

## Administration Guide

# Novell® Privileged User Manager

**2.2.2**

July 21, 2010

[www.novell.com](http://www.novell.com)



## Legal Notices

Novell, Inc., makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc., makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the [Novell International Trade Services Web page \(http://www.novell.com/info/exports/\)](http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 2009 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc., has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed on the [Novell Legal Patents Web page \(http://www.novell.com/company/legal/patents/\)](http://www.novell.com/company/legal/patents/) and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.  
404 Wyman Street, Suite 500  
Waltham, MA 02451  
U.S.A.  
[www.novell.com](http://www.novell.com)

*Online Documentation:* To access the latest online documentation for this and other Novell products, see the [Novell Documentation Web page \(http://www.novell.com/documentation\)](http://www.novell.com/documentation).

## **Novell Trademarks**

For Novell trademarks, see [the Novell Trademark and Service Mark list \(http://www.novell.com/company/legal/trademarks/tmlist.html\)](http://www.novell.com/company/legal/trademarks/tmlist.html).

## **Third-Party Materials**

All third-party trademarks are the property of their respective owners.



# Contents

<b>About This Guide</b>	<b>11</b>
<b>1 Welcome to the Framework</b>	<b>13</b>
1.1 Introduction to the Framework	13
1.1.1 System Requirements	13
1.1.2 Framework Manager	14
1.1.3 Framework Manager Console	15
1.1.4 Framework Agent	15
1.2 The Workspace Layout	15
1.2.1 Navigation Path	16
1.2.2 Navigation Pane	16
1.2.3 Task Pane	17
<b>2 Managing Package Distribution</b>	<b>19</b>
2.1 Downloading Packages to a Package Manager	19
2.1.1 Configuring the Package Manager	19
2.1.2 Adding Packages to the Package Manager	20
2.1.3 Checking for Updated Packages	20
2.1.4 Removing Packages	21
2.2 Managing the Workspace	21
2.2.1 Managing the Consoles	21
2.2.2 Adding a Console to the Framework Manager Console	21
2.2.3 Removing Consoles from the Framework Manager Console	21
2.2.4 Updating Consoles in the Framework Manager Console	22
<b>3 Managing Framework Hosts</b>	<b>23</b>
3.1 Managing Domains	23
3.1.1 Creating a Domain	23
3.1.2 Modifying a Domain	24
3.1.3 Deleting a Domain from the Framework	24
3.2 Managing Hosts	25
3.2.1 Adding a Host	25
3.2.2 Viewing Host Details	25
3.2.3 Modifying a Host	26
3.2.4 Moving a Host	27
3.2.5 Deleting a Host	28
3.2.6 Finding a Host	28
3.2.7 Privileged User Manager Databases	28
3.3 Monitoring Hosts	31
3.3.1 Viewing the Host Log	31
3.3.2 Modifying Log Settings	31
3.3.3 Example Rollover Script	32
3.3.4 System Alerts	33
3.3.5 Modifying Alert Settings	33
3.3.6 Viewing the Host Status	34
3.4 Managing Host Packages	35
3.4.1 Finding Packages on Hosts	35
3.4.2 Updating Packages for a Host	35

3.4.3	Rolling Back Packages	36
3.4.4	Committing Packages	36
3.4.5	Registering and Unregistering Packages for a Host	37
3.4.6	Installing Packages on a Host	37
3.4.7	Uninstalling Packages from a Host	38
3.4.8	Modifying Audit Settings for the Audit Manager Package	38
3.4.9	Configuring SMTP Settings for the Messaging Component Package	39
3.5	Increasing the Security When Accessing the Framework Manager Console	39
3.5.1	Requesting a Certificate for the Framework Manager Console	39
3.5.2	Installing a Certificate	40
3.5.3	Modifying the Connector	40
3.6	Troubleshooting	40
3.6.1	Promoting Managers When the Primary Manager Fails	41
3.6.2	Viewing Store and Forward Messages	41
3.6.3	Managing Low Disk Space	42
3.6.4	Restarting the Agent	43
3.6.5	Managing the Registry Cache	44
3.6.6	Time Synchronization	45
<b>4</b>	<b>Managing Framework Users and Groups</b>	<b>47</b>
4.1	Managing Users	47
4.1.1	Configuring Account Settings	47
4.1.2	Adding a Framework User	49
4.1.3	Modifying a Framework User	49
4.1.4	Removing a Framework User Group from a User	56
4.1.5	Deleting a Framework User	56
4.2	Managing Groups	56
4.2.1	Adding a Framework User Group	56
4.2.2	Modifying a Framework User Group	57
4.2.3	Configuring a Help Desk Group	57
4.2.4	Configuring Roles	58
4.2.5	Deleting a Framework User Group	62
4.3	Deploying the Access Control Module	62
4.4	Changing a Framework User's Password	63
<b>5</b>	<b>Command Control</b>	<b>65</b>
5.1	How Does Command Control Work?	65
5.2	Integrating Command Control into User Environments	66
5.2.1	Using usrun with a Command	67
5.2.2	Using rush for Privileged Sessions	68
5.2.3	Using crush for Complete Session Capture	70
5.2.4	Using rush for Complete Session Control	71
5.2.5	Using Shell Scripts	73
5.3	Importing Command Control Configuration Data	73
5.3.1	Importing Command Control Settings	73
5.3.2	Exporting Command Control Settings	74
5.3.3	Importing Command Control Samples	74
5.4	Command Control Transactions	75
5.4.1	Enabling Transactions and Configuring Settings	75
5.4.2	Making Command Control Configuration Changes with Transactions Enabled	75
5.4.3	Committing a Transaction	76
5.5	Configuring Command Control	76
5.5.1	Defining Audit Settings	77
5.5.2	Backing Up and Restoring	78
5.5.3	Finding a Reference	79

5.5.4	Defining Custom Attributes	79
5.5.5	Functions	79
5.5.6	Adding a Category	81
5.5.7	Deleting a Category	81
5.6	Rules	81
5.6.1	Adding a Rule	83
5.6.2	Modifying a Rule	83
5.6.3	Setting Conditions for a Rule	84
5.6.4	Removing Conditions for a Rule	85
5.6.5	Configuring Script Arguments and Entities for a Rule	85
5.6.6	Assigning a Script to a Rule	86
5.6.7	Removing Script Arguments and Entities	86
5.6.8	Removing a Script from a Rule	86
5.6.9	Populating LDAP Group Scripts	87
5.6.10	Finding a Rule	89
5.6.11	Moving a Rule	89
5.6.12	Copying a Rule	89
5.6.13	Linking a Rule	90
5.6.14	Deleting a Rule	90
5.6.15	Viewing Pseudocode	90
5.7	Command Control Groups	91
5.7.1	User Groups	91
5.7.2	Host Groups	93
5.7.3	Adding an Account Group	94
5.7.4	Modifying an Account Group	95
5.7.5	Deleting an Account Group	95
5.7.6	Copying a Group	96
5.7.7	Moving a Group	96
5.7.8	Enhanced Access Control	96
5.8	Commands	98
5.8.1	Adding a Command	98
5.8.2	Modifying a Command	99
5.8.3	Setting the Command Risk	101
5.8.4	Removing a Command Risk	101
5.8.5	Copying a Command	101
5.8.6	Moving a Command	102
5.8.7	Deleting a Command	102
5.8.8	Importing Sample Commands	102
5.9	Scripts	102
5.9.1	Adding a Script	103
5.9.2	Modifying a Script	103
5.9.3	Copying a Script	103
5.9.4	Moving a Script	104
5.9.5	Deleting a Script	104
5.9.6	Sample Scripts	104
5.10	Access Times	107
5.10.1	Adding an Access Time	107
5.10.2	Modifying an Access Time	108
5.10.3	Copying an Access Time	108
5.10.4	Moving an Access Time	108
5.10.5	Deleting an Access Time	109
5.11	Command Control Reports	109
5.11.1	Adding a Command Control Report	109
5.11.2	Modifying a Command Control Report	110
5.11.3	Copying a Command Control Report	110
5.11.4	Moving a Command Control Report	111
5.11.5	Deleting a Command Control Report	111
5.12	Test Suites	111

5.12.1	Adding a Test Suite	111
5.12.2	Adding or Modifying a Test Case	112
5.12.3	Running a Test Suite	113
5.12.4	Viewing a Test Suite	114
5.12.5	Modifying a Test Suite	114
5.12.6	Deleting a Test Case	114
5.12.7	Deleting a Test Suite	115
5.12.8	Importing a Test Suite	115
5.12.9	Exporting a Test Suite	115
5.13	Deploying Command Control	116
5.13.1	Command Control Modules	116
5.13.2	Auditing Modules	116
5.13.3	Compliance Auditor Modules	116
5.13.4	Installing Command Control	116
<b>6</b>	<b>Managing Audit Reports</b>	<b>119</b>
6.1	Audit Settings	119
6.2	Encryption Settings	120
6.3	Syslog Settings	120
6.4	Command Control Reports	121
6.4.1	Adding a Report	121
6.4.2	Viewing Report Data	122
6.4.3	Filtering the Viewable Records	122
6.4.4	Modifying General Report Information	124
6.4.5	Selecting Log Files	124
6.4.6	Replaying Keystrokes	124
6.4.7	Removing a Report	125
6.4.8	Generating an Activity Report	125
<b>7</b>	<b>Compliance Auditor</b>	<b>127</b>
7.1	Controlling Access to the Compliance Auditor	127
7.2	Compliance Audit Rules	128
7.2.1	Adding or Modifying an Audit Rule	128
7.3	Compliance Audit Reports	129
7.3.1	Adding or Modifying an Audit Report	130
7.3.2	Sample Command Control Report Template	131
7.3.3	Deleting a Report	135
7.4	Compliance Auditor Records	135
7.4.1	Viewing a Compliance Audit Record	136
7.4.2	Viewing and Editing a Command Control Keystroke Report	137
7.4.3	Viewing a Change Management Audit Record	138
7.4.4	Viewing a Report Audit Record	138
7.4.5	Editing an Audit Record	138
7.4.6	Archiving Records	139
7.4.7	Managing Archived Records	139
7.5	Access Control Levels	140
7.5.1	Adding or Modifying a User ACL	140
7.5.2	Deleting a User ACL	140
7.6	Deploying the Compliance Auditor	141
<b>8</b>	<b>Load Balancing and Failover</b>	<b>143</b>
8.1	Failover	143
8.2	Load Balancing	144

<b>9</b>	<b>Command Control Components</b>	<b>147</b>
<b>10</b>	<b>Command Line Options</b>	<b>149</b>
10.1	The unifi Options	149
10.2	Command Control Options	150
10.2.1	Importing and Exporting Command Control Settings	150
10.2.2	Backing Up and Restoring a Command Control Configuration	151
10.2.3	Running Test Suites	153
10.3	Package Distribution Options	153
10.4	Package Manager Options	154
10.5	Registry Agent Options	154
10.5.1	Registering an Agent	154
10.5.2	Finding a Primary Manager Package	155
10.5.3	Agent Status	155
10.6	Registry Manager Options	155
10.7	Compliance Auditor Options	156
10.7.1	Exporting and Importing Compliance Auditor Settings	156
10.7.2	Managing Compliance Auditor Records	156



# About This Guide

This Administration Guide explains how to use the Framework Manager to control and audit superuser access to Linux and UNIX machines.

- ♦ Chapter 1, “Welcome to the Framework,” on page 13
- ♦ Chapter 2, “Managing Package Distribution,” on page 19
- ♦ Chapter 3, “Managing Framework Hosts,” on page 23
- ♦ Chapter 4, “Managing Framework Users and Groups,” on page 47
- ♦ Chapter 5, “Command Control,” on page 65
- ♦ Chapter 6, “Managing Audit Reports,” on page 119
- ♦ Chapter 7, “Compliance Auditor,” on page 127
- ♦ Chapter 8, “Load Balancing and Failover,” on page 143
- ♦ Chapter 9, “Command Control Components,” on page 147
- ♦ Chapter 10, “Command Line Options,” on page 149

## Audience

This guide is intended for users who manage the Privileged User Manager product.

## Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation, or go to [www.novell.com/documentation/feedback.html](http://www.novell.com/documentation/feedback.html) and enter your comments there.

## Documentation Updates

For the most recent version of the *Privileged User Administration Guide*, visit the [Privileged User Manager Web site \(http://www.novell.com/documentation/privilegedusermanager22\)](http://www.novell.com/documentation/privilegedusermanager22).

## Additional Documentation

*Privileged User Manager Getting Started Guide* ([http://www.novell.com/documentation/privilegedusermanager22/npum\\_install/data/index.html](http://www.novell.com/documentation/privilegedusermanager22/npum_install/data/index.html))



# Welcome to the Framework

# 1

Welcome to the Novell Privileged User Manager, an application that enables immediate central distribution, updates, and management of the tools and services required for a safe and secure operating environment in today's and tomorrow's IT infrastructure.

Privileged User Manager provides complete control of restricted commands on UNIX and Linux platforms, plus fully configurable auditing.

- ♦ [Section 1.1, "Introduction to the Framework," on page 13](#)
- ♦ [Section 1.2, "The Workspace Layout," on page 15](#)

## 1.1 Introduction to the Framework

Novell Privileged User Manager uses a Framework as the base layer to provide an easy-to-use enterprise architecture into which Privileged User Manager modules are added to create the necessary problem-solving functionality. The Framework has several key features:

- ♦ It provides the core functionality needed to implement secure, enterprise-wide services.
- ♦ It provides services such as secure and authenticated communication among the components.
- ♦ It provides integrated databases and logging.
- ♦ It allows the deployment of Privileged User Manager modules to Framework hosts to implement new functionality.
- ♦ With each module that is installed, an additional console is added to the main Framework Manager console to allow access to new administration functionality.

The Framework is made up of three primary components:

- ♦ [Section 1.1.1, "System Requirements," on page 13](#)
- ♦ [Section 1.1.2, "Framework Manager," on page 14](#)
- ♦ [Section 1.1.3, "Framework Manager Console," on page 15](#)
- ♦ [Section 1.1.4, "Framework Agent," on page 15](#)

### 1.1.1 System Requirements

Recommended system requirements specify the minimum prerequisites to run Framework Agent and Framework Manager.

#### Framework Agent Requirements

The minimum requirements for the Framework Agent is as below:

- ♦ 300 MHz (RISC), 1 GHz (CISC) processor
- ♦ 50 MB additional RAM space
- ♦ 100 MB additional hard disk space

## Framework Manager Requirements

The minimum requirements for the Framework Manager is as below:

- ♦ 1 GHz or more (RISC), 2 GHz or more (CISC) processor
- ♦ 250 MB additional RAM space
- ♦ 150 MB additional hard disk space
- ♦ Hard disk space for Audit Storage

---

**NOTE:** Approximate additional space calculation for Audit Storage = (250 KB) X (number of users) X (average sessions per day, which is usually 8).

---

### 1.1.2 Framework Manager

The Framework Manager is the server component of the Framework. It provides a centralized registry, enabling services and administration of the entire Framework from any single point on the enterprise network.

The Framework Manager is administered through the Framework Manager console, using a suitable Web browser with the Adobe Flash Player.

The manager modules are installed on the Framework Manager by default. The modules can also be distributed to other Framework hosts to provide load balancing and failover for the Framework. If there are multiple occurrences of the same type of manager installed on the Framework, they operate in primary and backup roles. Updates to the data controlled by each group of like managers are only updated at the primary manager.

The default manager modules are:

- ♦ **Administration Manager (admin):** Provides the functionality for the Web-based user interface. Framework consoles can be installed on the Administration Agent and used to control product features.
- ♦ **Access Manager (auth):** Maintains a list of Framework user accounts and provides authentication services for the Framework. It needs to be installed with a local Registry Manager in order to create a secure user authentication token.
- ♦ **Audit Manager (audit):** Maintains the repository for all auditing information collected by the Framework.
- ♦ **Command Control Manager (cmdctrl):** Maintains the rule configurations and is responsible for validating user command requests.
- ♦ **Compliance Auditor (secaudit):** Collects samples of all audit data for compliance regulations.
- ♦ **Messaging Component (msgagnt):** Provides the transport mechanism and interacts with e-mail servers to provide reporting functionality.
- ♦ **Package Manager (pkgman):** Manages a repository for Framework packages.
- ♦ **Registry Manager (registry):** Maintains a database of all Framework hosts and modules. Provides certificate-based registration features for the hosts.
- ♦ **Syslog Emitter (syslogemit):** Provides logging of audit information to a syslog server.

### 1.1.3 Framework Manager Console

The Framework Manager console is the default user interface for the Framework. It allows configuration and management of the Framework through a graphical user interface.



For a description of this console, see [Section 1.2, “The Workspace Layout,”](#) on page 15.

### 1.1.4 Framework Agent

The Framework Agent is the client component of the Framework. It is responsible for receiving and carrying out instructions from the Framework Manager on all hosts. The following Framework Agent packages are installed on all Framework hosts:

- ◆ **Registry Agent (regclnt):** Provides a local cached lookup for module locations. The Registry Agent queries the Registry Manager when local cached information is not available or isn't fresh.
- ◆ **Distribution Agent (distrib):** Provides the interface to control the installation and removal of the packages in the Framework. It has methods to install and remove and to list available and updatable packages. The Distribution Agent retrieves packages from the local Package Managers.
- ◆ **Store and Forward Agent (strfwd):** Provides a store and forward mechanism for guaranteed delivery of messages. It is used for various core features such as replication of the manager databases.
- ◆ **Command Control Agent (rexec):** Enables the Framework to control and audit user commands.

