security Server.

- 8: Select the encryption key from the **Key** scroll-list that you want to apply to the currently encrypted data.

Encryption keys must already be generated for them to be displayed in the drop-down menu. For information about generating encryption keys, see [“Using Encryption Keys” on page 7-115](#).

A default key, `clear_key`, is provided with the Security Server. This key indicates that clear text (no encryption) is applied to the resource. `clear_key` is used for debug purposes or to rekey encrypted data into clear text, effectively restoring the resource to an unencrypted state. Select `clear_key` if you want to decrypt the resources.

- 9: Save the policy.

- 10: Run the `dataxform` utility.



Do not change an encryption key if any files that use that key are open or being accessed in cache. A key change takes effect immediately; however, open files and cached data continue to use the old key. Changing keys while protected files and cached data are in use is unpredictable and can corrupt data.

The `dataxform` utility and its configuration is described in [“Changing encryption keys” on page 10-183](#).

Changing encryption keys

The Security Server provides a utility to change the encryption keys that are applied to files in a guard point. This utility can be used to:

- Change from one key to another key. Do this to update the keys on a regular basis to improve data security.
- Remove encryption. Do this to decrypt files so they are usable outside of Encryption Expert guard points.
- Apply encryption. Do this to files that were not encrypted because they were in a directory prior to turning it into a guard point.

The process of changing between encryption keys is referred to as “rekeying”.

If the number of files to be rekeyed is not great, a user with the appropriate access permission can simply copy the files out of the guard point to an unguarded directory. Files copied in this manner are saved in an unencrypted form and can be used by anyone. To reapply encryption, copy the files into a guard point. The files will be encrypted with the key in the policy as they are being copied. This approach does not require policy changes.

A utility is provided to rekey large numbers of files automatically without having to copy files into and out of guard points, though it does require policy changes. The `dataxform` utility is executed in a UNIX terminal window on a File System Agent-enabled host. It can be applied to individual files in a guard point or all the files in a guard point.

The utility can be run manually or automatically:

- For manual execution, create a policy that configures the *Security Rules*, *Key Selection Rules*, and *Data Transformation Rules* tabs of the Policy Composer. Then run `dataxform`.
- For automatic execution, create and apply a rekey policy as you would for manual execution. Also, create a special configuration file that provides the input parameters needed to run `dataxform`. Data transformation begins immediately when the special configuration file is copied into the guard point or the rekey policy is applied.

An Encryption Expert policy provides two types of encryption keys that can be applied to a guard point. The first is a regular AES encryption key. The second is a special key, `clear_key`, that is used in the data transformation process to encrypt or decrypt files. `clear_key` is applied to decrypt encrypted files or to encrypt unencrypted files. `clear_key` is not used to rekey files from one AES key to another.



Run `dataxform` as the root or administrative user because data migration copies data out of and back into guard points. The data must remain static during the migration process to avoid the risk of corrupting the data. Disable any services that write into the directories being migrated during data transformation.

Using dataxform

Location

`dataxform` is located at:

- `./agent/vmd/bin/dataxform`

Add the path for the directory that contains `dataxform` to the user `$PATH` environment variable so you can run `dataxform` without having to specify the full path.

Syntax

```
dataxform --version
dataxform --rekey_supported -gp dir
dataxform --rekey --gp dir [--nq]
dataxform --recovery --file_list file --gp dir [--nq]
dataxform --rekey_list --file_list file --gp dir
```

Option	Parameter	Description
--check_links	none	Scans for hard links in the guard point. Use the <code>--gp</code> option to specify the guard point to scan. No files are re-keyed. This operation only generates a list of hard-link files.
--cleanup	none	Removes the <code>dataxform_auto_lock</code> , <code>dataxform_files</code> , and <code>dataxform_status</code> files from a guard point. The files were generated by a previous <code>dataxform</code> session and can affect the next <code>dataxform</code> session. Include the <code>--gp</code> option to specify the guard point that contains the files to be removed.

--dir_recovery	directory name	Specifies the directory in which to write the dataxform status file. If this option is not specified, the progress and status files are written to <code>/var/log/eet/eet</code> on UNIX systems and <code>\Documents and Settings\All Users\Application Data\IBM\DBTOOLS\LUWEncryptionExpert\Agent\log</code> on Windows systems. The name of the output file created by <code>--dir_recovery</code> is <code>dataxform_status-<i>gp</i></code> , where <i>gp</i> is the underscore-separated path to the guard point. For example, if the path to the guard point is <code>/home/apps/apps1/data</code> , then the output file name is <code>dataxform_status-_home_apps_apps1_data</code> . If this option is not specified, the status file will be named <code>dataxform_status</code> . The dataxform file, <code>dataxform_auto_lock</code> , is always written to the guard point.
--file_list	file name	Used with <code>--recovery</code> to specify the output file name, or with <code>--rekey_list</code> to specify the input file name. <code>--file_list</code> takes one argument, <i>file</i> .
--gp	guard point	Specifies the full path to the guard point directory to process.
--nq	none	The <code>--nq</code> (“no queries”) option performs rekeying without user prompts. Without this option, dataxform prompts the user to verify that s/he wants to continue with the rekey operation.
--preserve_access_time	none	Directs dataxform to leave the pre-transformation last-accessed time of files intact during the rekey process. This option is useful when the last-accessed time is used to trigger file backups.
--preserve_modified_time	none	Directs dataxform to leave the pre-transformation last-modified time of files intact during the rekey process. This option is useful when the last-modified time is used to trigger file backups.
--print_stat	none	Displays the time it takes dataxform to complete each phase of the transformation process.

--recovery	none	Use the <code>--recovery</code> option to rollback a partially completed <code>dataxform</code> session. This option generates a list of the files in the guard point that were transformed during a previous rekey session. Use the <code>--gp</code> option to specify the guard point. Use with the <code>--file_list</code> option to specify the output file name. If the <code>--file_list</code> option is not specified, the output file is named <code>"dataxform_files"</code> and it is placed in the guard point.
--rekey	none	Rekeys all the files in the guard point specified by the <code>--gp</code> option. Use with <code>--dir_recovery</code> option to specify the directory in which to generate the status file. If the <code>--dir_recovery</code> option is specified, the status file is named <code>"dataxform_status_gp"</code> , where <i>gp</i> is the underscore-separated path to the guard point. If <code>--dir_recovery</code> is not specified, the status file is named <code>"dataxform_status"</code> and it is placed in the guard point.
--rekey_list	none	The <code>--rekey_list</code> option directs <code>dataxform</code> to rekey only the files specified by the <code>--file_list</code> option. The <code>--rekey_list</code> option also requires the <code>--gp</code> option. If the <code>--file_list</code> option is not specified, <code>dataxform</code> will look for a file named <code>"dataxform_files"</code> in the guard point to use as the source of the filenames to transform. If it cannot find <code>dataxform_files</code> , <code>dataxform</code> will abort with an error message.
--rekey_supported	none	The <code>--rekey_supported</code> option checks a guard point to determine if the guard point is appropriately configured for rekeying. The <code>--gp</code> option specifies the guard point to check. No files are re-keyed. This operation only indicates if the specified guard point is appropriately configured to proceed with rekeying.
--scan	none	The <code>--scan</code> option counts the number of files and directories in a guard point. Use the <code>--gp</code> option to specify the guard point.
--status_interval	seconds	The <code>--status_interval</code> option sets the time interval (in seconds) at which data transformation status messages are sent to the Security Server. The default is 300 seconds (5 minutes).

--thd	integer between 1 and 32, inclusive	The <code>--thd</code> option sets the number of threads that dataxform can use to transform files. A “thread” equates to a file. The more threads specified, the more files that are transformed concurrently. You may process up to 32 files concurrently. The default number of threads for a dataxform session is 16. Threads are numbered 0 through 31, but the values you enter are 1 through 32. If you enter 0, it indicates all 32 threads. If messages like “access denied” are displayed or files are being skipped, try reducing the number of threads. If the errors still occur after the number of threads is set to 1, the errors are not due to processes colliding, and there is something wrong with the files or policy permissions.
--version	none	Displays dataxform version information.

Examples

To display dataxform version information:

```
# dataxform --version
Build version: 1.1.3.0
Build ID: 1.1.3.0-Linux-2.6-x86-64-Build0084
Build Date: 2008-07-11 14:07:35 -0700 (PDT)
#
```

To verify that a guard point is eligible for rekeying:

The `rekey_supported` option checks if the directory specified by the `--gp` option is guarded with a rekey policy. The rekey policy is configured in the *Security Rules* and *Data Transformation Rules* tabs of a policy.

```
# dataxform --rekey_supported --gp /home/apps/appsl/data/
Attempting to contact the server in order to obtain setup
information...
Successfully contacted server.
Checking if data transform is supported for guard point
/home/apps/appsl/data/.
Data transformation is supported on /home/apps/appsl/data/
#
```

To check for hard-linked files inside the guard point:

```
# dataxform --check_links --gp /home/apps/appsl/data/
Attempting to contact the server in order to obtain setup
information...
Successfully contacted server.
About to perform the requested data transform operation.
-- Be sure to back up your data.
-- Please do not access files in the guard point during the
transform process.
```

```

-- Please do not attempt to terminate the application
Do you wish to continue (y/n)?y
Starting data transform of /home/apps/apps1/data/.
Status information for directories in guard point.
Directory /home/apps/apps1/data/ contains 108 files (3 KB).
Scanning 108 files (3 KB) in 0 directories for guard point
/home/apps/apps1/data/.
Scanning for hard links in directory /home/apps/apps1/data/.
Finished check for hard links in /home/apps/apps1/data/.
Found 2 references to hard link.
Hardlink reference 1 /home/apps/apps1/data/file_99.dat.
Hardlink reference 2 /home/apps/apps1/data/file_A.dat.
#

```

To run dataxform with a file list to remove encryption:

Note that the file list, `dx_fileList.txt`, contains only one file, `file_9.dat`. The list is small for brevity sake.

```

# dataxform --rekey_list --file_list dx_fileList.txt
--gp /home/apps/apps1/data
Attempting to contact the server in order to obtain setup
information...
Successfully contacted server.
Checking if data transform is supported for guard point
/home/apps/apps1/data.
Data transformation is supported on /home/apps/apps1/data
About to perform the requested data transform operation.
-- Be sure to back up your data.
-- Please do not access files in the guard point during the
transform process.
-- Please do not attempt to terminate the application
Do you wish to continue (y/n)?y
Starting data transform of /home/apps/apps1/data for files listed
in dx_fileList.txt.
Created thread ID 0x7b38d0 (0).
Created thread ID 0x7b0790 (1).
Created thread ID 0x74e870 (2).
...
Created thread ID 0x7b5ea0 (14).
Created thread ID 0x7b62a0 (15).
Starting data transform of /home/apps/apps1/data/file_9.dat.
Finished data transform of /home/apps/apps1/data/file_9.dat.
Data transform of 1 of 1 files listed in dx_fileList.txt for
guard point /home/apps/apps1/data.
#

```

Running multiple dataxform sessions

You can execute multiple dataxform sessions on the same host as long as each dataxform session is executed in a different guard point. The sequencing varies between manual and automatic dataxform execution. If you execute multiple instances of dataxform manually, the instances will run in parallel, meaning that you can rekey multiple guard points simultaneously. If you execute multiple instances of dataxform automatically, the instances will run in serial.

Preparing for dataxform

Regardless of the mode of execution, manual or automatic, you must configure a policy for dataxform to use to rekey a guard point. This policy is used to specify the current encryption key and the key that you would like to change to.

This section describes how to configure a simple rekey policy. For additional information see [“Configuring encryption in file system security rules” on page 9-175](#).

Create Security Rules

Configure two or three rules in the *Security Rules* tab.

- The first rule protects the files used to perform dataxform. The first rule comprises a resource set that applies to all files that start with “dataxform_”, an action that applies to any method of access, and an effect that permits any access to dataxform files. It does not apply key encryption so the dataxform files are not encrypted or decrypted along with the other files. In the example, this is set to **Resource**=dxFiles, **Action**=all_ops, and **Effect**=permit audit. The dxFiles resource contains one member, “dataxform_*”.
- The second rule comprises an action that applies to key operations only, and an effect that permits key operations and enables encryption. In the example, this is set to **Action**=key_op and **Effect**=permit apply_key audit.
- The third rule is optional. It denies access to all other users to provide added security. It comprises an action that applies to any method of access and an effect that denies any access to the guard point. This rule is not shown in the example, but would be set to **Action**=all_ops and **Effect**=deny.




Figure 10.1: Configuring the *Security Rules* tab

The example shows a simple rekey policy configuration. The rekey policy does not apply any access security. It applies encryption to everything accessed in the guard point, except

to files that begin with “dataxform_”. Your policy will most likely be different. In an actual production environment, your policy should limit access to just dataxform administrators, while denying access to all other users.

Specify the current encryption key

Configure the *Key Selection Rules* tab. Select the encryption key that is currently used to encrypt files. Specify the correct encryption key. The transformed files will be garbage if the wrong key is used. If files in the guard point are not encrypted, but you want to encrypt them, select the `clear_key` key in this tab. In this example, the symmetric key that is currently being applied to files in the guard point is `aes128_cached`.




Figure 10.2: Configuring the *Key Selection Rules* tab

Specify the new encryption key

Configure the *Data Transformation Rules* tab. Select the encryption key to which to apply to the files. To rekey so encryption is not applied, use the `clear_key` key. In this example, the selected key is `clear_key`, indicating that the files are to be decrypted. The dataxform utility will transform the contents of files in the guard point from their current encrypted form into an unencrypted form.




Figure 10.3: Configuring the *Data Transformation Rules* tab with `clear_key`
Save the policy.

Manual dataxform



To execute `dataxform` manually, log onto the host system as the administrative user to avoid any file system permission issues in the guard points. The guard point must remain static during the transformation process to avoid the risk of file corruption.

To run manual `dataxform`:

- 1: Backup data in the guard point.
This step is not required but it ensures a means to restore a guard point if anything should go wrong.
- 2: If specific files are to be rekeyed, create a file list now.
A file list is a text file that comprises the full path name of each file to be transformed. Enter one file path per line. If a file list is not specified, `dataxform` will rekey all the files in the guard point.
- 3: Log into the Management Console as a user with `Server` role permissions.
- 4: Open the *Guard FS* tab of the host with the guard point to be transformed.
The applied policies and guard points of the host are displayed.
- 5: Disable the guard point that is currently in effect.
 - a: Select the **Select** radio button for the guard point.
 - b: Click **Disable**.
- 6: (Optional) Enter the `df` command on the host system repeatedly until the `secfs` mount for the guard point disappears, or execute the “`secfsd -status guard`” command repeatedly until the guard point is no longer displayed.
- 7: Apply the rekey policy to the now disabled guard point.
See “[Preparing for dataxform](#)” on page 10-189 for rekey policy details.



If the Security Server does not detect a file named `dataxform_auto_config` in the guard point when the rekey policy is applied to the guard point, the Security Server will check for the file every 2 minutes until the file is detected, or until the rekey policy is removed from the guard point, or until `dataxform` is executed manually.

- 8: (Optional) Enter the `df` command on the host system repeatedly until the `secfs` mount for the guard point reappears, or execute the “`secfsd -status guard`” command repeatedly until the guard point is displayed.
- 9: Execute the `dataxform` command with the desired options on the host system.
For example,


```
#dataxform --rekey_list --file_list dx_fileList.txt
--gp /home/apps/appsl/data --dir_recovery /root
```
- 10: (Optional) Monitor `dataxform` progress on the host system.


```
# tail -f /var/log/eet/eetdxf_usr.log
```
- 11: Wait until `dataxform` completes.
- 12: Disable or delete the rekey policy.
- 13: Create a new policy, or modify the previous policy.

- a: Configure the policy to use the new encryption key in the *Key Selection Rules* tab.
- b: Leave the *Data Transformation Rules* tab empty.
- c: If encryption is not going to be applied, leave both the *Key Selection Rules* and *Data Transformation Rules* tabs empty.
- d: Remove `key_op` from the **Action** fields of rules in the *Security Rules* tab.
- e: Do not apply a policy that is configured for encryption to a directory that contains unencrypted files because, when `apply_key` is configured, the unencrypted files are encrypted when they are accessed. The data will be unusable.

14: Apply the policy to the guard point.

Automatic dataxform



(Windows only) UNIX and Windows users can run both run manual and automatic dataxform. However, after Windows users run automatic dataxform, the folder for the guard point remains in a “busy” state and it cannot be unguarded. Windows users must reboot the system after performing automatic dataxform in order to reset the state of the guard point and unguard the folder.



(Windows only) Do not run automatic dataxform in guard points that are on network drives. Automatic dataxform runs as the local “system” account. This account operates only on the file system and cannot operate across the network. To transform files on a network drive, log onto the system as a user with privileges to access the network drive and to modify the guard point to be transformed, and then run dataxform manually.

Automatic transformation is the automatic rekeying of files that occurs when a rekey policy is applied and a uniquely named file is detected in the guard point. Rekey policy configuration is described in “[Preparing for dataxform](#)” on page 10-189. The uniquely named file is `dataxform_auto_config` and it is placed in the top-level directory of a guard point.

The `dataxform_auto_config` file comprises one or two lines. The first line is mandatory. It lists the dataxform version to use. For now, always set the first line to “version=1”. The second line is optional. The second line lists additional dataxform parameters. By default, dataxform executes the `--rekey` and `--gp` options. These options will rekey all the files in the guard point. Do not enter the `--rekey_list`, `--file_list`, or `--gp` options in the `dataxform_auto_config` file. An example `dataxform_auto_config` file is shown below:

```
version=1
--thd 4 --preserve_modified_time
```

To run automatic dataxform:

- 1: Backup the guard point.
This step is not required but it allows you to recover if anything should go wrong.
- 2: Create a `dataxform_auto_config` file and place it in the top-level directory of the guard point.
- 3: Log into the Management Console as a user with `Server` role permissions.

- 4: Open the *Guard FS* tab of the host with the guard point to be transformed.
The applied policies and guard points of the host are displayed.
- 5: Disable the guard point that is currently in effect.
 - a: Select the **Select** radio button for the guard point.
 - b: Click **Disable**.
- 6: (Optional) Enter the `df` command on the host system repeatedly until the `secfs` mount for the guard point disappears, or execute the `secfsd -status guard` command repeatedly until the guard point is no longer displayed.
- 7: Apply the rekey policy to the now disabled guard point.
See [“Preparing for dataxform” on page 10-189](#) for rekey policy details.
Data transformation should start immediately after the rekey policy is applied.



If the Security Server does not detect a file named `dataxform_auto_config` in the guard point when the rekey policy is applied to the guard point, the Security Server will check for the file every 2 minutes until the file is detected, or until the rekey policy is removed from the guard point, or until `dataxform` is executed manually.

- 8: (Optional) Monitor dataxform progress on the host system.

```
# tail -f /var/log/eet/eetdxf_usr.log
```
- 9: Wait until dataxform completes.
- 10: Disable or delete the rekey policy.
- 11: Delete the `dataxform_auto_config` file.
- 12: Create a new policy, or modify the previous policy.
 - a: Configure the policy to use the new encryption key in the *Key Selection Rules* tab.
 - b: Leave the *Data Transformation Rules* tab empty.
 - c: If encryption is not going to be applied, leave both the *Key Selection Rules* and *Data Transformation Rules* tabs empty.
 - d: Remove `key_op` from the **Action** fields of rules in the *Security Rules* tab.
 - e: Do not apply a policy that is configured for encryption to a directory that contains unencrypted files because, when `apply_key` is configured, the unencrypted files are encrypted when they are accessed. The files will become unusable.
- 13: Apply the policy to the guard point.

Monitoring dataxform

Dataxform activity is recorded in three places:

- Security Server log
 - `/var/log/eet/eetdxf_usr.log`
 - `dataxform_*` files
-

Using the Security Server log

During a rekey operation, the Security Server logs both the current encryption key and the new encryption key. In [Figure 10.4](#), each file being processed is listed twice. The first entry shows the current encryption key, `aes128_cached`, and the second entry shows the newly applied encryption key, `clear_key`.

DashboardAdministratorsHostsKeysPoliciesHigh AvailabilityLogConfiguration

Logs

Log Filter

Log Type

All

Search Source

Jump To Page

Search Message

dataxform

Last Refreshed:

03/11/08 12:02 PM PDT

Refresh

Purge Debug LogsExport LogsDelete Logs

Total: 1732

Page 1 of 3

ID	Time	Severity	Source	Message
173560	2008-03-11 10:50:12.338	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_203.dat] Key[clear_key] Effect[PERMIT Code (1M)]
173559	2008-03-11 10:50:12.338	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_203.dat] Key[aes128_cached] Effect[PERMIT Code (1M)]
173558	2008-03-11 10:50:12.338	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_203.dat] Key[clear_key] Effect[PERMIT Code (1M)]
173557	2008-03-11 10:50:12.338	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_203.dat] Key[aes128_cached] Effect[PERMIT Code (1M)]
173556	2008-03-11 10:50:12.337	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_202.dat] Key[clear_key] Effect[PERMIT Code (1M)]
173555	2008-03-11 10:50:12.337	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_202.dat] Key[aes128_cached] Effect[PERMIT Code (1M)]
173554	2008-03-11 10:50:12.337	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_202.dat] Key[clear_key] Effect[PERMIT Code (1M)]
173553	2008-03-11 10:50:12.337	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_202.dat] Key[aes128_cached] Effect[PERMIT Code (1M)]
173552	2008-03-11 10:50:12.336	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_201.dat] Key[clear_key] Effect[PERMIT Code (1M)]
173551	2008-03-11 10:50:12.336	I	elbrus	EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128] User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\] Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[/home/apps/apps1/data/file_201.dat] Key[aes128_cached] Effect[PERMIT Code (1M)]

Figure 10.4: Displaying dataxform activity in the Security Server log

(UNIX only) There should be entries in the Security Server log like the following for each file that is rekeyed when running automatic dataxform:

```
EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128]
User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\]
Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/
dataxform] Action[key_op] Res[/home/apps/apps1/data/file_9.dat]
Key[clear_key] Effect[PERMIT Code (1M)]
EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128]
User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\]
Process [/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/
dataxform] Action[key_op] Res[/home/apps/apps1/data/file_9.dat]
Key[aes128_cached] Effect[PERMIT Code (1M)]
```

One log entry indicates the previous encryption key and the other indicates the new encryption key. Looking at the preceding example, the previous key for `file_9.dat` was `aes128_cached` and the new key is `clear_key`.

Ensure that there are no errors in the Security Server log that have to do with dataxform. Look for errors that contain strings like “denied” and “failed”. For example,

```
EET2601I: [SecFS, 0] AUDITPolicy[reKeyAES128]
User[billyBob,uid=1002,gid=10\wheel\] Process [/opt/IBM/DB2TOOLS/
LUWEncryptionExpert/agent/vmd/bin/dataxform] Action[key_op] Res[ /
```

```
home/apps/appsl/data/file_46.dat] Key[aes128_cached] Effect[DENIED
Code (1A,1U)]
```

This log entry indicates that an unauthorized user attempted to run `dataxform`.

Using `eetdx*_` files.

Data transformation activity is recorded and prepared for transmission to the Security Server. One file that records transformation activity is updated dynamically and should be checked first should any problems arise. You may want to use the “`tail -f`” command on the file to monitor the data transformation process in real-time.

The data transformation log file resides in `/var/log/eet` on the host system.

The base name of the data transformation log file is `eetdx*_usr.log`, where `usr` is the name of the user executing `dataxform`. For example, `eetdx*_root.log`. When the log file reaches the maximum capacity, it is appended with the current date and time, and a new, empty `eetdx*_usr.log` file is created. The naming convention for cycled log files is:

```
eetdx*_usr.log.YYY-MM-DD-HH-MM-SS
```

For example,

```
eetdx*_root.log <!-- the current dataxform log file
eetdx*_root.log.2008-03-11-11-50-18
eetdx*_root.log.2008-03-10-15-57-59
eetdx*_root.log.2008-03-10-15-31-04
```

The following `eetdx*_root.log` excerpt corresponds to the Security Server log example shown in [Figure 10.4](#). Always wait for the “Data transform for guard point *path* complete.” message before making any changes to the guard point.

```
2008-03-11 11:50:12.336 [VMD] [INFO ] [13554] [EET4310I] Starting
data transform of /home/apps/appsl/data/file_201.dat.
2008-03-11 11:50:12.336 [VMD] [INFO ] [13554] [EET4311I] Finished
data transform of /home/apps/appsl/data/file_201.dat.
2008-03-11 11:50:12.337 [VMD] [INFO ] [13554] [EET4310I] Starting
data transform of /home/apps/appsl/data/file_202.dat.
2008-03-11 11:50:12.337 [VMD] [INFO ] [13554] [EET4311I] Finished
data transform of /home/apps/appsl/data/file_202.dat.
2008-03-11 11:50:12.338 [VMD] [INFO ] [13554] [EET4310I] Starting
data transform of /home/apps/appsl/data/file_203.dat.
2008-03-11 11:50:12.338 [VMD] [INFO ] [13554] [EET4311I] Finished
data transform of /home/apps/appsl/data/file_203.dat.
2008-03-11 11:50:12.338 [VMD] [INFO ] [13554] [EET4380I]
Transformed 104 files (2 KB) of 104 files (2 KB) for guard point /
home/apps/appsl/data.
2008-03-11 11:50:13.263 [VMD] [INFO ] [13554] [EET4361I] Data
transform for guard point /home/apps/appsl/data complete.
```

Dataxform logging is tied to VMD logging. The size of the `eetdx*_usr.log` file is set by the **Maximum File Size** text-entry box on the *FS Log* tab. The default maximum `eetdx*_usr.log` file is 1 MB.

Using `dataxform_*` files

dataxform_auto_config – This is a user-generated file that comprises one or two lines and is placed in the guard point to initiate automatic data transformation. This file is described in “Automatic dataxform” on [page 10-192](#).

dataxform_auto_lock – This two-line file is generated by dataxform and placed in the guard point during manual and automatic data transformation. The first line is the name of the host system on which dataxform is being executed. The second line indicates the dataxform status. Status is “in progress” or “done”. An example of a complete and successful dataxform run is:

```
# cat dataxform_auto_lock
  elbrus
  done
#
```

dataxform_status – By default, this multi-line file is generated by dataxform during manual and automatic data transformation, and placed in /var/log/eet/eet on UNIX systems and \Documents and Settings\All Users\Application Data\IBM\DBTOOLS\LUWEncryptionExpert\Agent\log on Windows systems.

It is used to track the status of the current or last execution of dataxform. The dataxform_status file contains the following entries (the “Line x” text is just for descriptive purposes only. It does not appear in the actual file.)

```
Line 1  version=1
Line 2  status=done
Line 3  current=/home/apps/apps1/data/file_93.dat
Line 4  0 in-progress files
Line 5  0 error files
Line 6  0 skipped files
```

```
Line 1  The file format version number
Line 2  The dataxform exit code: “done” or “in progress”
Line 3  The last file transformed during the dataxform session
Line 4  The number of files in progress when dataxform exited
Line 5  The number of files that caused an error
Line 6  The number of files that dataxform skipped
```

If any of the file counts are greater than zero, the names of the relevant files will be listed in the status file.

Do not delete this file if dataxform ends abruptly. If you want to re-execute dataxform, it will read the status file to determine where to resume rekeying files.

An example of a dataxform status file is:

```
# cat dataxform_status
  version=1
  status=done
  current=/home/apps/apps1/data/file_93.dat
  0 in-progress files
  0 error files
  1 skipped files
  /home/apps/apps1/data/file_99.dat
#
```

Backing up and restoring DB2 databases

Encryption Expert is easy to use and almost invisible during the backup and restore process. Applying Encryption Expert protection involves adding an extra argument on the DB2 command line when making backups. Nothing extra is required to restore the backup. The command syntax to restore an Encryption Expert-protected backup is the same as restoring any other backup.