## 1.2 The Workspace Layout

The Framework Manager console consists of three areas: a navigation path, a navigation pane, and a task pane.



- ◆ [Section 1.2.1, “Navigation Path,” on page 16](#)
- ◆ [Section 1.2.2, “Navigation Pane,” on page 16](#)
- ◆ [Section 1.2.3, “Task Pane,” on page 17](#)

## 1.2.1 Navigation Path

The navigation path near the top center of the screen shows the current position in the Framework Manager console.



Click an item on the navigation path for quick access to a given navigation pane. For example, to return to the home page, click *Home*.

## 1.2.2 Navigation Pane

The navigation pane on the right of the screen provides the current administrative facilities, consisting of icons, data grids, and forms.

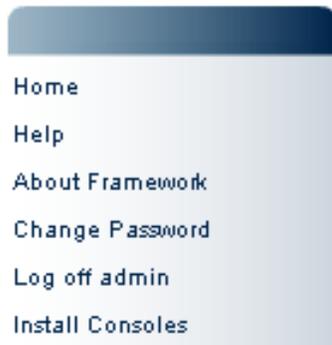


In the navigation pane, you have access to six administrative consoles:

- ♦ **Compliance Auditor:** Proactive auditing tool that pulls events from the event logs for analysis, according to predefined rules. It pulls filtered audit events at hourly, daily, weekly or monthly intervals. This enables auditors to view prefiltered security transactions, play back recordings of user activity, and record notes for compliance purposes. In an era of increasing regulatory compliance, the ability to supply demonstrable audit compliance at any time provides a more secure system and reduces audit risk. For more information, see [Chapter 7, “Compliance Auditor,” on page 127](#).
- ♦ **Framework User Manager:** Manages users who log in to the Framework Manager through role-based grouping. For more information, see [Chapter 4, “Managing Framework Users and Groups,” on page 47](#).
- ♦ **Hosts:** Centrally manages Privileged User Manager installation and updates, load-balancing, redundancy of resources, and host alerts. For more information, see [Chapter 3, “Managing Framework Hosts,” on page 23](#).
- ♦ **Reporting:** Provides easy access and search capability for event logs and allows you review and color-code user keystroke activity through the Command Risk Analysis Engine. For more information, see [Chapter 6, “Managing Audit Reports,” on page 119](#).
- ♦ **Command Control:** Uses an intuitive graphical interface to manage security policies for privilege management. For more information, see [“Command Control” on page 65](#).
- ♦ **Package Manager:** Lets you easily update any Privileged User Manager application. For more information, see [Chapter 2, “Managing Package Distribution,” on page 19](#).

### 1.2.3 Task Pane

The task pane on the left of the screen contains options that are applicable to the current Framework Manager console display.



The items in the top frame change, depending upon what is selected in the navigation pane.

# Managing Package Distribution

# 2

- ♦ [Section 2.1, “Downloading Packages to a Package Manager,” on page 19](#)
- ♦ [Section 2.2, “Managing the Workspace,” on page 21](#)

## 2.1 Downloading Packages to a Package Manager

To update packages to your Framework hosts, you must first download them to a Package Manager, then update individual hosts or domains with the packages.

There are two options for downloading packages to a Package Manager:

- ♦ You can download packages directly from the Novell Update Server.
- ♦ You can download packages from the Novell Update Server onto a local server, then download packages from this local server to your Framework hosts.

You must configure the Package Manager to access the server you require.

- ♦ [Section 2.1.1, “Configuring the Package Manager,” on page 19](#)
- ♦ [Section 2.1.2, “Adding Packages to the Package Manager,” on page 20](#)
- ♦ [Section 2.1.3, “Checking for Updated Packages,” on page 20](#)
- ♦ [Section 2.1.4, “Removing Packages,” on page 21](#)

### 2.1.1 Configuring the Package Manager

You must supply the Package Manager with a location for downloading the packages before you can add packages for distribution.

- 1 Click *Package Manager* on the home page of the console.
- 2 Click *Settings* in the task pane.
- 3 (Conditional) To use a local server:
  - 3a Select *Local Package Manager*.
  - 3b Fill in the following fields:
    - Host name:** Specify the DNS name of the host.
    - Port:** Specify the communication port. The default is 29120.The local server is a Framework host that has been configured to store the packages.
- 4 (Conditional) To use an update server:
  - 4a Select *Novell Update Server*.
  - 4b Specify the username and password for your Novell Customer Center account.
  - 4c To view the update server information, select *Advanced Settings*.
    - ♦ When *Packages* is selected, the following URL is configured:  
`https://nu.novell.com:443/PUM/packages`

Use this URL to download update packages.

- ♦ When *Eval* is selected, the following URL is configured:

`https://nu.novell.com:443/PUM/eval`

Use this URL to download evaluation packages.

5 Click *Finish*.

6 Continue with [Section 2.1.2, “Adding Packages to the Package Manager,”](#) on page 20.

## 2.1.2 Adding Packages to the Package Manager

If you have configured the Package Manager to use an update server (see [Section 2.1.1, “Configuring the Package Manager,”](#) on page 19), you use the following procedure to add packages to the Package Manager. If you downloaded the packages to a directory, you can use a command line option to add the packages. See [Section 10.3, “Package Distribution Options,”](#) on page 153.

1 Click *Package Manager* on the home page of the console.

2 Click *Add Packages* in the task pane.

3 Set the *Package Filter* options:

**Versions:** Select 2.2.

**Platforms:** Select the operating system, then use the arrow to display and select the platforms.

**Types:** Select the types of packages you want to add (consoles, modules, interfaces, patches).

**Components:** Select the components (command control, framework, miscellaneous)

4 Select the packages from the list of available packages.

To select multiple packages, press the Ctrl key and select the packages one at a time, or press the Shift key to select a consecutive list of packages. To select all packages, use Ctrl+A.

5 Click *Next* to start downloading.

6 Click *Finish*.

7 To install these packages on your hosts, continue with [Section 3.4.2, “Updating Packages for a Host,”](#) on page 35.

8 To use this Package Manager as a local server for downloading packages, configure other Package Managers with the DNS name of this host.

## 2.1.3 Checking for Updated Packages

After you have configured the package manager for a group of packages, you can use the *Check for Updates* option to see if any updates are available.

1 Click *Package Manager* on the home page of the console.

2 Click *Check for Updates* in the task pane.

The navigation pane displays updated packages available for download.

3 Select the packages from the list of available packages.

To select multiple packages, press the Ctrl key and select the packages one at a time, or press the Shift key to select a consecutive list of packages. To select all packages, use Ctrl+A.

4 Click *Next* to start downloading.

- 5 Click *Finish*.
- 6 If updates were available, continue with [Section 3.4.2, “Updating Packages for a Host,” on page 35](#) to install these packages on your hosts.

## 2.1.4 Removing Packages

- 1 Click *Package Manager* on the home page of the console.
- 2 In the list of available packages, select the packages you want to remove  
To select multiple packages, the Ctrl key and select the packages one at a time, or the Shift key to select a consecutive list of packages. To select all packages, use Ctrl+A.
- 3 Click *Remove Packages* in the task pane.
- 4 Click *Next* to start removing packages.
- 5 Click *Finish*.

## 2.2 Managing the Workspace

- ♦ [Section 2.2.1, “Managing the Consoles,” on page 21](#)
- ♦ [Section 2.2.2, “Adding a Console to the Framework Manager Console,” on page 21](#)
- ♦ [Section 2.2.3, “Removing Consoles from the Framework Manager Console,” on page 21](#)
- ♦ [Section 2.2.4, “Updating Consoles in the Framework Manager Console,” on page 22](#)

### 2.2.1 Managing the Consoles

The Framework Manager console can be extended by installing console packages. Console packages provide the administrative and reporting panes for the Privileged User Manager modules.

Console packages must be downloaded to the Package Manager before they become available for installation.

### 2.2.2 Adding a Console to the Framework Manager Console

- 1 Click *Install Consoles* on the home page of the console.  
If the list is empty, you have installed all available consoles.
- 2 In the list of available consoles, select the consoles you want to add.  
To select multiple consoles, press the Ctrl key and select the consoles one at a time, or press the Shift key to select a consecutive list of consoles. To select all consoles, use Ctrl+A.
- 3 Click *Next* to start installing.
- 4 Review the list of installed consoles.
- 5 Click *Finish*.

### 2.2.3 Removing Consoles from the Framework Manager Console

- 1 Click *Uninstall Consoles* on the home page of the console.

**2** In the list of available consoles, select the consoles you want to remove.

To select multiple consoles, press the Ctrl key and select the consoles one at a time, or press the Shift key to select a consecutive list of consoles. To select all consoles, use Ctrl+A.

**3** Click *Next* to start uninstalling.

**4** Review the list of removed consoles.

**5** Click *Finish*.

## 2.2.4 Updating Consoles in the Framework Manager Console

When updated console packages become available on your Package Manager, you can update the console packages installed on your Framework Manager console.

**1** Click *Update Consoles* in the task pane.

The navigation pane displays updated consoles available for deployment.

**2** In the list of available consoles, select the consoles you want to update.

To select multiple consoles, press the Ctrl key and select the consoles one at a time, or press the Shift key to select a consecutive list of consoles. To select all consoles, use Ctrl+A.

**3** Click *Next* to start installing.

**4** Review the list of updated consoles.

**5** Click *Finish*.

---

**NOTE:** After updating a console, you must shut down and reopen the Framework Manager console to see the changes.

---

# Managing Framework Hosts

# 3

The Hosts console provides a hierarchical view of all currently defined hosts. Each host machine on which you have installed managers and agents must be added to the Framework Manager console through the Hosts console. Hosts are identified to the Framework Manager console by a unique agent name that is used to register the manager or agent after installation.

Hosts are added to domains, and domains allow you to organize your hosts into a tree structure. If Privileged User Manager is configured to manage 40, 100, or 400 hosts, the hierarchical structure allows for easy access to a particular host.

- ◆ [Section 3.1, “Managing Domains,” on page 23](#)
- ◆ [Section 3.2, “Managing Hosts,” on page 25](#)
- ◆ [Section 3.3, “Monitoring Hosts,” on page 31](#)
- ◆ [Section 3.4, “Managing Host Packages,” on page 35](#)
- ◆ [Section 3.5, “Increasing the Security When Accessing the Framework Manager Console,” on page 39](#)
- ◆ [Section 3.6, “Troubleshooting,” on page 40](#)

## 3.1 Managing Domains

Privileged User Manager provides load balancing and failover capabilities, based on the hierarchical structure of your hosts. Before organizing your hosts into domains and subdomains, refer to [Chapter 8, “Load Balancing and Failover,” on page 143](#) for information on how these features work.

- ◆ [Section 3.1.1, “Creating a Domain,” on page 23](#)
- ◆ [Section 3.1.2, “Modifying a Domain,” on page 24](#)
- ◆ [Section 3.1.3, “Deleting a Domain from the Framework,” on page 24](#)

### 3.1.1 Creating a Domain

When you install Privileged User Manager, a top-level domain called Hosts is automatically created for you. To rename this domain, see [Section 3.1.2, “Modifying a Domain,” on page 24](#). Under this top-level domain, you can create subdomains.

- 1 Click *Hosts* on the home page of the console.
- 2 To add a new subdomain to an existing domain, select the existing domain.
- 3 Click *Add Domain* in the task pane.
- 4 Specify a domain name.
- 5 Click *Finish*.
- 6 Select from the following tasks:
  - ◆ To configure the domain, continue with [Section 3.1.2, “Modifying a Domain,” on page 24](#).

- ♦ To add hosts to the domain, continue with [Section 3.2.1, “Adding a Host,” on page 25.](#)
- ♦ To move existing hosts to this domain, continue with [Section 3.2.4, “Moving a Host,” on page 27.](#)

## 3.1.2 Modifying a Domain

Use this page to modify the domain name and to modify encryption options. The encryption settings apply to all hosts within the domain, unless you modify the host encryption settings. Host settings overwrite domain settings.

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the domain you want to modify.
- 3 In the task pane, click *Modify Domain*.
- 4 Configure the following options:

**Domain name:** Specify a new domain name.

**Key configuration:** Select this option to enable configuration of the encryption key and encryption of the databases stored on the hosts in this domain.

**Host Key rollover (days):** Specify how many days the host key can be used before generating a new key for the hosts in this domain.

**DB Key rollover (days):** Specify how many days the database key can be used before generating a new key for the hosts in this domain.

**Encrypt:** Select the databases you want to encrypt for the hosts in this domain.

Use care in selecting the databases you enable for encryption. Encrypting the data can affect performance. Novell<sup>®</sup> recommends the following:

- ♦ `auth.db` because it contains usernames
- ♦ `registry.db` because it contains the hosts.
- ♦ `cmdctrl.db` because it contains command control rules with usernames and hosts.

For a brief description of databases, see [Section 3.2.7, “Privileged User Manager Databases,” on page 28.](#)

The encryption of auditing data should be enabled from the Reporting console. See [Section 6.1, “Audit Settings,” on page 119.](#)

- 5 Click *Finish*.

## 3.1.3 Deleting a Domain from the Framework

You cannot delete a domain if it contains any hosts. You must delete or move the hosts first. See [Section 3.2.4, “Moving a Host,” on page 27.](#)

---

**IMPORTANT:** This action cannot be undone.

---

- 1 Click *Hosts* on the home page of the console.  
The navigation pane displays the current hierarchy for your Framework.
- 2 In the navigation pane, select the domain you want to delete.

- 3 In the task pane, click *Delete Domain*.
- 4 Click *Finish*.

## 3.2 Managing Hosts

A host is created for each machine you want to manage with Privileged User Manager. You need to create the host before you can register the agents that you install on the machines. If you have created a text file that lists the domains that you have created and the hosts that you want in each domain, you can cut and paste from the text file to add the hosts to the domains.

- ♦ [Section 3.2.1, “Adding a Host,” on page 25](#)
- ♦ [Section 3.2.2, “Viewing Host Details,” on page 25](#)
- ♦ [Section 3.2.3, “Modifying a Host,” on page 26](#)
- ♦ [Section 3.2.4, “Moving a Host,” on page 27](#)
- ♦ [Section 3.2.5, “Deleting a Host,” on page 28](#)
- ♦ [Section 3.2.6, “Finding a Host,” on page 28](#)
- ♦ [Section 3.2.7, “Privileged User Manager Databases,” on page 28](#)

### 3.2.1 Adding a Host

- 1 Click *Hosts* on the home page of the console.
- 2 Select the domain for the new host.
- 3 Click *Add Hosts* from the task pane.
- 4 In the text box, specify the agent names for the hosts.

You can type the names one at a time, using one name per line, or you can paste a list of names. When you add a host to the Framework Manager, the name does not need to relate to the existing DNS name used to locate the host on your network.

- 5 Click *Next*.

A list of agents names is displayed.

- 6 Click *Finish*.

The status of the host is unregistered until the agent is installed and registered on the host machine. For instructions on this process, see “[Installing and Registering a Framework Agent](#)”.

### 3.2.2 Viewing Host Details

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the domain containing the hosts whose details you want to view.
- 3 Click the arrow ► next to the domain icon to display the hosts on the left side of the navigation pane.
- 4 Click the host icon to display the host details and status.

Field	Description
Agent name	The agent name configured for this host.
Host name	The name of the host. This is either a resolvable DNS name or the IP address.
Host port	The port the host is using for Privileged User Manager communication.
Platform	The operating system on the host.
Processor	The type of processor on the host.
OS Version	The version of the kernel running on the host.
Agent version	The version of the agent software that the host is running.
System time	The current date and time that the host is configured for, displayed in UTC.  Use this time to verify that the agent's time is synchronized with the other hosts.
Service uptime	The number of days, hours, minutes, and seconds the agent has been running since the last start up.
Active sessions	The number of connections currently open between the agent and any other agent, including itself.
Active tasks	The number of internal tasks that the agent is running at any one time.
Installation path	The directory location of the installed agent software.
Disk space	The total amount of available disk space, the amount of free disk space, and the percentage of disk space in use.
Memory (approx)	The amount of memory (heap) currently being used by the agent to store its data.  This is the virtual data segment size minus the thread stack and the statically initialized data (because this is a constant value) as returned by the <code>sbrk</code> system call.
Registration	The licensing state of the software, either licensed or unlicensed.
Status	The status of the host: online, offline, unregistered.

**5** Click the arrow ► next to the host icon to display the *Packages* icon.

**6** Click *Packages* to view details of the packages installed on this host.

### 3.2.3 Modifying a Host

**1** Click *Hosts* on the home page of the console.

**2** In the navigation pane, select the host you want to modify.

**3** In the task pane, click *Modify Host*.

**4** Modify the general details:

**Agent name:** Specify a display name for this agent.

The agent name does not need to relate to the existing DNS name used to locate the host on your network.

**Description:** Add a description. This description is displayed next to the agent name in the hierarchical view.

**Host name:** Specify the DNS name used to locate the host on your network or the IP address of the host.

**Port:** Displays the port that was specified when the agent was registered.

- 5 Modify the encryption options. When these settings are modified for an individual host, the host settings overwrite the settings specified for the domain.

**Key configuration:** Select this option to enable configuration of the encryption key.

**Host Key rollover (days):** Specify how many days the host key can be used before generating a new key.

**DB Key rollover (days):** Specify how many days the database key can be used before generating a new key.

**Encrypt:** Select the databases you want to encrypt.

Use care in selecting the databases you enable for encryption. Encrypting the data can affect performance. Novell recommends the following:

- ♦ auth.db because it contain usernames
- ♦ registry.db because it contains the hostnames.
- ♦ cmdctrl.db because it contains command control rules with usernames and hostnames.