Before attempting to back up and restore Encryption Expert-protected data, it is assumed that you have already configured both the Encryption Expert Database Backup Agent and Security Server, and configured a policy that allows backup and restore operations.

Backing up a DB2 database

Depending upon the Encryption Expert policy for the database, you can encrypt, compress, or encrypt and compress, the data that is being backed up.

To make an encrypted backup, a compressed backup, or both, include the Encryption Expert shared library on the DB2 command line. The shared library starts the agent, which then queries the Security Server for policy constraints. If the policy allows, the backup is performed in the manner that is determined by the policy. The shared library is included in the Encryption Expert metadata for database backup. When the backup is restored, the shared library is extracted and used by DB2 to activate the Encryption Expert agent. You do not need to specify the shared library on the command to restore the backup. However, if the shared library is included on the restore command line, it has precedence over the shared library in the metadata.

The following figure shows the primary phases of a DB2 backup operation when you include the Encryption Expert library on the command line. The phases are:

- Issuing a DB2 command to make a backup with the Encryption Expert library (1)
- Activating the agent by including the Encryption Expert library on the DB2 command line, which then contacts the Security Server with information about the agent host and database (2)
- The Security Server randomly generating an asymmetric encryption key, which itself is encrypted with all the public keys in the assigned key group (3)
- The Security Server returning the FEK and Encryption Expert metadata (4)
- The agent encrypting the data payload with the asymmetric key to create the backup image (5)
- The backup image and the backup metadata (with the Encryption Expert library) being placed in the backup (6)




Figure 11.1: Encryption Expert for DB2 backup process

The size of the backup image is affected by several factors. One obvious factor is whether or not you are using compression. If you are using compression, the backup image should be smaller than the original image. If you are not using compression, the backup image will be larger than the original image because of the metadata that is included with the backup. If you lower the backup buffer size, the backup image will be even larger.

A header is included in each buffer. The header contains DB2 metadata and Encryption Expert metadata. The header can require an additional 64 to 100 bytes per buffer. If you reduce the buffer size, the backup image must comprise more buffers and consequently more space is used. Even more space might be required when dealing with page-aligned data. It is recommended that you use the default buffer size.

To back up an offline database with Encryption Expert encryption, compression, or both:

- 1: Log onto the agent host as the DB2 database administrative user.

You must log onto the agent host and perform the backup on the agent host because you must include the Encryption Expert library on the command line, and, if you are performing a remote backup, the library would be unavailable on the remote host.

The user in this example is `db2inst1`.

- 2: If you connected to the database, disconnect before initiating the backup.

The database cannot be active during backup.

```
# db2 connect reset
```

- 3: (Optional) Open another terminal window and run the `tail` command on the agent log file to monitor backup activity.

```
# tail -f /var/log/eet/eetdb2_usr.log
```

where, *user* is the name of the database instance user that is performing the backup or restore operation. For example, `db2inst1`. Only the latest log file is named `eetdb2_usr.log`. When the log file reaches the maximum allowed size, it is renamed `eetdb2_usr.log.date-num` and a new `eetdb2_usr.log` is created.

- 4: Run the DB2 backup command as you would normally, but include the additional arguments shown in the following example:

```
# db2 backup db sample to /db2back compress comprlib /
opt/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2/lib/
libeetdb2.so
```

where, *sample* is the name of the database and `/db2back` is the target destination for the backup image. `compress` and `comprlib` are literals to be entered exactly as shown. The path argument is the full-path to the Encryption Expert agent shared library object, `libeetdb2.so`. The default library path is shown. If you installed the agent software elsewhere, modify the path accordingly.

- 5: Wait for DB2 to complete and to display the following message.

```
Backup successful. The timestamp for this backup
image is : YYYYMMDDHHMMSS
```

If this message is not displayed, an error has occurred and you must check the logs to determine the cause.

- 6: (Optional) Verify that the database was successfully backed up.

Run the `db2ckbp` command with the `-cl` option, followed by the full path to the Security Server shared library object and the name of the database backup.

The library you specify will be the same library you specified to make the backup in step 4 above. For example:

On UNIX:

```
# db2ckbkp -cl /opt/IBM/DB2TOOLS/LUWEncryptionExpert/  
agent/db2/lib/libeetdb2.so  
ACTDB.0.db2inst9.NODE0000.CATN0000.20070816085836.001  
[1] Buffers processed: #####  
Image Verification Complete - successful.  
#
```

On Windows, open the DB2 Command Line Processor and enter:

```
$ db2ckbkp -cl C:\Program  
Files\IBM\DB2TOOLS\LUWEncryptionExpert\Agent\db2\bin\eetdb  
2.dll  
ACTDB.0.db2inst9.NODE0000.CATN0000.20070816085836.001
```

How Encryption Expert Database Backup Agent encrypts backups

A sophisticated key management scheme is built into the Encryption Expert Database Backup Agent software. This section describes how the keys interact to provide iron-clad protection.

The two primary encryption key components are:

- Backup Encryption Key (BEK) component
- Public and Private Key Encrypting Keys (PubKEK and PrivKEK) component

The BEK encrypts the data payload. The PubKEK and PrivKEK encrypt and decrypt the BEK. These components are described in the following table:

Abbrev	Full name	Description
BEK	Backup Encryption Key	Symmetric AES256 key that is used to encrypt and decrypt the data stream content. That is, the data in the DB2 backup is encrypted with this type of key. The BEK is automatically generated by the Encryption Expert. Each BEK is unique to each backup; however, in the case of parallel backups, all instances in a parallel backup are considered to be the same instance, so they all receive the same AES key. An encrypted form of this key is stored in the backup header. No local copy is maintained.
PubKEK	Public Key Encrypting Key	Public key part of the RSA asymmetric key-pair. The PubKEK is used in Encryption Expert to encrypt the AES256 Backup Encryption Key (BEK) that is stored in the backup header. The BEK is encrypted with the Public Key (PubKEK) of each key in the key group that is defined in a policy and assigned to the agent. PubKEKs are stored in the Security Server database.

Abbrev	Full name	Description
PrivKEK	Private Key Encrypting Key	Private key part of the RSA asymmetric key-pair. The PrivKEK is used in Encryption Expert to decrypt the BEK that was encrypted with the corresponding PubKEKs. PrivKEKs are stored in the Encryption Expert Security Server database.

The PubKEKs and PrivKEKs are paired with SHA-2 hash code values to conceal the key and to ensure that the correct key is applied to the backup or restore operation. The SHA-2 function takes a long string of any length, such as the 2048-bit PubKEK, and produces a fixed-length string as output, which is sometimes called a message digest or a digital fingerprint. The hash code is paired with the PubKEK and written to the backup. When the backup is restored, the hash code is used to identify the PubKEK that is needed to decrypt the BEK and the backup payload.

As shown in Figure 11.2, the BEK is encrypted with all the Public Keys (PubKEK) that are authorized to decrypt the backup data. Each PubKEK(BEK) is paired with the SHA-2 hash code of the PubKEK. The resulting list of (SHA-2(PubKEK(BEK))) keys is placed into the backup header. The Encryption Expert Security Server uses the SHA-2 hash codes to locate the matching private keys with which to decrypt the BEK.




Figure 11.2: Data encryption process

The public keys, and their corresponding hash codes, are derived from a key group that you created during agent host configuration.

Restoring a DB2 database

The only requirements for restoring a backup is to restore it on a system that is configured as an Encryption Expert Database Backup Agent and that agent has access to the appropriate private keys. If the system is not configured with Database Backup Agent software, the restore attempt will fail and DB2 will return an error message saying that it is unable to process `libeetdb2.so` or `eetdb2.dll`.

A nominal degradation in restoration performance will occur because of the overhead associated with data decryption.

The following diagram shows the phases for restoring a Encryption Expert backup. The phases comprise:

- A DB2 command to restore a backup (1)
- DB2 accesses the backup metadata and extracts the Encryption Expert shared library
- The Encryption Expert agent gets the backup header (2) and sends it, along with system parameters, to the Encryption Expert Security Server.
- The Encryption Expert Security Server evaluates the policy for the agent and, if allowed by the policy, the server uses a private key from the data store to decrypt the BEK from the header (3)
- The Encryption Expert Security Server returns the decrypted BEK to the agent system
- The agent system uses the key in the decrypted BEK to decrypt the data in the backup image (4)




Figure 11.3: Encryption Expert for DB2 restore process

The DB2 Database Backup Agent library is specified on the command line to make backups. When the backup is restored, the shared library is extracted from the backup metadata and is used by DB2 to activate the Encryption Expert agent. You do not need to specify the shared library on the command to restore a backup because the library is

embedded in the backup. However, if the shared library is included on the restore command line, it has precedence over the shared library in the metadata. If you are restoring an older backup, or encounter problems restoring a backup, include the library in the restore command line.

The syntax of a DB2 Database Backup Agent restore command is:

```
db2 restore db dbname from source_loc taken at time [to
target_loc] [compress comprlib libpath]
```

To restore an offline database backup with Encryption Expert:

- 1: Log onto the agent host as the DB2 database administrative user.

The user in this example is `db2inst1`.

- 2: (Optional) Open another terminal window and run the `tail` command on the agent log file to monitor restoration activity.

```
# tail -f /var/log/eet/eetdb2_db2inst1.log
```

where, `db2inst1` in `eetdb2_db2inst1.log` is the database user who is restoring the backup.

- 3: Run the DB2 restore command.

The following is a simple example:

```
# db2 restore db sample from /db2back taken at
20080220210650
```

where, `sample` is the name of the database, `/db2back` is the source location for the backup image, and `20070214155104` is the timestamp that uniquely identifies the backup instance.

You can also specify the DB2 Database Backup Agent library on the restore command line. Should you have any problems restoring a backup in the manner shown above, try again but this time include the DB2 Database Backup Agent library. The following example restores the same backup but this time it specifies a target destination and the library:

```
db2 restore db sample from /db2back taken at 20080220210650
to /db2data compress comprlib /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/agent/db2/libeetdb2.so
```

- 4: Wait for DB2 to complete and to display the following message.

```
The RESTORE DATABASE command completed successfully.
```

If this message is not displayed, an error condition has occurred and you must check the logs to determine the cause.

Working with online DB2 data

Encryption Expert supports backup and restore operations on both offline data-at-rest and online active databases. An offline backup or restore operation requires an exclusive connection to the database.

To back up and restore the `sample` database:

- 1: Configure DB2 to support online active backups.

```
db2 -v "update db cfg for sample using logretain recovery"
```

- 2: Get out of the "backup pending" state.

```
db2 backup db sample to /data1/db2back
```

- 3: Back up the database:

On UNIX:

```
db2 backup db sample online to /data1/db2back compress  
comprlib /opt/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/  
db2/lib/libeetdb2.so
```

On Windows, open the DB2 Command Line Processor and enter:

```
$ backup database sample to C:\data1\db2back compress  
comprlib C:\Program  
Files\IBM\DB2TOOLS\LUWEncryptionExpert\Agent\db2\bin\eetdb  
2.dll
```

- 4: Wait for the table space backups to complete.
 - 5: Restore the table spaces:

```
db2 restore db sample from /data1/db2back
```
 - 6: Roll forward to the end of the logs:

```
db2 rollforward db sample to end of logs and stop
```
-

Backing up and restoring IDS databases

This chapter describes how to configure the IDS Database Backup Agent to backup and restore IDS databases using the `ontape` and `onbar` commands.

Policies must be configured and applied before using the IDS Database Backup Agent. Database Backup Agent policies are described in [“Configuring Database Backup Agent policies” on page 8-134](#).

The IDS Database Backup Agent is a filter that is applied between the `ontape` or `onbar` command performing the operation and the IDS database. This filter supports all the features and functions normally associated with the IDS command, including third-party storage managers, such as the IBM Tivoli Storage Manager (TSM) or the Veritas Storage Manager.




Figure 12.1: IDS Database Backup Agent in an IDS configuration

The Encryption Expert IDS Database Backup Agent is flexible, requires little additional configuration, and is easy use. In addition to minor configuration changes, you must create a short attributes file that is read when the Database Backup Agent collects information for policy evaluation.

Information that is unique to a backup or restore operation can be provided to the agent through the attributes file. The attributes file is used to process policies and refine the policy match criteria. For example, if the policy names a specific IDS server, specify the same IDS server name in the attributes file. You can create multiple attributes files and reuse them to perform iterative backup and restore operations. Also, the attributes files can be configured on-the-fly during backup and restore operations with user environment variables .

The IDS Database Backup Agent reads data from stdin, its parameters from an attributes file, and writes compressed and/or encrypted data to stdout. Errors are reported on stderr. On an Encryption Expert-enabled system, the IDS Database Backup Agent:

- Reads the attributes file
- Contacts the Security Server
- Passes attributes file contents to the Security Server
- If restoring a backup, the agent also passes backup metadata to the server
- The IDS agent is sent one of the following:
 - Request denied
 - Backup request granted, plus AES key to encrypt data, and RSA public key groups to encrypt the AES key
 - Restore request granted, plus the server returns the AES key with which to decrypt the backup

Requirements

See [“Before you begin” on page 3-29](#).

In addition:

- The system hosting the IDS Database Backup Agent must run RedHat Linux 4 Update 3 or higher.
- Ensure that the correct libraries are installed on the system that is hosting the IDS Database Backup Agent:

The required libraries are:

```
libaio-0.3.105
libgcc-3.4.5
libstdc++-3.4.5
ncurses-5.4
pam-0.77
```

You can use the `rpm -qa` command and `grep` the output to determine currently installed libraries. For example:

```
# rpm -qa | grep libaio
libaio-0.3.105-2
libaio-devel-0.3.105-2
libaio-0.3.105-2
#
```

- Verify that the system hosting the IDS Database Backup Agent is running Informix IDS Dynamic Server V11.10 FC2 on Linux EM64.
- Check that the IDS processes are running before making a database backup. Execute the `ps` command and `grep` for “oninit”. Start the IDS database with the `oninit` command if it is not running.
- Provide sufficient disk space to perform backup operations. Two to four Gigabytes should be sufficient.
- The `informix` group is generated during IDS Database Backup Agent installation. Check that the primary group for the `informix` user is also `informix`. For example:

```
# ls -aslF /home/informix
total 76
```

```

4 drwx----- 5 informix informix 4096 Aug 18 13:03 ./
8 drwxr-xr-x 5 root      root      4096 Aug 16 11:06 ../
8 -rw----- 1 informix informix 7104 Aug 18 14:43 .bash_history
4 -rw-r--r-- 1 informix informix   24 Aug 15 17:59 .bash_logout
4 -rw-r--r-- 1 informix informix  191 Aug 15 17:59 .bash_profile

```

Configuring the IDS onconfig file

The `ontape` command does not support a storage manager. The `ontape` command supports STDIO. This allows you to specify the Database Backup Agent filter on the command line. The `onbar` command does not support STDIO, so you must specify the filter in an IDS `onconfig` configuration file and execute the command through a storage manager, such as the Tivoli Storage Manager (TSM) or Informix Storage Manager (ISM). The `onconfig` configuration file specifies parameters for executing the Database Backup Agent filter with the attributes file.

To configure the IDS `onconfig` file:

- 1: Create two IDS attributes files, one for backup operations and one for restore operations.

These are described in [“Configuring the IDS Database Backup Agent attributes file” on page 12-208](#).

- 2: Edit the `onconfig` file.

For example:

```
vi $INFORMIXDIR/etc/onconfig.demo_on
```

- 3: Change, or add, the `BACKUP_FILTER` and `RESTORE_FILTER` variables so that backup and restore operations wake the agent and pass it the respective IDS attributes file.

For example:

```

BACKUP_FILTER '/opt/IBM/DB2TOOLS/LUWEncryptionExpert/
agent/ids/bin/eetidsagt -f /home/informix/
backup_attributes.cfg'
RESTORE_FILTER '/opt/IBM/DB2TOOLS/LUWEncryptionExpert/
agent/ids/bin/eetidsagt -f /home/informix/
restore_attributes.cfg'

```

- 4: Save and exit the file.

Configuring the IDS Database Backup Agent attributes file

The attributes file contains configuration information that pertains to a particular backup or restore operation. The attributes file provides unique, job-specific information to use in policy evaluation. The agent passes this information to the Security Server. The Security Server evaluates the policy, using the variables in the attributes file. The Security Server then permits or denies the requested operation. If it permits the operation, the Security Server provides the agent the encryption keys needed to complete the operation.

At a minimum, you should configure two attributes files, one for backup operations and one for restore operations.

The syntax of an attributes file is:

- `#` is an optional comment that starts after the occurrence of the `#`
- Expressions are the form: `name=value`

- The attribute name is a text string without whitespaces.
- The attribute name and its value are case sensitive. Attribute names must be upper-case. Lower-case attribute names are ignored.
- An equal sign (=) separates the attribute name from the attribute value.
- Whitespace is stripped from the beginning and end of the value.

The following table lists valid attributes and their values:

Attribute		Value	Required /Optional
BLOCKSIZE	=	This attribute specifies the size of the data chunks that the Database Backup Agent reads from stdin to process read and write operations. This variable is not related to, or have any effect on, disk or tape blocking. You can optimize operational throughput by adjusting the block size. The attribute value is an integer expressed in kilobyte units (KB). The minimum value is 1 and the maximum is 1835008. The default is 128.	Optional

Attribute		Value	Required /Optional
ENV	=	This attribute specifies the environment variables that, if configured for the user that is performing the IDS backup or restore operation, will be used to set additional match criteria. The attribute value is a comma-separated list of environment variables. Acceptable variables are INFORMIXSERVER, SERVERNUM, and BLOCKSIZE. If the user or application submitting the backup or restore request has environment variables that you want to use, include the variable names in the ENV attribute. For example, the user or application submitting the backup or restore request can have the following variables set in their user environment: <pre>BLOCKSIZE=256 INFORMIXSERVER=qa_db01</pre> To use these variables, enter a line like the following in the attributes file: <pre>ENV=BLOCKSIZE, INFORMIXSERVER</pre> The environment variable values will be passed along to the Security Server as if they had been entered into the attributes file. You can specify environment variables in the ENV list but not configure them in the user environment. The ENV attribute skips unconfigured environment variables. If the same variable, but with different values, is defined in the attributes file and in the user environment, the value in the user environment has precedence. For example, if the SERVERNUM attribute is set, and also, SERVERNUM is specified in the ENV attribute, the value of the user \$SERVERNUM variable is used. Valid separators for the ENV list are a comma, space, and tab.	Optional

Attribute		Value	Required /Optional
INFORMIXSERVER	=	This attribute specifies the name of the database server. This attribute corresponds to the Server Name text-entry box in the <i>Backup Rule</i> window for the policy. If Server Name is not set, then you can enter any name you like as the value of this attribute because it will be ignored.	Required
OPERATION	=	This attribute specifies the operation that is to occur. The value is either “read” (restore) or “write” (backup).	Required
SERVERNUM	=	This attribute specifies the IDS server number. This attribute corresponds to the Server Number text-entry box in the <i>Backup Rule</i> window for the policy. If Server Number is not set, then you can enter any number you like as the value of this attribute because it will be ignored.	Required
<i>additional conditions</i>	=	These are other, arbitrary attributes to be included and passed to the Security Server as additional match criteria for evaluating a backup request. Specify a <i>name=value</i> expression by itself on a line in the attributes file and specify the same <i>name:value</i> pair in the Additional Conditions portion of the backup rule of a policy. If the two match, then the criteria is satisfied and backup permission can be granted. As in the backup rule, you can specify up to four additional conditions. The name and value are case-sensitive. If additional conditions are not set in the policy, the additional conditions set in the attributes file will be ignored.	Optional

The following is an example of configuring additional conditions. Additional conditions provide more data to use to evaluate a backup request. If the policy is configured as shown below:

Additional Conditions

Name 1	compress_encrypt	Value 1	c1e2
Name 2	group	Value 2	backup_admins
Name 3	favorite_color	Value 3	green
Name 4		Value 4	

Figure 12.2: Additional Conditions

then include the following lines in the attributes file:

```
compress_encrypt=c1e2
group=backup_admins
favorite_color=green
```

(Note that both the names and values are case-sensitive.)

The following is a sample backup attributes file:

```
# Filename: backup_attributes.cfg
OPERATION=write
BLOCKSIZE=128
SERVERNUM=3
INFORMIXSERVER=wrongServer
ENV=BLOCKSIZE, INFORMIXSERVER
compress_encrypt=c1e2
group=backup_admins
favorite_color=green
```

The INFORMIXSERVER attribute in the sample above specifies the wrong server; however, the server specified in the user environment will be used because INFORMIXSERVER is also specified in the ENV attribute. If the INFORMIXSERVER variable is not set in the user environment, the value of the INFORMIXSERVER attribute will be used instead.

The following is a sample restore attributes file:

```
# Filename: restore_attributes.cfg
OPERATION=read
BLOCKSIZE=256
INFORMIXSERVER=demo_on
SERVERNUM=1
```

Setting the IDS environment

Set your IDS environment variables as you normally would. For example, you can configure something like the following in the informix user `.bashrc` file, and then source the `.bashrc` file. This example is setup to administer the `demo_on` database. Your environment configuration will be different.

```
INFORMIXTERM=terminfo
export INFORMIXTERM
INFORMIXSERVER=demo_on
ONCONFIG=onconfig.demo_on
INFORMIXDIR="/opt/IBM/informix"
INFORMIXSQLHOSTS="/opt/IBM/informix/etc/
sqlhosts.demo_on"
PATH=$PATH:/opt/IBM/informix/bin
export INFORMIXSERVER INFORMIXDIR ONCONFIG
INFORMIXSQLHOSTS PATH
```

Making an IDS backup

Depending upon the Encryption Expert policy, you can encrypt, compress, or encrypt and compress IDS databases. To make a backup, run the IDS backup command with a Database Backup Agent filter. This is represented as “1” in [Figure 12.3](#).

The agent forwards the backup request and configuration data to the Security Server (2).

The Security Server evaluates the request and data and, if all the criteria match, the Security Server allows the operation (3). It creates the unique AES key information that will be used to encrypt the backup data. It then encrypts that unique AES key information with the RSA public keys in a key group. The Security Server returns this data to the agent system (4).

The agent uses the AES key to encrypt the data from the IDS database. The agent sends the encrypted data, along with the appropriate metadata, to the storage device (5).




Figure 12.3: Encryption Expert for IDS backup process

Check the progress of backup and restore operation. You can check the `/var/log/eet/eetids_user.log` file. For example, `eetids_informix.log`. Also, check the *Logs* window in the Management Console.



Do not backup an IDS database to an existing file. During a backup operation, IDS creates a zero-size database file and then it processes the EE filter. If the EE filter denies the operation, IDS aborts the operation and the zero-size database file is left on the system. The zero-size database file is unusable. If a backup operation is initiated that overwrites an

existing database backup, and the backup operation is denied, the backup file will be empty and unusable.

Making an IDS database backup using ontape

The `ontape` command can be executed to support both command-line and configuration file filter parameters. Both ways of running `ontape` are described in this section.

Using ontape without onconfig

A simple syntax for using the `ontape` command without using the `onconfig` configuration file is: :

```
ontape options | eetidsagt -f attrs > dbBkup
```

where,

- *options* are `ontape` command-line options. At a minimum, specify the type of backup (logical, physical, full) desired and the STDIO output destination.
- *eetidsagt* is the full path to the `eetidsagt` filter.
- *attrs* is the full path to the attributes file with which to moderate the backup process.
- *dbBkup* is the full path to the backup file to be generated.

The `ontape` command begins execution and forwards the write request to the Encryption Expert agent file, `eetidsagt`. If the policy permits, the `ontape` output is encrypted and/or compressed, and written to *dbBkup*.

For example:

```
# ontape -v -s -L 0 -t STDIO | /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/agent/ids/bin/eetidsagt -f /home/
informix/backup_attributes.cfg > /data1/eetback/
eet_L0.bkp
```

To make an encrypted and/or compressed IDS backup:

- 1: Configure an IDS policy in the Management Console.
- 2: Set the appropriate user environment.
- 3: Configure the attributes file.
- 4: Execute the `ontape` command and include the `eetidsagt` filter.

For example:

```
# ontape -v -s -L 0 -t STDIO | /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/agent/ids/bin/eetidsagt -f /home/
informix/backup_attributes.cfg > /data1/eetback/
eet_L0.bkp
10 percent done.
20 percent done.
30 percent done.
100 percent done.
```

Please label this tape as number 1 in the arc tape sequence.

This tape contains the following logical logs:

```
11
```

Program over.

Using ontape with onconfig

You can execute the `ontape` command with a minimum of arguments when you configure the `onconfig` file.

To run `ontape` with `onconfig`:

- 1: Log in as the `informix` user.
- 2: Edit the `onconfig` file for the database server to use the `EEagent`.
See [“Configuring the IDS onconfig file” on page 12-208](#).
- 3: Set the user environment.
- 4: Configure the attributes file.
- 5: Run `ontape`.

For example:

```
# ontape -s -L 0
Using the backup and restore filter /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/agent/ids/bin/eetidsagt -f /home/
informix/backup_attributes.cfg.
10 percent done.
20 percent done.
30 percent done.
100 percent done.
File created: /idsback/vmlinux101_0_L0

Please label this tape as number 1 in the arc tape
sequence.
This tape contains the following logical logs:

12

Program over.
#
```

Making an IDS database backup using onbar

To backup an IDS database using `onbar`:

- 1: Log in as the `informix` user.
- 2: Check that the ISM is running.
`onbar` requires a Storage Manager (SM). Start it if it is not running.
- 3: Edit the `onconfig` file for the database server.
See [“Configuring the IDS onconfig file” on page 12-208](#).
- 4: Set the user environment.
- 5: Configure the attributes file with parameters for `EE` the agent to use. .
- 6: Run `onbar`.

For example:

```
# onbar -b -w
```

- 7: Verify the results.

You can:

- Run `onbar -m 99`

- View the log file `/opt/IBM/informix/bar_act.log`
- View the log file in `/var/log/eet`
- View the *Logs* window in the Management Console

For example:

```
# onbar -m 99
Onbar activity log file: /opt/IBM/informix/bar_act.log
2008-07-22 12:04:47 25557 25555 /opt/IBM/informix/bin/
onbar_d -b -w
2008-07-22 12:04:48 25557 25555 Archive started on
rootdbs
2008-07-22 12:04:48 25557 25555 Begin level 0 backup
rootdbs.
2008-07-22 12:04:48 25557 25555 Starting Filter /opt/
IBM/DB2TOOLS/LUWEncryptionExpert/agent/ids/bin/
eetidsagt -f /home/informix/backup_attributes.cfg.
2008-07-22 12:04:48 25559 25557 Successfully connected
to Storage Manager.
2008-07-22 12:06:20 25559 25557 child process blocks
read = 21
2008-07-22 12:06:20 25559 25557 The child process for
the backup and restore filter is terminating with exit
code 0.
2008-07-22 12:06:50 25557 25555 Successfully connected
to Storage Manager.
2008-07-22 12:06:50 25557 25555 Completed level 0
backup rootdbs (Storage Manager copy ID: 1273869633 0).
2008-07-22 12:06:51 25557 25555 Archive on rootdbs
Completed (Requested Level 0).
2008-07-22 12:06:52 25557 25555 /opt/IBM/informix/bin/
onbar_d complete, returning 0 (0x00)
2008-07-22 12:23:16 25734 25732 /opt/IBM/informix/bin/
onbar_d -m 99
#
```

Restoring an IDS backup

The only requirement for restoring a backup is to restore it on a system that is configured as an Encryption Expert agent, and which has access to a Security Server with the appropriate keys. If the system is not configured with agent software, the restore may complete but the data can be encrypted and unusable. Also, IDS can return a message like the following: `/opt/IBM/informix/bin/onbar_d complete, returning 147 (0x93)`.

Restoring an IDS database backup using ontape

The format of the `ontape` restore command line is:

```
cat dbBkup | eetidsagt -f attrs | ontape -r -t STDIO -v
```

where,

- *dbBkup* is the full path to the IDS database backup that is to be restored.
- *eetidsagt* is the full path to the eetidsagt filter.
- *attrs* is the full path to the attributes file with which to moderate the restore process.
- If the policy permits, the results are passed to the *ontape* command for execution. At a minimum, specify the type of restoration (logical, physical, full) desired and the STDIO input source.

For example:

```
# cat /data1/idsback/eet_L0.bkp | /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/agent/ids/bin/eetidsagt -f /home/
informix/restore_attributes.cfg | ontape -r -t STDIO -v
```

To restore a backup:

- 1: As the *informix* user, open a terminal window on the Database Backup Agent system.
- 2: Set the user environment.
- 3: Configure the attributes file.
- 4: Shutdown the *oninit* processes before restoring a backup.

```
# onmode -ky
```

Though the policy would otherwise permit a restore request, restoration can fail and the following message can be displayed if you do not shutdown the *oninit* processes:

```
Physical restore failed - Cannot warm restore ROOT
DBspace
```

- 5: Enter the database restore command line.

```
# cat /data1/idsback/eet_L0.bkp | /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/agent/ids/bin/eetidsagt -f /home/
informix/restore_attributes.cfg | ontape -r -t STDIO -v
```

The *ontape* results are displayed on the terminal:

```
Archive Tape Information
Tape type:          Archive Backup Tape
Online version: IBM Informix Dynamic Server Version
11.10.FC2
Archive date:       Mon Jul 21 17:04:45 2008
User id:            informix
Terminal id:        /dev/pts/1
Archive level:      0
Tape device:        STDIO
Tape blocksize (in k): 32
Tape size (in k):   0
Tape number in series: 1

...
Initialization Time      07/15/2008 13:39:43
System Page Size         2048
Version                  15
Index Page Logging       OFF
Archive CheckPoint Time  07/21/2008 16:45:15
```

```

Dbspaces
number  flags  fchunk nchunks flags owner      name
1       40001  1      1        N  B  informix  rootdbs
...
/opt/IBM/informix/bin/onmode -sy
Program over.
#

```

After the database is brought online (i.e., `onmode -m`), the `oninit` processes are automatically restarted and the database is operational once again.

Restoring an IDS database backup using onbar

To restore a database using onbar:

- 1: As the `informix` user, open a terminal window on the Database Backup Agent system.
- 2: Set the user environment.
- 3: Configure the attributes file.
- 4: Edit the `onconfig` file for the database server.

See [“Configuring the IDS onconfig file” on page 12-208](#).

- 5: Shutdown the `oninit` processes before performing a cold restore.

```
# onmode -kuy
```

Even if the policy would otherwise permit a restore request, restoration will fail and the following message will be displayed if you do not shutdown the `oninit` processes:

```
Physical restore failed - Cannot warm restore ROOT
DBspace
```

- 6: Run the `onbar` command.

For example, to restore an entire archive:

```
onbar -r -w
```

- 7: Verify the results.

You can:

- Run `onbar -m 99`
 - View the log file `/opt/IBM/informix/bar_act.log`
 - View the log file in `/var/log/eet`
 - View the *Logs* window in the Management Console
-

Configuring high availability

High availability is a Security Server configuration comprised of two or more Security Server installations. One installation is designated as the primary server and the others are the failover servers. All Security Server/Management Console activity occurs on the primary server. Changes and updates to the primary server are copied to the failover servers at set intervals. If the primary server fails to respond to agent requests, the agent automatically switches to the first failover server. If that is unavailable, the agent tries the next failover server. And so on.

You can configure as many failover servers as you wish; however, the practical limit is four. More than four failover servers can result in increased network traffic and degrade Security Server performance.



Only users with the Admin role can configure High Availability. See [“Configuring users” on page 5-69](#) for details.

A failover server is a mirror of the primary Security Server. They both contain the same keys, policies, and configurations. An agent can submit a backup or restore request to any Security Server in a high availability configuration and achieve the same results. You are prompted during installation to indicate if the current Security Server installation is the primary server or a failover server.

Displaying the high availability configuration

To display the current high availability configuration:

- 1: Log into the Management Console on the primary Security Server as a user with `Server` role permissions.
- 2: Select **High Availability** in the menu bar.

The *High Availability Servers* window opens. It displays the primary server and, if configured, the failover servers.



You should never attempt to make configuration changes to a failover Security Server because the changes will be lost each time the primary server updates the failover server.

Adding a failover server

To add a failover server to the Security Server database using the fewest steps possible:

1: Log into the Management Console as a user with **Server** role permissions.

2: Select **High Availability** in the menu bar.

The *High Availability Servers* window opens.

3: Click **Add**.


The *Add Server* window opens.

4: Enter the name to identify the failover Security Server in the **Server Name** text-entry box.

This is the same name that will be used to identify the server during the certificate generation phase of failover Security Server software installation. It will also be the CN name in the failover server certificate.

5: Click **Ok**.

The *High Availability Servers* window displays the specified server with a **Role** of **Failover**. The **Registered** checkbox is not checked because the failover server has not been registered and integrated with the primary server.



High Availability Servers			
Total Servers: 2			
Add Delete			
Selected	Name	Role	Registered
☒	vmlinux01	Primary	☒
☐	vmlinux02	Failover	☐
Add Delete			

Figure 13.1: Configured but not registered failover server

6: Install the failover Security Server software, as described in [“Installing and configuring the Encryption Expert failover server” on page 3-38](#).

Failover Security Server software installation should complete successfully and return “Encryption Expert successfully configured to communicate with Security Server on *hostname/FQDN*.” If it does not, most likely there is a mismatch between the server information configured in step 4 above and the server information configured during failover server software installation.

You can manually configure server certificates later if the certificate generation and exchange phase of server software installation fails.

7: Check the *High Availability Servers* window for the new server.

The **Registered** checkbox should be enabled, indicating that the failover server and the primary server can now communicate and exchange data with each other.

8: Enable database replication with the `configure_replication` utility.

This is described in [“Configuring SQL replication” on page 13-222](#).

9: Add the failover server information to the `./db2/etc/agent.conf` file, the `./ids/etc/agent.conf` file, and/or the `./vmd/etc/agent.conf` file on each agent system. See step 4 on [page 101](#) for an example of configuring the respective `agent.conf` files.

Replicating the Encryption Expert Security Server database

You can use SQL replication to automatically synchronize live Security Servers and ensure that database changes are correctly downloaded to the failover servers. The failover servers each run a security database that is automatically updated by the primary server at set intervals. This ensures that the failover servers, when activated, work with the same keys, policies, and host configurations as the primary server.

Four utilities are provided to configure and maintain a failover server with database replication. They are:

- `configure_replication`
- `start_replication`
- `stop_replication`
- `cleanup_replication`

Run the utilities as the Encryption Expert administrative user (for example, `voradmin`) on the primary server. The utilities change only the primary server. That is, they modify the primary server database schema and start some processes on the primary server. The failover servers are not modified at all. Only after replication is activated, and the update interval is satisfied, that the primary server schema changes are applied to the failover servers schema.

Only the server configuration database (CGSSDB) is replicated. The audit log database (LOGDB), the audit log configuration, and the Java Message Service database (JMSDB), are not part of replication and are not included in the server configuration transfer.

Before you configure Security Server database replication:

- You must have network access between the primary and failover servers.
- You must know the passwords for the Encryption Expert administrative user for all the servers in the high availability configuration.
- You will need the root password to restart the server configuration database (CGSSDB) on the primary server.
- Be prepared to restart the primary server.

SQL replication of the Security Server database

Each Encryption Expert agent maintains a list of the server URLs that it can access. The first URL in the list is always the primary server. The second and remaining URLs are the failover servers. If communication with the primary server fails, the agent tries the next URL in the list.

SQL replication architecture is used to update each failover Security Server database with the primary Security Server database.

The primary server acts as the Certificate Authority and creates signing certificates. The primary server must be installed and configured with its own certificates before you can configure agent and failover-server certificates. You should back up both the primary and failover Security Servers after configuring the two to communicate with each other. The backups can be used to revert the SQL replication setup to a known good state if subsequent changes critically affect the failover server.

It is highly recommended that you back up the primary Security Server database after every major configuration is performed, including replication. It is also recommended that you backup new and unreplicated failover Security Servers.

This section describes:

- how to initially configure SQL replication
- how to back up the SQL replication configuration of the primary and failover Security Servers
- how to delete a failover server configuration
- how to restore the SQL replication configuration of primary and failover Security Servers
- what to do in the event of a failure

Configuring SQL replication

Multiple failover Security Servers can be configured with one primary Security Server.

To configure Security Server database replication:

- 1: If you have not done so already, install the failover servers, as described in [“Installing and configuring the Encryption Expert failover server” on page 3-38](#).
- 2: It is recommended that you back up the primary server database with `./server/bin/backup_config` before proceeding. See [“Backing up the Security Server configuration” on page 16-295](#) for details.
- 3: Log onto the primary Security Server as the administrative user (for example, `voradmin`).
- 4: Execute the utility that prepares the primary server for replication.


```
# ./server/bin/configure_replication
...
Do you wish to proceed? (y|n)[n]: y
```
- 5: Enter `y` to continue.

The utility prompts you to enter and re-enter the Encryption Expert administrative user (for example, `voradmin`) password for the primary Security Server you are currently logged onto.

```
Enter the admin (voradmin) password for this Security
Server:
```

- 6: Enter the password for the local Encryption Expert administrative user.

The utility checks the Security Server database for configured, but unregistered, failover Security Servers. It then lists the available failover servers and prompts you to select one of them. For example:

```
You currently have no failover servers configured for
replication. We will configure the first failover server.
Failover Security Server host are: vmlinux02
The following failover server are available to be
configured for replication:
[1]: vmlinux02
Choose the failover server to be configured [1]:
```

- 7: Enter the number for the failover server to be configured.

The utility prompts you to enter and re-enter the Encryption Expert administrative user (`voradmin`) password for the selected failover Security Server.

- 8: Enter the password for the failover server administrative user, `voradmin`.

This allows the primary Security Server to administer the failover Security Server database.

The utility registers the failover server and restarts the datastore.

- 9: Switch to the root user.

```
# su root
```

- 10: Restart the primary server using the `cgss restart` command.

```
# /etc/init.d/cgss restart
```

The failover server has been added to the local data store and the primary server is restarted in order to use the new configuration. The `cgss restart` command restarts the embedded database, the Web application server, and automatically starts the processes necessary for database replication on the primary server.

- 11: Verify that the database replication processes are running on the primary server.

There should be two processes running, `asncap` and `asnapply`.

```
# ps aux | grep asn
voradmin 11857 ...
    asncap capture_server=CGSSDB
    CAPTURE_PATH=/home/voradmin
voradmin 11867 ...
    asnapply control_server=CGSSDB
    apply_qual=CGSSDB_AQ pwdfile=asnpwd.aut
    apply_path=/home/voradmin
#
```

`asncap` captures changes to the primary server database. `asnapply` applies those changes to each failover server database. `asnapply` does nothing if the capture tables are empty. That is, if the primary server hasn't been changed since the last failover server update, `asnapply` will wait until the next scheduled update and check again for changes.

The log files for the `asncap` and `asnapply` processes are `user.CGSSDB.ASN.CAP.log` and `user.CGSSDB.CGSSDB_AQ.APP.log`, respectively. The default `user` is `voradmin`. These logs are stored in `/home/voradmin`. Refer to [“voradmin.CGSSDB.ASN.CAP.log” on page 14-245](#) and [“voradmin.CGSSDB.CGSSDB_AQ.APP.log” on page 14-246](#) for examples and descriptions of these logs.

- 12: If the replication processes are not running, trying starting them manually. On the primary server, execute the utility that activates the primary server for replication.

```
# ./server/bin/internal/start_replication
```

This utility:

- refreshes the primary server replication configuration
- starts a process that monitors and captures database changes on the primary server
- starts a process that updates the failover server database with recent primary server database changes at set intervals



You should back up the SQL replication configuration immediately after successfully installing and configuring Security Server software on both the primary and failover Security Servers. See [“Backing up SQL replication configurations” on page 13-224](#).



If you want to stop database replication, use the `stop_replication` utility to stop the processes, stop tracking changes to the primary server database, and stop sending updates to all the failover servers.

The utilities are described more fully below.

`./server/bin/cleanup_replication`

This utility removes failover Security Server authentication credentials from the primary Security Server. This utility is used before running `configure_replication`. If failover server credentials exist when `configure_replication` is run, `configure_replication` will fail because it does not overwrite existing credentials. Run `cleanup_replication` on the primary Security Server as the administrative user (i.e., `voradmin`).

`./server/bin/configure_replication`

This utility is run on the primary server by the Encryption Expert administrative user. It modifies the schema and inserts data to support replication. Nothing is done to the failover server schema. Back up the primary server database with `backup_config` before executing this utility. The effect of this utility can be undone by restoring the backup with `restore_config`. You must restart the Security Server application for the configuration changes to take effect.

See also [“Backing up the Security Server configuration” on page 16-295](#).

`./server/bin/internal/start_replication`

This utility activates database replication. Run it on the primary server as the administrative user (for example, `voradmin`). It starts processes to capture changes (`asncap`) and apply changes (`asnapply`) on the primary server. Nothing is run on the failover server. However, when the configured wait interval passes, the failover server schema is updated if any configuration changes had occurred on the primary server since the previous update. Note that on configured primary servers, the `cgss` utility automatically calls `start_replication`.

`./server/bin/internal/stop_replication`

This utility deactivates database replication. Run it on the primary server as the administrative user (for example, `voradmin`). It stops the processes that capture changes (`asncap`) and apply those changes (`asnapply`) on the primary server. Nothing is done on, or to, the failover server. The failover server just stops getting updates. Execute `start_replication` to resume server replication. Execute `cleanup_replication` and `configure_replication` again only if the failover server had changed.

Backing up SQL replication configurations

For purposes of maintenance or SQL replication reconfiguration, back up the primary server when SQL replication is working.

Back up the failover server configuration immediately after successfully installing the Security Server software but before starting SQL replication. This backup is a “zero-state” image of the failover server database and it will be required should the primary server and failover server go out of sync.

To back up SQL replication configurations:

- 1: Log into the primary server as the Security Server administrator. The default administrator is `voradmin`.
-

- 2: Stop server replication.

```
# /opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/bin/
internal/stop_replication
```

Perform this step to back up the primary server SQL replication configuration only. Skip this step if you are backing up the failover server SQL replication configuration.

- 3: Log into the primary or failover server as root.

- 4: Run the server database backup utility.

```
/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/bin/
backup_config
```

See “[Backing up the Security Server configuration](#)” on page 16-295 for details.

- 5: Log into the primary server as the Security Server administrator.

- 6: Restart replication on the primary server.

```
# /opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/bin/
internal/start_replication
```

Deleting a failover server

You can delete a failover server from the Encryption Expert database and leave the actual failover server software intact on the host system. You will delete the failover server from the Encryption Expert database and then regenerate it when updating primary-failover certificates and when updating failover server software.

To delete a failover server from the Encryption Expert database:

- 1: Log into the Management Console as a user with `server` role permissions.

- 2: Select **High Availability** in the menu bar.

The *High Availability Servers* window opens.

- 3: Enable the failover server selection checkbox.

The selection checkbox is located in the **Select** column next to the failover server name.

- 4: Click **Delete**.

You are prompted to verify the deletion.

- 5: Click **Ok**.

- 6: Run `cleanup_replication` to remove the last traces of the failover server configuration.

What to do in the event of a failure

This section describes what to do if you experience a failure. The topics covered in this section are:

- Primary Security Server failure
 - Failover Security Server failure
 - Primary/failover Security Server databases go out of sync
-

Primary Security Server failure

In this scenario, the primary Security Server has failed. To recover, you will need to configure a new primary Security Server. To do this:

- 1: Install a new Security Server.

During installation:

- Set the server role to **Primary**.
- If you are installing the server software on a different system, and if possible, use the same host name as the failed primary server.
- If you are installing the server software on a different system, update the DNS server with the new primary server information.

If you cannot reuse the host name, you will have to change the configuration and regenerate the certificates for each agent that uses the primary server. Also, client systems usually cache DNS information that can sometimes go stale and require flushing.

- 2: Execute the `./server/bin/restore_config` utility with the most recent primary server backup.
- 3: Execute the `./server/bin/re_gen_cert` utility to regenerate server certificate.
- 4: Use the `cleanup_replication` tool on the primary Security Server.

You might observe some failure at this step depending on the time backup was performed in previous primary server.

- 5: Run `configure_replication` on the primary Security Server and specify the newly installed Security Server for failover
- 6: Restart the primary Security Server using `"/etc/init.d/cgss restart"`. This will automatically activate replication.
- 7: On each agent, re-generate certificates: run `"setup_eet_agent_cert_db2 clean"`; then run `setup_eet_agent_cert_db2`.

Before running the `setup_eet_agent_cert_db2` utility on each agent to re-register, click on each of the hosts in management console and uncheck "Communication Enabled", otherwise the agent registration process will fail.

- 8: Restart the primary Security Server using `"/etc/init.d/cgss restart"`. This will automatically activate replication.

Failover Security Server failure

In this scenario, the failover Security Server getting the automated SQL replication updates has failed. To recover, you can setup another failover Security Server.

- 1: Stop replication process on primary Security Server by running the `stop_replication` utility.
 - 2: Go to High Availability section of primary Server management console. Delete the existing failover server and add a new failover server.
 - 3: Install a new Security Server; choose its role as Failover.
 - 4: Use the `cleanup_replication` tool on the primary Security Server.
 - 5: Run the `configure_replication` tool on the primary Security Server and specify the newly installed Security Server for failover
-

- 6: Restart the primary Security Server using `"/etc/init.d/cgss restart"`. This will automatically activate replication.

There are two scenarios for installing a failover Security Server on a new host:

- Keep the new failover server host name same as the previous failover server host name.
 - Update the DNS server with the IP of new failover server and keep the same host name. This approach has an advantage that agent configuration does not require any modification. However, clients usually cache DNS information and sometimes that information can become stale.
- Use a new failover server host name.
 - This requires that all agent configurations are updated with this new failover server host name and new certificates generated.

If the Security Server databases go out of sync

In this scenario, the failover Security Server and primary Security Server databases are out of sync. This could be caused by either a network glitch, or by making a configuration change directly on the failover server.

To recover from this:

- 1: Stop replication on the primary server using the `stop_replication` utility. Note that `stop_replication` stops ALL replication. You cannot stop replication for a specific failover server.
 - 2: Restore the backup of the empty failover server database using the `restore_config` utility and providing the path to the failover server database backup file.

This step is optional. However, replication goes in one direction, from the primary server to the failover server. If the failover server configuration is to be completely in sync with the primary, restore the backup that you had made earlier of the clean, newly installed failover server. For example:

```
# restore_config -f /tmp/Backup_2008_06_18_1218.tar
```
 - 3: Run `synchronize_replication` on the primary Security Server and specify the failover server host name and `voradmin` password. This utility will synchronize the primary Security Server and failover Security Server databases.
 - 4: Restart the primary Security Server using `"/etc/init.d/cgss restart"`. This will automatically activate replication.
-

Managing Encryption Expert logs

System events are recorded in various log files, and they are displayed in the Security Server management interface. Several log files are generated during software installation and while using the Security Server. You can view logs files directly in an editor, or you can view select log entries in the Security Server management interface.

You can display the Security Server log and configure the events that it displays in the management interface. Displaying and configuring the Security Server log are described in this section. Viewing the individual log files is described in [“Using log files” on page 14-237](#).



Log entries are usually sorted by their timestamps because this sorting method is the only way to determine event sequence. It is important to synchronize the clocks of all the systems that host the Security Server and the agents that access that server. If the system clocks have a tendency to diverge from the actual time, configure the systems to synchronize with an NTP server.

The Security Server management interface is used to:

- Display log entries for both Security Server and agent events
- Configure Security Server logging parameters
- Configure agent logging parameters

Displaying the Security Server log

The Security Server log displays:

- Attempts to make a database backup that includes the Encryption Expert shared library option, `libeetdb2.so`, `eetdb2.dll`, or the `eetidsagt` application. If the appropriate command-line options are not included in the backup request, it is considered to be a regular database operation, the Security Server is not accessed, and no entry is generated in the Security Server log. Attempts to restore a backup are logged unless you are restoring a non-Encryption Expert-protected database backup. Refer to the regular DB2 or IDS log to view non-Encryption Expert backup and restore operations.
 - Any Security Server configuration change that occurs. Only an indication that a change has occurred is provided. Specific information about the changes is not provided.
-

sort by

Log Filter

Log Type: All

Jump To Page: 1

Last Refreshed: 2008-07-18 11:43:02.149

Search Source:

Search Message:

Refresh

Purge Debug Logs Export Logs Delete Logs

Total: 26

Page 1 of 1

ID	Time	Severity	Source	Message
2619	2008-07-18 10:46:49.238	I	vmlinux01	EET0446: Administrator secadmin deleted SysLog Host localhost.
2618	2008-07-18 10:36:04.226	I	vmlinux01	EET0016: The Security Server permitted a backup request from host vmlinux01 for server name(demo_on), Object ID(0), Server Number(1). The policy used is (allowAll_ids), the rule used is (backupRule01).
2617	2008-07-18 10:35:55.536	I	vmlinux01	EET0150: The Security Server matched policy (allowAll_ids) and rule (backupRule01) for the BACKUP request from host (vmlinux01), the server name demo_on, Object ID 0 and Server Number 1.
2616	2008-07-18 10:35:55.415	I	vmlinux01	EET0020: The Security Server received a backup request from host vmlinux01 for server name(demo_on), Object ID(0), Server Number(1).
2615	2008-07-18 10:27:03.777	E	vmlinux01	EET2602: [SecFS, 0] ALARM Policy[allowAdminFileOpsOnly_fs] User[root,uid=0,gid=0[root,bin,daemon,sys,adm,disk,wheel/] Process[/bin/cat] Action[read_file] Res[/opt/apps/apps1/doc/file1.txt] Key[None] Effect[DENIED Code (1U,2M)]
2614	2008-07-18 10:29:55.414	I	vmlinux01	EET0016: The Security Server permitted a backup request from host vmlinux01 for server name(demo_on), Object ID(0), Server Number(1). The policy used is (allowAll_ids), the rule used is (backupRule01).
2613	2008-07-18 10:29:55.407	I	vmlinux01	EET0150: The Security Server matched policy (allowAll_ids) and rule (backupRule01) for the BACKUP request from host (vmlinux01), the server name demo_on, Object ID 0 and Server Number 1.
2612	2008-07-18 10:29:55.348	I	vmlinux01	EET0020: The Security Server received a backup request from host vmlinux01 for server name(demo_on), Object ID(0), Server Number(1).
2611	2008-07-18 10:24:46.54	E	vmlinux01	EET2602: [SecFS, 0] ALARM Policy[allowAdminFileOpsOnly_fs] User[root,uid=0,gid=0[root,bin,daemon,sys,adm,disk,wheel/] Process[/bin/cat] Action[read_file] Res[/opt/apps/apps1/doc/file1.txt] Key[None] Effect[DENIED Code (1U,2M)]
2610	2008-07-18 10:24:44.132	I	vmlinux01	EET2601: [SecFS, 0] AUDIT Policy[allowAllOps_fs] User[root,uid=0,gid=0[root,bin,daemon,sys,adm,disk,wheel/] Process[/bin/cat] Action[read_file] Res[/opt/apps/apps2/doc/file1.txt] Key[None] Effect[PERMIT Code (1M)]
2609	2008-07-18 10:24:41.631	E	vmlinux01	EET2602: [SecFS, 0] ALARM Policy[denyAllOps_fs] User[root,uid=0,gid=0[root,bin,daemon,sys,adm,disk,wheel/] Process[/bin/cat] Action[read_file] Res[/opt/apps/apps3/doc/file1.txt] Key[None] Effect[DENIED Code (1M)]
2608	2008-07-18 10:27:10.339	I	vmlinux01	EET0445: Administrator secadmin added SysLog Host localhost.
2607	2008-07-18 10:27:04.231	I	vmlinux01	EET0446: Administrator secadmin deleted SysLog Host ss2.i.vormetric.com.
2606	2008-07-18 10:47:49.246	I	vmlinux01	EET0016: The Security Server permitted a backup request from host vmlinux01 for server name(demo_on), Object ID(0), Server Number(1). The policy used is (allowAll_ids), the rule used is (backupRule01).

Figure 14.1: Management Console log

To display the Security Server log:

- 1: Log into the Management Console as a user with **Server** role permissions.
- 2: Select **Log** in the menu bar.

The *Logs* window opens. The *Logs* window is comprised of the following five columns:

Column	Description
ID	Entries are sequentially numbered in the same order in which the Security Server enters them into the log database.
Time	The time at which the event occurred. Timestamps are in the form <i>YYYY-MM-DD HH:MM:SS.mm</i> , where <i>Y</i> =year, <i>M</i> =month, <i>D</i> =day, <i>H</i> =hour, <i>M</i> =minute, <i>S</i> =second, and <i>m</i> =millisecond. For example, 2008-02-30 16:37:07.83. Timestamps are relative to the Security Server for Database Backup Agents and relative to the agent for File System Agents. If the Security Server and File System Agents are in different timezones, time order will not match the event sequences.

Column	Description
Severity	Severity levels are D(ebug), I(nfo), W(arn), E(rror), and F(atal). Severity is configurable and only messages that match the severity level are entered in the log. Log levels are cumulative, so each level includes the levels below it. For example, FATAL logs only fatal messages, whereas WARN logs warning messages, and includes ERROR messages and FATAL messages. Details about the log levels are provided in “Log message levels” on page 14-237 .
Source	The name of the host on which the event takes place.
Message	The message associated with the event. For example, “EET0282I: Administrator “bc” updated online policy “fs_policy01.” or “EET0006I: The Security Server permitted a backup request from host vmlinux02 for dbinstance(db2inst1), dbalias(SAMPLE), dbpartition(0). The policy used is (db2_policy01), the rule used is (backupRule01).”

See [“Analyzing log entries” on page 14-233](#) for a description of log entries and how to interpret them.

Setting general display attributes

You can sort the entries displayed in the *Logs* window based upon ID, Time, Severity, or Source. Click a column label, followed by **Refresh**, to display Security Server log entries based upon the selected criteria. Do it again to reverse the order based.

By default, entries are displayed based upon their ID number, with the highest numbers displayed at the top of the window and the smallest at the bottom. If you click **ID**, the entries with the lowest numbers will be displayed at the top and the highest at the bottom. The little green arrow next to the column label indicates the sort order for that column. An upward pointing arrow indicates an ascending order. In the case of ID numbers, an upward pointing arrow indicates that the display starts with the smallest number and ascends to the highest number. In the case of time values, an upward pointing arrow indicates that time values are displayed starting with the oldest time and ascending to the most recent time. The sort criteria in [Figure 14.1: “Management Console log”](#) displays log entries based upon the event timestamp, and they are displayed in descending order, with the most recent at the top of the display.

You can display specific log entries in addition to sorting the entries. The table below lists display fields that you can set.

Field	Description
Log Type	Limit the log entries displayed to just server entries, just agent entries, or both. You may choose Agent , Server , or All from the pop-up menu. All log entries are displayed by default.

Field	Description
Search Source	Searches the currently displayed log entries for messages generated by the specified host and displays only those entries generated by the specified host. Enter full host names, as displayed in the Source column.
Search Message	Searches the currently displayed entries in the Message column for all occurrences of the specified string. All entries that contain the specified string are displayed.
Jump to Page	Displays the log entries of a particular page. Enter an integer between 1 and the total number of pages in the Management Console log.
Last Refreshed	Indicator as to the last time the log page was refreshed.

Click the **Refresh** button to apply the display settings.