For a brief description of databases, see [Section 3.2.7, “Privileged User Manager Databases,” on page 28](#).

The encryption of auditing data should be enabled from the Reporting console. See [Section 6.1, “Audit Settings,” on page 119](#).

- 6 Click *Finish*.

### 3.2.4 Moving a Host

You can move hosts among the domains.

- 1 Click *Hosts* on the home page of the console.

The navigation pane displays the current hierarchy for your Framework.

- 2 In the navigation pane, select the domain containing the hosts you want to move. The hosts in that domain are displayed on the right side of the navigation pane.
- 3 Select the hosts you want to move.

To select multiple hosts, press the Ctrl key and select the hosts one at a time, or press the Shift key to select a consecutive list of hosts. To select all hosts in a domain, use Ctrl+A.

- 4 Drag the hosts to the new domain.

## 3.2.5 Deleting a Host

---

**IMPORTANT:** This action cannot be undone.

---

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the domain containing the hosts you want to delete. The hosts in that domain are displayed on the right side of the navigation pane.
- 3 Select the hosts you want to delete.  
To select multiple hosts, press the Ctrl key and select the hosts one at a time, or press the Shift key to select a consecutive list of hosts. To select all hosts in a domain, use Ctrl+A.
- 4 In the task pane, click *Delete Host*.  
The selected hosts are listed.
- 5 Click *Finish*.

## 3.2.6 Finding a Host

- 1 Click *Hosts* on the home page of the console.
- 2 Click *Hosts* or a domain name.
- 3 In the task pane, click *Find Host*.
- 4 In the *Agent name* field, specify the name of the host you are looking for.  
You can use the wildcard characters \* and ?. For example, entering h\* finds all hosts with agent names beginning with h. This field is case sensitive.
- 5 Click *Find*.
- 6 To go to a host's details, double-click its agent name.
- 7 To return the Hosts page, click *Close*.

## 3.2.7 Privileged User Manager Databases

The following databases are created on the Framework Manager console machine.

Location of the database files for SLES compatible PUM is `/var/opt/novell/npum/service/local/`. Location of the database files for other UNIX platforms is `/opt/novell/npum/service/local/`, which is shown in the table below.

Database and Standard Location	Description
admin.db <code>/opt/novell/npum/service/local/admin/</code>	Not used.
admin.ldb <code>/opt/novell/npum/service/local/admin/</code>	Not used.
audit.db <code>/opt/novell/npum/service/local/audit/</code>	Contains all configured report definitions and the settings for roll over.

Database and Standard Location	Description
audit.ldb /opt/novell/npum/service/local/audit/	Contains role history and the metadata for audit logs.
auth.db /opt/novell/npum/service/local/auth/	Contains fully replicated authorization data including user details and settings for access to the Framework Manager console.
auth.ldb /opt/novell/npum/service/local/auth/	Not used.
cmdctrl.db /opt/novell/npum/service/local/cmdctrl/	Contains rules and configuration for Command Control.
cmdctrl.ldb /opt/novell/npum/service/local/cmdctrl/	Not used.
cmdctrl.db /opt/novell/npum/service/local/audit/	Contains the current audit log.
cmdctrl<timestamp>.ldb /opt/novell/npum/service/local/audit/	Contains the archived audit log.
distrib.db /opt/novell/npum/service/local/distrib/	Not used.
distrib.ldb /opt/novell/npum/service/local/distrib/	Not used.
msgagnt.db /opt/novell/npum/service/local/msgagnt/	Contains SMTP configuration information.
msgagnt.ldb /opt/novell/npum/service/local/msgagnt/	Not used.
pkgman.db /opt/novell/npum/service/local/pkgman/	Contains the metadata for the packages stored locally on the package manager for deployment to the Framework. It also contains replicated settings for the location to download packages from.
pkgman.ldb /opt/novell/npum/service/local/pkgman/	Contains local data, such as the location of the local package repository.

<b>Database and Standard Location</b>	<b>Description</b>
regclnt.db  /opt/novell/npum/service/local/ registry/	Not used.
regclnt.ldb  /opt/novell/npum/service/local/ registry/	Local cache of the registry.db.
registry.db  /opt/novell/npum/service/local/ registry/	Contains public keys, hostnames, and access ports.
registry.ldb  /opt/novell/npum/service/local/ registry/	Contains data used to manage registry agent caches.
rexec.db  /opt/novell/npum/service/local/rexec/	Not used.
rexec.ldb  /opt/novell/npum/service/local/rexec/	Not used.
secaudit.db  /opt/novell/npum/service/local/ secaudit/	Contains rule configuration and audit data.
secaudit.ldb  /opt/novell/npum/service/local/ secaudit/	Not used.
strfwd.db  /opt/novell/npum/service/local/ strfwd/	Not used.
strfwd.ldb  /opt/novell/npum/service/local/ strfwd/	Contains messages intended for and received from other hosts connected to the Framework until they are acknowledged or processed by the appropriate local module.
syslogemit.db  /opt/novell/npum/service/local/ syslogemit/	Contains configuration information.
syslogemit.ldb  /opt/novell/npum/service/local/ syslogemit/	Not used.

## 3.3 Monitoring Hosts

Privileged User Manager maintains a log file for each host, can be configured to send alerts to the Framework Manager console when errors occur, and allows you to view the status of each host:

- ◆ [Section 3.3.1, “Viewing the Host Log,” on page 31](#)
- ◆ [Section 3.3.2, “Modifying Log Settings,” on page 31](#)
- ◆ [Section 3.3.3, “Example Rollover Script,” on page 32](#)
- ◆ [Section 3.3.4, “System Alerts,” on page 33](#)
- ◆ [Section 3.3.5, “Modifying Alert Settings,” on page 33](#)
- ◆ [Section 3.3.6, “Viewing the Host Status,” on page 34](#)

### 3.3.1 Viewing the Host Log

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the host whose details you want to view.
- 3 In the task pane, click *View Host Log*.
- 4 Select the information you want to view:

**Log level:** Set the level of information you want to see on the screen.

- ◆ *Error* displays only Error messages.
- ◆ *Warning* displays Warning and Error messages.
- ◆ *Information* displays Information, Warning, and Error messages.

**Refresh (secs):** Set the interval in seconds between screen refreshes. You can select intervals from 1 to 60 seconds

**Maximum cached log messages:** Set the maximum number of log messages to display on the screen. You can view from 10 to 1000 messages.

- 5 Click the *Pause* check box to pause the screen display.
- 6 Click the *Clear* button to clear the screen display.
- 7 Click *Close* to return to the Framework hierarchy view.

### 3.3.2 Modifying Log Settings

You can modify log settings for all hosts, all hosts in a domain, or an individual host by using the *Domain Log Settings* or *Host Log Settings* options.

- 1 Click *Hosts* on the home page of the console.
- 2 To modify the log settings for all hosts in a domain, select the domain. To modify the log settings for an individual host, select the host in the navigation pane.
- 3 Click *Domain Log Settings* or *Host Log Settings* in the task pane, then modify the following settings:

**Filename:** Specify the filename and location of the log file.

**Level:** Set the level of information you need:

- ◆ *Error* displays only Error messages.

- ♦ *Warning* displays Warning and Error messages.
- ♦ *Info* displays Information, Warning, and Error messages.
- ♦ *Debug* displays Debug, Information, Warning, and Error messages.
- ♦ *Trace* displays Trace, Debug, Information, Warning, and Error messages.

The *Debug* and *Trace* settings are primarily for the use of Novell Support.

**Show all tasks:** Click *Show all tasks* to have the log show all tasks. The *Show all tasks* option is primarily for the use of Novell Support.

**Rollover:** Select the rollover point from the drop-down list to specify when the log file is overwritten with new information. If the maximum size set for the log file is reached, the log file is overwritten regardless of this setting.

**Max Size (MB):** Select the maximum size of the log file from the drop-down list to specify when the log is overwritten with new information.

**Rollover Script:** Enter a Perl script to be executed at the rollover point. For a sample script, see [“Example Rollover Script” on page 32](#).

- 4 Click *Next* to apply the changes.

If the changes are applied successfully, a green box  is shown next to the agent name.

If the changes are not applied successfully (for example, if the host is not online), a red box  is shown next to the agent name.

- 5 Click *Close*.

### 3.3.3 Example Rollover Script

This is an example of a Perl script that can be called at the rollover point for the host log file. The script compresses the old `unifid.log` and then removes any log files that are more than 30 days old.

```
use File::Basename;
system("/usr/bin/gzip $LOG_FILE");
my $log_root = dirname($LOG_FILE);
opendir(LOGDIR, $log_root);
# Find all the compressed log files
my @log_files = map { $_->[1] }
                 map { [ $_, "$log_root/$_" ] }
                 grep { /\.gz$/ }
                 readdir(LOGDIR);

closedir(LOGDIR);
# Delete all log files older than 30 days
my $time = time();
foreach my $log (@log_files) {
    my ($mtime) = (stat($log))[9];
    my $age = int((($time - $mtime) / 3600) / 24);
    $ctx->log_info("Checking $log $age");
    next unless $age > 30;

    $ctx->log_info("Deleting $log ($age days old)");
    unlink $log;
}
```

### 3.3.4 System Alerts

The System Alerts page shows system status alert messages from all hosts in your Framework. The page shows the time of the alert, the host that originated the alert, the type of alert, and information about the alert.

You can define the level of system alerts you want by using the *Domain Alert Settings* or *Host Alert Settings* options.

The existence of system alerts is indicated by a flashing Framework icon in the bottom right corner of the screen.

- 1 Click the icon to display the System Alerts page.
- 2 To clear the existing alerts, click *Finish*.
- 3 To close the System Alerts page without clearing the existing alerts, click *Cancel*.

The Framework icon continues to flash.

### 3.3.5 Modifying Alert Settings

You can configure your Framework hosts to generate system status alert messages when specific events occur, such as when the agent exceeds a specified memory usage. If a system alert is triggered, the Framework icon in the bottom right corner of the screen flashes. To view the system alerts, click the icon (see “[System Alerts](#)” on page 33 for details).

You can modify alert settings for all hosts, all hosts in a domain, or an individual host by using the *Domain Alert Settings* or *Host Alert Settings* options.

To modify alert settings:

- 1 Click *Hosts* on the home page of the console.
- 2 To modify alert settings for all hosts in a domain, select *Hosts* or the name of a domain. To modify alert settings for an individual host, select the host in the navigation pane.
- 3 In the task pane, click *Domain Alert Settings* or *Host Alert Settings*.

Changes made to a domain’s alert settings override the current settings for individual hosts in that domain. However, subsequent changes made to an individual host’s alert settings override the current domain alert settings on that host.

- 4 Configure the following options:

**Alert on log level:** Select the level of log information needed to trigger an alert. For example, if you want alerts to be triggered when error messages occur in the log, select *Error*. The *Warning* option includes *Warning* and *Error* messages. The *Info* option includes *Info*, *Warning*, and *Error* messages. Select *Never* to switch this setting off.

**Alert log filter:** Define a specific message you want to trigger alerts, or part of a message with wildcard symbols \*. You can use regular expressions in this field by selecting the *Regular expression* check box and specifying your regular expression.

This setting is independent of the setting in *Alert on log level*.

**Time offset (mins):** Specify the time offset in minutes when you want to trigger an alert. An alert is triggered if a host’s time setting differs from the time setting of the Primary Registry Manager by this number of minutes. Time offsets can cause problems because certificates are time-based. The UTC (Universal Time Coordinated) value is used.

**Pending messages (mins):** Specify the interval for when you want to trigger an alert. An alert is triggered if an event has been in the queue of store and forward messages for this number of minutes.

**Maximum memory (MB):** Specify the amount of memory in MB that you want as the threshold for an alert. An alert is triggered if a host is using more than this amount of memory.

**Minimum disk space (MB):** Specify the minimum amount in MB of disk space that you want as the threshold of an alert. An alert is triggered if a host has less than this amount of disk space remaining in the default installation location.

**Expired certificate:** Select this option to cause an alert to be triggered when an agent's certificate expires.

- 5 Click *Next* to apply the settings to the hosts.

If the settings are applied successfully, the indicator next to the hostname is green .

If the settings are not applied successfully (for example, if the host is offline), the indicator is red .

- 6 Click *Close*.

If any of your settings cause an alert to be triggered, the Framework icon flashes.

### 3.3.6 Viewing the Host Status

The *Host Status* option allows you to view the current status of all your hosts, or all the hosts in a domain, on one page.

- 1 Click *Hosts* on the home page of the console.
- 2 Select a domain.
- 3 Click *Host Status* in the task pane.

The status for each host is displayed, as shown in the following table, with a summary at the bottom of the screen.

---

	The host is online.
	There is a status problem with the host; for example, the host's time offset exceeds the defined level (see <a href="#">Step 5</a> ). Click the arrow to the left of the green box to display status messages.
	The host is offline.
	The host is unregistered.

---

- 4 Use the *Online*, *Offline* and *Unregistered* check boxes to select the hosts you want to view.

If you have a long list of hosts, deselect the *Auto scroll* check box to stop the automatic scrolling.

- 5 (Optional) Change the filter settings from the default values and select *Restart* to check the status again. The filters available are:

**Maximum time offset (minutes):** The difference in system time between the host and the Primary Registry Manager. If the time offset exceeds the value in this field, a warning indicator is displayed.

**Minimum disk space (MB):** If the available disk space on the host machine goes below the value in this field, a warning indicator is displayed.

**Maximum Memory (MB):** If the memory used by the host exceeds the value in this field, a warning indicator is displayed.

**6** To view a host's details, double-click the host or click *Close* to return to the hierarchical view.

To use a command line option to view the status, see [Section 10.5.3, "Agent Status," on page 155](#).

## 3.4 Managing Host Packages

- ♦ [Section 3.4.1, "Finding Packages on Hosts," on page 35](#)
- ♦ [Section 3.4.2, "Updating Packages for a Host," on page 35](#)
- ♦ [Section 3.4.3, "Rolling Back Packages," on page 36](#)
- ♦ [Section 3.4.4, "Committing Packages," on page 36](#)
- ♦ [Section 3.4.5, "Registering and Unregistering Packages for a Host," on page 37](#)
- ♦ [Section 3.4.6, "Installing Packages on a Host," on page 37](#)
- ♦ [Section 3.4.7, "Uninstalling Packages from a Host," on page 38](#)
- ♦ [Section 3.4.8, "Modifying Audit Settings for the Audit Manager Package," on page 38](#)
- ♦ [Section 3.4.9, "Configuring SMTP Settings for the Messaging Component Package," on page 39](#)

### 3.4.1 Finding Packages on Hosts

You can search through all your hosts to find where a specific package is installed, or find whether the package is installed at all.

- 1** Click *Hosts* on the home page of the console.
- 2** Click *Hosts* or a domain name.
- 3** In the task pane, click *Find Package*.
- 4** From the *Package* drop-down list, select the package you are looking for.
- 5** If you want to find out which hosts do not have this package installed, select the *package not installed* check box. If you want to find out where the package is installed, deselect the *package not installed* check box.
- 6** Click *Find*. A list of agents where the package is installed or not installed is displayed.
- 7** To view a host's details, double-click the agent name. Otherwise, to return to the Hosts page, click *Close*.

### 3.4.2 Updating Packages for a Host

When updated packages become available on your local Package Manager, you can update the packages installed on individual hosts or for all hosts in a domain. The distribution agent performs a number of checks to ensure that the host has enough disk space to extract and install the package. If there is insufficient space, the package is not updated.

- 1** Click *Hosts* on the home page of the console.

- 2 In the navigation pane, select either the host or the domain where you want to update the packages.
- 3 Click *Update Domain Packages* or *Update Packages* in the task pane.  
The console checks for updates on your Package Manager and displays any updated packages available for download.
- 4 Select the packages from the list of available packages.  
To select multiple packages, press the Ctrl key and select the packages one at a time, or press the Shift key to select a consecutive list of packages. To select all packages, use Ctrl+A.
- 5 (Optional) If you want to create a backup of the current installed packages, select *Create backup*.  
If necessary, you can use the *Rollback Packages* option to roll back to the backed-up packages. Only the last backup is stored. Creating backups requires extra space.
- 6 Click *Next* to start downloading the selected packages.
- 7 Click *Finish*.

### 3.4.3 Rolling Back Packages

If you chose to create a backup when updating packages for an individual host or for a domain, you can roll back to the last backup.

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the domain or the host where you want to roll back the packages.
- 3 Click *Rollback Domain Packages* or *Rollback Packages* in the task pane.
- 4 Select the packages from the list of available backed-up packages.  
To select multiple packages, press the Ctrl key and select the packages one at a time, or press the Shift key to select a consecutive list of packages. To select all packages, use Ctrl+A.
- 5 Click *Next* to start rolling back to the previously backed-up packages.
- 6 Click *Finish*.

### 3.4.4 Committing Packages

When packages are updated and the *Create backup* option is enabled, a backup of the current package is created and stored in a backup directory in the working package directory. This allows you to roll back to the previous level if the current package does not perform correctly in your environment. (See [Section 3.4.3, “Rolling Back Packages,” on page 36.](#))

If the current package does perform correctly in your environment, you can commit the package, which frees up disk space by deleting the files in the backup directory. If your hosts have limited disk space, Novell recommends that you commit the packages on all hosts before performing the next update.