The following table lists additional actions that can be applied to the log data:

Button	Description
Refresh	Click this button to display the latest log entries. You must click Refresh to: -- display log entries that may have been generated since you first opened the <i>Log</i> window -- apply a display criteria, such as Log Type , Search Source , Search Message , and Number of Messages
Purge Debug Logs	Click this button to remove all DEBUG level messages from the <i>Log</i> window and the Security Server database.
Export Logs	Saves the entire Security Server log database to a file, regardless of the number or type of log entries currently being displayed. It is recommended that you export your logs and back them up with some regularity. The log data is saved as ASCII text. The file is saved to the <code>/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/server/default/exported_logs</code> directory. The file name is: <code>CGSS_LOG_VIEW_UNTIL_YYYY-MM-DD-HH.MM.SS.XXXXXX.csv</code> . For example, <code>CGSS_LOG_VIEW_UNTIL_2008-04-25-22.31.32.703000.csv</code> .
Delete Logs	Removes the entire Security Server log database, regardless of the number or type of log entries currently being displayed. It is recommended that you export the log data before deleting it.

VCR-style navigation buttons are displayed in the *Logs* window. These buttons are always displayed. Use these buttons to advance between pages. The VCR-style buttons are:

	First. Display the first group of log entries in the <i>Logs</i> window. The number of log entries in each group is determined by the log display attribute defined in the <i>General</i> tab of the <i>Preferences</i> window.
	Previous. Display the previous group of log entries in the <i>Logs</i> window.
	Next. Display the next group of log entries in the <i>Logs</i> window.
	Last. Display the last group of log entries in the <i>Logs</i> window.

Configuring default log attributes

The *Server* tab in the *Preference* window displays the current Security Server logging configuration.

- See [“Setting Security Server logging preferences” on page 15-253](#) for details about setting the default Security Server logging attributes.
- See [“Setting DB2 Database Backup Agent logging preferences” on page 15-255](#) for details about setting the default Database Backup Agent logging attributes.
- See [“Setting DB2 Database Backup Agent logging preferences” on page 15-255](#) for details about setting the default File System Agent logging attributes.

Analyzing log entries

Setting log levels and general attributes for all hosts are described in [“Setting preferences” on page 15-249](#).

Setting log levels and general attributes for a specific host or host group are described in [“Setting preferences” on page 15-249](#).

[Figure 14.2: “Management Console log entries”](#) shows a typical Management Console log. It shows both File System Agent and Database Backup Agent activity.

View All Total: 200				
Purge Debug Logs Export Logs Delete Logs				
ID	Time	Severity	Source	Message
100000001229	2007-11-02 12:56:04.747	I	vmlinux02	EET4537I: [db2inst1, 0, SAMPLE, 24969] ACCESS GRANTED on request to backup database SAMPLE, on hostname vmlinux02. Instance name is db2inst1, partition 0, catnode 0. The timestamp is 2007-11-02 12:55:37, and the DB2 version is 8020500. The request protocol version is 4.0.0.
100000001228	2007-11-02 12:55:54.725	I	vmlinux02	EET4553I: [db2inst1, 0, SAMPLE, 24969] Internal initialization successful.
100000001224	2007-11-02 12:27:57.819	I	vmlinux02	EET2602I: [SecFS, 0] AUDIT: Policy[permitAll] User[root,uid=0,gid=0] Process[/bin/l] Access[read_dir] Res[/] Effect [PERMIT Code (1M)]
100000001223	2007-11-02 12:27:57.796	I	vmlinux02	EET2602I: [SecFS, 0] AUDIT: Policy[permitAll] User[root,uid=0,gid=0] Process[/bin/l] Access[read_dir] Res[/] Effect [PERMIT Code (1M)]
100000001222	2007-11-02 12:27:56.831	I	vmlinux02	EET2602I: [SecFS, 0] AUDIT: Policy[dailyMaint_policy] User[root,uid=0,gid=0] Process[/bin/l] Access[read_dir] Res[/] Effect [DENIED Code (1U,2M)]
100000001232	2007-11-02 09:31:59.445	I	vmlinux01	EET0006I: The Security Server permitted a backup request from host vmlinux02 for dbinstance(db2inst1), dbalias(SAMPLE), dbpartition(0). The policy used is (policy01), the rule used is (backupRule01).
100000001231	2007-11-02 09:31:59.43	I	vmlinux01	EET0006I: The Security Server permitted a backup request from host vmlinux02 for dbinstance(db2inst1), dbalias(SAMPLE), dbpartition(0). The policy used is (policy01), the rule used is (backupRule01).
100000001230	2007-11-02 09:31:59.412	I	vmlinux01	EET0144I: The Security Server matched policy (policy01) and rule (backupRule01) for the BACKUP request from host (vmlinux02), the database instance db2inst1, alias SAMPLE and partition 0.
100000001227	2007-11-02 09:31:48.616	I	vmlinux01	EET0144I: The Security Server matched policy (policy01) and rule (backupRule01) for the BACKUP request from host (vmlinux02), the database instance db2inst1, alias SAMPLE and partition 0.
100000001226	2007-11-02 09:31:44.819	I	vmlinux01	EET0010I: The Security Server received a backup request from host vmlinux02 for dbinstance(db2inst1), dbalias(SAMPLE), dbpartition(0).
100000001225	2007-11-02 09:31:44.415	I	vmlinux01	EET0010I: The Security Server received a backup request from host vmlinux02 for dbinstance(db2inst1), dbalias(SAMPLE), dbpartition(0).
100000001221	2007-11-02 08:37:44.412	I	vmlinux01	EET0267I: Administrator "bcobain" logged in at 2007-11-02 08:37:38.967.
100000001220	2007-11-02 08:31:29.416	I	vmlinux01	EET0116I: The security server has successfully initialized the logging system to log messages from hosts. The server can now log messages from hosts.

Figure 14.2: Management Console log entries

The format and content of log entries for File System Agents and Database Backup Agents are different and are described separately below.

Analyzing Database Backup Agent log entries

The messages in the log file are descriptive and should be self-evident. Some messages of note are:

- 2008-02-02 15:02:22.770 [COM] [DEBUG] [11198] Successful Agent init with server https://10.2.0.37:8443/communication-server/SSControllerDB2

This message indicates that the agent can make a network connection to the server. Every backup and restore operation must include a similar message. If the operation does not, the server configuration on the agent is wrong or the Security Server is not accessible on the network.

- 2008-02-15 15:28:11.528 [POL] [ERROR] [14614] [ET0003E] ACCESS DENIED on request to back up database SAMPLE, on hostname vmlinux02. Instance name is db2inst1, partition 0, catnode 0. The timestamp is 2008-02-15 15:28:08, and the DB2 version is 8020500. The request protocol version is 4.0.0.

This message indicates a successful connection to the Security Server and recognition by the Security Server. The requested operation was denied based upon the policy. Check the policy for this agent on the Security Server.

- 2008-02-17 16:10:26.044 [POL] [INFO] [24164] [EET0005I]
ACCESS GRANTED while attempting to restore database FOO,
on hostname vmlinux02. Instance name is db2inst1,
partition 0, catnode 0. The timestamp is 2006-12-17
16:10:25, and the DB2 version is 8020000. The request
protocol version is 4.0.0.

This message indicates that the agent is granted permission to complete the requested restore operation. It indicates only that permission is granted. It does not indicate that the operation completed successfully.

- 2008-02-02 11:03:40.089 [POL] [INFO] [2741] [EET4576I]
Restore completed successfully.

This message indicates that the restore operation completed successfully.

- 2008-02-07 17:11:03.626 [DB2] [SEVERE] [12078] [EET2029S]
Backup data request failed: access denied or a related
cause.

This message is generated by DB2 and indicates that a restore operation failed. The agent successfully contacted the server, but the Security Server denied permission to complete the operation. Check the policy for this agent to ensure that the agent has permission to perform the requested operation.

- 2008-02-02 15:07:16.691 [COM] [SEVERE] [11384] [EET4515S]
There was a communication failure with the Security
Server. The server returned a fault: Error 26; "SOAP-
ENV:Server"; Subcode: "[no subcode]"; Reason: "SSL
error"; Detail: "Can't read certificate key file".

2008-02-02 15:07:16.350 [COM] [SEVERE] [32417] [EET0019F]
SSL context initialization failure: error 8.

2008-02-02 15:07:16.353 [COM] [SEVERE] [32417] [EET0007F]
Security Server connection failure.

These three messages are part of a SSL authentication failure between the agent and the server. The agent certificate file, `agent.pem`, does not exist in the `./agent/pem` directory. Regenerate the agent SSL certificate with `setup_eet_agent_cert_db2` and try again. You might need to run `setup_eet_agent_cert_db2 clean` first.

Attempts to perform a backup may cause DB2 to return a *reason code* of 1. For example: `SQL2062N An error occurred while accessing media "libeetdb2.so". Reason code: "1".` A reason code of 1 doesn't help you to troubleshoot an issue; however, the other log entries adjacent to this entry should describe the problem.

Analyzing File System Agent log entries

The general format of a File System Agent log entry is:

```
EET2602I: [SecFS, 0] AUDIT: Policy[policyName?] User[userID?]  
Process[command?] Access[whatIsItDoing?]  
Res[whatIsItDoingItTo?] Effect[allowOrDeny? Code  
(whatMatched?)]
```

where:

- `Policy[]` indicates the name of the policy that is being used to evaluate the access attempt.
- `User[]` identifies the system user attempting to access data in the guard point. It typically displays the user name, user ID, and group ID.
- `Process[]` indicates the command, script, or utility being executed.
- `Access[]` indicates what is being attempted. Access may be `read_dir`, `remove_file`, `write_file_attr`, `write_app`, `create_file`, etc. These correspond to the **Access** methods that you configure in the policy. `read_dir` corresponds to `d_rd`. `remove_file` corresponds to `f_rm`. And so on.
- `Res[]` indicates the object being accessed by `Process[]`.
- `EFFECT[]` indicates the rule that matched and, based upon that rule, whether or not the Security Server grants access. Rule matching is described below. Access states may be either `PERMIT` or `DENIED`.

For example:

```
EET2602I: [SecFS, 0] AUDIT: Policy[dailyMaint_policy]
User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\]
Process[/bin/ls] Access[read_dir] Res[/] Effect[DENIED Code
(1U,2M)]
```

The format of a rule match is:

intchar

int is an integer representing the security rule being used or violated. Security rules are numbered sequentially in the Policy Composer window, beginning at the top of the scroll-list down to the bottom.

char is an uppercase letter indicating the thing that is using or violating the policy.

Character Code	Description
A	The Action component of a security rule failed to match.
M	All security rule components match and, unless overridden, the Effect for that security rule is applied.
P	The Process component of a security rule failed to match.
R	The Resource component of a security rule failed to match.
T	The time specified in the When component of a security rule failed to match.
U	The User component of a security rule failed to match.

For example, the following match codes indicate:

- **1R** – Mismatch in Resource for Security Rule 1.
- **3U** – Mismatch in User for Security Rule 3.
- **4A** – Mismatch in Action for Security Rule 4.
- **2M** – All components matched for Security Rule 2. Since all the rules matched, Security Rule 2 will be used and no other rules will be evaluated.

Log message levels

The detail and extent of information logged is determined by the selected log level. The agent supports five log levels. In sequence they are:

- `DEBUG`
- `INFO`
- `WARN`
- `ERROR`
- `FATAL`

Log levels are cumulative. The level that you select not only generates log entries for events that occur at that level, but all the levels below. For example, the `WARN` level also includes events that occur on the `ERROR` and `FATAL` levels.

Severity	Description
DEBUG	The <code>DEBUG</code> level designates fine-grained informational events that are targeted towards support engineers and developers.
INFO	The <code>INFO</code> level designates informational messages that highlight the progress of the application at coarse-grained level.
WARN	The <code>WARN</code> level designates potentially harmful situations.
ERROR	The <code>ERROR</code> level designates error events that might still allow the application to continue running.
FATAL	The <code>FATAL</code> level designates very severe error events that will presumably lead the application to abort.

Using log files

Check the log files to verify the successful installation and configuration of the Encryption Expert software, to determine why a backup or restore operation failed, or to monitor Encryption Expert activity.

A logged event falls into one of the following categories:

- Operational status. The result of any significant action performed by an Encryption Expert agent or Security Server is logged.
- Administrative activity. The result of any maintenance or administrative activity on the Security Server is logged (for example, a key has been created or exported).
- System status. The result of any system errors are logged (for example, if the database connection is interrupted).
- A policy-specified audit. If the result of a policy evaluation specifies that it should be audited, then a suitable message is logged.

Several logs files are provided. Each serves a different purpose. The log files are:

- `db2setup.log` - Generated during the DB2 installation portion of Encryption Expert installation or after installing DB2 by itself.
- `eet_jboss.log` - Updated when the Encryption Expert Security Server is stopped (for example, `/etc/init.d/cgss stop`), started, or restarted. View this log file when you start or stop the Security Server.

Some files, such as `server.log` and `cgss.log` are large, and large amounts of data are written to them as processes and operations are performed. To better spot log entries and monitor Encryption Expert status as events are occurring, execute the `tail` command with the `-f` option on the target file. For example:

```
# tail -f cgss.log
```

Log files can become large and should be backed up and the older, inactive log files should be purged at regular intervals. Log files can get large enough to exceed partition capacity and disrupt logging. Log file space should be considered on the Security Server, and the agents that access that server, because exceeding the available space on one can affect the other. For example, when the server LOGDB database partition becomes full, the Security Server refuses to accept additional agent log entries. The log data is then written locally on the agent system. Also, the client system `/tmp` directory can become full, which makes the `/tmp` directory unusable to anyone.

You can delete the Encryption Expert log files in `/tmp`, and inactive Encryption Expert logs in other directories. Active logs are log files that being currently written to and updated by Encryption Expert processes. Inactive logs are logs that have been filled to capacity and then closed. The name of the closed log file is the original name usually appended with the date and some random numbers. For example, the name of the active agent log is `eetdb2_usr.log`. When it reaches the configured capacity, it is made inactive and renamed to `eetdb2_usr.log.YYYY-MM-DD-XX-XX-XX`. Do not attempt to manually modify or remove active logs. Doing so can render the Security Server inoperable. Use the Management Console interface to configure server and agent logs. Backup and delete inactive logs with some regularity to maximize available hard disk space.



(Windows only) The `\Documents and Settings\All Users\Application Data` folder is hidden by default. Encryption Expert stores agent logs, configuration data, and certificates under that folder. If you cannot browse the folder for your platform, enable the **Show hidden files and folders** radio button in the `Folder Options` menu to view the folder and its contents.

Encryption Expert File System Agent log files

The agent logs are the first places to check when communication between the Security Server and agent system fails. Also, you may want to check these logs after setting up a new agent or changing the agent configuration.

eetvmd_usr.log

(UNIX) /var/log/eet/eetvmd_usr.log
 (Windows) \Documents and Settings\All Users\Application Data\IBM\DBTOOLS\LUWEncryptionExpert\Agent\log\eetvmd_usr.log

eetvmd_usr.log contains the File System Agent transactions for a specific user. Since vmd is started and run by the root user, the filename is always eetvmd_root.log. Transactions consist of a record of vmd actions, such as starting the vmd daemon and setting up communication links with the Security Server. It also contains log messages that are normally displayed in the Management Console log but for some reason cannot be uploaded to the Security Server. An excerpt of this file is shown below.

```
2008-03-30 16:59:30.812 [COM] [DEBUG] [2311] Initialized server
structure to https://vmlinux01:8443
2008-03-30 16:59:30.867 [COM] [SEVERE] [2311] [EET4515S] There was
a communication failure with the Security Server. The server
returned a fault: SOAP 1.1; "SOAP-ENV:Client"; Subcode: "[no
subcode]"; Reason: "Connection refused"; Detail: "connect failed in
tcp_connect()".
2008-03-30 16:59:30.874 [COM] [SEVERE] [2311] [EET4583F] Error
initializing agent.
...
...
2008-03-30 17:43:53.841 [POL] [INFO ] [2311] [EET4553I] Internal
initialization successful.
2008-03-30 17:43:53.848 [COM] [DEBUG] [2311] Initialize Agent VMD
succeeded
2008-03-30 17:43:53.853 [COM] [DEBUG] [2311] Server communication
established
2008-03-30 17:43:54.880 [COM] [INFO ] [2311] [EET3640I] STARTED on
linux PORT 7024 VERSION: 4.0.0.0
2008-03-30 19:21:09.684 [COM] [DEBUG] [8739] Locale set to default:
en_US
2008-03-30 19:21:09.791 [COM] [INFO ] [8739] [EET3773I] The
VMD is starting. It was invoked at the command line like this: "/
usr/bin/vmd"
```

messages (UNIX only)

/var/log/messages

messages is a syslog-generated file. It contains standard syslog entries. It contains kernel entries for enabling/disabling the logger, memory usage, CPU usage, system calls, device initialization, etc. It also contains log entries that would otherwise be displayed in the Management Console log but, for some reason, cannot be uploaded to the Security Server.

An excerpt of /var/log/messages with Encryption Expert-related messages is shown below:

```
Oct 30 13:36:29 vmlinux02 kernel: vmfs_switch_sdk_logger: Switching
to upload-messages mode
```

```
Oct 30 13:37:26 vmlinux02 sshd(pam_unix)[28241]: authentication
failure; logname = uid=0 euid=0 tty=ssh ruser= rhost=10.2.60.25
user=bcobain
Oct 30 13:37:27 vmlinux02 kernel: Login for user(bcobain, 1001, 10/
[wheel]) Failed with PAM_AUTH_ERR
Nov  1 14:49:11 vmlinux02 kernel: [AUD] [INFO ] [EET2602I] AUDIT:
Policy[permitAll] User[bcobain,uid=1001,gid=10\wheel\] Process[/
bin/ls] Access[read_dir] Res[/] Effect[PERMIT Code (1M)]
```

secfs.log

The `secfs.log` file contains kernel-related messages, and the `secfsd.log` file contains process-related messages. The `secfs.log` file is only generated on AIX systems. The `secfs.log` file is maintained in the `./agent/secfs/tmp` directory. It is used instead of `syslog` to log kernel messages. The same log messages are placed in both `/var/log/messages` and `secfs.log`. The `secfs.log` file is archived at 32MB and renamed to `secfs.log.archive`. Only one archive file is maintained.

SECFS_Logger_secfs_log.XXX

```
/var/log/eet/SECFS_Logger_secfs_log.XXX
```

Template file for the header that is prepended to `vmd_upload` files.

On Windows systems, this is `\Documents and Settings\All`

`Users\Application`

```
Data\IBM\DBTOOLS\LUWEEncryptionExpert\Agent\log\SECFS_upload_
host.num
```

```
<logfile>
<agenttype>FS</agenttype>
<version>1</version>
<sdkversion>4.0.0</sdkversion>
<pid>0</pid>
<hostname>vmlinux02</hostname>
<tag>588477390</tag>
<component>SecFS</component>
<subtag>236323513</subtag>
```

secfsd.log

```
/opt/IBM/DB2TOOLS/LUWEEncryptionExpert/agent/secfs/tmp/secfsd.log
\Documents and Settings\All Users\Application
Data\IBM\DBTOOLS\LUWEEncryptionExpert\Agent\log\secfsd.log
```

The `secfs.log` file contains kernel-related messages, and the `secfsd.log` file contains process-related messages. `secfsd.log` contains a record of guard point mounts and guard point dismounts (guard points are mounted file systems). Entries are added to this file when you add and remove guard points, as well as when you reboot the agent system. An excerpt is shown below.

```
/usr/bin/secfsd[2250]: Information: Tue Oct 30 23:48:11 2007 mount
cmd (mount -t secfs2 "/home/bob/dir1" "/home/bob/dir1" -o
"sa=23,lock=1,type=1,dir=/home/bob/dir1")
```

See also [“secfsd” on page 289](#).

statusfile

```
/opt/IBM/DB2TOOLS/LUWEEncryptionExpert/agent/secfs/tmp/statusfile
```


\Program
Files\IBM\DB2TOOLS\LUWEncryptionExpert\agent\secfs\tmp\statusfile
statusfile is a current record of the local File System Agent configuration. View this file after updating the File System Agent configuration on the Security Server to verify that the changes have actually been applied. This file should always be checked when the configuration of the File System Agent is in question. This file:

- lists each guard point and guard point properties, such as the lock status, protection status, and guard point directory
- lists the names of applied policies
- lists the logging information that is captured
- lists where captured log information is sent
- lists hosts settings

You can also display the file timestamp to see when the agent was last updated.



This file is deleted each time the File System Agent configuration is updated. You must manually regenerate it using the “secfsd -status” command. If you want to keep records of File System Agent configuration changes, either copy the statusfile to a different name, or run “vmsec status” and tee the output to a different file.

You may view the file contents using an ASCII display command, such as cat. An excerpt of this file is shown below:

```
# cat ./secfs/tmp/statusfile
FILE_FORMAT=1
GP_1_Policy=1
GP_1_Dir=/home/apps/app2
GP_1_lock=1
GP_1_type=1
GP_1_gtype=local
GP_1_opt=gtype=0,policy=1,lock=1,type=1,dir=/home/apps/app2/
GP_1_config_state=guarded
GP_1_prev_config_state=guarded
GP_1_status=guarded
...
TOTAL_GP=2
coreguard_locked=false
system_locked=false
logger_upload_url=https://vmlinux01:8444/upload/logupload
logger_cert_dir=/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/
pem
hostname_for_logging=vmlinux02
policy_name_1=auditOnly
policy_type_1=regular
...
logger_threshold_CGS=INFO
logger_threshold_XFM=WARN
logger_threshold_CGA=INFO
logger_threshold_POL=INFO
...
AUTHBIN_1=|authenticator|/bin/su
C2D8AE0175363C649976A109288A91CFE7F717B1
```

```
AUTHBIN_2=|authenticator|/usr/sbin/sshd
4FD810A4F2E9577E48475596AD30A61E98298A92
```

```
...
```

vmd_upload_hostname.XXX

```
/var/log/eet/vmd_upload_hostname.XXX
\Documents and Settings\All Users\Application
Data\IBM\DBTOOLS\LUWEncryptionExpert\Agent\log
```

vmd_upload_hostname.XXX lists the log uploaded to the Security Server.

Encryption Expert Security Server log files

Encryption Expert Security Server logs are logs that reside on the Security Server. The primary log is viewed in the *Logs* window of the Security Server Web interface. This log is generally the first log that you check to diagnose server problems. Check the Security Server log after making or restoring a database backup. Look for entries like “Backup Request for SAMPLE from host vmlinux02 is allowed.” and “Backup/Restore completed successfully.” Messages like “Backup data request failed: access denied or a related cause.” indicate a problem has occurred and some debugging on your part is required.

badlog.log

Log data that cannot be successfully parsed are referred to as “bad logs”. The `badlog.log` file contains any log data from any agent that is intended for display in the Security Server Web interface but cannot be displayed because it cannot be parsed due to format irregularities. Each attempt by an agent to upload a log file to the server that fails is placed in the `badlogs` directory as a unique file. The server will continue to accept uploaded logs from the agent regardless of the number of failed attempts.

These badlog files originate on the agent in the `/var/log/eet` directory and removed from the agent after they are successfully uploaded to the server.

The log file is `vmdb2logup_hostname_dbInstName_num` and is placed in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/server/default/auditlog/badlogs`. `hostname` is the name of the agent system uploading log data. `dbInstName` is the database user instance on the agent system. `num` is a large number to uniquely and sequentially name badlog files. For example, `vmdb2logup_aix32005_db2inst9_1011720.814`.

boot.log

A record of the JBoss startup process is recorded in `boot.log`. This file lists JBoss parameters and their values. Such as the Java version, the Encryption Expert administrative user that started the JBoss server, plus all the components that the JBoss server runs. Check this log file if the JBoss application server cannot be started.

This file is located in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/server/default/log`.

eetdb2_jboss.log

The `eetdb2_jboss.log` file contains information that is related to starting and stopping the JBoss Web application server. This file is generated when the `/etc/init.d/cgss` command is used to start and stop JBoss. Check this log file for problems that are related to JBoss, such as when you are unable to initiate a Management Console session.

This file is located in `/tmp`.

Example `eetdb2_jboss.log` excerpts:

```
11:44:08,763 INFO [STDOUT] CGSS_HOME=/opt/IBM/DB2TOOLS/
LUWEEncryptionExpert/server/jboss-4.0.4.GA/server/default
11:44:09,796 INFO [EARDeployer] Started J2EE application: file:/
opt/IBM/DB2TOOLS/LUWEEncryptionExpert/server/jboss-4.0.4.GA/
server/default/deploy/cgss.ear
11:44:10,573 INFO [Http11BaseProtocol] Starting Coyote HTTP/1.1 on
http-localhost%2F127.0.0.1-8080
11:44:12,055 INFO [Server] JBoss (MX MicroKernel) [4.0.4.GA
(build: CVSTag=JBoss_4_0_4_GA date=200605151000)] Started in
3m:8s:440ms
```

cgss.log

The `cgss.log` file contains a record of the events that comprise the BEK generation process for an agent requesting to make a backup, as well as the names of uploaded audit files. This file does not contain events that pertain to restore operations. Check this file if the agent fails to back up a database, even though agent/server authentication is correctly configured and the policy for this agent permits the backup operation.

This file is located in `/opt/IBM/DB2TOOLS/LUWEEncryptionExpert/server/jboss-4.0.4.GA/server/default/log`.

Example `cgss.log` excerpts:

```
2008-02-15 11:32:14,209 INFO
[com.widgets.server.crypto.CryptoUtil] Entering generateBek.
2008-02-15 11:32:14,228 INFO
[com.widgets.server.crypto.CryptoUtil] Exiting generateBek.
2008-02-15 11:32:28,637 INFO
[com.widgets.server.upload.LogUploadServlet] HEADER INFO: filename:
vmdb2logup_vmlinux02_db2inst1_10882.42
2008-02-15 11:32:28,854 INFO
[com.widgets.server.upload.LogUploadServlet] The message file /
opt/IBM/DB2TOOLS/LUWEEncryptionExpert/server/jboss-4.0.4.GA/
server/default/auditlog/vmdb2logup_vmlinux02_db2inst1_10882.42 has
been sent for processing.
2008-02-15 11:32:29,258 INFO
[com.widgets.server.upload.MessageProcessor] The uploaded log file
/opt/IBM/DB2TOOLS/LUWEEncryptionExpert/server/jboss-4.0.4.GA/
server/default/auditlog/vmdb2logup_vmlinux02_db2inst1_10882.42 has
been processed successfully.
```

eet_server_install_YYYY-MM-DD.log

The `eet_server_install_YYYY-MM-DD.log` file lists the JBoss files installed in `/opt/IBM/DB2TOOLS/LUWEEncryptionExpert/server/jboss-4.0.4.GA`, status of the DB2 database manager, database server, name of the user with administrative access (for example, `voradmin`), database alias, and the commands used to create the Security Server database schema. At the end it lists the status of the DB2 and JBoss servers.

db2setup.log

The `db2setup.log` file is generated while installing the embedded DB2 package that comes with the Security Server. Check this log file if you suspect that the embedded Security Server database was not correctly installed.

This file is located in `/tmp`.

panic.log

The `panic.log` file is generated when the Security Server Web application server is unable to successfully write to the Security Server embedded database, which happens when the database server unexpectedly shuts-down or when you execute a shut down command like `db2stop force`. If this occurs, the same entries that would have been written to the database are written to the panic log.

The default `panic.log` directory is:

```
./server/jboss-4.0.4.GA/server/default/auditlog/paniclog
```

The panic log location or name cannot be configured. There is no limit to the number of entries that can be placed in the panic log. The panic log can continue to grow until the system runs out of free disk space.