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the domain or the host where you want to commit packages.
- 3 Click *Commit Domain Packages* or *Commit Packages* in the task pane.
- 4 Select the packages from the list of available packages.

To select multiple packages, press the Ctrl key and select the packages one at a time, or press the Shift key to select a consecutive list of packages. To select all packages, use Ctrl+A.

- 5 Click *Next* to start the commit process.
- 6 Click *Finish*.

### 3.4.5 Registering and Unregistering Packages for a Host

If you want to stop a package from functioning without removing it completely, you can unregister it. You can then register it again later if necessary. Packages are automatically registered when you add them, so you only need to register them if you have previously unregistered them.

Registered packages are shown with a green check mark: .

Unregistered packages are shown with a red exclamation mark: .

To register or unregister a package for a host:

- 1 Click *Hosts* on the home page of the console.  
The navigation pane displays the current hierarchy for your Framework.
- 2 In the navigation pane, select the domain where you want to register or unregister packages.
- 3 Select the host.
- 4 With the host's packages displayed, select the packages you want to register or unregister.  
To select multiple packages, press the Ctrl key and select the packages one at a time, or press the Shift key to select a consecutive list of packages. To select all packages, use Ctrl+A.
- 5 Click *Register Package* or *Unregister Package* in the task pane.

---

**NOTE:** The Framework Manager console does not refresh automatically. To check whether your packages have been successfully registered or unregistered, go to another screen and then return to the list of packages.

---

### 3.4.6 Installing Packages on a Host

Ensure that the packages have been downloaded to the Framework Package Manager by viewing packages available to deploy. See [Section 2.1, "Downloading Packages to a Package Manager," on page 19](#).

The distribution agent performs a number of checks to ensure that the host has enough disk space to extract and install the package. If there is insufficient space, the package is not installed.

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the domain where you want to install the packages.
- 3 Select the host.
- 4 With the host's *Packages* icon selected, click *Install Packages* in the task pane.
- 5 Select the packages from the list of available packages.

To select multiple packages, press the Ctrl key and select the packages one at a time, or press the Shift key to select a consecutive list of packages. To select all packages, use Ctrl+A.

6 Click *Next* to start installing the selected packages.

7 Click *Finish*.

To use a command line option to install packages on hosts, see [Section 10.4, “Package Manager Options,” on page 154](#).

### 3.4.7 Uninstalling Packages from a Host

You do not need to uninstall a package to disable it. You can disable it by unregistering the package. See [Section 3.4.5, “Registering and Unregistering Packages for a Host,” on page 37](#).

To uninstall a package:

1 Click *Hosts* on the home page of the console.

2 In the navigation pane, select the domain where you want to uninstall packages.

3 Select the host.

4 With the host’s *Packages* icon selected, select the packages from the list of installed packages.

To select multiple packages, press the Ctrl key and select the packages one at a time, or press the Shift key to select a consecutive list of packages. To select all packages, use Ctrl+A.

5 Click *Uninstall Packages* in the task pane.

6 Click *Next* to start uninstalling the selected packages.

7 Click *Finish*.

To use a command line option to install packages on hosts, see [Section 10.4, “Package Manager Options,” on page 154](#).

### 3.4.8 Modifying Audit Settings for the Audit Manager Package

The databases containing audited data from command control (`cmdndctrl.db`) can be placed in an alternative location. The administration audit files (`audit.db` and `audit.ldb`) and `log.msqs` are still stored in the default location `/opt/novell/service/local/audit` or `C:\Program Files\Novell\npum\service\local\audit`, but these files are relatively small.

To define an alternative location for the audit databases:

1 Click *Hosts* on the home page of the console.

2 In the navigation pane, select the host with the Audit Manager installed.

3 With the host’s packages displayed, select the *Audit Manager (audit)* package.

4 Click *Audit Settings* in the task pane.

5 In the *Audit Path* field, specify the location for the audit databases.

6 Click *Finish*.

### 3.4.9 Configuring SMTP Settings for the Messaging Component Package

The *SMTP Settings* option allows you to provide details of your e-mail server so reports such as the compliance auditor reports and custom command control reports can be automatically e-mailed to the necessary personnel.

To configure the SMTP settings:

- 1 Click *Hosts* on the home page of the console.
- 2 Select the host where the Compliance Auditor and Messaging Component are installed.
- 3 Click *Packages* to view details of the packages installed on this host.
- 4 Select the *Messaging Component (msgagnt)*.
- 5 Click *SMTP Settings* in the task pane.
- 6 Configure the following fields:
  - SMTP Host:** Specify the IP address of your e-mail server.
  - SMTP Port:** Specify the port of your e-mail server.
  - SMTP Domain:** If you are using a Lotus Notes server, specify the name of your SMTP domain.
- 7 Click *Finish*.

## 3.5 Increasing the Security When Accessing the Framework Manager Console

The following options increase the security of the communication between the browsers and the Framework Manager console. These configuration options do not affect the communications between the Framework managers and the Framework agents.

- ♦ [Section 3.5.1, “Requesting a Certificate for the Framework Manager Console,” on page 39](#)
- ♦ [Section 3.5.2, “Installing a Certificate,” on page 40](#)
- ♦ [Section 3.5.3, “Modifying the Connector,” on page 40](#)

### 3.5.1 Requesting a Certificate for the Framework Manager Console

For added security, you can install a certificate to use when accessing the Framework Manager console. To access this option, you need to select the specific packages on the host you want to protect. You must then complete a certificate request form, send it to your chosen certification authority, and then install the certificate you receive.

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the host that you want to protect with a certificate.
- 3 With the host’s packages displayed, select the *Administration Agent (admin)* for the Framework Manager console.
- 4 Click *Request Certificate* in the task pane.
- 5 Specify the necessary details as described in your chosen certification authority documentation.

**6** Click *Finish*.

The text for your certificate request is displayed in the text area.

- 7** Copy the certificate request into an e-mail and send it to your chosen certification authority.
- 8** When you receive the certificate from your certification authority, install it as described in [“Installing a Certificate” on page 40](#).

### 3.5.2 Installing a Certificate

When you have received the requested certificate from your certification authority:

- 1** Copy the certificate to the machine you use to access the Framework Manager console.
- 2** Click *Hosts* on the home page of the console.
- 3** In the navigation pane, select the host where you want to install the certificate.
- 4** With the host’s packages displayed, select the package for which you have requested the certificate.
- 5** Click *Install Certificate* in the task pane.
- 6** Paste the certificate into the text area.
- 7** Click *Finish*.

### 3.5.3 Modifying the Connector

You can modify the way you connect to the Framework Manager console. You can define which interface card and port to use, and increase the security of the connection by using SSL.

To access this option, you need to select the specific packages on the host you want to modify.

- 1** Click *Hosts* on the home page of the console.
- 2** In the navigation pane, select the host where you want to modify the connector.
- 3** With the host’s packages displayed, select the *Administration Manager (admin)* for the Framework Manager console.
- 4** Click *Modify Connector* in the task pane.
- 5** Make the changes you want:
  - ♦ Define the address of a specific interface card
  - ♦ Define which port to use
  - ♦ Select the *SSL* check box if you want to use SSL.
- 6** Click *Finish*.

## 3.6 Troubleshooting

- ♦ [Section 3.6.1, “Promoting Managers When the Primary Manager Fails,” on page 41](#)
- ♦ [Section 3.6.2, “Viewing Store and Forward Messages,” on page 41](#)
- ♦ [Section 3.6.3, “Managing Low Disk Space,” on page 42](#)
- ♦ [Section 3.6.4, “Restarting the Agent,” on page 43](#)

- ♦ [Section 3.6.5, “Managing the Registry Cache,” on page 44](#)
- ♦ [Section 3.6.6, “Time Synchronization,” on page 45](#)

### 3.6.1 Promoting Managers When the Primary Manager Fails

If you have multiple Framework Managers deployed, the first manager installed is defined as the primary manager by default, and its packages are defined as primary. Manager packages on all other manager hosts act as backups. If your primary manager becomes unavailable, you can select single or multiple manager packages on a host to be promoted to primary status.

The Framework continues to function when the primary manager is unavailable, but no changes can be made to the Framework. Changes can only be written to the databases on the primary manager, which are then replicated to the backup managers. The only exception to this is the audit database. Each audit agent is responsible for sending its audit messages to each audit manager. This ensures that audit data is not lost.

Novell recommends having one agent designated as a complete mirror of your primary manager. In event of a total failure of the primary manager, you can log into the backup console and promote it to primary status with no disruption of Privileged User Manager services.

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the host where you want to promote a manager.
- 3 With the host’s packages displayed, select the manager packages you want to promote.  
To select multiple manager packages, press the Ctrl key and select the packages one at a time, or press the Shift key to select a consecutive list of manager packages.
- 4 Click *Promote Manager* in the task pane.
- 5 Review the list of manager packages you have selected.
- 6 Click *Finish*.
- 7 View the host’s packages again and verify that the *Status* of the promoted manager packages has changed to Primary.

To use a command line option to promote backup host to primary status, see [Section 10.6, “Registry Manager Options,” on page 155](#).

### 3.6.2 Viewing Store and Forward Messages

Messages from one host to another are stored if the sending host cannot communicate with the receiving host, and forwarded when the communication link is restored. You can view these messages and delete them if you do not need them.

You can use this feature to analyze a host and to discover whether it is having problems contacting a particular host. This problem usually occurs when a host is down or when a DNS name for a host name cannot be resolved.

- 1 Click *Hosts* on the home page of the console.  
The navigation pane displays the current hierarchy for your Framework.
- 2 Select the host for which you want to view store and forward messages.
- 3 Click the host’s *Packages* icon (select the arrow next to the host’s name to display it).

4 Select *View Messages* in the task pane.

If any stored messages exist, they are displayed. Information about the message is shown, including the time the message was sent, the host the message was being sent to, the module that sent the message, the type of message (method), the number of failed attempts at sending the message, and the next scheduled attempt to send the message, if any.

5 To attempt to send one or more messages again, select the messages and select *Retry*.

6 To delete one or more messages, select the messages and click *Delete*.

7 To refresh the screen, click *Refresh*.

8 Click *Close*.

### 3.6.3 Managing Low Disk Space

In previous releases of Privileged User Manager, `usrun` sessions were terminated with an auditing error when the server ran out of disk space. For long term running processes, this is not the ideal solution.

You can now use Command Control scripts to slow down or freeze input/output for the following conditions:

- ♦ In `usrun` sessions when disk space is low.
- ♦ When the store and forward process cannot contact an audit manager and its queue size is increasing.

You can control what happens under these conditions by configuring the following attributes:

**disk\_min\_free (default: 1MB):** Minimum free disk space. When free disk space goes below this level, the action defined in `backoff_action` is applied.

- ♦ If the `backoff_action` is `block`, the audit message is paused until disk space becomes available.
- ♦ If the `backoff_action` is `fail`, the request to store the audit message fails with an error, and the user session is terminated.
- ♦ If the `backoff_action` is `allow`, the session is unaffected.

**disk\_wm\_free (default: 2MB):** Free disk space watermark. When the free disk space goes below this level, the delay defined in `backoff_delay` is applied between each audited message.

**queue\_max\_size (default: 100MB):** Maximum queue size. When the queue size goes above this level, the action defined in `backoff_action` is applied.

- ♦ If the `backoff_action` is `block`, the audit message is paused until the queue size reduces below this level.
- ♦ If `backoff_action` is `fail`, the request to store the audit message fails with an error, and the user session is terminated.
- ♦ If `backoff_action` is `allow`, the session is unaffected.

**queue\_wm\_size (default: 20MB):** Queue size watermark. When the queue size goes above this level, the delay defined in `backoff_delay` is applied between each audited message.

**backoff\_divisor (default: 1):** Provides the ability to increase the delay as the disk space reduces or the queue size increases. The delay is calculated by dividing the range between the `disk_wm_free` and `disk_min_free` (or `queue_max_size` and `queue_wm_size`) by the `backoff_divisor` and then applying the delay for each increment.

**backoff\_delay (default: 500ms):** Time in milliseconds to delay the audit request.

**backoff\_action (default: block):** Either block, fail, or allow.

The following Command Control script illustrates how to change these settings:

```
my $t=$meta->child("Audit");
$t=$meta->add_node("Audit") if(! $t);
$t->arg("disk_min_free","10");
$t->arg("disk_wm_free","20");
return 1;
```

This script sets the `disk_min_free` attribute to 10MB and the `disk_wm_free` attribute to 20 MB. You can assign this script to any rule or you can assign it to a rule at the top of the tree that all commands pass through.

You should create an emergency policy that allows administrators to access the machine when disk space is low or the store-and-forward queue size is large. Such a script would look similar to the following:

```
my $t=$meta->child("Audit");
$t=$meta->add_node("Audit") if(! $t);
$t->arg("disk_min_free","0");
$t->arg("disk_wm_free","0");
$t->arg("queue_max_size","0");
$t->arg("queue_wm_size","0");
$t->arg("backoff_action","allow");
return 1;
```

You can assign this script to any rule or you can assign it to a rule at the top of the tree that all commands pass through

### 3.6.4 Restarting the Agent

If you are having problems, Novell Support might ask you to restart an agent.

- 1 Click *Hosts* on the home page of the console.
- 2 In the navigation pane, select the host on which you want to restart the agent.  
To select multiple hosts in a domain, select the domain, then press the Ctrl key and select the hosts one at a time, or press the Shift key to select a consecutive list of hosts. To select all hosts in a domain, use Ctrl+A.
- 3 Click *Restart Agent* in the task pane.
- 4 Select the type of restart you want to perform, as advised by Novell Support.  
**Soft restart:** Reloads the module libraries and resets the service uptime.  
**Hard restart:** Restarts the daemon, reloads all modules, and resets the service uptime.
- 5 Click *Finish*.

## 3.6.5 Managing the Registry Cache

The registry cache is held by the Registry Agent on each host, and it contains a list of the packages deployed on each host in your Framework. This list is a copy of part of the information held by the Registry Manager, and it enables Framework components to locate and communicate with each other, according to their position in the hierarchy created when you add domains and hosts to your Framework. Agents send requests to managers in the immediate subdomain, and if a request is unsuccessful, they try a manager higher up in the hierarchy. See [Chapter 8, “Load Balancing and Failover,”](#) on page 143 for details.

- ◆ You can view the registry cache to check hosts in your Framework to see if a specific manager or agent module is installed, and check the order in the Framework hierarchy according to the hosts the modules are installed on. See [“Viewing the Registry Cache”](#) on page 44.
- ◆ If the registry cache becomes out-of-date, communication problems can occur. To fix this, try clearing the registry cache on the Registry Agent to allow it to be updated by the Registry Manager. See [“Clearing the Registry Cache”](#) on page 44.

### Viewing the Registry Cache

When viewing the registry cache, you can use the stale cache (the default option). The cache is considered stale if it has not been updated by the Registry Manager for 2 hours, and this is usually adequate. If you deselect the *Use Stale Cache* check box, the information is provided by the Registry Manager.

- 1 Click *Hosts* on the home page of the console.

The navigation pane displays the current hierarchy for your Framework.

- 2 Select the host for which you want to view the registry cache.
- 3 Click the host's *Packages* icon (click the arrow next to the host's name to display it).
- 4 Click *View Cache* in the task pane.
- 5 From the drop-down list, select the package you want to look up in the registry cache.
- 6 If you want to view the latest information from the Registry Manager, deselect the *Use stale cache* check box, then click *Lookup*.

Details of the hosts where the module is installed are displayed in order according to their position in the Framework hierarchy. Information shown includes the Framework agent name, IP address, port number, and whether the host has the primary manager component installed (indicated by 1 in the *Primary* column) or not (indicated by 0).

- 7 (Optional) To clear the registry cache, click *Clear Cache*.

This marks the cache as stale, and it is automatically updated by the Registry Manager. You can also clear the cache by using the *Clear Cache* option in the task pane.

- 8 Click *Close*.

### Clearing the Registry Cache

Novell Support might advise you to try clearing the registry cache if you have communication problems among Privileged User Manager components. The registry cache is held by the Registry Agent and contains a list of manager and agents in your Framework, copied from the Registry Manager. See [“Managing the Registry Cache”](#) on page 44 for more details.

- 1 Click *Hosts* on the home page of the console.

The navigation pane displays the current hierarchy for your Framework.

- 2 Select the host for which you want to clear the registry cache.
- 3 Click the host's *Packages* icon (click the arrow next to the host's name to display it).
- 4 Click *Clear Cache* in the task pane.

The registry cache is marked as stale and is updated by the Registry Manager. You can also clear the registry cache by using the *View Cache* option (see [“Viewing the Registry Cache” on page 44](#)).

### 3.6.6 Time Synchronization

All agents should be configured to use a Network Time Protocol server. Agents must have their time synchronized with the primary registry manager so that the time difference is less than two hours.

If the time difference is greater than two hours, the agent can appear offline and command control requests can fail.



# Managing Framework Users and Groups

# 4

Privileged User Manager provides comprehensive user management facilities to control access to the Framework consoles. The admin user created when the Framework is initially installed belongs to the admin group, which has full access to all installed consoles and can perform all tasks. You can use this user account to create additional user accounts and groups through the Framework User Manager console, which is part of the Access Control module. Role-based authorization is used to determine which user groups can access specific consoles and perform specific tasks.