For example, you can see entries like the following in the `cgss.log` file. These messages indicate that the server is unable to write to the audit log in the database:

```
2008-05-07 23:09:22,378 ERROR
[com.widgets.server.logger.AuditLogLet] Error while creating a
Prepared Statement for audit log[IBM][JDBC Driver] CLI0601E Invalid
statement handle or statement is closed. SQLSTATE=S1000
2008-05-07 23:09:22,378 ERROR
[com.widgets.server.logger.AuditLogger] Prepared Statement from
AuditLogLet is null
2008-05-07 23:09:22,378 ERROR
[com.widgets.server.logger.AuditLogger] Audit log storage to
database failed[IBM][CLI Driver] CLI0106E Connection is closed.
SQLSTATE=08003
2008-05-07 23:09:22,378 ERROR
[com.widgets.server.logger.AuditLogger] Calling PanicLogger to
write to file
```

The `panic.log` file contains log entries that are identical to those entries that would otherwise be displayed in the Web browser. For example:

```
I 2008-05-08 06:05:07.379 POL7407I The Security Server matched rule
(c1e1) to the BACKUP request for host aix55050, database instance
xxinst, alias xxalias and partition 0. But the conditions do not
match. The rule will not be applied
I 2008-05-08 06:09:22.378 COM5313I The Security Server denied a
restore request from host aix55050 for dbinstance(db2inst1),
dbalias(ABSOLUT), dbpartition(0). There were no applicable policy
or applicable rule found.
```

server.log

The `server.log` file contains details about agent backup and restore requests, connection status, Management Console interaction, Java exceptions, JBoss start and stop processes, and more. This file contains diverse information and should be checked for almost any problem that is related to the Security Server. Sometimes it is easier to grep a specific error level, like WARN, INFO, or DEBUG, than it is to view the entire file.

This file is located in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/server/default/log`.

The following is an example `server.log` excerpt:

```
2008-05-15 15:04:13,469 INFO [STDOUT] CGSS_HOME=/opt/IBM/
DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/server/
default
```

```

2008-05-15 15:04:13,475 INFO [STDOUT] CGSS_HOME set to: "/opt/
IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/server/
default/"
2008-05-15 15:04:13,668 INFO [STDOUT] PolicyDAOImpl
findApplicablePolicies policies that match the host with
name:vmlinux02
2008-05-15 15:04:13,758 INFO [STDOUT] PolicyDAOImpl
findApplicablePolicies policies that match the host groups:
2008-05-15 15:04:13,760 INFO [STDOUT] PolicyDAOImpl
com.widgets.server.dao.policy.PolicyDB2DTO@cc60cc6[id->1; name-
>vmlinux02policy; enabled->true; updateTime->2008-05-15
15:03:01.11instanceName->; dbAlias->; partitionNum->-1;
relevancyScore->0]
2008-05-15 15:04:13,790 INFO [STDOUT] PolicyDAOImpl
findApplicablePolicies policies that apply to all hosts:
2008-05-15 15:04:13,791 INFO [STDOUT] PolicyDAOImpl
com.widgets.server.dao.policy.PolicyDB2DTO@cc60cc6[id->1; name-
>vmlinux02policy; enabled->true; updateTime->2008-05-15
15:03:01.11instanceName->; dbAlias->; partitionNum->-1;
relevancyScore->0]
2008-05-15 15:04:13,855 INFO [STDOUT] RuleDB2DTO toRule condition
size: 0
2008-05-15 15:04:13,856 INFO [STDOUT] RuleDB2DTO toRule
encryptionAlgorithm: AES128
2008-05-15 15:04:13,886 INFO [STDOUT] RuleDB2DTO toRule condition
size: 0
2008-05-15 15:04:13,888 INFO [STDOUT] PolicyDAOImpl
findApplicablePolicies final results:
2008-05-15 15:04:13,889 INFO [STDOUT] PolicyDAOImpl
com.widgets.server.policy.Policy@60e460e4[id->1; name-
>vmlinux02policy; order->0; enabled->true; updateTime->2008-05-15
15:03:01.11]
2008-05-15 15:04:13,891 INFO [STDOUT] Evaluate
Rule:vmlinux02BkupRule

```

vmdb2logup

The `vmdb2logup_host_usr_num.num` file is a temporary file that is generated when a configured host submits a backup or restore request to the Security Server. This file contains agent logs that are uploaded to, and very briefly stored on, the Security Server auditlog directory before it is loaded into the log viewer database.

This file is located in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/server/default/auditlog`.

The actual file name can be determined by checking the `cgss.log` file. An example file name is `vmdb2logup_vmlinux02_db2inst1_13331.53`.

voradmin.CGSSDB.ASN.CAP.log

The `asncap` process monitors and records database changes to a log file. `asncap`, watches tables on the primary server, like `CGSS_HOST`, and copies changes to corresponding capture tables, like `CDCGSS_HOST` in the following example.

```

# tail /home/voradmin/voradmin.CGSSDB.ASN.CAP.log
2008-04-25-13.48.37.725557 <pruneCDNormal> ASN0105I CAPTURE "ASN"
: "PruneThread". "21" rows have been pruned from the table
"VORADMIN"."CDCGSS_HOST" at "Fri Apr 25 13:48:37 2008".
2008-04-25-13.48.39.035023 <pruneCDNormal> ASN0105I CAPTURE "ASN"
: "PruneThread". "504" rows have been pruned from the table

```

```

"VORADMIN"."CDCGSS_LOGGER_APPENDER_MAPPING" at "Fri Apr 25 13:48:39
2008".
2008-04-25-13.48.39.544099 <pruneCDNormal> ASN0105I CAPTURE "ASN"
: "PruneThread". "252" rows have been pruned from the table
"VORADMIN"."CDCGSS_LOGGER_CONFIG" at "Fri Apr 25 13:48:39 2008".
2008-04-25-13.48.39.711860 <pruneUOW> ASN0105I CAPTURE "ASN" :
"PruneThread". "21" rows have been pruned from the table
"ASN"."IBMSNAP_UOW" at "Fri Apr 25 13:48:39 2008".
2008-04-25-13.53.37.515120 <pruneCDNormal> ASN0105I CAPTURE "ASN"
: "PruneThread". "42" rows have been pruned from the table
"VORADMIN"."CDCGSS_APPENDER_CONFIG" at "Fri Apr 25 13:53:37 2008".
2008-04-25-13.53.37.895529 <pruneCDNormal> ASN0105I CAPTURE "ASN"
: "PruneThread". "189" rows have been pruned from the table
"VORADMIN"."CDCGSS_APPENDER_PROPERTY" at "Fri Apr 25 13:53:37
2008".
2008-04-25-13.53.37.944006 <pruneCDNormal> ASN0105I CAPTURE "ASN"
: "PruneThread". "7" rows have been pruned from the table
"VORADMIN"."CDCGSS_HOST" at "Fri Apr 25 13:53:37 2008".
2008-04-25-13.53.38.407467 <pruneCDNormal> ASN0105I CAPTURE "ASN"
: "PruneThread". "168" rows have been pruned from the table
"VORADMIN"."CDCGSS_LOGGER_APPENDER_MAPPING" at "Fri Apr 25 13:53:38
2008".
2008-04-25-13.53.38.679975 <pruneCDNormal> ASN0105I CAPTURE "ASN"
: "PruneThread". "84" rows have been pruned from the table
"VORADMIN"."CDCGSS_LOGGER_CONFIG" at "Fri Apr 25 13:53:38 2008".
2008-04-25-13.53.38.898837 <pruneUOW> ASN0105I CAPTURE "ASN" :
"PruneThread". "7" rows have been pruned from the table
"ASN"."IBMSNAP_UOW" at "Fri Apr 25 13:53:38 2008".
#

```

voradmin.CGSSDB.CGSSDB_AQ.APP.log

The asnaply process used to upload database changes to a failover server at set intervals logs to:

```

# tail /home/voradmin/voradmin.CGSSDB.CGSSDB_AQ.APP.log
2008-04-25-14.57.00.657141 <CPREST(01/00)> ASN1044I APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "1" minutes and "0" seconds.
2008-04-25-14.58.01.179182 <CPREST(01/00)> ASN1044I APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "0" minutes and "59" seconds.
2008-04-25-14.59.00.724051 <CPREST(01/00)> ASN1044I APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "1" minutes and "0" seconds.
2008-04-25-15.01.44.904505 <CPREST(01/00)> ASN1044I APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "0" minutes and "59" seconds.
2008-04-25-15.02.44.653848 <CPREST(01/00)> ASN1044I APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "0" minutes and "16" seconds.
2008-04-25-15.03.01.199875 <CPREST(01/00)> ASN1044I APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "0" minutes and "59" seconds.
2008-04-25-15.04.00.838227 <CPREST(01/00)> ASN1044I APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "1" minutes and "0" seconds.

```

```

2008-04-25-15.05.02.004614 <CPREST(01/00)> ASN1044I  APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "0" minutes and "59" seconds.
2008-04-25-15.06.01.558036 <CPREST(01/00)> ASN1044I  APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "0" minutes and "59" seconds.
2008-04-25-15.07.01.089212 <CPREST(01/00)> ASN1044I  APPLY
"CGSSDB_AQ" : "WorkerThread". The Apply program will become
inactive for "0" minutes and "59" seconds.
#

```

Database Backup Agent log files

The primary log is displayed in the *Logs* window of the Management Console. Check this log for activity and status information. If problems occur that do not appear in the *Logs* window, check the log file on the agent to diagnose problems. Check the *Logs* window after making or restoring a database backup. Look for entries like “Backup Request for *db* from host *host* is allowed.” and “Backup/Restore completed successfully.” Messages like “Backup data request failed: access denied or a related cause.” indicate a problem has occurred and some debugging on your part is required.

Check the agent log when a backup or restore operation fails, after setting up a new agent, or after changing an existing agent configuration. A communication error between the server and agent can occur for many reasons: the network goes down, the agent network is improperly configured, certificates expire or are modified, and so on. If there is a communication error between the server and agent, agent log messages might not be entered in the *Logs* window.

The agent log is not generated until a backup or restore operation is executed with the Encryption Expert agent library or filter. Only the backup and restore operations that use the Encryption Expert library or filter are recorded in this file.

The agent log file names are `eetdb2_usr.log`, for the DB2 Database Backup Agent, `eetids_usr.log`, for the IDS Database Backup Agent: where, *usr* is the login name of the user who is executing the backup or restore operation. For example, if the database user is `db2inst1` running the DB2 Encryption Expert software, the log file name will be `eetdb2_db2inst1.log`.

The agent log file is stored in two places. The log directory in the Encryption Expert installation hierarchy is symbolically linked to a directory in the system log hierarchy. The files that they contain are identical so it does not matter which log file you view: The locations are:

- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2/log/eetdb2_usr.log`
- `/var/log/eet/eetdb2_usr.log`
- `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/ids/log/eetids_usr.log`
- `/var/log/eet/eetids_usr.log`

Hint: Run “`tail -f`” on the log file to spot events between the agent and server as they occur. For example log file entries and descriptions as to what they mean, see [“Analyzing Database Backup Agent log entries” on page 14-234](#). Some example log entries are presented below.

- The log file displays character encoding and network activity:


```

2008-07-16 18:46:56.722 [COM] [DEBUG] [24662] Locale set
to default: en_US

```

```
2008-07-16 18:46:56.749 [COM] [DEBUG] [24662] Unable to
determine fully qualified hostname. Using vmlinux101
from gethostname
2008-07-16 18:46:56.751 [COM] [DEBUG] [24662] Server
URLs: https://vmlinux01
2008-07-16 18:46:57.123 [COM] [INFO ] [24662] [EET4601I]
Successful Agent init with Security Server (https://
vmlinux01:8443/communication-server/SSControllerDB2).
```

- The log file displays agent information:

```
2008-07-16 18:49:27.603 [COM] [DEBUG] [24662] Agent
build version: 1.1.3.0
2008-07-16 18:49:27.603 [COM] [DEBUG] [24662] Agent
build ID: BUILD_U0084
2008-07-16 18:49:27.603 [COM] [DEBUG] [24662] Agent
build date: 2008-06-23 20:04:51 -0700 (PDT)
```

- The log file displays IDS policy evaluation:

```
2008-07-16 18:49:27.885 [POL] [INFO ] [24662] [EET4615I]
ACCESS GRANTED on request to perform backup for hostname
vmlinux101. Server name is demo_on, server number 1. The
timestamp is 2008-07-16 18:49:27. The request protocol
version is 4.0.0.
2008-07-16 18:49:29.608 [POL] [INFO ] [24662] [EET4575I]
Backup completed successfully.
```

```
2008-07-16 19:05:07.135 [POL] [ERROR] [24791] [EET4613E]
ACCESS DENIED on request to perform backup for hostname
vmlinux101. Server name is demo_on, server number 1. The
timestamp is 2008-07-16 19:05:07. The request protocol
version is 4.0.0.
2008-07-16 19:05:07.136 [IDS] [SEVERE] [24791]
[EET2328S] Backup data request failed: access denied or
a related cause.
```

The `/var/log/eet` directory can also contain additional log files in the form of `vmdb2logup_hostName_dbInstName_num`. For example, `vmdb2logup_vmlinux02_db2inst1_20838.3`. These log files are generated when the agent is unable to successfully upload log information to the server. See [“badlog.log” on page 14-242](#) for a description as to how these files are used.

Setting preferences

You can set the following preferences in the Management Console Web interface.

- General preferences that allow you to set the number of objects displayed based upon the object type. For example, you can set a preference that displays all configured policies on one Web page, rather than just 20 per page.
- Password preferences that specify how long a password must be, the types of characters that a password must contain, and password duration. Password preferences are also used to configure how the Security Server responds to repeated failed login attempts.
- Security Server preferences that allow you to set log maintenance parameters on the Security Server system. For example, you can set the interval to wait before moving agent log entries from temporary buffers on the Security Server to the Security Server log database, and consequently to the log viewer.
- DB2 Database Backup Agent preferences that allow you to set how the Security Server maintains DB2-specific log data. For example, you can set the maximum number of log entries the agent can upload at one time to the Security Server.
- IDS Database Backup Agent preferences that allow you to set how the Security Server maintains IDS-specific log data. For example, you can set the maximum number of log entries the agent can upload at one time to the Security Server.
- File System Agent preference that allow you to set how the Security Server maintains file system-specific log data. For example, you can set the interval at which the agent uploads log data to the Security Server.

Setting general preferences

To set general Security Server log preferences:

- 1: Log into the Management Console as a user with `Server` role permissions.
- 2: Select **Configuration->General Preferences** in the menu bar.
The *General Preference* window opens to the *General* tab.
- 3: Change the values displayed in the attribute text-entry boxes or scroll-list.
The following table lists and describes attributes and their values.

Category	Parameter	Description
Administrator Page Settings	View All Administrators	Checkbox to display all the configured administrators in the <i>Administrators</i> window on one page. Disabled by default.
	Number of Administrators Per Page	Set the maximum number of administrators in the <i>Administrators</i> window to display on one page. Navigation buttons are displayed in the <i>Administrators</i> window to move between the pages. The default is 20.
Host Page Settings	View All Hosts	Checkbox to display all the configured hosts in the <i>Hosts</i> window on one page. Disabled by default.
	Number of Hosts Per Page	Set the maximum number of hosts in the <i>Hosts</i> window to display on one page. Navigation buttons are displayed in the <i>Hosts</i> window to move between the pages. The default is 20.
Policy Page Settings	View All Policies	Checkbox to display all the configured policies in the <i>Policies</i> window on one page. Disabled by default.
	Number of Policies Per Page	Set the maximum number of policies in the <i>Policies</i> window to display on one page. Navigation buttons are displayed in the <i>Policies</i> window to move between the pages. The default is 20.
Key Page Settings	View All Keys	Checkbox to display all the configured keys in the <i>Keys</i> window on one page. Disabled by default.
	Number of Keys Per Page	Set the maximum number of keys in the <i>Keys</i> window to display on one page. Navigation buttons are displayed in the <i>Keys</i> window to move between the pages. The default is 20.
	View All Key Groups	Checkbox to display all the configured key groups in the <i>Key Groups</i> window on one page. Disabled by default.
	Number of Key Groups Per Page	Set the maximum number of key groups in the <i>Key Groups</i> window to display on one page. Navigation buttons are displayed in the <i>Key Groups</i> window to move between the pages. The default is 20.
Log Page Settings	Number of Log Messages Per Page	Set the maximum number of log entries to display on one page. Navigation buttons are displayed in the <i>Logs</i> window to move between the pages. The default is 20.
Management Console Timeout	Management Console Timeout	Set the interval of inactivity allowed before automatically logging users out of the Management Console Web session. Unsaved changes are discarded. Choices are 5 minutes, 20 minutes, 1 hour, 2 hours, and 8 hours. The default is 1 hour.

- 4: Click **Apply** to finalize the changes.

Setting password preferences

User passwords are a vital part of a good security system. A Management Console user password, at a minimum, can contain standard ASCII alphabet characters (i.e., a-z, A-Z) and integers (i.e., 0-9). Blank spaces are also permitted. The individual elements in this combination of characters, integers, and blanks cannot occur in adjacent sequence. That is, a password cannot contain two instances of the same element if they are next to each other. For example, `mississippi` will not be accepted, but `misSisSipPi` will.

Additional restraints can be applied that require all new passwords to contain at least one uppercase alphabet character, at least one special character, and the minimum number of characters that must be used.

To set Security Server password preferences:

- 1: Log into the Management Console as a user with `Server` role permissions.
- 2: Select **Configuration->General Preferences**.
The *General Preferences* window opens.
- 3: Select the *Password* tab.
- 4: Change the values displayed in the attribute text-entry boxes or scroll-list.

The following table lists and describes attributes and their values.

Category	Parameter	Description
Password Characteristics	Password Duration	Passwords expire after the number of days set by an administrative user. The password expiration interval is applied globally to each user account. If the user does not change the password prior to expiration, the user must reset the password immediately the next time s/he attempts to log in; otherwise the Management Console will not start. The expiration interval is an integer between 7 and 365. The default is 90. Password Duration must be set to a value greater than Password Expiration Notification. For example, if Password Duration is set to 90 days, then the Password Expiration Notification must be set to 89 days or less.
	Password History	The Security Server maintains a password history. You cannot use the same password more than once per the set limit. The default is 4. You may set this value to 0 to permit reuse of the current password.
	Minimum Password Length	Sets the minimum number of characters, including blank spaces, that must be in the password. The minimum password length is an integer between 8 and whatever your operating system can handle. The default is 8.
	Password Expiration Notification	Sets the number of days prior to password expiration at which to begin telling the user that their password is about to expire. Users are notified of the impending expiration at Management Console login. The notification interval is an integer between 6 and 31. The default is 31.
Password Complexity	Require Uppercase	When enabled, requires at least one uppercase alphabet character (i.e., A-Z) in the user password. This checkbox is enabled by default.
	Require Numbers	When enabled, requires at least one integer (i.e., 0-9) in the user password. This checkbox is enabled by default.
	Require Special Characters: !@#\$%^&*(){}[]	When enabled, requires at least one of the special characters listed in the left column in the user password. This checkbox is enabled by default.

Category	Parameter	Description
Account Lockout	Maximum Number of Login Tries	Set the maximum number of unsuccessful login attempts before disabling the Management Console Web interface for a set interval of time. The Management Console becomes inoperable and ignores further login attempts by any user for the specified interval. The default number of tries allowed is 3.
	User Lockout Time	The interval to wait before re-enabling the Management Console Web interface and allowing users to login. The default is 30 minutes.

5: Click **Apply** to finalize the changes.

Setting Security Server logging preferences

The *Server* tab displays aspects of the current Security Server configuration. It allows you to configure logs and log contents. Currently, you can configure five attributes.

To set Security Server log preferences:

- 1: Log into the Management Console as a user with *Server* role permissions.
- 2: Select **Configuration->Log Preferences**.
The *Log Preferences* window opens.
- 3: Select the *Server* tab.
- 4: Change the values displayed in the attribute text-entry boxes or scroll-list.

The following table lists and describes attributes and their values.

Attribute	Description
Logging Level	Sets the severity level at which entries are placed in the <code>cgss.log</code> file, and consequently displayed in the <i>Log</i> window. If syslog support is enabled, Logging Level also sets the log entries sent to the syslog server. The choices are <code>DEBUG</code> , <code>INFO</code> , <code>WARN</code> , <code>ERROR</code> , and <code>FATAL</code> . Log levels are cumulative, so each level includes the levels below it. For example, <code>FATAL</code> logs only fatal errors, whereas <code>WARN</code> logs warnings, and includes <code>ERROR</code> and <code>FATAL</code> conditions. Avoid using low levels for a sustained period because the logs can consume disk space quickly. The default is <code>INFO</code> .

Attribute	Description
Log Upload DB Retry	The interval to wait before resuming the transfer of agent log data that had been uploaded, and is stored in system files, into the log viewer database after a failure, such as after losing the connection to the database. The default is 30 seconds.
Log Buffer Size	The maximum number of entries to place in the Security Server log. When this limit is reached, or when Log Buffer Flush Time has elapsed, the entries are moved to the log viewer database. The default is 100 entries.
Log Buffer Flush Time	The interval to wait before moving log entries in the server log buffers to the log viewer database. The default is 15 seconds.
Update Host Frequency	The interval the Security Server waits to push policy and configuration updates to all File System Agent hosts. The default interval between data pushes is 30 seconds. Little data is pushed to each File System Agent host; however, you should increase the interval as you increase the number of hosts so as not to degrade network performance. If a File System Agent host is not available online, the Security Server will attempt to push updates to the non-responsive host at oscillating intervals. The intervals are short at first, and the delay in between increases with each attempt, up to one hour intervals. The Security Server starts the cycle over again after reaching one hour intervals. The Security Server continues this sequence indefinitely or until the host becomes accessible on the network and accepts the update.

Attribute	Description
Default Host Communication Port	The number of the port on the Security Server and on the File System Agent through which the two exchange File System Agent policy-related data. When you change the Default Host Communication Port number, the new port number is applied to all new hosts that are added after the configuration change is made. Existing File System Agent hosts are unaffected. The change is visible in the Communication Port text-entry box that is displayed in the <i>General</i> tab of each new host. Note that if you change the Communication Port number for an existing host, you must also restart the File System Agent process that runs on that host. Restarting the File System Agent refreshes the configuration data used by the File System Agent process, vmd. Execute one of the following commands on the File System Agent host to restart the File System Agent: On Linux: <pre>service secfs restart</pre> On AIX: <pre>/etc/rc.d/rc2.d/S99secfs restart</pre> The default File System Agent communication port number is 7024.

- 5: The **Restore to Default** button is used to return attribute values back to the factory default.
- 6: Click **Apply** to finalize the changes.

Setting DB2 Database Backup Agent logging preferences

You can configure the DB2 Database Backup Agent information that is entered in the Management Console log. You can configure the agent information globally, in which all the DB2 Database Backup Agents added after the configuration change inherit the agent log attributes, but all the existing DB2 Database Backup Agent log configurations remain intact. Or, you can configure log attributes for individual agents. This section describes global agent log configuration.

The default agent log file is `eedb2_usr.log`, and it is described in [“Database Backup Agent log files” on page 14-247](#).

To set the logging preferences for new DB2 Database Backup Agents:

- 1: Log into the Management Console as a user with `Server` role permissions.
- 2: Select **Configuration->Log Preferences**.
The *Log Preferences* window opens.
- 3: Select the *DB2 Log* tab.
- 4: Change values, enable features, and select message levels.

The following table lists and describes attributes and their values.

Attribute	Description
Message Types	
DB2 (Interaction with DB2 database)	Enters messages that are related to the agent and DB2 server interaction in the agent logs. Log to File and Upload to Server are enabled by default. The default log message level is <code>WARN</code> .
Policy (Evaluation of Policies)	Enters messages that are related to policy evaluation in the agent log. Log to File records policy events in log files on the agent. Upload to Server uploads the same information to the Security Server log viewer database. Log to File and Upload to Server are enabled by default. Set the log message level. Log message levels are described in “Log message levels” on page 14-237 . The default log message level is <code>INFO</code> .
Transformation (Encryption and Compression)	Enters messages that are related to data transformation in the agent logs. Log to File and Upload to Server are enabled by default. The default log message level is <code>WARN</code> .
File Logging Settings	
Maximum File Size (bytes)	The agent start a new, empty log file when the specified limit is exceeded. The default is 1000000 bytes.
Delete Old Log Files	Enable this checkbox to delete old agent logs. This checkbox works in conjunction with the Number of Old Log Files to Keep text-entry box. For example, enable this checkbox and enter 3 as the Number of Old Log Files to Keep value. After 3 logs are generated, the first log, <code>log1</code> , is deleted and a new log, <code>log4</code> , is created. If you do not enable this checkbox, log files will continue to accumulate in the server database and you will have to remove them manually.
Number of Old Log Files to Keep	Specifies the maximum number of agent log files to leave in the server database. This text-entry box is only displayed when the Delete Old Log Files checkbox is enabled. The default is 5.
Upload Logging Settings	

Attribute	Description
Maximum Number of Messages to Upload At Once	A network gating feature to limit the number of messages sent to the Security Server at one time. When the specified number of log entries is reached, those entries are uploaded to the Security Server. The default is 1000.
Upload Messages At Least Every (seconds)	The maximum interval to wait before the agent is to upload messages to the Security Server. Use this attribute to update the log viewer even when the Maximum Number of Messages to Upload At Once has not been reached. You can lower the interval if there is little agent activity. The default is 10 seconds.
Upload Messages At Most Every (seconds)	The minimum interval to wait before the agent is to upload messages to the Security Server. You can increase the interval if there is considerable agent activity, so the agents do not flood the network with log messages. The default is 1.
Normal Time Out (seconds)	The maximum interval of time the agent is to wait for the Security Server to acknowledge a backup or restore request and upload related message data. If the agent cannot connect to the Security Server within the specified interval, the agent will try again after the interval configured by the Upload Messages At Least Every attribute. The default is 2 seconds.
Shutdown Time Out (seconds)	The maximum interval of time the agent is to wait for the Security Server to acknowledge job completion and upload related message data. If the agent is unable to upload the log messages within the specified interval, they are left on the agent system. The agent will resend the messages at the beginning of the next job. The default is 30 seconds.
Duplicate Message Suppression Settings	
Threshold	See “Limiting the number of duplicate log entries” on page 15-265 for a description of the Duplicates column and the Threshold and Interval text-entry boxes.
Interval	See “Limiting the number of duplicate log entries” on page 15-265 for a description of the Duplicates column and the Threshold and Interval text-entry boxes.

- 5: Click **Apply** to finalize the changes.

Setting IDS Database Backup Agent logging preferences

You can configure the IDS Database Backup Agent information that is entered in the Management Console log. You can configure the agent information globally, in which all the IDS agents added after the configuration change inherit the agent log attributes, but all the existing IDS agent log configurations remain intact. Or, you can configure log attributes for individual agents. This section describes global agent log configuration. For information about individual agent log configuration, see [“Changing host configuration” on page 3-61](#).

The default agent log file is `eetids_usr.log` and it is described in [“Database Backup Agent log files” on page 14-247](#).

To set the logging preferences for new IDS Database Backup Agents:

- 1: Log into the Management Console as a user with `Server` role permissions.
- 2: Select **Configuration->Log Preferences**.
The *Log Preferences* window opens.
- 3: Select the *IDS Log* tab.
- 4: Change values, enable features, and select message levels.

The following table lists and describes attributes and their values.

Attribute	Description
Message Types	
Informix (Interaction with IDS database)	Enters messages that are related to the agent and IDS server interaction in the agent logs. Log to File and Upload to Server are enabled by default. The default log message level is <code>WARN</code> .
Policy (Evaluation of Policies)	Enters messages that are related to policy evaluation in the agent log. Log to File records policy events in log files on the agent. Upload to Server uploads the same information to the Security Server log viewer database. Log to File and Upload to Server are enabled by default. Set the log message level. Log message levels are described in “Log message levels” on page 14-237 . The default log message level is <code>INFO</code> .
Transformation (Encryption and Compression)	Enters messages that are related to data transformation in the agent logs. Log to File and Upload to Server are enabled by default. The default log message level is <code>WARN</code> .
File Logging Settings	

Attribute	Description
Maximum File Size (bytes)	The agent start a new, empty log file when the specified limit is exceeded. The default is 1000000 bytes.
Delete Old Log Files	Enable this checkbox to delete old agent logs. This checkbox works in conjunction with the Number of Old Log Files to Keep text-entry box. For example, enable this checkbox and enter 3 as the Number of Old Log Files to Keep value. After 3 logs are generated, the first log, log1, is deleted and a new log, log4, is created. If you do not enable this checkbox, log files will continue to accumulate in the server database and you will have to remove them manually.
Number of Old Log Files to Keep	Specifies the maximum number of agent log files to leave in the server database. This text-entry box is only displayed when the Delete Old Log Files checkbox is enabled. The default is 5.
Upload Logging Settings	
Maximum Number of Messages to Upload At Once	A network gating feature to limit the number of messages sent to the Security Server at one time. When the specified number of log entries is reached, those entries are uploaded to the Security Server. The default is 1000.
Upload Messages At Least Every (seconds)	The maximum interval to wait before the agent is to upload messages to the Security Server. Use this attribute to update the log viewer even when the Maximum Number of Messages to Upload At Once has not been reached. You can lower the interval if there is little agent activity. The default is 10 seconds.
Upload Messages At Most Every (seconds)	The minimum interval to wait before the agent is to upload messages to the Security Server. You can increase the interval if there is considerable agent activity, so the agents do not flood the network with log messages. The default is 1.

Attribute	Description
Normal Time Out (seconds)	The maximum interval of time the agent is to wait for the Security Server to acknowledge a backup or restore request and upload related message data. If the agent cannot connect to the Security Server within the specified interval, the agent will try again after the interval configured by the Upload Messages At Least Every attribute. The default is 2 seconds.
Shutdown Time Out (seconds)	The maximum interval of time the agent is to wait for the Security Server to acknowledge job completion and upload related message data. If the agent is unable to upload the log messages within the specified interval, they are left on the agent system. The agent will resend the messages at the beginning of the next job. The default is 30 seconds.
Duplicate Message Suppression Settings	
Threshold	See “Limiting the number of duplicate log entries” on page 15-265 for a description of the Duplicates column and the Threshold and Interval text-entry boxes.
Interval	See “Limiting the number of duplicate log entries” on page 15-265 for a description of the Duplicates column and the Threshold and Interval text-entry boxes.

5: Click **Apply** to finalize the changes.

Setting File System Agent logging preferences

You can configure the File System Agent process information that is entered in the Management Console log. You can configure the process information globally, in which all the File System Agents that are added after the configuration change inherit the log attributes, while all current file system configurations remain intact. Or, you can configure log attributes for individual File System Agent installations. This section describes global File System Agent log configuration.

The default File System Agent log file is `eetvmd_usr.log` and it is described in [“Encryption Expert File System Agent log files” on page 14-239](#).

To set the logging preferences for new File System Agent installations:

- 1: Log into the Management Console as a user with `Server` role permissions.
- 2: Select **Configuration->Log Preferences**.
The *Log Preferences* window opens.
- 3: Select the *FS Log* tab.

- 4: Change values, enable features, and select message levels.

The following table lists and describes attributes and their values.

Attribute	Description
Message Types	
Management Service	This option logs: * Communication between File System Agent and Security Server * Dataxform session activity Enable the Log to File checkbox to trap File System Agent messages on the host that pertain to the network activity between the File System Agent and Security Server. Display these messages by executing <code>dmesg</code> or by executing <code>tail</code> on the <code>/var/log/messages</code> file. These messages can also be sent to syslog. These are samples of the types of messages you can expect to see when the Log to File checkbox is enabled: <pre>2008-03-26 13:58:26.024 [VMD] [INFO] [29390] [EET3794I] Query security config succeeded. 2008-03-26 13:59:08.014 [VMD] [INFO] [29390] [EET3781I] Successfully received and implemented a new security configuration. (Note that Management Service messages contain the string "[VMD]").) Log to File and Upload to Server are enabled by default. The default log message level is INFO.</pre>

Attribute	Description
Policy Evaluation	This option logs: * Audit messages from policy evaluation * Deny - set to level "ERROR" * Allow - set to level "INFO" * Incidental system errors during policy evaluation Enable the Log to File checkbox to trap File System Agent messages on the host that pertain to policy evaluation and which can be displayed by executing <code>dmesg</code> or by executing <code>tail</code> on the <code>/var/log/messages</code> file. These messages can also be sent to syslog. These are samples of the types of messages you can expect to see when the Log to File checkbox is enabled: <pre>Mar 21 20:18:15 vmlinux05 kernel: [CGP] [ERROR] [EET2602E] ALARM Policy[denyAllO User[root,uid=0 (User Not Authenticated)] Process[/home/apps/apps2/bin/ xclock] Action[read_file] Res[/home/apps/ apps2/bin/xclock] Key[None] Effect[DENIED Code (1M)] Mar 21 20:30:50 vmlinux05 kernel: [CGP] [INFO] [EET2601I] AUDIT Policy[allowAllOps_fs] User[root,uid=0 (User Not Authenticated)] Process[/home/apps/apps1/ bin/bc] Action[read_file] Res[/home/apps/apps1/ bin/bc] Key[aes128_cached] Effect[PERMIT Code (1M)]</pre> Set the log message level. Log message levels are described in “Log message levels” on page 14-237. The default log message level is <code>ERROR</code>.
System Administration	This option logs extreme errors having to do with the File System Agent system, such as insufficient system memory. Other examples as to what triggers these messages are unavailable at this time. Messages of this type are extremely rare. (Note that Policy Evaluation messages contain the string “[CGS]”.) The default log message level is <code>ERROR</code>.

Attribute	Description
Security Administration	This option logs: * Guard/unguard (not mounting) messages * Policy version updates to a guard point * SecFS up/down * SecFS messages related to VMD issues * Most catastrophic and debug messages Enable the Log to File checkbox to trap File System Agent messages on the host that pertain to security administration and which can be displayed by executing <code>dmesg</code> or by executing <code>tail</code> on the <code>/var/log/messages</code> file. These messages can also be sent to syslog. These are samples of the types of messages you can expect to see when the Log to File checkbox is enabled: <pre>Mar 26 12:32:35 vmlinux05 kernel: [CGA] [INFO] [EET3001I] EVENT: Policy denyAllOps_fs took effect or was modified Mar 26 12:32:35 vmlinux05 kernel: [CGA] [INFO] [EET3001I] EVENT: Successfully guarded /home/apps/apps3/etc (Note that Security Administration messages contain the string "[CGA]").) The default log message level is INFO.</pre>
File Logging Settings	
Maximum File Size (bytes)	The agent start a new, empty log file when the specified limit is exceeded. The default is 1000000 bytes.
Delete Old Log Files	Enable this checkbox to delete old FS agent logs. This checkbox works in conjunction with the Number of Old Log Files to Keep text-entry box. For example, enable this checkbox and enter 3 as the Number of Old Log Files to Keep value. After 3 logs are generated, the first log, <code>log1</code>, is deleted and a new log, <code>log4</code>, is created. If you do not enable this checkbox, log files will continue to accumulate in the server database and you will have to remove them manually.
Number of Old Log Files to Keep	Specifies the maximum number of agent log files to leave in the server database. This text-entry box is only displayed when the Delete Old Log Files checkbox is enabled. The default is 5.
Upload Logging Settings	

Attribute	Description
Maximum Number of Messages to Upload At Once	A network gating feature to limit the number of messages sent to the Security Server at one time. When the specified number of log entries is reached, those entries are uploaded to the Security Server. The default is 1000.
Upload Messages At Least Every (seconds)	The maximum interval to wait before the agent is to upload messages to the Security Server. Use this attribute to update the log viewer even when the Maximum Number of Messages to Upload At Once has not been reached. You can lower the interval if there is little agent activity. The default is 10 seconds.
Upload Messages At Most Every (seconds)	The minimum interval to wait before the agent is to upload messages to the Security Server. You can increase the interval if there is considerable agent activity, so the agents do not flood the network with log messages. The default is 1.
Normal Time Out (seconds)	The maximum interval of time the agent is to wait for the Security Server to acknowledge a backup or restore request and upload related message data. If the agent cannot connect to the Security Server within the specified interval, the agent will try again after the interval configured by the Upload Messages At Least Every attribute. The default is 2 seconds.
Shutdown Time Out (seconds)	The maximum interval of time the agent is to wait for the Security Server to acknowledge job completion and upload related message data. If the agent is unable to upload the log messages within the specified interval, they are left on the agent system. The agent will resend the messages at the beginning of the next job. The default is 30 seconds.
Duplicate Message Suppression Settings	
Threshold	See “Limiting the number of duplicate log entries” on page 15-265 for a description of the Duplicates column and the Threshold and Interval text-entry boxes.
Interval	See “Limiting the number of duplicate log entries” on page 15-265 for a description of the Duplicates column and the Threshold and Interval text-entry boxes.

5: Click **Apply** to finalize the changes.

File System Agent logs can accumulate quickly and consume a lot of disk space. In addition to setting File System Agent preferences, occasionally monitor the available space of the file system for the `/var/log/eet` directory. When the free disk space of the file system reaches a minimum threshold, the File System Agent begins to issue warning messages. These messages are displayed in the Management Console *Logs* window.

If the disk fills, and the agent cannot create a log entry, the agent will stop operation. This applies to both File System Agents and Database Backup Agents. In the case of Database Backup Agents, if the agent cannot write to the log, the agent ignores the backup or restore request, and, after some time-out interval, the DB2 or IDS server displays an error message. You must create some free space for the agents to resume operation. In the case of File System Agents, you must also restart the File System Agent processes using the `secfs restart` command.

The File System Agent begins to send warning messages when the amount of free disk space drops to 512 MB, and additional messages are sent with each successive drop to 256 MB, 128 MB, 64 MB, and then to 32 MB. Once the available space drops to 32 MB, one message is sent every ten minutes.

Messages like the following are sent to the Security Server when space drops to 512 MB and less:

```
EET4620I=The space available for log files in /var/log/eet is low
(62 MB).
```

The message changes when space drops below the 32 MB threshold. For example:

```
EET4621W=The space available for log files in /var/log/eet is
dangerously low (27 MB).
```

Gating is performed to minimize the number of messages that are sent when free space oscillates above and below a threshold. A message is not sent when the available space increases above a threshold and then decreases below that threshold, unless the amount of increase is more than half of the higher threshold. For example a warning is issued when free space drops to 240 MB. This is triggered by a drop below the 256 MB threshold. The next threshold is now 128 MB. No additional messages are sent as long as the free space remains between 384 MB (i.e., $256 + 256/2$) and 128 MB.

Limiting the number of duplicate log entries

The *DB2 Log*, *IDS Log*, and *FS Log* tabs allow you to configure the number of duplicate messages of a specific type that agents send to the Security Server. This feature allows you to minimize the number and frequency of redundant log messages. An agent can be configured to:

- send every message
- send just the first message in a series of sequential and duplicate messages
- send up to a specified number of duplicate messages when they occur within a configured time period

Use the scroll-lists in the **Duplicates** column to set controls that limit the number of messages sent to the Security Server.

Level	Description
Allow	All duplicate messages of the corresponding Message Type are captured and displayed in the log. The default is Allow .
Drop	Messages of the corresponding Message Type that are repeated in sequence will be displayed once, followed by the number of times the same message occurred. This setting ignores the Threshold and Interval settings.
Suppress	Messages of the corresponding Message Type will follow the configured Threshold as to how many times duplicate messages are sent to the Security Server during the given Interval .

The main difference between **Drop** and **Suppress** is that **Drop** requires duplicate messages to occur in sequence, one after the other. **Suppress** does not take into account whether duplicate messages are in sequence or not. With **Suppress**, the agent counts the number of duplicate messages within an allocated period. The time count begins with the first message generated. Only the specified number of duplicate messages are sent to the Security Server during the specified interval.

The **Duplicate Message Suppression Settings** options, **Threshold** and **Interval**, are integers that set the maximum number of times that the same message can appear in the log within a specified amount of time. By default, no messages are dropped.

Set the **Duplicates** scroll-lists to **Suppress**. The default **Threshold** is 5 occurrences and the default **Interval** is 600 seconds. With these settings, the first message starts the interval count, after sending a total of 5 duplicate messages to the Security Server, all duplicate messages that occur within the 10 minute interval are discarded on the agent system. Once the time from the first message to the Interval is exceeded, the agent begins the cycle again, sending the first 5 duplicate messages and discarding the rest during the next interval.

Setting syslog logging preferences

The messages sent to the Management Console *Logs* window can also be sent to a remote syslog server or to the local `/var/log/messages` file.

To set the preferences for syslog messages and enable logging:

- 1: Verify that the syslog server is operational and accessible from the Security Server system.

It is usually sufficient to ping the syslog server and run `ps` to check the `syslog` process on the syslog server system.

If you are going to send the messages to `localhost`, check that the `syslog` process on the local system is accepting messages. You may need to restart `syslogd` with the “-r” argument.
- 2: Log into the Management Console as a user with `Server` role permissions.

- 3: Select **Configuration->Log Preferences**.

The *Log Preferences* window opens to the *Server* tab.

- 4: Set the **Logging Level** property.

The level you select affects the number of messages that are displayed in the Management Console *Logs* window and the number of messages sent to the syslog server.

Redundant syslog failure messages are filtered so that only one out of every fifty is sent to `/var/log/messages` and the Management Console *Logs* window. All the redundant syslog failure messages are sent when the level is set to `DEBUG`.

- 5: Click **Apply**.

- 6: Select the *Syslog* tab.

- 7: Click **Add**.

The *Add Syslog Server* window opens.

- 8: Enter the host name or FQDN of the syslog server in the **Server Name** text-entry box.

- 9: Click **Ok**.

The *Log Preferences* window opens.

- 10: Enable the **Syslog Enabled** checkbox.

- 11: Click **Apply**.

- 12: Perform some task on an agent system that normally generates a syslog entry, such as backing up an IDS database or accessing a guard point.

- 13: Check the `/var/log/messages` file on the syslog server for Security Server log entries.

Maintaining Encryption Expert

Verifying agent to server SSL authentication

There is a simple and quick way to ensure mutual SSL communication and authentication between an agent and its server without actually having to run a DB2 backup or restore job. The `commtest` utility mimics a DB2 backup request and attempts to connect to the server.

The `commtest` utility is located in the DB2 agent `./bin` directory, and it takes one argument: the path to the Encryption Expert library file. For example, on UNIX, assuming you are in the `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2` directory, the utility syntax is:

```
# ./bin/commtest ./lib/libeetdb2.so
```

On Windows, assuming you are in the `\Program Files\IBM\DB2TOOLS\LUWEncryptionExpert\agent\db2\bin` folder, the utility syntax is:

```
$ .\commtest.exe .\eetdb2.dll
```

The utility should return a “Success” message. For example:

```
Attempting to open shared object at .\db2\bin\eetdb2.dll
Success!
$
```

The `commtest` utility submits a backup request to the Security Server as the user `CT_inst` and the database alias `CT_dbal`. If **Instance** and **Alias** are not defined in the policy, the utility should return a “Success” message. The test will fail if the **Instance** and **Alias** names are configured in the policy for the host. The test will also fail if the agent does not have the correct server certificate or there is no network connection to the server. If the test fails, check the Security Server log first then the local agent log.

If the test appears in the Security Server log, it indicates that the Database Backup Agent can communicate across SSL to the server, and it is just a policy mismatch that prevented successful completion. For example, the following log entry indicates that the server and agent can communicate with each other; however, since the `commtest` user instance (`CT_inst`) is not recognized as a valid user instance in the policy, policy matching will fail.

```
EET0007I: The Security Server denied a backup request
from host vmlinux02 for dbinstance(CT_inst),
dbalias(CT_dbal), dbpartition(0). The policy used is
(allowAdminsOnly), the rule used is (backupRule05).
```

If the test fails, and there is no log entry in the Security Server, it indicates that the Database Backup Agent cannot connect to, or be authenticated by, the Security Server. Check the local log in `/var/log/eet`. For example, the following message indicates that the agent certificate cannot be read.

```
[EET4515S] There was a communication failure with the
Security Server. The server returned a fault: Error 26;
"SOAP-ENV:Server"; Subcode: "[no subcode]"; Reason: "SSL
error"; Detail: "Can't read certificate key file".
```

Check the `./agent/db2/pem` directory for a missing or expired `agent-cert.pem` file. Clean out and regenerate agent certificates if necessary.

The following messages indicate that the agent cannot establish a TCP connection to the server:

```
2007-09-13 16:06:38.455 [COM] [SEVERE] [26335]
[EET4515S] There was a communication failure with the
Security Server. The server returned a fault: SOAP 1.1;
"SOAP-ENV:Client"; Subcode: "[no subcode]"; Reason: "No
route to host"; Detail: "connect failed in
tcp_connect()".
2007-09-13 16:06:38.458 [COM] [SEVERE] [26335]
[EET4527S] The Agent could not obtain configuration
information. The Agent failed to initialize. Verify
connectivity to the Security Server and retry the
operation.
```

Check the network connection to the Security Server system or check the Security Server network configuration. Check that you can ping the server, or log onto the server, from the agent system.

Determining the Encryption Expert Database Backup Agent version

The `buildinfo` command displays the Database Backup Agent version information. The command takes one argument, the agent library.

- On DB2 Database Backup Agents (UNIX), the agent library is `./lib/libeetdb2.so`.
- On DB2 Database Backup Agents (Windows), the agent library is `.\bin\eetdb2.dll`.
- On IDS agents, the library is built into the IDS executable `./bin/eetidsagt`.
- `buildinfo` provides information about the Database Backup Agent. For information about the File System Agent, use a command like “`secfsd -version`” or “`vmd -v`” instead.
- For Windows hosts, use the Encryption Expert tray instead.

To display the Database Backup Agent version:

- 1: Change to the Database Backup Agent directory.
`cd /opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent`
- 2: (DB2 only) Enter the `buildinfo` command, followed by the path to the agent library. (On Windows, the library is `.\db2\bin\eetdb2.dll`)
`./db2/bin/buildinfo ./db2/lib/libeetdb2.so`

MODULE	KEY	VALUE
DB2AGENT	BuildId	BUILD_U0064

```

DB2AGENT      Version      1.1.3.0
DB2AGENT      BuildDate    2008-03-27 15:34:08 -0700 (PDT)
SDK           BuildId      BUILD_U0064
SDK           Version      1.1.3.0
SDK           BuildDate    2008-03-27 15:33:58 -0700 (PDT)
#

```

The SDK BuildId row displays the build number. In this example, the build number is 0054.

- 3: (IDS only) Enter the buildinfo command, followed by the path to the Encryption Expert executable for IDS.

```

# ./ids/bin/buildinfo ./ids/bin/eetidsagt
MODULE      KEY      VALUE
IDSAGENT    BuildId    BUILD_U0062
IDSAGENT    Version    1.1.3.0
IDSAGENT    BuildDate  2008-03-14 12:40:56 -0700 (PDT)
SDK         BuildId    BUILD_U0062
SDK         Version    1.1.3.0
SDK         BuildDate  2008-03-14 12:40:42 -0700 (PDT)
#

```

Determining the Encryption Expert File System Agent version

The vmd command displays the File System Agent version information. This command is located in /opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/bin/vmd. A symbolic link to this file from /usr/bin/vmd is created when the agent software is installed.

The vmd command without any arguments starts the vmd process. This process is required for File System Agent protection. Use the ps command to verify that the vmd process is running.

To display the File System Agent version:

- On UNIX agents, execute the vmd command with the -v argument:

```

# which vmd
/usr/bin/vmd
# vmd -v
1.1.2.0-Linux-2.6-x86-64-Build0048
2008-03-13 11:25:03 -0700 (PDT)
#

```

- For Windows hosts, use the Encryption Expert tray instead.
- On Windows agents, execute the vmd command in a command window with the -v argument:

```

$ which vmd
\Program
Files\IBM\DB2TOOLS\LUWEncryptionExpert\agent\vmd\bin\vmd
$ vmd -v
1.1.3.0
0072
2008-05-29 16:38:00 -0700 (PDT)

```

Helpful utilities

Server utilities:

These utilities are located in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/bin`.

backup_config

The `backup_config` utility creates a backup of the Security Server database. See [“Backing up the Security Server configuration” on page 16-295](#) for usage information.

cleanup_replication

The `cleanup_replication` utility cleans up the Security Server replication configuration. This utility should be run on the primary server only. This utility performs database operations, hence only the `voradmin` user can run this utility. This utility does not take any arguments.

To run `cleanup_replication`:

- 1: Log on to the primary Security Server as `voradmin`.
- 2: Execute the certificates cleanup utility without any arguments.

For example:

```
# ./bin/cleanup_replication
```

The utility displays a brief message and prompts you to proceed.

```
We will now cleanup the Security Server replication on this
machine ...
```

```
Do you wish to proceed? (y|n)[n]:
```

- 3: Enter “y” to continue.

The utility lists all the configured failover servers and prompts you to choose one of them.

```
Stopping replication...done.
```

```
The following security servers are currently configured for
replication:
```

```
[1]: vmlinux02
```

```
[2]: vmlinux03
```

```
Choose the failover server for replication cleanup [1]:
```

- 4: Enter the number of the failover server to be cleaned.

The utility processes the failover server and completes.

```
Cleaning up replication configuration...DB20000I
```

```
The SQL command completed successfully.
```

```
DB20000I The SQL command completed successfully.
```

```
DB20000I The SQL command completed successfully.
```

```
DB20000I The SQL command completed successfully.
```

```
Cleaning up replication configuration in data store....done
```

```
#
```

configure_replication

The `configure_replication` utility configures the Security Server replication processes. Run this utility on the primary server only. This utility performs database operations, so only the `voradmin` user can run it.

This utility prompts for:

- The password for the voradmin user on the primary server
- The hostname or FQDN of the failover server
- The password for the voradmin user on the failover server

An example of how to use the `configure_replication` utility is provided in [“Configuring SQL replication” on page 13-222](#).

You must restart the Security Server for the changes to take affect.

deinstall_eet_server

The `deinstall_eet_server` utility removes all Security Server software from the system. See [“Removing the Encryption Expert Security Server” on page 3-54](#) for usage information.

import_v1120_config

The `import_v1120_config` utility imports a version 1.1.2 server configuration into a version 1.1.3 server configuration. The system must be running version 1.1.3 Security Server to use this utility.



Import the v1.1.2 Security Server configuration immediately after installing the v1.1.3 Security Server because this procedure completely overwrites the v1.1.3 Security Server configuration, and changes made to the v1.1.3 Security Server are discarded.

To import a v1.1.2 Security Server configuration into a v1.1.3 Security Server:

- 1: On the v1.1.2 Security Server system, backup the Security Server configuration.

```
# ./bin/backup_config
```

Since an optional directory path was not added to the end of the example command, the default staging area, `/tmp`, will be used.

Details for using `backup_config` are provided in [“Backing up the Security Server configuration” on page 16-295](#)

- 2: Copy the backup file onto the v1.1.3 Security Server system.
- 3: Restore the backup file on the v1.1.3 Security Server system.

For example,

```
# ./bin/import_v1120_config -f /tmp/
Backup_2008_03_17_0959.tar
```

The utility prompts you to stop Security Server processes so that it can restore the Security Server configuration in a static environment.

The Security Server needs to shutdown before the restore can proceed. All the current data will be over-written with the backup data.

Are you OK to proceed? (yes|no)[no]:

- a: Enter `yes` to temporarily shutdown the Security Server and proceed with restoring the Security Server database.

The utility prompts you to add the users that are in the backup to the current Security Server configuration.

Restore users from backup? (yes|no):

- b: Enter `yes` to add the users in the backup to the current Security Server configuration or `no` to discard the users in the backup.
-

Note that, if you add the users from the backup, and there are similarly named users already in the current Security Server configuration, any changes made to the user configurations since the backup are discarded and the old passwords and roles are reapplied.

The utility prompts you to enter the same password that was used to make the backup.

Please provide the password the configuration backup was protected with.

enter aes-256-cbc decryption password:

c: Enter the password that was used to make the backup.

The utility restores the backup and restarts the Security Server.

- 4: Open the Management Console on the v1.1.3 Security Server.
- 5: Disable all the **Registration Allowed** checkboxes for each agent configured on each host.

Click **Apply** or **Ok** to apply the configuration changes.

This step removes the agent certificates from the Security Server.

- 6: Enable the **Registration Allowed** checkboxes for each agent to be configured on each host.

Click **Apply** or **Ok** to apply the configuration changes.

This step allows the exchange of certificates between the Security Server and the respective agents on a host.

- 7: Remove the agent certificates of each agent from each host.

The certificate removal utilities are:

```
(FS) # ./vmd/bin/register_host clean
(DB2) # ./db2/bin/setup_eet_agent_cert_db2 clean
(IDS) # ./ids/bin/setup_eet_agent_cert_ids clean
```

- 8: Regenerate agent certificates and reregister the host agents with the v1.1.3 Security Server.

The certificate generation and reregistration utilities are:

```
(FS) # ./vmd/bin/register_host
(DB2) # ./db2/bin/setup_eet_agent_cert_db2
(IDS) # ./ids/bin/setup_eet_agent_cert_ids
```

logdb_controller

The `logdb_controller` utility controls the growth of tables in the LOGDB database. There are five tables in the LOGDB database. They are `cgss_server_log`, `cgss_agent_log_fs`, `cgss_agent_log_db2`, `cgss_agent_log_ids`, and `cgss_log_view`. `logdb_controller` is automatically setup as a cron job by the installer and it wakes up every 15 minutes to check the size of these five tables.

Since LOGDB is setup automatically as a cronjob, the user does not need to interact directly with the utility. However, if the end user needs to change certain parameters, the cronjob can be changed.

This utility performs database operation, hence only `voradmin` user can run this utility or cronjob.

This utility takes three parameters in the following order to operate:

- Number of logs to be retained in database: 100000 or other number
- Logs to be exported or not: <export|noexport>
- Path where exported log files will be saved as a tar.gz file

The defaults are:

- Number of logs to be retained is 100000
- Logs to be exported
- Path to export, server's etc directory /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/server/etc

The utility logdb_controller controls, or limits, the indefinite growth in the number of records in the LOGDB database. cron job is run every 15 minutes. The utility goes and checks all four tables. There is an internal threshold number which is set as one fourth of the input parameters.

Default Input parameter for Number of log records to be retained is 100000.

Threshold level is set at 125,000.

The utility first goes to one table, for example cgss_agent_log_fs, and checks the minimum LOG_ID and maximum LOG_ID. If the difference of min. and max. ID is less than the threshold number, it skips the operation. If the difference is greater than threshold, it first exports the logs to a default (or changeable) directory and then prunes the records. If it fails to export for some reason, it does not prune. Exported files are in db2 ixf and expmsgs format. These files are tarred and then gzipped to conserve disk space. These are saved with a date and timestamp for ease of identity and to maintain uniqueness of filename.

A different path can be specified to store the exported logs tar.gz file. The utility checks if this path is writable by the voradmin user, if it is not, it aborts the operation.

Configuring logdb_controller with cronjob

To see if the cronjob is set up with every cgss start

```
# su voradmin
$ crontab -l
```

This will the following list after cronjob is running

```
0,15,30,45 * * * * . .bashrc; /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/server/bin/internal/
logdb_controller > /tmp/logdb_controller.log 2>&1
```

The cronjob will automatically terminate when cgss is stopped.