This section describes the following Framework User Manager tasks:

- ♦ [Section 4.1, “Managing Users,” on page 47](#)
- ♦ [Section 4.2, “Managing Groups,” on page 56](#)
- ♦ [Section 4.3, “Deploying the Access Control Module,” on page 62](#)
- ♦ [Section 4.4, “Changing a Framework User’s Password,” on page 63](#)

## 4.1 Managing Users

When you add a new user, the user cannot access any of the Framework consoles until the user is added to a group that contains a role allowing the appropriate access. For example, if you want a user to be able to access only the Compliance Auditor console, you must create a group and configure the appropriate Compliance Auditor roles, then create the user and add the user to the group.

You can create additional users with the same access as the admin user by adding them to the admin group, or create your own group with access to all modules and roles. You can also configure these additional users to be superusers. Only users who belong to a group with the “super” role can view and administer superusers.

- ♦ [Section 4.1.1, “Configuring Account Settings,” on page 47](#)
- ♦ [Section 4.1.2, “Adding a Framework User,” on page 49](#)
- ♦ [Section 4.1.3, “Modifying a Framework User,” on page 49](#)
- ♦ [Section 4.1.4, “Removing a Framework User Group from a User,” on page 56](#)
- ♦ [Section 4.1.5, “Deleting a Framework User,” on page 56](#)

### 4.1.1 Configuring Account Settings

The *Account Settings* option allows you to set the default values for user settings such as minimum password length. When you add a new user, these default settings apply, but they can be overridden for individual users by modifying the individual account settings.

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Click *Users* in the navigation pane.
- 3 Click *Account Settings* in the task pane.

**4** Configure the following account options:

**Inactivity timeout (minutes):** Specify the number of minutes that users can be inactive before logging them out of the Framework Manager console.

**Account lockout:** Specify the number of times a user can enter the wrong password before being locked out. You can re-enable the user's account by using the *Modify User* option and clearing the *Disabled* check box. You can reset the user's password by using the *Modify User* option.

**Inactive days (disable):** Specify the number of days that a user's account can be inactive before it is disabled. You can reactivate the user's account by using the *Modify User* option and using the *Reactivate* account check box in the *Account* section.

**Inactive days (delete):** Specify the number of days a user's account can be inactive before it is deleted.

**Password lifetime (days):** Specify the number of days a user's password can be used before it expires and the user is prompted to change the password.

**Minimum password length:** Specify the minimum number of characters that must be used in a user's password.

**Password history:** Specify the number of unique passwords that a user must use before being allowed to reuse an old password.

**Minimum alpha:** Specify the minimum number of alphabetic characters that must be used in a user's password.

**Minimum numerics:** Specify the minimum number of numeric characters that must be used in a user's password.

**Cache native passwords:** Enable this option if you want the Framework Manager passwords updated with LDAP passwords. When you set up a mapping for users with an LDAP server, the Framework Manager password is updated to match the LDAP password with each successful login. (For information on setting up an LDAP mapping, see [“Modify User: Native Maps” on page 54.](#))

If a user never successfully logs into the LDAP server, the local password is never updated, and the user can use the local Framework Manager password to log in.

If this option is disabled, the local Framework Manager passwords are never updated with the LDAP passwords. Users can attempt an LDAP login, and if that fails, they can log in locally with their Framework Manager passwords.

**5** Configure the following help desk attributes. These attributes control the functionality of the Help Desk role and determine the actions that can be performed by a help desk user. For information about creating a group to use these attributes, see [Section 4.2.3, “Configuring a Help Desk Group,” on page 57.](#)

**Disabled:** Allows the help desk user to enable and disable user accounts.

**Password:** Allows the help desk user to change an existing password.

**Change at next login:** Allows the help desk user to determine whether the user is forced to change the password on the next login.

**Last changed:** Displays the last time the password was changed and allows the help desk user to reset it to the current date and time.

**Bad logons:** Displays the number of bad logins and allows the help desk user to reset the count.

**Last bad logon:** Displays the time and date of the last bad login and allows the help desk user to reset it to the current date and time.

**Last logon:** Displays the last successful login of the user and allows the help desk user to reactivate the account.

**Group membership:** Allows the help desk user to assign the user to non-administrative accounts.

6 Click *Finish*.

## 4.1.2 Adding a Framework User

When you add a new Framework user:

- ♦ The user's account is set up according to the default values defined in the *Account Settings* option. You can change these settings for individual users by using the *Modify User* option.
- ♦ The user's password is set to expire immediately so he or she is prompted to change it on the first login to the Framework Manager console. You can change this setting for individual users by using the *Modify User* option.
- ♦ The user cannot access any of the Framework consoles until you have added the user to a group with the required roles defined. For more information, see [Section 4.1.3, "Modifying a Framework User," on page 49](#) and [Section 4.2.4, "Configuring Roles," on page 58](#).

To add a new Framework user:

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Click *Users* in the navigation pane.
- 3 Click *Add User* in the task pane.
- 4 Specify a name for the user in the *Username* field.
- 5 Specify a password for the user in the *Password* field.  
The password must comply with the default account settings for the Framework.
- 6 Click *Finish*.
- 7 To configure additional settings for the user's account, continue with [Section 4.1.3, "Modifying a Framework User," on page 49](#).

## 4.1.3 Modifying a Framework User

The *Modify User* option allows you to override the default account settings for an individual user, and also provides a number of additional configuration settings and tasks, including resetting a user's password and assigning a user to a group.

To modify a Framework user account:

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Click *Users* in the navigation pane.
- 3 In the left pane, select the user account you want to modify.
- 4 Click *Modify User* in the task pane.
- 5 Change the settings as desired:

**Disabled:** Select this option to disable the user's account.

**Comment:** Specify a short comment in the text box.

**Description:** Specify a detailed description in the text box.

- 6 To configure additional options, select the section you want:

**Password:** Allows you to reset the user's password and configure other password settings. For specific instructions and additional options, see [“Modify User: Password” on page 50](#).

**Password validation:** Allows you to define the minimum number of alphabetic and numeric characters required in the user's password. For specific instructions and additional options, see [“Modify User: Password Validation” on page 51](#).

**Account:** Allows you to configure the user as a superuser, provides information about the user's account, and provides other account configuration options. For specific instructions and additional options, see [“Modify User: Account” on page 51](#).

**Account Details:** Allows you to enter personal information for the user, including Staff ID and contact details. For specific instructions and additional options, see [“Modify User: Account Details” on page 52](#).

**Host Access Control:** Allows you to control where the user can access the console from. For specific instructions and additional options, see [“Modify User: Host Access Control” on page 53](#).

**Native Maps:** Allows you to map the Framework user account to a user account on a UNIX platform or on an LDAP server. For specific instructions and additional options, see [“Modify User: Native Maps” on page 54](#).

**Logon Script:** Allows you to define a Perl logon script for the user. For specific instructions and additional options, see [“Modify User: Logon Script” on page 55](#).

**Groups:** Allows you to add the user to one or more groups. For specific instructions and additional options, see [“Modify User: Groups” on page 55](#).

- 7 When you have completed your changes, click *Finish*.

## Modify User: Password

To set password options for a Framework user:

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Select the user account you want to modify, then click *Modify User*.
- 3 Click *Password*.
- 4 Change the options as desired:

**Password:** To reset the user's password, type the new password and retype it in the *Confirm* field.

---

**NOTE:** The password must comply with the default account settings for the Framework, and comply with individual user settings defined by using this option and the *Password Validation* option.

---

**Expired:** Select the *Expired* check box to expire the user's current password immediately, forcing the user to change it on the next login.

**Last changed:** Indicates when the password was last changed by the user, or, if the password has not yet been changed by the user, indicates when the user and password were created.

**Reset password age:** Select the *Reset password age* check box to reset the age of the password to zero. The user can use the password for the full number of days defined in *Password lifetime (days)* (see [Section 4.1.1, “Configuring Account Settings,”](#) on page 47), or in the *Maximum age* field if it has been configured.

**Minimum length:** Override the default account settings by specifying the minimum number of characters you require in a user’s password.

**Maximum age:** Override the default account settings by specifying the number of days before a user’s password expires, prompting the user to change the password.

**History:** Override the default account settings by specifying the number of unique passwords that a user must use before being allowed to reuse an old password.

- 5 Click *Finish* or select another option.

## Modify User: Password Validation

To set password validation options for a Framework user:

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Select the user account you want to modify, then click *Modify User*.
- 3 Click *Password validation*.
- 4 To override the default account settings for this user, select the appropriate check box and set the required values as follows:

**Min alpha characters:** Specify the minimum number of alphabetic characters you require in the user’s password.

**Min numeric characters:** Specify the minimum number of numeric characters you require in the user’s password.

- 5 Click *Finish* or select another option.

## Modify User: Account

To set account options for a Framework user:

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Select the user account you want to modify, then click *Modify User*.
- 3 Click *Account*.
- 4 Change the options as desired:

**Super user:** Select the *Super user* check box to make this user a superuser.

---

**NOTE:** The *Super user* option is available only if you are logged in as a superuser. Superusers can be viewed and administered only by users belonging to a group with the super role defined for the auth module.

---

**Last bad logon:** The last time the user failed to log on successfully.

**Last logon:** Indicates when the user last logged in to the Framework Manager console.

**Reactivate account:** Select the *Reactivate account* check box to re-enable a user’s account that has been locked through bad logons.

**Disable inactive days:** Override the default account settings by specifying the number of days the user's account can be inactive before it is disabled. You can reactivate the user's account by using the *Reactivate* account option described above.

**Delete inactive days:** Override the default account settings by specifying the number of days the user's account can be inactive before it is deleted.

**Inactivity logout mins:** Override the default account settings by specifying the number of minutes the user can be inactive before the user is logged out of the Framework Manager console.

**Bad logons:** The number of times the user has failed to log on successfully since the last successful logon.

**Reset bad logon count:** Resets the number of unsuccessful logons to zero.

**Lockout:** Override the default account settings by specifying the number of times the user can enter the wrong password before being locked out. You can re-enable the user's account by clearing the *Disabled* check box in the main *Modify User* section. You can reset the user's password in the *Password* section.

**Message of the day:** Override the default account settings by specifying a message to be displayed to the user after a successful logon.

- 5 Click *Finish* or select another option.

## Modify User: Account Details

To set personal account details for a Framework user:

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Select the user account you want to modify, then click *Modify User*.
- 3 Click *Account Details*.
- 4 To set the following options, select the appropriate check box and specify the text:

**Staff ID:** Specify the user's staff ID, for example, the user's unique company identifier.

**Display name:** Specify a display name for the user, for example, the user's full name. If a name is defined here it can be automatically entered as the *Manager Name* in Account Group and User Group definitions for Command Control by selecting the manager's Framework user name (see [“Modifying an Account Group” on page 95](#) and [“Modifying a User Group” on page 92](#)). It can also be used in Compliance Auditor reports (see [Section 7.3.1, “Adding or Modifying an Audit Report,” on page 130](#)).

**Email address:** Specify the user's e-mail address. If an e-mail address is defined here, it can also be used in Command Control (see [“Modifying an Account Group” on page 95](#) and [“Modifying a User Group” on page 92](#)) and in the Compliance Auditor (see [Section 7.3.1, “Adding or Modifying an Audit Report,” on page 130](#)).

**Telephone number:** Specify the user's telephone number. If a telephone number is defined here, it can also be used in Command Control (see [“Modifying an Account Group” on page 95](#) and [“Modifying a User Group” on page 92](#)) and in the Compliance Auditor (see [Section 7.3.1, “Adding or Modifying an Audit Report,” on page 130](#)).

- 5 Click *Finish* or select another option.

## Modify User: Host Access Control

You can control where the user can access a Framework Manager console from by defining a list of ports and hosts to which access is allowed, or a list of ports and hosts to which access is denied.

If you make no entries for this option, access is allowed from any location.

To control where the user can access the Framework Manager console from:

- 1** Click *Framework User Manager* on the home page of the console.
- 2** Select the user account you want to modify, then click *Modify User*.
- 3** Click *Host Access Control*.
- 4** (Optional) Define a list of locations from where the user is allowed to access the console, and deny access from all other locations:
  - 4a** If auditing is required, select the *Auditing* check box and use the drop-down list to select the events you want to be audited.
  - 4b** Select the *Host Access* check box.
  - 4c** Click the *Add* button below the *Host Access* list.
  - 4d** In the *Port Range* column, specify the required port number or range of port numbers. The following entries are allowed:

---

*	All ports
port	A single port, such as 80
port-port	A range of ports, such as 20-30
svcname	Resolves a service name to its port, such as HTTP

---

- 4e** In the *Host/IP Subnet* column, specify the required host definition. The following entries are allowed:

---

*	All hosts
ip address	A full IP address, such as 192.168.1.1
ip address-ip address	A range of IP addresses, such as 192.168.1.1-192.168.1.12
part ip address	Part of an IP address, such as 192.168.1
network/netmask	A network/netmask pair, such as 192.168.1.0/255.255.255.0
network/nnn CIDR	A network/nnn CIDR, such as 192.168.11.0/24
hostname	A hostname, such as dellsrv1.novell.com
domain	A domain name, such as *.novell.com

---

- 4f** In the *Allow* column, click the check box.
- 4g** Repeat [Step 4c](#) through [Step 4e](#) for any other required location definitions.

- 5** (Optional) Define a list of locations from which the user is denied access to the console, and allow access from all other locations:
  - 5a** If auditing is required, select the *Auditing* check box and use the drop-down list to select the events you want to be audited.
  - 5b** Select the *Host Access* check box.
  - 5c** Click the *Add* button below the Host Access list.
  - 5d** Specify the desired locations as described in [Step 4d](#) and [Step 4e](#) above.
  - 5e** To make this a deny entry, make sure the check box is not selected in the *Allow* column.
  - 5f** Repeat steps [Step 5c](#) and [Step 5e](#) for any other required location definitions.
- 6** Click *Finish* or select another option.

### Modify User: Native Maps

The *Native Maps* option allows you to map Framework User accounts to UNIX or Linux accounts and to LDAP accounts.

- ♦ [“UNIX or Linux Account Mapping” on page 54](#)
- ♦ [“LDAP Account Mapping” on page 55](#)

### UNIX or Linux Account Mapping

The Privilege User Manager Framework provides the ability to perform a number of functions from the command line. When using the command line, you are required to authenticate to the Framework. For example, the following command returns the status of all agents:

```
/opt/novell/npum/sbin/unifi -u admin regclnt status -a
```

The command contains a switch for the username (`-u admin`). When the command is executed, the user is prompted for a password.

You can use the *Native Maps* option to map a platform system user to a Privileged User Manager account. If you use an additional switch in the command line call, you are no longer required to provide authentication. A user with a native map can enter the following command:

```
/opt/novell/npum/sbin/unifi -n regclnt status -a
```

The native map plus the `-n` switch allows the command to be executed without prompting the user for a name or a password.

To add a native map for a UNIX or Linux user:

- 1** Click *Framework User Manager* on the home page of the console.
- 2** Select the user account you want to modify, then click *Modify User*.
- 3** Click *Native Maps*.
- 4** Click *Add*.
- 5** In the *User* column, specify the user’s name for the UNIX or Linux platform.
- 6** In the *Host* column, select the hostname for the UNIX or Linux platform.
- 7** Repeat [Step 4](#) through [Step 6](#) for any additional maps you require.
- 8** To edit a native map, select it and make the required changes.

- 9 To remove a native map, select it and click *Remove*.
- 10 Click *Finish* or select another option.

## LDAP Account Mapping

Native maps can be used to allow Framework Manager users to obtain their authentication credentials from an LDAP server. This allows the LDAP server to remain the authoritative source for user credentials and active accounts. If you want LDAP mapped users to be able to log in when the LDAP server is not available, see the *Cache native passwords* option in [Section 4.1.1, “Configuring Account Settings,” on page 47](#).

To configure an LDAP mapping:

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Select the user account you want to modify, then click *Modify User*.
- 3 Click *Native Maps*.
- 4 Click *Add*.
- 5 In the *User* column, specify the user’s fully qualified distinguished name. For example:  
`cn=plou,ou=development,o=novell`
- 6 In the *Host* column, specify the scheme (`ldap` or `ldaps`) and IP address of the LDAP server. Specify a port only if the LDAP server is not using the standard port for the scheme. For example:  
`ldaps://10.10.16.165`  
`ldaps://10.10.16.166:789`
- 7 Click *Finish* or select another option.

## Modify User: Logon Script

You can assign a Perl script to a user to be run when the user logs on to the Framework Manager console. For example, you could assign a script that causes an e-mail to be sent to a manager when the user logs on.

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Select the user account you want to modify, then click *Modify User*.
- 3 Click *Logon Script*.
- 4 Specify the logon script you require for this user. You can type the script or paste it from another document.
- 5 Click *Finish* or select another option.

## Modify User: Groups

To assign a Framework user to one or more groups:

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Select the user account you want to modify, then click *Modify User*.
- 3 Click *Groups*.
- 4 Select the check boxes for the groups you want this user to belong to.
- 5 Click *Finish* or select another option.

You can also assign a user to a group by using the *Modify Group* option, by dragging the user onto the group, or by dragging the group onto the user.

You can remove a user from a group by deselecting the check box for the required group. See [Section 4.1.4, “Removing a Framework User Group from a User,” on page 56](#) for other methods.

#### 4.1.4 Removing a Framework User Group from a User

There are several ways of removing a Framework user group from a Framework user’s account. You can modify the user, modify the group, or use the objects in the navigation pane.

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Select the group you want to remove from the user’s account.
- 3 In the right pane, select the user.
- 4 Click *Remove User* in the task pane. The user is removed.