Example to change cronjob

```
# su voradmin
$ crontab -e
```

The cronjob will display as shown below in the vi editor or whatever text editor is configured in the user profile

```
0,15,30,45 * * * * . .bashrc; /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/server/bin/internal/
logdb_controller > /tmp/logdb_controller.log 2>&1
```

To change the default number of logs to be retained from 100000 to 200000, to disable export logs, and to change the path to a different path for logs to be exported, use the following example. Please note that the order of arguments should be maintained properly.

```
0,15,30,45 * * * * . .bashrc; /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/server/bin/internal/
```

```
logdb_controller 200000 noexport /home/voradmin/
exportedlogs > /tmp/logdb_controller.log 2>&1
```



Disabling the exporting logs before deleting the logs should not be done at a production site. This feature is only for internal QA.

restore_config

The `restore_config` utility restores a Security Server database backup that was made with `backup_config`. To do a restore, check that there is nothing you want to keep before running restore on a server system because everything will be replaced by the data in the previous backup. When running this utility, the machine must have a Security Server installed.

- 1: Copy the backup file to a location on the machine where the server is installed, for example `/testDir/myBackup`.
- 2: Run: `restore_config -f tar_file stage_dir`
tar_file is the full path to the backup file. For example, `/testDir/myBackup/Backup_2008_04_22_1523.tar`
stage_dir is an optional temporary staging area, for example, `/tmp`. If you do not specify a temporary staging area, `/tmp` will be used.

An example of this command is shown below:

```
./bin/restore_config -f /tmp/Backup_2008_04_22_1523.tar
```

The utility prompts you to stop Security Server processes so that it can restore the Security Server configuration in a static environment.

The Security Server needs to shutdown before the restore can proceed. All the current data will be over-written with the backup data.

Are you OK to proceed? (yes|no)[no]:

- a: Enter `yes` to temporarily shutdown the Security Server and proceed with restoring the Security Server database.

The utility prompts you to add the users that are in the backup to the current Security Server configuration.

Restore users from backup? (yes|no):

- b: Enter `yes` to add the users in the backup to the current Security Server configuration or `no` to discard the users in the backup.

Note that, if you add the users from the backup, and there are similarly named users already in the current Security Server configuration, any changes made to the user configurations since the backup are discarded and the old passwords and roles are reapplied.

The utility prompts you to enter the same password that was used to make the backup.

Please provide the password the configuration backup was protected with.

enter aes-256-cbc decryption password:

- c: Enter the password that was used to make the backup.

The utility restores the backup and restarts the Security Server.

- 3: If you restore the backup on the same system on which you made the backup, and nothing has changed, you are done.
-

This example assumes this server is not the same as that you did the backup for.

- 4: Run `re_gen_cert` to regenerate the server certificate. This will stop the server. Run `service cgss start` to restart the server.
- 5: For each Database Backup Agent, do the following: For the corresponding host in the server, uncheck "register allowed" for each host to clean up the certificate for that host (through the Management Console) and recheck it so as the agent can register itself to the server. For each Database Backup Agent, on the Database Backup Agent side, run `setup_eet_agent_cert_db2 clean` to clean up the current certificate. And run `setup_eet_agent_cert_db2` to recreate the certificate.
- 6: For each agent, modify the URL in the `agent.conf` to point to the new server hostname.

The `agent.conf` file is located at:

(UNIX) `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/etc/agent.conf`

(UNIX) `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2/etc/agent.conf`

(Windows) `\Documents and Settings\All Users\Application Data\IBM\DBTOOLS\LUWEncryptionExpert\Agent\vmd\etc\agent.conf`

See step 4 in [“Configuring an agent to use Security Servers” on page 6-100](#) for details.

- 7: (UNIX) Restart the `vmd` process to refresh the File System Agent configuration.

a: Kill the `vmd` process if it is still running:

```
# ps -aef | grep vmd
# kill -9 procID
```

b: Start the `vmd` process.

```
# ./agent/vmd/bin/vmd
```

If you are unable to kill the process, disable the **FS Agent Locked** checkbox for the host in the Management Console.

- 8: (Windows) Restart the `vmd` process to refresh the File System Agent configuration.

a: Open a CMD window.

b: Enter “`net stop vmd`”.

```
The vmd service is stopping.
The vmd service was stopped successfully.
```

c: Enter “`net start vmd`”.

```
The vmd service is starting.
The vmd service was started successfully.
```

re_gen_cert_auth

The `re_gen_cert_auth` utility regenerates the Certificate Authority which locates on the primary Security Server.

The user can run this utility in one of the following situations:

- when the signer key is compromised
- when the signer certificate expires
- one of the fields of the signer certificate has changed

Only root can run this utility.

The utility does the following, in the following order:

- generates a new signer certificate
- deletes the old signer certificate from the keystore
- imports the new signer certificate into the keystore with alias `cgss_server_ca`
- generates a new certificate request from the existing server certificate (with alias `cgss_server_app`)
- signs the server certificate with new CA
- imports the new server certificate into the keystore under alias `cgss_server_app`

The following operations are needed after running the `re_gen_cert_auth` utility.

- Run `re_sign_failover_cert` from the failover server.
- For every agent registered with the Security Server, perform the following steps:
 - d: Disable agent registration from the Management Console to clean up agent certificate on the `$SERVER_SHORT`.
 - e: Enable the agent registration from the Management Console.
 - f: (UNIX only) Run `register_host clean` on the agent installation to clean up old File System Agent certificate files.
 - g: Run `register_host` on the agent installation.
 - h: Run `setup_eet_agent_cert_db2 clean` on the agent installation to clean up old Database Backup Agent certificate files.
 - i: Run `setup_eet_agent_cert_db2` on the agent installation.

`re_gen_failover_cert`

This utility regenerates the server certificate for the failover Security Server. Use `re_gen_failover_cert` only if failover server certificates are already configured. If the failover server is not configured with certificates, as can occur when failover server software installation fails during the certificate generation and exchange stage, run `setup_failover_cert` instead. `re_gen_failover_cert` will fail if there are no certificates and it will direct you to run `setup_failover_cert` instead.

The user can run `re_gen_failover_cert` this utility in one of the following situations:

- when the failover server key has been compromised
- when the failover server certificate has expired
- when the host name of the failover server changes
- when one of the certificate fields (such as organization, city, etc.) of the certificate has changed

Only root can run this utility.

The utility does the following, in the following order:

- generates a new key pair in the keystore
 - swaps the master key encryption to use the new key pair
 - deletes the old key pair in the keystore
 - generates a new certificate request based on the new key
 - gets the certificate request signed by the CA (Certificate Authority, located on the primary server)
 - imports the new server certificate back to the keystore
-

Server and agent communication is not affected by this change.

re_gen_cert

This utility regenerates the server certificate for the primary Security Server.

The user can run this utility in one of the following situations:

- when the server key has been compromised
- when the server certificate has expired
- when the host name of the primary server changes
- when one of the certificate fields (such as organization, city, etc.) of the certificate has changed
- when the primary server is restored to another box with a different host name

Only root can run this utility.

The utility does the following, in the following order:

- generates a new key pair in the keystore
- swaps the master key encryption to use the new key pair
- deletes the old key pair in the keystore
- generates a new certificate request based on the new key
- gets the certificate request signed by the CA (Certificate Authority)
- imports the new server certificate back to the keystore

All server and agent communication is not affected by this change

re_sign_cert

This utility re-signs the server certificate for the primary Security Server.

The user can run this utility when the server certificate has expired.

Only root can run this utility.

The utility does the following, in the following order:

- generates a new certificate request from the existing key pair in the keystore
- gets the certificate request signed by the CA (Certificate Authority)
- imports the new server certificate back to the keystore

All server and agent communication is not affected by this change.

re_sign_failover_cert

This utility re-signs the server certificate for the failover Security Server.

The user can run this utility when the failover server certificate has expired.

Only root can run this utility.

The utility does the following, in the following order:

- generates a new certificate request from the existing key pair in the keystore
- gets the certificate request signed by the CA (Certificate Authority, located on the primary server)
- imports the new server certificate back to the keystore

All server and agent communication is not affected by this change.

setup_failover_cert

The `setup_failover_cert` utility creates failover server certificates. It uses the Java keytool to create the signer and server certificates, then it places them in the keystore.

Use this utility if the failover server is not configured with certificates, as can occur when failover server software installation fails during the certificate generation and exchange stage. `setup_failover_cert` will fail if there are certificates already configured and it will direct you to run `re_gen_failover_cert` instead.

To generate failover server certificates:

- 1: Log onto the failover server as the root user.

The failover server must have network access to the primary server.

- 2: Execute the `./server/bin/setup_failover_cert` utility without any arguments.

For example:

```
# ./server/bin/setup_failover_cert
```

The utility prompts you for the failover server network name. It uses the host name as the default.

```
Enter the host name of this computer. This will be used by
Agents to talk to this Security Server.
```

```
Failover Security Server host name [vmlinux02]:
```

- 3: Enter the host name or FQDN of the failover Security Server.

This is the recognized network identity that the failover server will use to communicate with the primary server and agents. Enter the same name or FQDN that is configured for the failover server in the *High Availability Servers* window.

The utility prompts you for the primary server network name.

```
Enter the host name for the Primary Security Server.
```

```
Primary Security Server host name: vmlinux01
```

- 4: Enter the host name or FQDN of the primary Security Server.

This is the recognized network identity that the primary server uses to communicate with the failover server and agents.

Hint: If you are unsure of the primary server network identity, execute the `./server/bin/internal/check_host` utility on the primary server and enter the name exactly as shown.

The utility creates a Java keystore and certificates. It then prompts with a series of questions for information to identify the certificate issuer.

- 5: Enter information to identify the certificate issuer.

The information that you provide is displayed when the failover server signer-certificate is viewed. You are prompted to specify:

- Your organizational unit, which is frequently a department or group name
- Organization name, which is frequently the company name
- City or locality in which the organization is located
- State or province in which the organization is located
- The country in which the organization is located

After this information is entered, the installation utility generates keys and certificates. It then exchanges them with the primary server using SSH through a secure SSL channel.

If the failover server had attempted to exchange certificates with a Security Server at some time in the past, a private key file for SSH may already exist and the installation utility will ask to overwrite the existing file.

```
Generating the private key...done.
done.
```

```
We will now contact the Primary Security Server for the
generation of certificates.
/home/voradmin/.ssh/id_rsa already exists.
Overwrite (y/n)?
```

- 6: If you are prompted to overwrite the `id_rsa` file, enter “y”.

The `id_rsa` file contains the private server key and it must be updated with the new private key.

- 7: Enter the password for the `voradmin` user on the primary Security Server when prompted.

For example:

```
voradmin@vmlinux01's password:
```

The utility begins to exchange certificates and keys with the primary server after you enter the password.

```
Establishing an ssh tunnel to "vmlinux01"...done
Generated certificate for failover server on the Security
Server machine.
Copying failover server certificate files from Security
Server done.
ca.keystore                100% 4577      4.5KB/s   00:00
keystore.password          100%   76      0.1KB/s   00:00
Importing the signer certificate into the keystore...done.
Importing the server certificate into server's
keystore...done.
```

```
Certificate and Key generation complete.
```

- 8: Start replication on the primary Security Server.

At this stage, the failover Security Server is fully configured and ready for replication. You must manually start replication on the primary Security Server. See [“Configuring SQL replication” on page 13-222](#) for details.

- 9: For each agent, modify the URL in the `agent.conf` to point to the new server hostname.

The `agent.conf` file is located at:

```
(UNIX) /opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/etc/
agent.conf
```

```
(UNIX) /opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2/etc/agent.conf
```

```
(Windows) \Documents and Settings\All Users\Application
Data\IBM\DB2TOOLS\LUWEncryptionExpert\Agent\vmd\etc\age
nt.conf
```

See step 4 in [“Configuring an agent to use Security Servers” on page 6-100](#) for details.

- 10: (UNIX) Restart the `vmd` process to refresh the File System Agent configuration.

a: Kill the `vmd` process if it is still running:

```
# ps -aef | grep vmd
```

```
# kill -9 procID
```

b: Start the vmd process.

```
# ./agent/vmd/bin/vmd
```

If you are unable to kill the process, disable the **FS Agent Locked** checkbox for the host in the Management Console.

11: (Windows) Restart the vmd process to refresh the File System Agent configuration.

a: Open a CMD window.

b: Enter “net stop vmd”.

```
The vmd service is stopping.
```

```
The vmd service was stopped successfully.
```

c: Enter “net start vmd”.

```
The vmd service is starting.
```

```
The vmd service was started successfully.
```

start_replication

The `start_replication` utility starts the SQL replication of the primary Security Server to the failover Security Server. See [“Configuring SQL replication” on page 13-222](#) for usage information.

stop_replication

The `stop_replication` utility stops the SQL replication of the primary Security Server to the failover Security Servers. See [“Configuring SQL replication” on page 13-222](#) for usage information.

synchronize_replication

The `synchronize_replication` utility synchronizes the failover Security Server database with the primary Security Server database. Normally, the primary Security Server synchronizes databases with the failover Security Server at regular intervals, plus the failover Security Server queries the primary Security Server for database updates when it is restarted. If the primary server is unable to connect to the failover server, the primary server will keep trying at one minute intervals. However, the connection can become stale when, for example, there is a long network interruption between the primary and failover servers. Use the `synchronize_replication` utility on the primary server to refresh the connection and push database updates to the failover server.

Run the `synchronize_replication` utility when you notice that, after agents switch to a failover server, the failover server does not produce the same results as the primary server. Or, when you are changing the primary server configuration and the network drops unexpectedly.

To use this utility:

- You must be the `voradmin` user.
- Replication must already be configured and running on both the primary and failover Security Servers.
- This utility takes no arguments.

If the utility fails to complete, you will be directed to view a file in `/tmp` named `synchronize_failover_YYYY-MM-DD.log`. For example, `synchronize_failover_2008-07-15.log`.

See also [“If the Security Server databases go out of sync” on page 13-227](#).

To synchronize the failover Security Server with the primary Security Server:

- 1: Login to the primary Security Server as voradmin.

- 2: Execute the synchronization utility.

```
# ./bin/synchronize_replication
```

The synchronization utility prompts you to continue:

```
Security Server Replication Synchronization
We will now synchronize the Security Server failover server
database
Do you wish to proceed? (y|n)[n]:
```

- 3: Enter “y” to continue.

You are prompted to enter the administrator password for the primary Security Server:

```
Stopping replication...done.
Enter the admin (voradmin) password for this Security
Server:
```

- 4: Enter and re-enter the voradmin administrator password for the primary Security Server.

The synchronization utility lists all the configured failover Security Servers,

The following security servers are currently configured for replication:

```
[1]: vmlinux02
[2]: vmlinux03
```

Choose the failover server for replication synchronize[1]:

- 5: Enter the number for the failover server that you wish to synchronize with the primary server.

The synchronization utility displays database properties and then prompts you to enter the voradmin administrative password for the failover server.

```
Database Connection Information
Database server          = DB2/LINUX 9.1.3
SQL authorization ID     = VORADMIN
Local database alias     = CGSSDB
Enter the admin (voradmin) password for vmlinux02:
```

- 6: Enter and re-enter the voradmin administrator password for the failover Security Server.

The synchronization utility refreshes the primary and failover server configurations, and then exits.

```
Re-enter the password to verify:
Synchronizing failover server database...
Cleaning up residual manual refresh load/unload files ...
Performing refresh...Starting initial change capture
processes...
```

```
Database Connection Information
Database server          = DB2/LINUX 9.1.3
SQL authorization ID     = VORADMIN
Local database alias     = CGSSDB
Security Server Replication Configuration complete.
You must restart the Security Server by running
"/etc/init.d/cgss restart" as the installation user.
#
```

- 7: Open a terminal window on the primary server as the root user.

- 8: Restart the primary server.
`/etc/init.d/cgss restart`

Agent utilities:

buildinfo

Displays the installed Encryption Expert Database Backup Agent version. Ignore the `buildinfo` utility in the `./vmd/bin` directory. There is no corresponding `*.so` file for this utility to display the File System Agent version. See [“Determining the Encryption Expert Database Backup Agent version” on page 16-270](#).

For Windows hosts, use the Encryption Expert tray instead.

check_host

See [“Verifying your system identity” on page 16-300](#) for usage information.

For Windows hosts, use the Encryption Expert tray instead.

register_host

The `register_host` command looks and operates differently on UNIX and Windows systems. See [“Changing host certificates” on page 6-88](#) for usage information on UNIX systems. Windows usage information is presented below.

The `register_host.exe` utility is used to create certificate requests, exchange certificates between the Security Server and the host, and to register agents on the Security Server. After the host is registered, the agents can be configured and guard points applied or database backups performed. The Windows `register_host.exe` utility is run via a graphic interface.

Syntax

```
./register_host.exe SSname
```

where, *SSname* is an optional argument to provide the Security Server name now rather than having to enter it later in the graphic interface.

You can also double-click the executable to open the graphic interface.

Example

To create certificates for a new File System Agent installation or to update existing File System Agent certificates:

- 1: If you have not done so already, add the host to the Security Server database, configure the types of Encryption Expert agents that are allowed to communicate with the Security Server, and enable the **Registration Allowed** and **Communication Enabled** checkboxes for the host.
- 2: Select **Register Host** in the Encryption Expert tray.

The registration utility can also be started by entering `register_host.exe` without any arguments in a command window or double-click the utility in Windows Explorer.

The utility is located in

```
\Program Files\IBM\DB2TOOLS\LUWEncryptionExpert\Agent
./db2/bin/register_host.exe
./vmd/bin/register_host.exe
```

The *Register Host* window opens.




Figure 16.1: Sample *Register Host* startup window

- 3: Click **Next**.

The next *Register Host* window opens.

- 4: Enter the name of the local host in the **Host Name** text-entry box.

This must be the same name that was used to add the host to the Security Server database.

- 5: Enable the checkboxes for the types of agents to be registered with the Security Server.

You can choose to configure the certificates of individual agents. All File System and DB2 agents are enabled by default.




Figure 16.2: Specifying the host name and agents to be configured

- 6: Click **Next**.

The next *Register Host* window prompts for the name of the primary Security Server that is to service this host.

If you entered the `register_host` command without an argument, the window appears as shown in [Figure 16.3](#) below. If you entered the `register_host` command with the name of a Security Server, that name will be displayed in the **Enter server name** text-entry box.




Figure 16.3: Specifying the Security Server for the host

- 7: Enter the network name of the primary Security Server in the **Enter server name** text-entry box, if the Security Server name is not already displayed in the text-entry box.
- 8: Click **Register**.

The information that you entered is used to register the local host with the Security Server and activity status is displayed in the *Register Host* window.

If `register_host` is run on a system that is already registered with the Security Server, the *register_host* pop-up opens. The *register_host* pop-up warns you of the existing registration and prompts you to continue.




Figure 16.4: Replacing the existing agent certificate

- 9: If the *register_host* pop-up is open, click **OK** to create new host authentication credentials or **Cancel** to leave the existing credentials intact and exit the *register_host* utility.

This example assumes that you clicked **Ok**.




Figure 16.5: Approving File System Agent certificate

The *register_host* utility displays the File System Agent fingerprint for your approval.

10: Click Ok.

You are prompted to accept the signer CA certificate. A pop-up window displays the signer CA certificate fingerprint, which is also displayed in the Management Console *Dashboard* window.




Figure 16.6: Accepting the CA fingerprint

11: Verify the signer fingerprint.

Compare the fingerprint value in the pop-up window with the fingerprint value in the Management Console *Dashboard* window.

12: Assuming the fingerprints match, accept the certificate and click **Yes**.

If only a File System Agent is installed, certificate generation, certificate exchange, and host registration are complete. The *Register Host* window displays the success or failure of the registration process. An example is shown in [Figure 16.9](#).

If a Database Backup Agent is installed on the host system, and if the host system has certificates on it, the *register_host* pop-up opens. The *register_host* pop-up warns you of the existing registration and prompts you to continue.




Figure 16.7: Replacing the existing agent certificate

- 13: If the `register_host` pop-up is open, click **OK** to create new host authentication credentials or **Cancel** to leave the existing credentials intact and exit the `register_host` utility.

This example assumes that you clicked **Ok**.




Figure 16.8: Approving Database Backup Agent certificate

The `register_host` utility displays the Database Backup Agent fingerprint for your approval.

- 14: Click **OK**.

The host and agent registration process is complete. You compare the fingerprints displayed here with the fingerprints displayed in the Management Console *Edit Host* window for this host.




Figure 16.9: Successful host registration

15: Click **Finish**.

secfsd

On UNIX, the `secfsd` utility is located in `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/secfs/.sec/current/bin` and a symbolic link to this file is placed in `/usr/bin/secfsd`.

On Windows, the `secfsd` utility is located in `\Program Files\IBM\DB2TOOLS\LUWEncryptionExpert\agent\secfs\sec\bin\secfsd.exe`.

The `secfsd` utility displays the following attributes of the Encryption Expert File System Agent:

- guard points defined in the *Guard FS* tab
- authentication parameters defined in the *Host Settings* tab
- lock status set by enabling **FS Agent Locked** and **System Locked**
- Web destination and SSL certificate for uploading log entries
- policies applied in the *Guard FS* tab
- status of required processes (`secfsd` and `vmd`)
- `secfs` version

Syntax

```
secfsd [OPTIONS...]
    -debug on|off          show debug log messages
                           for agent activity (secfsd)
    -help                  show utility syntax
    -status                update secfs/tmp/statusfile
    -status guard [-v]     list guard points
    -status auth            list authentication settings
    -status lockstat       list system and agent locks
    -status logger         list logging details
    -status policy         list configured policies
```

```

-status pslist          list protected process
-version               list secfs version
cmd -c debug.level.on|off show debug log messages
                        for kernel activity (vmd)

```

Examples

To create or update the `./secfs/tmp/statusfile` file.

```

# secfsd -status
#

```

Note the `statusfile` file is deleted after an agent configuration change is applied in order to prevent you from looking at stale configuration data.

To display the guard point paths, applied policies, policy type, and guard status:

```

# secfsd -status guard
GuardPoint      Policy      Type      ConfigState  Status
-----
/home/apps/app2  auditOnly   local     guarded      guarded
/home/apps/app3/bin denyOneFile local     guarded      guarded

```

GuardPoint is the full path of the guard point. Policy is the name of the policy applied to the guard point. Type may be either `local` or `automount`, and it is configured in the *Guard FS* tab. ConfigState is the guard status of the guard point, as recognized by the Security Server. It may be `guarded` or `not guarded`. Status is the current guard status, as recognized by the File System Agent. The ConfigState and Status states may vary when, for example, you apply a guard point but someone is currently working in the guard point, so the policy cannot be applied at this time. In this case, the ConfigState would be `guarded` and the Status would be `not guarded`.

To display the same information in a different format, include the `-v` argument. For example:

```

# secfsd -status guard -v
GUARD POINT: 1
    Policy:      auditOnly
    Directory:    /home/apps/app2
    Type:         local
    Status:       guarded
    ConfigState:  guarded
GUARD POINT: 2
    Policy:      denyOneFile
    Directory:    /home/apps/app3/bin
    Type:         local
    Status:       guarded
    ConfigState:  guarded
#

```

Use the `auth` argument to `secfsd` to display the signature SHA-1 hash of each File System Agent host setting. For example:

```

# secfsd -status auth
|authenticator|/bin/su
C2D8AE0175363C649976A109288A91CFE7F717B1
|authenticator|/usr/sbin/sshd
4FD810A4F2E9577E48475596AD30A61E98298A92

```

```
|authenticator|/bin/login
      F926C5D4209CBFE9E0F8CDB1ABB9FAB9C0C91D10
|authenticator|/usr/bin/gdm-binary
      C4802C61912D41DAABEE4448EB3E624C13A4C42C
|authenticator|/usr/bin/kdm
      2EE31C9304A669E76CDB2B68093945A20B7A208F
#
```

The `auth` argument stands for “authentication”, as in user authentication for those processes that have a concept of a user identity, such as `telnet` and `ssh`. Configure these values in the *Host Settings* tab. The host setting and hash value is also displayed in `./secfs/tmp/statusfile` and `./secfs/.sec/conf/configuration/.authbin.res`.

To view the status of File System Agent locks:

```
# secfsd -status lockstat
  FS Agent Lock: false
  System Lock: false
#
```

The value is `true` if the respective lock is applied. The value is `false` if the respective lock is not applied. The status of `FS Agent Lock` must be `false` before attempting to upgrade or remove File System Agent software. `FS Agent Lock` corresponds to **FS Agent Locked** in the *Host* window. The status of `System Lock` must be `false` before upgrading or changing operating system files. `System Lock` corresponds to **System Locked** in the *Host* window.

To view the status of the Encryption Expert logger:

```
# secfsd -status logger
  Upload URL:
    https://vmlinux01:8444/upload/logupload
  Logger Certificate directory:
    /opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/vmd/pem
#
```

This command sequence returns the URL to which the logger sends log data. It also returns the directory that contains the File System Agent certificate. The certificate is used to authenticate the agent when it uploads the log data to the Security Server.

To display the policies that are applied to the File System Agent:

```
# secfsd -status policy
  Policy: auditOnly
  Type: regular
  Policy: denyOneFile
  Type: regular
#
```

To display Encryption Expert File System Agent processes:

```
# secfsd -status pslist
  Protected pid list:      739    731
# ps -aef | grep 739
  root  739  1  0 Nov14 ?   00:05:58 /opt/IBM/DB2TOOLS/
LUWEncryptionExpert/agent/vmd/bin/vmd
```

```
# ps -aef | grep 731
root 731 1 0 Nov14 ? 00:04:37 /opt/IBM/DB2TOOLS/
LUWEEncryptionExpert/agent/secfs/.sec/current/bin/secfsd
#
```

The example displays the process PID numbers for the `vmd` and `secfsd` processes. The `ps` commands show the processes for those PIDs.

To display the Encryption Expert File System Agent version information:

```
# secfsd -version
version: 4.0-Linux-2.6-x86-32-Build0020
#
```

To enable or disable `secfs` File System Agent debugging:

```
# secfsd cmd -c debug.auth.on
secfsd: ok
# secfsd cmd -c debug.auth.off
secfsd: ok
#
```

This can greatly increase the amount of log information. When you are done debugging the File System Agent installation, disable debug mode. You can achieve the same result by substituting `secfsd` in the command line with `vmsec`.

vmd

The `vmd` utility displays the File System Agent software version information.

The `vmd` utility is located in `/opt/IBM/DB2TOOLS/LUWEEncryptionExpert/agent/vmd/bin` and a symbolic link to this file is placed in `/usr/bin/vmd`.

Syntax

```
vmd [OPTIONS...]
    -h                show utility syntax
    -v                display File System Agent version
```

Examples

To display the installed File System Agent version:

```
# vmd -v
1.1.2.0
1.1.1.0-Linux-2.6-x86-64-Build0054
2008-03-27 16:00:29 -0700 (PDT)
#
```

vmsec

On UNIX, the `vmsec` utility is located in `/opt/IBM/DB2TOOLS/LUWEEncryptionExpert/agent/vmd/bin` and a symbolic link to this file is placed in `/usr/bin/vmsec`.

On Windows, the `vmsec` utility is located in `\Program Files\IBM\DBTOOLS\LUWEEncryptionExpert\agent\vmd\bin`.

Syntax

```
vmsec [OPTIONS...]
    checkinstall      shows vmd kernel status
    cmd -c cmd      capture debug-level messages
```

passwd	enter the host password
status	display kernel status
version	display File System Agent version

Examples

To display the File System Agent kernel run status:

```
# vmsec checkinstall
    The kernel component is installed and running.
#
```

This utility will tell you if the File System Agent is configured and running. If it is not running, you may need to start it manually.

To enter the File System Agent host password:

```
# vmsec passwd
    Please enter password:
    OK passwd
#
```

To enable debug-level logging:

```
# vmsec cmd -c debug.auth.on
```

After entering the Encryption Expert debug mode, use the `tail` command to view the system messages log:

```
# tail -f /var/adm/messages
```

Debug mode can generate a lot of data. Turn off debug mode when it is no longer necessary.