#### 4.1.5 Deleting a Framework User

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Click *Users* in the navigation pane.
- 3 In the left pane, select the user you want to delete.
- 4 Click *Delete User* in the task pane.
- 5 Click *Finish* to confirm the deletion.

## 4.2 Managing Groups

Framework users must be assigned to one or more groups with the appropriate roles defined before they can access any Framework consoles or perform any tasks.

- ♦ [Section 4.2.1, “Adding a Framework User Group,” on page 56](#)
- ♦ [Section 4.2.2, “Modifying a Framework User Group,” on page 57](#)
- ♦ [Section 4.2.3, “Configuring a Help Desk Group,” on page 57](#)
- ♦ [Section 4.2.4, “Configuring Roles,” on page 58](#)
- ♦ [Section 4.2.5, “Deleting a Framework User Group,” on page 62](#)

#### 4.2.1 Adding a Framework User Group

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Click *Groups* in the navigation pane.
- 3 Click *Add Group* in the task pane.
- 4 Specify a name for the group in the *Group name* field.
- 5 Click *Finish*.
- 6 To configure the group, continue with [Section 4.2.2, “Modifying a Framework User Group,” on page 57](#).

## 4.2.2 Modifying a Framework User Group

The *Modify Group* option allows you to:

- ◆ Add a comment describing the group
- ◆ Add users and subgroups to the group
- ◆ Define administrative roles for the group
- ◆ Specify an audit manager for the group.

To modify a Framework user group:

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Click *Groups* in the navigation pane.
- 3 In the left pane., select the group you want to modify
- 4 Click *Modify Group* in the task pane.
- 5 (Optional) In the *Comment* field, enter a comment.
- 6 In the *Members* section, select the users you want to be members of this group.  
You can also add a user to groups in the *Groups* section of the *Modify User* option, by dragging the user onto the group, or by dragging the group onto the user.  
You can remove users from the group by deselecting them here. See [Section 4.1.4, “Removing a Framework User Group from a User,”](#) on page 56 for other methods.
- 7 In the *Sub Groups* section, select the groups you want to be subgroups of this group.  
You can also add subgroups to groups by dragging the group onto the main group.
- 8 In the *Roles* section, configure the roles you require for this group of users according to the consoles you want them to be able to access and the tasks you want them to be able to perform. You must assign at least one role. See [Section 4.2.4, “Configuring Roles,”](#) on page 58 for more details.
- 9 In the *Audit Manager* section, specify the details of the group’s manager.
- 10 Click *Finish*.

## 4.2.3 Configuring a Help Desk Group

The help desk role allows a predefined set of attributes to be set on the Account Settings page so that users assigned to the help desk group can only manage the subset of user attributes.

To set up a help desk group:

- 1 Configure the attributes:
  - 1a Click *Framework User Manager* on the home page of the console.
  - 1b Click *Users > Account Settings*.
  - 1c Configure the *Helpdesk Attributes*.  
For information about these attributes, see [Section 4.1.1, “Configuring Account Settings,”](#) on page 47.
  - 1d Click *Finish*.

- 2 Create the group:
  - 2a Click *Groups > Add Group*.
  - 2b Specify a name for the group, then click *Finish*.
  - 2c Select the group you just created, then click *Modify Group*.
  - 2d In the *Members* option, select the users that you want to belong to the help desk group.
  - 2e In the *Roles* option, click *Add*, then add the following roles:

Module	Role
auth	console
auth	read
auth	helpdesk

- 3 Click *Finish*.

## 4.2.4 Configuring Roles

When you create a new Framework user group, you must assign at least one role to the group to allow the users in the group to access one or more Framework modules and perform tasks.

To allow access to all modules and tasks, you can define a role with Module set to \* and Role set to \*. This is how the default admin group containing the default admin user is initially configured.

To allow access only to specific modules and tasks, use the *Modify Group* option (see [Section 4.2.2, “Modifying a Framework User Group,” on page 57](#)) and define one or more roles according to the tables below:

- ♦ [“Framework User Roles” on page 58](#)
- ♦ [“Audit Report Roles” on page 59](#)
- ♦ [“Compliance Auditor Roles” on page 60](#)
- ♦ [“Host Roles” on page 61](#)
- ♦ [“Package Manager Roles” on page 61](#)
- ♦ [“Command Control Roles” on page 61](#)

### Framework User Roles

The following roles can be assigned to the authentication module in order to control access to the Framework User Manager console. Select from these roles when you are setting up a group to manage Framework Manager users and groups.

Module	Role	Allows users to
auth	act_settings	Modify account settings.
	admin	Add or delete users and groups, and assign users to groups.
	console	View the Framework User Manager console.

Module	Role	Allows users to
	helpdesk	Modify the user account settings. To change which attributes are available for modification, see <a href="#">Section 4.1.1, “Configuring Account Settings,”</a> on page 47.  For information on how to use this role to create a Help Desk group that can manage user passwords, see <a href="#">Section 4.2.3, “Configuring a Help Desk Group,”</a> on page 57.
	read	Read the auth database.  This role must be used with all other auth roles.
	role_admin	Add or remove roles.
	super	View and modify superusers, and view and modify groups with the super role defined.
	*	Perform all roles.

### Audit Report Roles

The following roles can be assigned to the auditing module in order to control access to the Reporting console. Select from these roles when you are setting up a group to manage the command control reports.

Module	Role	Allows users to
audit	read	Read the audit database.  This role must be used with all other audit roles.
	console	View the Reporting console.
	admin	Modify reporting settings.
	command	View Command Control reports.
	logon	View Account Logon reports.
	*	Perform all roles.
	write	Create new audit reports and adjust filter settings.
	report	Access reports with the report defined roles.
	<report defined>	Read and update the reports defined in the <i>General</i> tab of the Reporting console.  This role is only useful when used in conjunction with the report role.

You can use these Audit Report roles to create the following types of audit managers:

- ♦ **Administrator:** To allow the group to update all aspects of the auditing module, including encryption and rollover, the group needs to be assigned the following roles for the audit module:

admin

write  
read  
command  
console

- ♦ **Manager:** To allow the group to update all aspects of the auditing module, except encryption and rollover, the group needs to be assigned the following roles for the audit module:

write  
read  
command  
console

- ♦ **User:** To allow the group to read and update a specific report, the group needs to be assigned the following roles for the audit module:

command  
console  
report  
<report defined read>  
<report defined update>

If you want the group to have read-only privileges to the report, do not assign the <report defined update> role. Users with read-only rights to a report can view the report from the console, view the keystroke sessions within the report, and select which audit databases to view (see the *LogFiles* tab). Users who also have the update right can update the report's filter, its name, and its description.

Each report allows you to specify a read role and an update role. You need to remember those names and manually enter them here. The console does not provide any error checking, so you need to make sure to enter the correct name. For information on how to enable a report for a role, see [Section 6.4.4, "Modifying General Report Information," on page 124](#).

## Compliance Auditor Roles

The following roles can be assigned to the compliance auditing module in order to control access to the Compliance Auditor console. For a group to manage compliance auditing, the group also needs read roles to the auditing and authentication modules.

Module	Role	Allows users to
secaudit	console	View the Compliance Auditor console.
	audit	View and edit records.
	admin	Add and modify audit rules.
	*	Perform the console, audit, and admin roles.
	<audit role name >	Access the records collected by audit rules with this role defined in the <i>Audit Role</i> field on the Modify Audit Rule page. You can choose your own name for the role.  See <a href="#">Section 7.2.1, "Adding or Modifying an Audit Rule," on page 128</a> for details about configuring audit rules.

Module	Role	Allows users to
audit	read	View a keystroke replay.
auth	read	Extract user credentials, including name and e-mail address, from the auth database for use with reports.

### Host Roles

The following roles can be assigned to the host module in order to control access to the Hosts console. Select from the following roles when creating a group to manage the hosts.

Module	Role	Allows users to
unifi	info	Run the host status check by using the command line interface.  You must type the word <code>info</code> because it is not available in the drop-down list.
	admin	View the Hosts console and perform administrative actions.

### Package Manager Roles

The following role can be assigned to the package manager module in order to control access to the Package Manager console. When you are creating a group that you want to manage the distribution of updates to Privileged User Manager, select the following:

Module	Role	Allows users to
pkgman	admin	View, add, update, or remove packages.

### Command Control Roles

The following roles can be assigned to the command control module in order to control access to the Command Control console. Select from the following roles when you are creating a group that you want to manage and test the rules in the command control database.

Module	Role	Allows users to
cmdctrl	read	View the Command Control console and run test suites.
	write	Modify the command control database. Users with this role cannot cancel other users' transactions or modify audit or transaction settings.  Must be used in conjunction with the cmdctrl read role.
	admin	Modify the Command Control database, including canceling other users' transactions and modifying audit and transaction settings.
	*	Perform all roles.

Module	Role	Allows users to
auth	read	Extract user credentials, including name and e-mail address, from the auth database into the account and user group definitions. Used in conjunction with the cmdctrl write (with read) and admin roles.

## 4.2.5 Deleting a Framework User Group

- 1 Click *Framework User Manager* on the home page of the console.
- 2 Click *Groups* in the navigation pane.
- 3 In the left pane, select the group you want to delete.
- 4 Click *Delete Group* in the task pane.
- 5 Click *Finish* to confirm the deletion.

## 4.3 Deploying the Access Control Module

When you install Privileged User Manager, the Framework User Manager console is installed with the other required modules. If you want to manage Framework users from other hosts, you need to deploy the Access Control modules on these hosts. You should always have at least one host that contains a backup of this console.

The Access Control module consists of the following packages:

- ♦ **Access Manager (auth):** Holds the Framework user account information and controls access to the Framework modules.
- ♦ **Access Control Console:** Required for configuring Framework users and groups. It is installed into the Framework Manager console as the *Framework User Manager* console.

The Access Control module has the following dependencies:

- ♦ The Access Manager package is shown as an available package only on hosts that have the Registry Manager (registry) package installed.
- ♦ The Access Control Console can only be deployed on hosts that have the Administration Manager (admin) package installed.

To deploy the Access Control modules on another host:

- 1 Download the following packages to your local Package Manager:
  - ♦ Access Manager
  - ♦ Access Control Console
  - ♦ Registry Manager
  - ♦ Administration Manager

See [Chapter 2, “Managing Package Distribution,” on page 19](#) for details.

- 2 Install the Registry Manager on the host you want to be the Access Manager, then install the Access Manager on the same host.

This can be on any operating system, including Windows. See [Section 3.4.6, “Installing Packages on a Host,” on page 37](#) for details. The packages can be deployed to as many hosts as you need in order to build an environment with load balancing and failover.

- 3 Install the Administration Manager on the same host or a different host.

It can be deployed to as many hosts as you need in order to build an environment with load balancing and failover.

- 4 Install the Access Control Console on a host where the Administration Manager is installed. See [Section 2.2.2, “Adding a Console to the Framework Manager Console,” on page 21](#) for details.

The Access Control module is now deployed and ready to use.

## 4.4 Changing a Framework User’s Password

Framework users can change their own passwords by using the *Change Password* option, which is always available in the task pane. If a Framework user belongs to a group with the appropriate auth role defined (see [Section 4.2.4, “Configuring Roles,” on page 58](#)), he or she can also change other users’ passwords by using the *Modify User* option.

To change your own password:

- 1 Click *Change Password* in the task pane.
- 2 In the *Old Password* field, specify your current password.
- 3 In the *New Password* field, specify your new password and confirm it in the *Confirm Password* field.

Your password must comply with the default *Account Settings* for the Framework, and comply with individual user settings defined by using the *Modify User* option.

- 4 Click *Finish*.



The Command Control feature provides UNIX and Linux users with controlled access to privileged commands in a secure manner across the enterprise. Command Control enables the complete lockdown of user privilege by providing rules to determine the commands that are authorized to be run, and a powerful account delegation feature that removes the need for common access to the `root` account.

Command Control provides centralized logging of activity across all platforms, and enables the selective capture of session activity for any user, to the keystroke level, which can be viewed through the Compliance Auditor and reporting features.

Additional features include external scripting that provides the ability to authenticate via third-party security databases or applications, and comprehensive test suite tools that allow the administrator to model and test new rule combinations before committing them to production use.

- ◆ [Section 5.1, “How Does Command Control Work?,” on page 65](#)
- ◆ [Section 5.2, “Integrating Command Control into User Environments,” on page 66](#)
- ◆ [Section 5.3, “Importing Command Control Configuration Data,” on page 73](#)
- ◆ [Section 5.4, “Command Control Transactions,” on page 75](#)
- ◆ [Section 5.5, “Configuring Command Control,” on page 76](#)
- ◆ [Section 5.6, “Rules,” on page 81](#)
- ◆ [Section 5.7, “Command Control Groups,” on page 91](#)
- ◆ [Section 5.8, “Commands,” on page 98](#)
- ◆ [Section 5.9, “Scripts,” on page 102](#)
- ◆ [Section 5.10, “Access Times,” on page 107](#)
- ◆ [Section 5.11, “Command Control Reports,” on page 109](#)
- ◆ [Section 5.12, “Test Suites,” on page 111](#)
- ◆ [Section 5.13, “Deploying Command Control,” on page 116](#)

## 5.1 How Does Command Control Work?

UNIX/Linux commands are passed to the Command Control through scripts, commands, or replacement shells, using the methods described in [Section 5.2, “Integrating Command Control into User Environments,” on page 66](#).

When a command is received, Command Control uses the following to evaluate the command:

- ◆ The command is validated against configured rule criteria such as submit user, submit host, run host requested, date/time, and the command name itself.
- ◆ Any Perl scripts associated with the rules are executed, such as setting environment variables.

If the evaluation authorizes the command to run:

- ♦ The command is executed on the requested run host unless a matching rule specifies a run host. When a rule specifies a run host, the command is executed on that host. Rule values always overwrite values in the command request.
- ♦ The command is executed as the requested run user unless a matching rule specifies a run user. When a rule specifies a run user, the command is executed as that user. Rule values always overwrite values in the command request.

If the evaluation returns unauthorized, the command is not executed and the reason for the failure is returned.

## 5.2 Integrating Command Control into User Environments

Privileged User Manager provides a number of shells and functions that are installed as part of the Command Control Agent.

- ♦ The `rush` and `crush` shells, which are based on the Korn shell (`ksh`)
- ♦ The `usrun` command, which provides for command execution as a privileged user.

These shells and functions allow you to integrate Command Control into the UNIX and Linux user environments.

Name of the shells differ in the SLES environment:

UNIX or Linux Shell Name	SLES Shell Name
<code>crush</code>	<code>cpcksh</code>
<code>rrush</code>	<code>rpcksh</code>
<code>rush</code>	<code>pcksh</code>

Crush is normally used to audit users who do not need any additional privileges. With Novell Privileged User Manager, you can change a user's login shell to `crush` (`/usr/bin/crush`), then configure a Command Control rule to authorize `crush` and enable session capture. When the users log in, the commands are captured and audited through Novell Privileged User Manager. Auditing is transparent to the users, and the users perceive no change.

Rush is normally used to grant privileges, such as the ability to run commands as `root`. A user logs in as a non-privileged user, then issues a `usrun rush` command to access the root `ksh` shell. Privileged User Manager can be configured to perform a complete session capture of everything the user does as `root`. In addition, you can configure rules that set up command risks for specific commands. You can also configure rules that prevent the execution of specific commands.

Depending upon the users logging in to your host machines, you can create rules for one or more of the following scenarios:

- ♦ [Section 5.2.1, "Using `usrun` with a Command," on page 67](#)
- ♦ [Section 5.2.2, "Using `rush` for Privileged Sessions," on page 68](#)
- ♦ [Section 5.2.3, "Using `crush` for Complete Session Capture," on page 70](#)

- ♦ [Section 5.2.4, “Using rush for Complete Session Control,” on page 71](#)
- ♦ [Section 5.2.5, “Using Shell Scripts,” on page 73](#)

## 5.2.1 Using usrun with a Command

Type `usrun` before any command to pass the command to the Command Control Manager for authorization.

The `usrun` function can be used with the following options:

```
usrun [-b] [-p] [-t] [-x] [-u <user>] [-h <host>] <command>
```

Option	Description
-b	Puts the execution of the command into the background.
-p	Provides a pipe compatibility option for competitive products. It is only used for a competitive swap-out.
-t	Provides a test command option that tests the specified command against the rule structure. A yes or no is printed to the screen, indicating whether the command would be accepted or not.
-x	Enables the X11 forwarding option.
-u <user>	Specifies the user you want the command to run as, although this can be overwritten by the Command Control rules.
-h <host>	Specifies the host you want the command run on, although this can be overwritten by the Command Control rules. For <host> you can use either the hostname of the computer or the agent name specified in the Hosts console.
<command>	Specifies the command to pass to the Command Control Manager.

For example, to provide administrators access to the `passwd` command so they can change user passwords, you must define a Command Control rule to authorize the `passwd` command for the appropriate administrators.

- 1 Click *Command Control* on the home page of the console.
- 2 Add a password command:
  - 2a Click *Commands*, then click *Add Command* in the task pane.
  - 2b Specify a name, then click *Finish*.
  - 2c Select your new command, then click *Modify Command*.
  - 2d Fill in the following fields:
 

**Description:** Explain the purpose of this command. Specify something similar to the following:

Allows a user to submit a `usrun passwd` command to change account passwords.