To turn off debug mode:

```
# vmsec cmd -c debug.auth.off
```

The following example shows executing a command, starting debug mode logging, and then executing the same command a second time:

```
#tail -f /var/log/messages
Nov 20 13:30:32 vmlinux02 kernel: [CGP] [INFO ] [EET2601I]
Policy[denyOneFile]
User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\]
Process[/home/apps/app3/bin/xclock] Action[read_file] Res[/home/
apps/app3/bin/xclock] Key[None] Effect[PERMIT Code (1R,2M)]
Nov 20 13:30:32 vmlinux02 last message repeated 2 times

Nov 20 13:30:44 vmlinux02 kernel: vmfs_cmd: setting debug (auth.on)
on
Nov 20 13:30:48 vmlinux02 kernel: vm_proc_get_attr: GetAuth ret
gen: 7 uid: 0 this_uid: 0 pid: 15392
Nov 20 13:30:48 vmlinux02 kernel: vm_proc_get_attr: id.uid: 0
count: 7 pid: 15392
Nov 20 13:30:48 vmlinux02 kernel: usr_find: root
Nov 20 13:30:48 vmlinux02 kernel: usr_inf_get: UID(0)
Nov 20 13:30:48 vmlinux02 kernel: [CGP] [INFO ] [EET2601I]
Policy[denyOneFile]
User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\]
Process[/home/apps/app3/bin/xclock] Action[read_file] Res[/home/
apps/app3/bin/xclock] Key[None] Effect[PERMIT Code (1R,2M)]
Nov 20 13:30:48 vmlinux02 kernel: vm_proc_get_attr: GetAuth ret
gen: 7 uid: 0 this_uid: 0 pid: 15392
```

```

Nov 20 13:30:48 vmlinux02 kernel: vm_proc_get_attr: id.uid: 0
count: 7 pid: 15392
Nov 20 13:30:48 vmlinux02 kernel: usr_find: root
Nov 20 13:30:48 vmlinux02 kernel: usr_inf_get: UID(0)
Nov 20 13:30:48 vmlinux02 kernel: [CGP] [INFO ] [EET2601I]
Policy[denyOneFile]
User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\]
Process[/home/apps/app3/bin/xclock] Action[read_file] Res[/home/
apps/app3/bin/xclock] Key[None] Effect[PERMIT Code (1R,2M)]
Nov 20 13:30:48 vmlinux02 kernel: vm_proc_get_attr: GetAuth ret
gen: 7 uid: 0 this_uid: 0 pid: 15392
Nov 20 13:30:48 vmlinux02 kernel: vm_proc_get_attr: id.uid: 0
count: 7 pid: 15392
Nov 20 13:30:48 vmlinux02 kernel: usr_find: root
Nov 20 13:30:48 vmlinux02 kernel: usr_inf_get: UID(0)
Nov 20 13:30:48 vmlinux02 kernel: [CGP] [INFO ] [EET2601I]
Policy[denyOneFile]
User[root,uid=0,gid=0\root,bin,daemon,sys,adm,disk,wheel\]
Process[/home/apps/app3/bin/xclock] Action[read_file] Res[/home/
apps/app3/bin/xclock] Key[None] Effect[PERMIT Code (1R,2M)]
Nov 20 13:30:54 vmlinux02 kernel: vmfs_cmd: setting debug (auth.off)
off

```

To display the kernel status:

```
# vmsec status
```

This command returns the contents of the `statusfile` file. See [“statusfile” on page 14-240](#).

To display the File System Agent build information:

```

# vmsec version
1.1.2.0
1.1.1.0-Linux-2.6-x86-64-Build0054
2008-03-27 16:00:33 -0700 (PDT)
#

```

Encryption Expert database maintenance

The Encryption Expert DB2 databases are named CGSSDB and LOGDB. CGSSDB contains all the policies, host configurations, and keys that are used by the Encryption Expert interface. You should not need to modify this database directly; however, you should back up the database before making major changes or upgrading the Encryption Expert Security Server installation. Doing so will give you a means to restore the Security Server configuration, if the change or upgrade produces undesirable results. If you remove the embedded database or the system crashes, and you do not have backups of the embedded database, you cannot restore DB2 backups.

LOGDB is the embedded Security Server database that contains all the log entries that are displayed in the Security Server log window. You can backup and restore this database. Note that when restoring any database, the current database is wholly replaced by the backup database.

Also, back up the `./server/pem` and `./server/etc` directories. The contents of `./server/pem` are required for agent and server authentication. The contents of `./server/etc` are required to restore previous backups.

Backing up the Security Server configuration

The `backup_config` utility creates a secure backup of the embedded Security Server databases. This backup can be used to restore the hosts, encryption keys, and policies of a Security Server for software crash recovery or system changes. Only the two embedded databases and the audit log are backed-up. Agent/server certificates are not included in the backup.

A backup can be restored only on a system that is running the same Security Server build and version numbers as the system on which it was made. This utility is also intended to upgrade to a new build or version, unless the README file indicate otherwise; as in the case of a database structure change that makes an old database incompatible with a new database schema.

The utility creates an encrypted tar file in `/tmp` named `Backup_YYYY_MM_DD_RRRR.tar`, where *Y* is the year, *M* the month, and *D* is the day. *R* is a random number that is used to differentiate between other backups made on the same day. You must use the companion restore utility, `restore_config`, with the password to extract the contents of this file.

To back up a Security Server configuration:

- 1: As the root user, execute the `backup_config` utility.

```
# ./server/bin/backup_config
```

The utility initializes itself, then prompts you twice to enter a password with which to secure the backup file.

```
Please provide a password to protect the configuration backup:
```

```
Please enter password again to confirm:
```

- 2: Enter an alphanumeric string, up to 8 characters, to use to encrypt the backup.

You must use this same password to restore the backup.

The utility prompts you for permission to temporarily shut down the Security Server so it can copy the database offline.

```
The Security Server needs to shutdown before the backup can proceed.
```

```
Shutdown Security Server? (yes|no)[no]:
```

- 3: Enter `yes` to continue Security Server backup.

If replication is enabled on the server, the backup utility will also stop the replication processes.

The utility encrypts the tar file, removes all temporary files, and restarts the Security Server.

- 4: Note the name and location of the generated tar file.

```
...
Security Server configuration backup complete.
Backup file is /tmp/Backup_2008_04_22_1523.tar
```

```
Starting Security Server...
Starting Security Server done.
Starting Security Server Replication...done.
#
```

- 5: Copy the backup file to a secure location.
-

Do not rename the backup file. Copy the backup file exactly as-is. Also, note somewhere the password used to encrypt the backup file.

Restoring the Security Server configuration

The `restore_config` utility restores an encrypted Security Server backup that was generated by executing the `backup_config` utility.

A backup is a snapshot of a Security Server configuration at one point in time. When restored, the Security Server Web interface will contain and display the same information captured at the time the backup was made. After an upgrade, information like the Encryption Expert version and build numbers, still show the older values. (This is a display issue only. You are running the new Encryption Expert software. You must use other methods to determine the actual version and build numbers.)

If you are restoring a backup of a 1.1.2 Security Server, see also [“import_v1120_config” on page 16-273](#).

You are prompted for the encrypted tar file and password that were used to make the backup.

To restore a Security Server configuration:

- 1: Change to the Security Server `bin` directory.

```
cd /opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/bin/
```

- 2: Run: `restore_config -f tar_file stage_dir`

`tar_file` is the full path to the backup file. For example,
`/testDir/myBackup/Backup_2008_04_22_1523.tar`

`stage_dir` is an optional temporary staging area, for example, `/tmp`. If you do not specify a temporary staging area, `/tmp` will be used.

An example of this command is shown below:

```
./bin/restore_config -f /tmp/Backup_2008_04_22_1523.tar
```

The utility prompts you to stop Security Server processes so that it can restore the Security Server configuration in a static environment.

The Security Server needs to shutdown before the restore can proceed. All the current data will be over-written with the backup data.

Are you OK to proceed? (yes|no)[no]:

- 3: Enter `yes` to temporarily shutdown the Security Server and proceed with restoring the Security Server database.

The utility prompts you to add the users that are in the backup to the current Security Server configuration.

Restore users from backup? (yes|no):

- 4: Enter `yes` to add the users in the backup to the current Security Server configuration or `no` to discard the users in the backup.

Note that, if you add the users from the backup, and there are similarly named users already in the current Security Server configuration, any changes made to the user configurations since the backup are discarded and the old passwords and roles are reapplied.

The utility prompts you to enter the same password that was used to make the backup.

Please provide the password the configuration backup was protected with.

```
enter aes-256-cbc decryption password:
```

- 5: Enter the password that was used to make the backup.

The utility restores the backup and restarts the Security Server.

The following message indicates that the restoration is complete, the Security Server has been restarted, and, if this is a primary server, replication has been enabled.

```
Starting Security Server done.
```

- 6: Open the Security Server in a Web browser and verify the restoration.
- 7: If you restoring the database backup on the same system on which the backup was originally made, and the system configuration has not been changed, and you are using the same Security Server installation that was used to make the backup, you can continue to use the certificates in the `./server/pem` directory. You can also restore a backup of the certificate files to the `./server/pem` directory if the backup was made on the same system, etc.
- 8: If you cannot reuse existing server certificates for the Database Backup Agent, execute the `setup_eet_agent_cert_db2` utility with the “clean” argument on each Database Backup Agent to remove any existing certificates and related certificate files. This command is explained in [“Changing host certificates” on page 6-88](#).
- 9: If you cannot reuse existing server certificates for a UNIX File System Agent, execute the `register_host` utility with the “clean” argument on each File System Agent to remove any existing certificates and related certificate files. The equivalent Windows command opens a window in which to This command is explained in [“Changing host certificates” on page 6-88](#).
- 10: Execute the `setup_eet_agent_cert_db2` utility without any arguments on each Database Backup Agent system to regenerate the agent authentication certificates and reestablish agent/server communication.
- 11: Execute the `register_host` utility without any arguments on each File System Agent system to regenerate the agent authentication certificates and reestablish agent/server communication.

Planning file system policies

Writing Security Server policies is easy. The hard part is anticipating system behavior. When you think you are performing just one action, you may actually be performing several. And though these additional actions are transparent to you, the Encryption Expert Security Server catches them and applies policy rules to them. The policy is applied correctly but the end results may not be what you expect. This section provides several examples of behavior variations so you can create better and more predictable policies.

- Why doesn't the Security Server prevent a command from being executed when I have it configured in my security rule?

Some commands are not executed in the manner you expect. For example, you can write a policy to prevent the UNIX `ls` command from being executed. The policy is correctly applied to `/usr/bin/ls`, but the policy doesn't prevent it from being executed. This can be because you are not actually executing `/usr/bin/ls`, as expected, but a version of `ls` that is internally compiled into the shell you are running. If you enter the full-path to the command and it is not executed, it is safe to assume the internal version of the command is being used. Though Security Server

policies can be applied against the shell executable, you cannot apply similar enforcement to the internal processes it spawns.

- Why does a command work part way when it shouldn't work at all?

Some operating systems implement commands in multiple steps, some of which may be blocked by a Security Server policy. A good example is the `mv` command on older Linux systems. It used to be executed as a two-step process. It copies the data first, then attempts to remove the source data. If a policy is applied to prevent data deletion, but it allows the data to be read, the `mv` command copies the data to the move point then generate an error when it attempts to remove the source data.

- Why does an application hang when clearly there is no policy violation?

The best way to explain this one is with an example. You can create a policy with two encryption key rules. The first one is applied to "hidden" files (`. *`) and it is configured with a **Stored on Server** key (this key requires network access to the Security Server). The second encryption key rule is applied to configuration files that end in `.cfg`. This rule is configured with a **Cached on Host** key (this key does not require network access to the Security Server). Then, when you try to edit a file that ends in `.cfg`, the application appears to lock up.

Some operating systems create hidden files, like `.swp` or `.txt`, to perform simple operations like editing a file. When the operating system attempts to create a hidden file to cache the work being performed, the File System Agent software takes control and waits for the Security Server to supply the requisite encryption key. If there is no network access to the Security Server, the File System Agent software will wait indefinitely for the Security Server to become available.

- Why does an application return a policy violation when there is no violation of the security rule?

Again, it can be because of the way the operating system performs the task. Using a variation of the preceding example, you can create two security rules: one that allows you to edit data files and one that does not allow you to create hidden files (`. *`). If, in the process of editing a file, the system attempts to create a hidden cache file like `.swp`, permission is denied, and the application fails.

The most important rule in policy creation is to know how the host operating system behaves. You need to anticipate system behavior and always test new policies before putting them into a production environment. Enable tracking options like `Trace`, `Audit`, and `Warn Mode` to verify that the policy and the host operating system behave as expected.

Troubleshooting

Things to check

What you check and the depth of checking you perform is dependent upon the nature of your problem. The following is a list of general things you can check to determine the cause of a problem. Frequently, it is a small thing, such as a typo by entering an uppercase letter than a lowercase letter.

You've installed the Security Server and an agent, now what?

- Set the terminal I/O configuration to recognize backspace keys like `<Ctrl>-h` and `<Backspace>` (usually displayed on the keyboard as a left-pointing arrow). If this is not set, and you remove text using a backspace key during software installation, it will appear that you had deleted the text but, in fact, it is still there. Use the `stty`
-

command to set the `erase` parameter. The actual command that you enter depends upon your user shell. The following example shows how to set the `erase` parameter in a bash shell:

```
# read BACKSPACE
<!-- Press the <Backspace> key. -->
# stty erase $BACKSPACE
```

`<Ctrl>-u` works as expected.

- Verify that the secure shell daemon, `sshd`, is running on the server and is using port number 22. 22 is the default port number. If you can initiate an SSH session to the server, SSH works and no further checking is required. If it doesn't work, check the server system. For example,

```
# ps -aef | grep sshd
root      2301      1  0 Aug08 ?    00:00:00 /usr/sbin/sshd
# cat /etc/services | grep ssh
ssh        22/tcp      # SSH Remote Login Protocol
ssh        22/udp      # SSH Remote Login Protocol
#
```

If you do not see the `sshd` process or SSH ports, refer to your system documentation to configure and start SSH.

- Verify that system ports numbers 8443, 8444, and 8445 are open before installing Security Server software. The Security Server requires these ports for certificate exchange, logging, and Web access, respectively. These ports are not checked for availability during installation, so you have to do it manually. Since you are supposed to be running the Security Server on a dedicated system, these ports are normally open.

Use the `netstat` command on the system to host the Security Server to check that ports 8443, 8444, and 8445 are open. For example,

```
# netstat -tl --numeric-ports | grep 8443
```

If the port is open, nothing is returned. If the port is in use, the system returns a message like the following:

```
# netstat -tl --numeric-ports | grep 8443
tcp        0      0 *:8443          *:*             LISTEN
#
```

You can only determine that something is listening on the port, not what is listening. You have to determine what is listening by some other means.

- Another way to determine port usage, and also see who is using the port, is with the `lsof` command. For example, an excerpt is shown below.

```
# lsof -i
COMMAND      PID      USER     FD  TYPE DEVICE  NODE NAME
...
sshd          3244     root      3u   IPv6   7160  TCP  *:ssh (LISTEN)
sshd          5370     root      3u   IPv6  12907  TCP  vmlinux01-eth1:ssh-
>10.3.60.25:1046 (ESTABLISHED)
...
db2sysc      12240   voradmin    3u   IPv4  190079  TCP  *:db2c_voradmin (LISTEN)
java         14774   voradmin  134u  IPv6  196989  TCP  *:1098 (LISTEN)
java         14774   voradmin  136u  IPv6  197012  TCP  *:1099 (LISTEN)
java         14774   voradmin  138u  IPv6  197015  TCP  *:krb524 (LISTEN)
java         14774   voradmin  139u  IPv6  197024  TCP  *:4445 (LISTEN)
java         14774   voradmin  188u  IPv6  197293  TCP  *:webcache (LISTEN)
```

```

java      14774 voradmin  190u  IPv6 197294 TCP *:8443 (LISTEN)
java      14774 voradmin  191u  IPv6 197295 TCP *:8444 (LISTEN)
java      14774 voradmin  192u  IPv6 197296 TCP *:8445 (LISTEN)
java      14774 voradmin  217u  IPv6 216091 TCP vmlinux01:34782->vmlinux01:8093
(ESTABLISHED)
java      14774 voradmin  218u  IPv6 197390 TCP *:8093 (LISTEN)
java      14774 voradmin  221u  IPv6 216090 TCP vmlinux01:8093->vmlinux01:34782
(ESTABLISHED)
java      14774 voradmin  232u  IPv6 199021 TCP *:8009 (LISTEN)
#

```

- (Windows) If you had modified the network for the Windows host, restart the network. For example, in a command or terminal window, enter:


```
net stop vmd
net start vmd
```
- If your administrative terminal, the agent, or the Security Server system is behind a firewall, verify that you can access the SSH and Security Server ports through the firewall. Verify that you can access ports 22, 7024, 8443, 8444, and 8445. Check with your firewall documentation to determine port usage.
- Certificates may expire or you want to cycle through certificates at set intervals as a security precaution. You can manually regenerate locally-signed certificates. See [“Changing host certificates” on page 6-88](#).

Verifying your system identity

The identity of the system that is hosting the Encryption Expert Security Server must be appropriately configured in order to successfully install the software. If it is not, the software installation process may indicate that installation completed successfully, but the log files will contain ERROR messages and you will be unable to start a Web session with the Security Server. Sometimes this will cause an ‘internal error’ message to appear when you start a Security Server Web session.

Perform the following checks before installing Security Server software so you only have to install the software once. The information below is based on host names and static IP addresses. If you are using DNS, adjust the steps to suit your system and network setup.

Check that the system name is properly set. Run the `hostname` command. If it is not correct, or if it is set to ‘localhost’, set it using ‘`hostname name`’. The system host name is used for host registration, server start up, the Web interface, etc.

Check that the system IP address is properly set. Run the ‘`hostname -i`’ command. If it is not correct, or it returns ‘127.0.0.1’ (the localhost IP), set it to the IP address that the other systems will use to communicate with it from across the network.

If these two checks return the correct system information, the software should install without any problems. If they do not, continue with the following checks.

Run the `ifconfig` command and check that the internet addresses for the ethernet (`eth0`) and localhost (`lo`) interfaces are set correctly. If necessary, set the IP addresses using the `ip` command (e.g., `# ip addr add 10.2.0.37/16 dev eth0`).

You can also set interface IP addresses by editing the `ifcfg-eth0` and `ifcfg-lo` files in `/etc/sysconfig/network-scripts`.

Check that `/etc/sysconfig/network` is properly set. By default, Red Hat Linux sets the `HOSTNAME` parameter to `localhost.localdomain`. Set it to the name of the system (e.g., `HOSTNAME=vmlinux01`). For example:

```
# cat /etc/sysconfig/network
```

```
NETWORKING=yes
HOSTNAME=localhost.localdomain <---- wrong - reset to system name
#
```

Run `/etc/init.d/network` restart after you modify the system network, or reboot the system.

Try to start a Security Server Web session after you install the software. If you cannot, check the logs in `/tmp` and `/opt/IBM/DB2TOOLS/LUWEncryptionExpert/server/jboss-4.0.4.GA/server/default/log`. Grep for 'ERROR'. If there are a lot of errors having to do with being unable to open deployment files, unable to start voradmin, missing dependencies, failed connections to localhost:1099, etc., the most likely culprit is your system configuration.

After the Security Server is installed, and you are able to successfully initiate a Security Server Web session, install the Encryption Expert agent software. If it fails to install during the certificate generation and exchange stage, there is an additional utility that you can run to perform a quick check of the system identity. Run the `check_host` utility on Encryption Expert agents to verify that `/etc/hosts` is configured correctly. This utility returns the host name that the agent host uses to identify itself to the server. The following is an example of the `check_host` utility output:

```
[root@vmlinux02]# /opt/IBM/DB2TOOLS/LUWEncryptionExpert/agent/db2/
bin/check_host
    vmlinux02.widgets.com
#
```

If the hostname is not correct, or it is set to '127.0.0.1', or the displayed host name is something other than the actual host name, the `/etc/hosts` file or the system identity is incorrect. Correct these using the techniques described above.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new

editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation.

Actual results may vary. Users of this document should verify the applicable data for their specific environment. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

If you are viewing a softcopy of this information, the photographs and color illustrations may not appear.

TRADEMARKS

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both.

IBM

IBM logo

ibm.com

AIX

DB2

DB2 Universal Database

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.



Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Index

Symbols

|authenticator| 104, 107
|exempt| 104
|lock| 104
|trust| 105, 107
|trustfrom| 105, 107

A

access
 access permissions 146
access file 37, 48, 49
accessibility xvii
action
 access permissions 146
adding
 processes 159
 users 155
administrative roles 36
agent
 log 199
 version 270
agent.conf 48, 49, 101
anti-rootkit software, access permission 171
asnaply 223, 224, 246
asncap 223, 224, 245
audience xv

B

backup_config 222, 224, 225, 295
backups
 zero-state 42
before you begin 29, 32
browser
 browsing the file system 147
buildinfo

agent version, determining 270

C

Cached on Host 97
case sensitivity 133
cgss restart 38, 223
cgss.log 238
changing passwords 71
check_host 41, 280
check_hosts 301
cleanup_replication 221, 224, 272
clear_key 177, 179, 181, 183
cluster directory 93
commtest 269
compression 199
configure_replication 221, 222, 224
conventions xvii
creating
 resources 146
creating users 72
current key rules 179

D

database
 CGSSDB 221
 JMSDB 221
 LOGDB 221, 238
 server database, backup 295
dataxform
 apply_key, setting 169
 clear_key 179, 183
 current key rules 178
 examples 187
 key_op, setting 168

- keys, using 177
- location 184
- next key rules 179
- permit, setting 169
- syntax 184
- UNIX requirements 184

DB2

- backup command 199
- backup, making 198
- backup, restoring 201
- buffers 199
- compression 199
- conflicting software 32
- db2ckbp 199
- header size 199
- issues 32
- online backups 203
- process flow 198
- restore command 203
- rollforward option 204
- using with EE 198
- version installed 32

- db2ckbp 199

- db2setup.log 238

- deinstall_eet_agent_db2 54

- deinstall_eet_server 54

- directories used 33

E

- eet_jboss.log 238

- eet_prereq.zip 34

- effect 146, 169

- encryption

- See also*, keys

- Encryption Expert

- backup database size 199

- backups, restoring 199

- deleting, agent 54, 55

- deleting, server 54

- log, agent 247

- logs, server 242

- expired passwords 72

F

- failover server

- backups 42

- fcfg-eth0 300

- files

- access 37, 48, 49

- agent.conf 48, 49, 101

- boot.log 242

- cgss 38

- cgss.log 243

- eetdb2_jboss.log 242

- eetdb2_usr.log 203, 238

- installed, agent 47

- installed, server 37, 48

- libeetdb2.so 199

- server.log 244

- FS Agent Locked

- single-user mode 83

G

- gid 157

- global options

- warn mode 172

- group ID (gid) 155, 157

- group name (gname) variable 177

- guard point 147

H

- host groups

- changing configuration 114

- deleting 108

- displaying 108

- host protection

- selecting 173

- hosts

- fcfg-eth0 300

- hostname 300

- ifconfig 300

- localhost, setting 300

- locking/unlocking 85

- online 118

- password 97

- registered agents example 79

- Select All 76, 134

sysconfig 300
system locks 85
View All 76, 134

I

icons

Host 143

ifconfig 300

install_eet_prereq 34

install_eet_server.bin 34, 44, 52

internal error 300

J

JBoss

checking 244

startup log 242

K

key rules

defining 178, 179

deleting 176

variables, using 177

keyboard shortcuts xvii

keys

AES 137

Cached on Host 97

creating 118

displaying 116

offline storage 118

overview 115

Store on Server 98

L

libeetdb2.so 199

logs

active 33

agent log 199, 247

agent logging attributes, setting 255, 258, 260

asnapply 223, 224

asncap 223, 224

boot.log 242

cgss.log 238, 243

db2setup.log 238, 243

displaying 229

eet_jboss.log 238

eet_server_install 243

eetdb2_jboss.log 242

eetdb2_usr.log 199, 203, 238, 247

error levels, overview 237

inactive 33, 238

overview 229

panic.log 244

secfs.log 240

server logging attributes, setting 233, 253

server logs 242

server.log 238, 244

timestamps 229

vmdb2logup_host_usr_num.num 245

voradmin.CGSSDB.ASN.CAP.log 245

M

Management Console

accepting signer certificate 59

checking secure connection 61

Dashboard window 60

initial checklist 27

login window 59

starting 57

N

navigating the Policy Composer 144

netstat 299

network, restarting 301

new user passwords 73

next key rules 180

O

offline keys 97

online keys 98

P

password restrictions 70

passwords 97

Cached on Host keys 97

Store on Server keys 98

policies

additional conditions 137

attributes 132

compression 136

- creating 134
 - deleting 139
 - editing 139
 - navigating the Policy Composer 144
 - order dependence 134
 - overview 132, 134
 - setting time 136, 138, 163
- policy attributes 25
- Policy Composer
- Add New Resource Set window 147, 151, 156, 160
 - Add New Time Set window 163
 - Edit Resource Set Resource Name window 152
 - how to use 144
 - navigating 144
 - parts of 144
 - Remote Object Chooser 147, 173
 - Remote Object Chooser window 173
 - Resource Object window 147
 - Resource Set List window 151
- port numbers 28
- ports
- checking 299
 - required 299
- processes
- deleting process sets 162
 - editing process sets 162
- protection properties 97
- R**
- register_host 100
- related documents xvi
- remote file selection 173
- remote object selection 173
- removing locks 83
- replication
- configuration, backing up 224
 - utilities 221
- resources
- Add New Resource Set Window 147, 156, 160
 - adding 147
 - adding new members 147, 152, 156, 160
 - creating 146
 - resource list 147
 - specifying 149
 - specifying groups (gid) 155
 - UNIX restrictions 147
- restore_config 224, 295, 296
- rootkit, special considerations for 171
- S**
- secfs
- description 48
 - file type 48
 - location 49
 - process log 48
 - status 49
- secfs.log 240
- security rules
- access methods 167
 - access permissions 169
 - adding processes 159
 - default 171
 - excluding users 161, 166
 - when 163
 - who 155
- Security Server
- firewall configuration 28
 - TCP/IP port numbers 28
- sendmail
- allowed process 159
- server.log 238
- servers
- backup 223
 - installing failover 38
- setting
- trust 105
 - trustfrom 105
- setup_eet_agent_cert_db2 235, 297
- single-user mode 83
- special characters 147, 156, 160
- SQL
- backup 224
- SSH
- sshd 299
- ssh_setup 41
- sshd daemon 299
- SSL
-

- affect on policies 133
- case sensitivity 133
- certificate exchange 39

start directory 174

start_replication 221, 223, 224, 225

stop_replication 221, 224, 225

Store on Server 98

sudo 29

synchronize_replication 227, 282

syntax

- special characters 147, 156, 160

sysconfig 300

system clocks 29

system lock

- enabling and disabling 85

System Locked 84

T

time range

- when in Security Rule 146

time sets 163

Trace 97

trust 105

trustfrom 105

U

uid 155, 156

uname 155

user ID (uid) 156

user name (uname) variable 177

users

- adding new 72

- admin 36

- administrative roles 36

- changing passwords 71

- creating user definitions 155

- excluding 161, 166

- login with expired password 72

- new user password 73

- one admin at a time 57

- password restrictions 70

- uname 155

- user ID (uid) 155

- user roles 69

user set 155, 159

voradmin 222

who 155

V

VCR-style buttons 77, 134, 233

vmd

- description 48

- process 48

- process starting 49

vmmount 49

vmumount 49

voradmin 58, 69, 221

voradmin.CGSSDB.ASN.CAP.log 245

voradmin.CGSSDB.CGSSDB_AQ.APP.log 246

W

warn mode 172

Web browser

- troubleshooting 300

Web browsers, supported 57

when

- time range 163

windows

- Protection Properties 97

- Remote Object Chooser 173

X

X.509 digital certificates 33, 43

Z

zero-state backups 42



Program Number: 5724-R38

Printed in USA

SC19-1275-03