**Commands:** Specify the following command.

```
passwd *
```
  - 2e Click *Finish*.

- 3 Add an Account User Group for the password command:
  - 3a Click *Account Groups > User Groups*, then click *Add User Group* in the task pane.
  - 3b Specify a name, then click *Finish*.
  - 3c Select your password user group, then click *Modify User Group*.
  - 3d Fill in the following fields:
 

**Description:** Explain the purpose of this user group. Specify something similar to the following:

Defines the user accounts that can run the `usrun passwd` command to change account passwords.

**Users:** Specify the usernames of the users on your Linux and UNIX hosts that have your permission to use the `usrun passwd` command.
  - 3e Click *Finish*.
- 4 Add a rule for the password command:
  - 4a Click *Rules*, then click *Add Rule* in the task pane.
  - 4b Specify a name, then click *Finish*.
  - 4c Select your password command, then drag it to your password rule.
  - 4d Select your passwd user group, then drag it to your password rule.
  - 4e Select your password rule, then click *Modify Rule* in the task pane.
  - 4f Fill in the following fields:
 

**Description:** Explain the purpose of this rule and the users it matches. Specify something similar to the following:

Matches users who submit a `usrun passwd` command. It authorizes their session and sets the run user to `root`.

**Session Capture:** Select *Off*.

**Authorize:** Select *Yes*, then select *Stop from the drop-down menu*.

**Run User:** Specify `root`.
  - 4g Click *Finish*.

Repeat this process for each command you want your users to run with the `usrun` command.

## 5.2.2 Using `rush` for Privileged Sessions

To use the `rush` shell, the user logs in using a non-privileged account. Then the user uses the `usrun` command to gain privileged shell access to perform administrative functions. The privileged shell performs complete session capture and can optionally audit the actual commands executed for use with Command Risk.

For this scenario, you must enable session capture and define a Command Control rule to authorize `rush` for the appropriate users.

- 1 Click *Command Control* on the home page of the console.
- 2 Add a `rush` command:
  - 2a Click *Commands*, then click *Add Command* in the task pane.
  - 2b Specify a name, then click *Finish*.

- 2c** Select your new command, then click *Modify Command*.
- 2d** Fill in the following fields:
- Description:** Explain the purpose of this command. Specify something similar to the following:
- When a user submits a `usrun rush` command or a `usrun shell` command, the command is rewritten to `/usr/bin/rush`. The Command Control Audit level is set to 1, which enables an additional level of audit to use with the Command Risk.
- Rewrite:** Specify the following:
- ```
/usr/bin/rush -o audit 1
```
- Commands:** Specify the following commands, each on a separate line.
- ```
rush
shell
```
- 2e** Click *Finish*.
- 3** Add an Account User Group for the rush shell:
- 3a** Click *Account Groups > User Groups*, then click *Add User Group* in the task pane.
- 3b** Specify a name, then click *Finish*.
- 3c** Select your rush user group, then click *Modify User Group*.
- 3d** Fill in the following fields:
- Description:** Explain the purpose of this user group. Specify something similar to the following:
- Defines the user accounts that can run the `usrun rush` command to access root privileges.
- Users:** Specify the usernames of the users on your Linux and UNIX hosts that have your permission to use the `usrun rush` command.
- 3e** Click *Finish*.
- 4** Add a rule for the rush command:
- 4a** Click *Rules*, then click *Add Rule* in the task pane.
- 4b** Specify a name, then click *Finish*.
- 4c** Select your rush command, then drag it to your rush rule.
- 4d** Select your rush user group, then drag it to your rush rule.
- 4e** Select your rush rule, then click *Modify Rule* in the task pane.
- 4f** Fill in the following fields:
- Description:** Explain the purpose of this rule and the users it matches. Specify something similar to the following:
- Matches users who submit a `usrun rush` or `usrun shell` command. It authorizes their session and enables session capture as root. The command assigned to this rule also included a rewrite that enables the additional level of audit to be used in conjunction with the command risk list. For information on this feature, see [Section 5.8.3, “Setting the Command Risk,” on page 101](#).
- Session Capture:** Select *On*.
- Authorize:** Select *Yes*, then select *Stop from the drop-down menu*.

**Run User:** Specify `root`.

**4g** Click *Finish*.

- 5** (Conditional) If your users need different environment variables to run some of their privileged commands, you can use a script to set up the values.

By default, Command Control uses the environment variables of the executing user. If your users receive a “not found” message for a command, you need add environment variables to the rule. For configuration information, see [Section 5.9, “Scripts,” on page 102](#) and [“Modify Environment Script” on page 105](#).

You can also define illegal commands, including built-in shell commands, in a script assigned to the rule. For configuration information, see [Section 5.9, “Scripts,” on page 102](#) and [“rush Illegal Commands Script” on page 107](#).

### 5.2.3 Using crush for Complete Session Capture

This method of integration provides the most auditing functionality. By changing the user’s shell to the crush client instead of the rush client, Command Control can be configured to capture the user’s complete session, in addition to all other audit and control features.

When the user logs in to the server, the session is started with the crush client, which executes as a normal Korn shell. A request is sent to the Command Control Manager for authorization. You must define a crush rule that enables session capture, as described in the steps below. Functions and aliases that can replace normal system commands are read from `/etc/profile.rush`. When the user issues a command that needs privileges to run, it is executed through the Command Control system.

- 1** Use the tool provided in the UNIX or Linux environment to set the user login shell to

```
/usr/bin/crush
```

- 2** Add a crush command:

**2a** Click *Commands > Add Command*.

**2b** Specify a name (for example, `crush shell`), then click *Finish*.

**2c** Select the name of the crush command, then click *Modify Command*.

**2d** Fill in the following fields:

**Description:** Explain the purpose of the rule. Specify something similar to the following:

When a user’s shell is set to `/usr/bin/crush` and the user logs in, a Command Control request is sent with a submitting command of `-crush` to indicate login. The user’s login shell is rewritten to `/usr/bin/rush`. The Command Control Audit level is set to 1, which enables an additional level of audit to use with the Command Risk.

**Rewrite:** Specify the following:

```
/usr/bin/rush -o audit 1
```

**Commands:** Specify the following:

```
-crush
```

**2e** Click *Finish*.

- 3** Add a crush Account User Group:

**3a** Click *Account Groups > User Groups*, then click *Add User Group* in the task pane.

**3b** Specify a name, then click *Finish*.

- 3c** Select your crush user group, then click *Modify User Group*.
- 3d** Fill in the following fields:
  - Description:** Explain the purpose of this user group. Specify something similar to the following:  
Defines the user accounts that can use the crush command.
  - Users:** Specify the usernames of the users on your Linux and UNIX hosts that have your permission to use the `crush` command.
- 3e** Click *Finish*.
- 4** Add a crush rule:
  - 4a** Click *Rules > Add Rule*.
  - 4b** Specify a name, then click *Finish*.
  - 4c** Select your crush command, then drag it to your crush rule.
  - 4d** Select your crush user group, then drag it to your crush rule.
  - 4e** Select your crush script, then drag it to your crush rule.
  - 4f** Select your crush rule, then click *Modify Rule* in the task pane.
  - 4g** Fill in the following fields:
    - Description:** Explain the purpose of this rule. Specify something similar to the following:  
Authorizes the matching of submit users who have `/usr/bin/crush` as their defined login shell. It authorizes their session and enables session capture, when they are still running as their original login ID.
    - Session Capture:** Select *On*.
    - Authorize:** Select *Yes*, then select *Stop if authorized* and from the drop-down menu. These settings allow subsequent commands to be executed without authorization checks whenever the user has had one command authorized.
  - 4h** Click *Finish*.

## 5.2.4 Using `rush` for Complete Session Control

You can change the user's login shell to the `rush` client so that no authorization request is sent when the user logs on. This provides a method of integration that is invisible to the user.

The `rush` client executes as a normal Korn shell. Functions and aliases that replace normal system commands are read from `/etc/profile.rush`. When the user issues a command that needs privileges to run, it is authorized through the Framework.

- 1** Use the tool provided in the UNIX or Linux environment to set the users' shell to  
`/usr/bin/rush`
- 2** To ensure that configured commands are authorized at the Framework, add the following line to either the user's `.profile` file or to the central `profile.rush` file in the `/etc` directory on the appropriate UNIX or Linux servers:  
`set -o remote`

---

**IMPORTANT:** The `set -o remote` option forces all commands that are not built in to the user's shell to be authorized at the Framework. Commands for which a defined rule does not exist are not permitted to execute. To prevent all commands in the `profile.rush` file or `.profile` file from being passed to the Framework for authorization, add the `set -o remote` command at the end of the file.

---

**3** (Optional) Set the following additional options in the profile file:

---

Option	Description
<code>set -o host &lt;hostname&gt;</code>	Specifies that all authorized commands are executed on the defined host, if permitted.
<code>set -o user &lt;username&gt;</code>	Specifies that all authorized commands are executed as the defined user if permitted.
<code>set -o audit &lt;n&gt;</code>	Enables auditing, Set <n> to one of the following values: <ul style="list-style-type: none"><li>◆ <b>1:</b> Enables auditing of all commands that are not built into the user's shell.</li><li>◆ <b>2:</b> Enables auditing of all commands including commands that are built into the user's shell. This level of auditing can affect login times.</li></ul> After the audit value has been set, it cannot be changed. If it is turned on in the profile, the user cannot turn it off later.
<code>set -o ignoreperm</code>	Enables commands that have not been successfully authorized at the Framework to execute according to the local permissions in effect on the server where the command was issued.
<code>set -o test</code>	Allows typed commands to be checked to see if they would be accepted by the rule structure.  A yes or no output to screen indicates the result.  The <code>set -o test</code> option is normally used in conjunction with the <code>set -o remote</code> option.
<code>set -o test '\${}{'</code>	Returns the complete metadata result that is generated by the Command Control manager.

---

Rule definitions override the settings for user and host. If a successfully matched rule specifies a run user or a run host, this user or host is used to execute the command, and not those specified in the `set -o` commands.

You can use rule conditions to match the run user or run host to the username or hostname defined by using these commands (see [“Setting Conditions for a Rule” on page 84](#)), but if a run user or run host is defined in the rule configuration, these are the ones that are used.

You can define a list of illegal commands, including built-in shell commands, in a script assigned to a rule. Users using the rush shell cannot run these commands, even if they are `root`.

## 5.2.5 Using Shell Scripts

You can hide some of the complexities of the privileged command syntax from your users by using scripts and aliases to wrap privileged tasks. Using this technique, the end user can log in with their non-privileged account and use what appear to be standard commands to perform privileged tasks. Alternatively, you could create a script that provides a menu system to access a set of administrative tasks. With this method, the user would simply select options from the menu to perform their privileged tasks.

Either method requires a shell script that executes under the `rush` shell and performs remote authorization. For example:

```
#!/usr/bin/rush
set -o remote
passwd $*
```

This script executes the `rush` client, sets it to use Command Control, and executes the `passwd` command.

## 5.3 Importing Command Control Configuration Data

You can import a complete command control configuration database, including test suites, using the `Import Settings` option, or you can import test suites only, using the `Import Test Suites` option under `Test Suites`.

If you import a complete command control configuration database, all existing data is overwritten, including test suites. If you import test suites only, they are added to your existing configuration and do not overwrite your existing test suites.

- ♦ If you want to import a complete command control configuration database, see [Section 5.3.1, “Importing Command Control Settings,”](#) on page 73.
- ♦ If you want to import test suites only, see [Section 5.12.8, “Importing a Test Suite,”](#) on page 115.

### 5.3.1 Importing Command Control Settings

You can use the `Import Settings` option to restore a previously backed-up version of your Command Control configuration settings, or to import Command Control configuration settings from another Framework. You then use the `Export Settings` option to obtain configuration settings so you can paste them into a text document for backup or for use on another Framework.

---

**IMPORTANT:** This process overwrites your existing configuration settings.

---

- 1 Access the Command Control configuration settings you need and copy the whole configuration.
- 2 Click *Command Control* on the home page of the console.
- 3 Click *Import Settings* in the task pane.

- 4 Click in the text area, then use Ctrl+V to paste the copied settings, or right-click in the text area and click *Paste*.
- 5 Click *Finish*.

To use a command line option to import Command Control settings, see [Section 10.2.1, “Importing and Exporting Command Control Settings,” on page 150](#).

## 5.3.2 Exporting Command Control Settings

You can export your Command Control configuration settings to a text file for backup purposes, or for use in another Framework. You use the *Import Settings* option to restore the backed-up configuration settings, or to import the settings into another Framework.

---

**NOTE:** Novell recommends that you take frequent backups of your Command Control configuration settings.

---

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Export Settings* in the task pane.
- 3 Use Ctrl+A to select all your Command Control configuration settings, or right-click in the text window and click *Select All*.
- 4 Use Ctrl+C to copy the settings, or right-click in the text window and click *Copy*.
- 5 Paste the text into a text document and save it.
- 6 Click *Finish*.

To use a command line option to export Command Control settings, see [Section 10.2.1, “Importing and Exporting Command Control Settings,” on page 150](#).

## 5.3.3 Importing Command Control Samples

Novell has provided a set of sample commands and Perl scripts to assist you with configuring your Command Control rules.

To add these samples to your configuration:

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Import Samples* in the task pane.
- 3 Select the samples you want.  
To select multiple samples in a folder, display the samples, then press the Ctrl key and select the required samples one at a time, or press the Shift key and select a consecutive list of samples. You cannot import samples by selecting a folder.
- 4 Click *Finish*. The samples are added to the appropriate section of the configuration.

## 5.4 Command Control Transactions

Your Command Control database can be protected through the use of the Transactions feature, which automatically locks the database when you start making changes and prevents other Framework users from making any changes. You must then commit the transaction to save the changes and release the lock, and you are prompted by customized questions to provide information that can be viewed in the Compliance Auditor. You can cancel the transaction at any time.

To use this feature, you must first enable it and create a customized Commit Transactions page, then you can use the feature and commit the changes you have made.

- ♦ [Section 5.4.1, “Enabling Transactions and Configuring Settings,” on page 75](#)
- ♦ [Section 5.4.2, “Making Command Control Configuration Changes with Transactions Enabled,” on page 75](#)
- ♦ [Section 5.4.3, “Committing a Transaction,” on page 76](#)

### 5.4.1 Enabling Transactions and Configuring Settings

You can configure the Command Control Manager to require the Transactions feature to be used when configuring Command Control rules.

You can also configure your own Commit Transaction page to be used for committing a transaction. The data entered on the Commit Transaction page can be viewed in the Compliance Auditor.

To configure this feature:

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Transaction Settings* in the task pane.
- 3 Select the *Enable transactions* check box to enable the use of Command Control transactions.
- 4 Click *Add*.
- 5 Specify a name for the field you want to display when a user commits a transaction. For example, you might want to request the user’s name, so you could specify *Name*. A blank field called *Name* appears on the screen when a user commits a transaction.
- 6 Select *Text* if you want the user to enter one line of text, or select *TextArea* if you want the user to be able to enter several lines of text.
- 7 Select *required* if you want to force the user to enter text in this field. The *Finish* button on the *Commit Transaction* page does not become available until the user has entered text in this field.
- 8 Repeat [Step 4](#) through [Step 7](#) for any other fields you want to display when the user commits the transaction.
- 9 Select *Finish*.

### 5.4.2 Making Command Control Configuration Changes with Transactions Enabled

- 1 Click *Command Control* on the home page of the console.
- 2 Make the configuration changes you want.

A message appear next to *Command Control* in the navigation pane to indicate that the Command Control database is locked, by whom, and when it was locked.

- 3 Click *Command Control* in the navigation pane, then click *Commit Transaction*. Complete the fields as set up on the Transaction Settings page, then click *Finish*.

Alternatively, if you do not want to keep the changes you have made, select *Cancel Transaction* in the task pane and select *Yes* to confirm. Any changes you have made since the database was locked are removed.

### 5.4.3 Committing a Transaction

When you have finished changing your Command Control database, you must commit your transaction to save the changes and release the lock on the database. The Commit Transaction page can be customized to request whatever information you require when a transaction is committed (see “[Enabling Transactions and Configuring Settings](#)” on page 75 for details).

To commit a transaction:

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Commit Transaction* in the task pane.
- 3 To create a backup of the Command Control database before the pending changes are written to the database, enable the Create backup option, then specify a reason for the backup in the text box.
- 4 Complete the customized fields according to company policies.
- 5 Click *Finish*.

## 5.5 Configuring Command Control

Command Control uses rules to protect and control user commands. When configuring a rule, you need to set rule conditions to determine which rule or rules are processed, depending, for example, on the command submitted or the user who submitted it. You also need to define what processing to do if the rule conditions are matched.

The components you can define and configure for a rule are as follows:

- ♦ The rule itself. For configuration information, see [Section 5.6, “Rules,”](#) on page 81.
- ♦ Account groups, user groups, and host groups, which determine who matches the rule. For configuration information, see [Section 5.7, “Command Control Groups,”](#) on page 91.
- ♦ Commands. For configuration information, see [Section 5.8, “Commands,”](#) on page 98.
- ♦ Scripts for additional functionality. For configuration information, see [Section 5.9, “Scripts,”](#) on page 102.
- ♦ Access times to define specific times during which access is denied or granted. For configuration information, see [Section 5.10, “Access Times,”](#) on page 107.

---

**NOTE:** To enable access to the Command Control console for a Framework user and to control the level of access available, you must add the user to a group with the appropriate roles defined. See [Section 4.2.4, “Configuring Roles,”](#) on page 58 for details.

---

The following additional features are provided to assist you with Command Control configuration and management:

- ♦ [Section 5.5.1, “Defining Audit Settings,”](#) on page 77

- ♦ Section 5.5.2, “Backing Up and Restoring,” on page 78
- ♦ Section 5.5.3, “Finding a Reference,” on page 79
- ♦ Section 5.5.4, “Defining Custom Attributes,” on page 79
- ♦ Section 5.5.5, “Functions,” on page 79
- ♦ Section 5.5.6, “Adding a Category,” on page 81
- ♦ Section 5.5.7, “Deleting a Category,” on page 81
- ♦ Section 5.3, “Importing Command Control Configuration Data,” on page 73
- ♦ Section 5.4, “Command Control Transactions,” on page 75
- ♦ Section 5.12, “Test Suites,” on page 111

## 5.5.1 Defining Audit Settings

All Command Control audit records contain the following information:

- ♦ Submit details such as the submitting username, hostname, and primary group.
- ♦ Target details such as the run username and the run hostname.
- ♦ Command details, which include the original command requested and the actual command run.
- ♦ Authorization status, either yes or no.
- ♦ Session capture status, either yes or no.
- ♦ Audit ID, which is the unique ID used to group audit events for the user’s session.
- ♦ Codeset, which is the character encoding used for localization.
- ♦ Terminal details such as tty name, terminal dimensions, and type.

The *Audit Settings* option allows you to modify this default record and add the following:

- ♦ Encryption of sensitive password data in keystroke capture reports along with a password that allows authorized Framework administrators to decrypt it.
- ♦ Additional options that can be audited for each record.

To define audit settings:

- 1** Click *Command Control* on the home page of the console.
- 2** Click *Audit Settings* in the task pane.
- 3** Configure the Password keystroke settings:
  - 3a** Select the *Password filter* check box.
  - 3b** In the *Password filter* text box, specify the text that is used to prompt users for their passwords.  
 For example, if your systems request a user’s password by using the word `Password`, specify `Password` in this field. If your systems use `password`, enter `password` in this field. If your systems use either, enter `assword` in this field. This ensures that the password the user enters in response to this prompt is encrypted in reports.
  - 3c** Select the *Encryption password* check box.
  - 3d** In the *Encryption password* text box, specify the password to be used to decrypt the sensitive password data in the report.

This password must be entered on the *Command Control Keystroke Report* page to decrypt the password data.

- 3e Specify the password again in the *Confirm password* text box.
- 4 (Optional) Select from the following check boxes to add more information to the audit record:
  - Command:** Complete information about the command being run, including the actual filename and arguments.
  - Host:** Information about the submitting host
  - Environment:** Complete list of the environment variables passed to the executed command.
  - Local time:** The time on the machine that submitted the request.
  - Cwd:** Details about the current working directory where the command was executed.
  - Options:** Details about the various process control options for executing the command.
  - Run Account:** Information about the account that is used to execute the command.
  - Process:** Details about the process that submitted the request.
  - Jobs:** The job control setting that were passed to the executed command.
  - Passwd:** Details of the `/etc/passwd` entry for the user submitting the request.
  - Groups:** The group membership details for the executed command.
  - Logon:** The logon time and source for the user submitting the request.
- 5 Click *Finish*.

## 5.5.2 Backing Up and Restoring

The backup option allows you to create snapshots of the command control database and restore these snapshots at future date. You can back up and restore from the Framework Manager console, but you need to use the command line to remove a backed-up snapshot. For information about the command line options, see [Section 10.2.2, “Backing Up and Restoring a Command Control Configuration,” on page 151](#).

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Backup and Restore*.
- 3 To back up the database, specify a reason for the backup, then click *Backup*.
- 4 To restore a previous version of the database, select the version, then click *Restore*.  
The current version is overwritten by the selected version.
- 5 Click *Close*.

The following information is recorded for each backed-up version:

**Date:** The date and time the backup was performed.

**Administrator:** The user that performed the backup.

**Reason:** The reason for performing the backup. This is optional information, but recommended.

### 5.5.3 Finding a Reference

The *Find References* option allows you to find where a specific account group, user group, host group, command, script, or access time is referenced in the database. For example, you could use this option to find out which account group or groups a specific user group belongs to.

- 1 Click *Command Control* on the home page of the console.
- 2 Select the entity for which you want to find references.
- 3 Click *Find References* in the task pane. The groups or rules in which the entity is referenced are displayed.
- 4 To go to one of the listed groups or rules, double-click it, or to return to the navigation pane, click *Close*.

### 5.5.4 Defining Custom Attributes

Custom attributes can be defined for account groups, user groups, host groups, commands, and access times to provide additional parameters for use in scripts. For example, you could set an expiration date as a custom attribute for a user group, check for this date in your script, then expire the user group when the date is reached.

To define custom attributes:

- 1 Click *Command Control* on the home page of the console.
- 2 Select the entity you want to add custom attributes to.
- 3 Click *Custom Attributes* in the task pane.
- 4 Click *Add*.
- 5 In the *Name* field, specify the name of the custom attribute, such as `Expiration date`.
- 6 In the *Value* field, specify the value for the attribute, such as the date you want the entity to expire.
- 7 Repeat [Step 4](#) through [Step 6](#) for any other custom attributes you want to add.
- 8 Click *Finish*.

### 5.5.5 Functions

The `udsh` command invokes commands on a set of hosts. It concurrently issues a *Command Control* request for each host that is specified and returns the output from all the hosts, formatted so that command results from all hosts can be managed.

- ♦ [“Syntax” on page 79](#)
- ♦ [“Options” on page 80](#)
- ♦ [“Keywords” on page 80](#)

#### Syntax

```
udsh [-bcdqv] [-t <timeout>] [-l <user>] [-f <num>] [-w <host>, <host  
wildcard>] [-g <hostgrp>, <hostgrp wildcard>] [cmd ...]
```

## Options

The following options can be specified only on the command line:

**Table 5-1** *udsh Options*

Option	Description
-b	Do not break lines to column width when displaying output.
-c	Do not remove the host from the list if the command fails.
-d	Add a time stamp to the displayed output.
-f <num>	Specify the maximum number of concurrent processes to run.
-g <hostgrp>, <hostgrp wildcard>	Specify the Command Control host groups to retrieve the list of agents to run the command on. Wildcards must be properly escaped. For example to run udsh against all host groups that begin with ho, enter the following:  -g ho\*
-l <user>	Specify the user to run the command as.
-q	Quiet. Do not display output.
-t <timeout>	Specify the timeout in seconds for the command to complete on each host.
-v	Verbose output.
-w <host>, <host wildcard>	Specify the agents to run the command on. Wildcards must be properly escaped. For example, to run udsh against all hosts that begin with host1, enter the following:  -w host1\*

If a command is not specified, the user is placed at a command prompt. Each entry run from this prompt is run separately on each host. If readline(3) is available, command line editing and history are provided.

## Keywords

There are various macros that can be specified in the command to substitute keywords when the command is run on the remote host. For example, the following command uses the `${rhost}$` keyword. It performs a `usrun echo` command of the remote host name on all agents that have a command control agent deployed:

```
udsh -w \* /bin/echo '${rhost}$'
```

**Table 5-2** *udsh Keywords*

Keyword	Description
<code>\${uid}\$</code>	Calling user's UID

Keyword	Description
<code>\${gid}\$</code>	Calling user's primary group ID
<code>\${gecos}\$</code>	Calling user's gecos
<code>\${home}\$</code>	Calling user's home directory
<code>\${shell}\$</code>	Calling user's shell
<code>\${cwd}\$</code>	Calling user's current working directory
<code>\${lhost}\$</code>	Local hostname
<code>\${rhost}\$</code>	Remote hostname
<code>\${pid}\$</code>	PID of the individual <code>udsh</code> call
<code>\${ppid}\$</code>	PID of the <code>udsh</code>

### 5.5.6 Adding a Category

You can use the appropriate *Add Category* option to group your account groups, user groups, host groups, commands, scripts, and access times into categories for ease of use and maintenance.

- 1 Click *Command Control* on the home page of the console.
- 2 Select the section to which you want to add a category. You can also add subcategories to existing categories.
- 3 Click the *Add Category* option in the task pane.
- 4 Specify a name for the category.
- 5 Click *Finish*.

### 5.5.7 Deleting a Category

Before deleting a category, you must delete or move the items and subcategories that it contains.

- 1 Click *Command Control* on the home page of the console.
- 2 Select the category you want to delete.
- 3 Click the *Delete Category* option in the task pane. The category is deleted.

## 5.6 Rules

Rules provide the means by which you can control commands. Commands can be authorized to run, or not authorized to run, by setting rule conditions based on different criteria:

- ♦ The command being submitted
- ♦ The user and host submitting the command
- ♦ The user and host assigned to run the command
- ♦ The time the command is submitted
- ♦ The contents of Perl scripts you have defined.

See [“Setting Conditions for a Rule” on page 84](#) for details.

If a rule's conditions are met, there are a number of options you can set to determine how the rule processes the command. You can configure a rule to:

- ◆ Display a message to the user submitting the command
- ◆ Capture the user session for reporting and auditing purposes
- ◆ Authorize or not authorize the command to be run
- ◆ Specify what further rule processing to do. The rule can specify that the processing of additional rules ends by using the stop conditions (*Stop*, *Stop if authorized*, *Stop if unauthorized*).

When the Framework Manager receives a command request, the evaluation starts at the top of the rule tree. Even when a request matches a rule, the evaluation continues until a rule has a stop condition or the rule tree has been processed.

You can also:

- ◆ Specify the user and host to run the command
- ◆ Set a risk level for use with keystroke reports
- ◆ Assign an audit group to the rule for use with the Compliance Auditor.

See [“Modifying a Rule” on page 83](#) for details.

You can also create and assign Perl scripts to the rule to provide additional functionality. See [“Adding a Script” on page 103](#) and [“Assigning a Script to a Rule” on page 86](#) for details.

---

**NOTE:** If you are using a different user (run user) to run an authorized command than the user who submitted the command (submit user), by default the submit user's environment variables are used for the run user. If you want to use the environment variables associated with the run user, you can add a script to your rule containing the following text:

```
$meta->get_params("Job")->arg("job_default_env",0);  
return 1;
```

- 
- ◆ [Section 5.6.1, “Adding a Rule,” on page 83](#)
  - ◆ [Section 5.6.2, “Modifying a Rule,” on page 83](#)
  - ◆ [Section 5.6.3, “Setting Conditions for a Rule,” on page 84](#)
  - ◆ [Section 5.6.4, “Removing Conditions for a Rule,” on page 85](#)
  - ◆ [Section 5.6.5, “Configuring Script Arguments and Entities for a Rule,” on page 85](#)
  - ◆ [Section 5.6.6, “Assigning a Script to a Rule,” on page 86](#)
  - ◆ [Section 5.6.7, “Removing Script Arguments and Entities,” on page 86](#)
  - ◆ [Section 5.6.8, “Removing a Script from a Rule,” on page 86](#)
  - ◆ [Section 5.6.9, “Populating LDAP Group Scripts,” on page 87](#)
  - ◆ [Section 5.6.10, “Finding a Rule,” on page 89](#)
  - ◆ [Section 5.6.11, “Moving a Rule,” on page 89](#)
  - ◆ [Section 5.6.12, “Copying a Rule,” on page 89](#)
  - ◆ [Section 5.6.13, “Linking a Rule,” on page 90](#)

- ♦ [Section 5.6.14, “Deleting a Rule,” on page 90](#)
- ♦ [Section 5.6.15, “Viewing Pseudocode,” on page 90](#)

## 5.6.1 Adding a Rule

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 To add a rule at the top level, click *Add Rule* in the task pane. To add a rule as a child of another rule, select the rule and click *Add Rule* in the task pane.
- 4 Specify a name for the rule.
- 5 Click *Finish*. The new rule is added.
- 6 To configure the rule, select the rule, then click *Modify Rule* in the task pane.  
For configuration information, see [Section 5.6.2, “Modifying a Rule,” on page 83](#).
- 7 Move the rule to the correct position according to the order in which you want to process your rules.

When a user issues a command under Command Control, the following rule processing takes place:

- ♦ The conditions set for the first rule in the hierarchy are checked.
- ♦ If there is a match, the rule is processed. Depending on how the rule is configured, processing of additional rules takes place or stops. If rule processing is not stopped, the next rule for which conditions are checked is the child of this rule. Rule checking and processing continues until it is stopped by a rule, or until all appropriate rules have been processed.
- ♦ If there is no match, the conditions for the next rule at the same hierarchical level as the first rule are checked, and this continues until a match is found. Rule processing then takes place as described above.

You can change the default order of rule processing on the *Modify Rule* screen, or by using scripts. See [“Modifying a Script” on page 103](#).

## 5.6.2 Modifying a Rule

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Select the rule you want to modify.
- 4 Click *Modify Rule* in the task pane.
- 5 Make the changes you want:
  - Name:** Change the name of the rule.
  - Disabled:** To disable the rule, select the *Disabled* box. A disabled rule is dimmed.
  - Description:** Specify a description of the rule.
  - User Message:** Specify a user message to be displayed to the user when this rule is processed, before any commands are run.

**Session Capture:** Select either *On* or *Off*. Setting *Session Capture* to *On* allows the Audit Manager to perform keystroke logging for the rule. To view a captured session from a Command Control report, an Auditing Manager and the Reporting Console must be installed.

**Authorize:** Select either *Yes* or *No*, depending on whether you want the command protected by the rule to be authorized or not authorized if the rule conditions are met.

Define what happens next by using the drop-down list as follows:

- ♦ **Blank:** The next rule in the hierarchy is checked.
- ♦ **Stop:** No more rules are checked for the command.
- ♦ **Return:** The next rule to be checked is up one level in the hierarchy from the current rule.
- ♦ **Stop if authorized:** If *Authorize* is set to *Yes*, no more rules are checked for the command.
- ♦ **Stop if unauthorized:** If *Authorize* is set to *No*, no more rules are checked for the command.

**Run User:** Define a run user by selecting the name of the user you want to run this command (this overrides any username defined through a set command).

**Run Host:** Define a run host by selecting the name of the host on which you want to run this command (this overrides any hostname defined through a set command).

**Risk Level:** Set a *Risk Level* of 0 to 99. This option allows you to set a value representing the relative risk of a rule when using the rush or crush clients with the session auditing option (see [Section 5.2, “Integrating Command Control into User Environments,” on page 66](#)). When viewing a Command Control Keystroke Report, you see commands controlled by rules with different risk values represented in different colors.

**Audit Group:** Define an *Audit Group*. This setting is for use in Compliance Auditor reports.

- 6 Click *Finish*. The settings you have defined for the rule are displayed in the console.

### 5.6.3 Setting Conditions for a Rule

You can set a number of conditions for a rule to determine whether the rule is processed or not. For example, you can set a particular command as a condition, and only process the rule if a user enters that command.

There are two ways of setting conditions for a rule:

- ♦ Dragging an entity onto the rule.
- ♦ Using the *Edit Condition* option, as described in the steps below.

---

**NOTE:** When you drag an entity onto a rule, you might need to edit the condition to ensure that the condition logic is what you want. If you want to use a script in rule conditions, you must set it to Conditional first (see [“Modifying a Script” on page 103](#)).

---

To set conditions by using the *Edit Condition* option:

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Select the rule for which you want to set conditions.
- 4 Select the currently defined condition in the right pane. If you have not yet defined a condition, this is *Match All*.

- 5 Select *Edit Condition* in the task pane.
- 6 In the *Add Condition* drop-down list, select the type of condition you want. The condition is displayed on the screen.
- 7 Set the condition to the value and logic you want. For example, if you set a condition to match a run user to a user group:
  - 7a Change *user* (submit user) to *run user*.
  - 7b Leave the logic setting as *IN*.
  - 7c Select the user group you require from the user group drop-down list.
- 8 Repeat [Step 6](#) and [Step 7](#) for any other conditions you want. Set the condition logic as necessary.
 

You can use parentheses to group conditions according to the necessary logic by selecting the parentheses ( ) entry from the *Add Condition* drop-down list. The opening and closing parentheses are displayed.

  - 8a Select the opening parenthesis.
  - 8b Select the condition type you want to place inside the parentheses and set it as necessary.
  - 8c Select the opening parenthesis again.
  - 8d Select another condition type to place inside the parentheses and set it as necessary.
  - 8e If necessary, change OR to AND.
  - 8f Repeat [Step 8d](#) through [Step 8f](#) for any other conditions you require inside this set of parentheses. You can also place parentheses within parentheses.
- 9 Click *Finish*.

#### 5.6.4 Removing Conditions for a Rule

You can remove all the conditions for a rule, or you can remove individual conditions.

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Use the arrow to display the rules and select the rule for which you want to remove conditions.
- 4 Select the currently defined condition in the right pane.
- 5 To remove all conditions, click *Remove Conditions* in the task pane, then click *Yes*.  
The rule condition is returned to *Match All*.
- 6 To remove individual conditions, click *Edit Condition* in the task pane, select the condition to remove, then click *Finish*.

#### 5.6.5 Configuring Script Arguments and Entities for a Rule

You can configure script arguments and entities for the scripts assigned to a rule before or after assigning the scripts. You can define only one set of arguments and entities, which applies to all scripts assigned to a rule.

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Select the rule for which you want to add script arguments.

- 4 Click *Script Arguments* in the task pane.
- 5 Click *Add*.
- 6 In the *Name* field, specify a name for the argument.
- 7 In the *Value* field, specify a value for the argument.
- 8 To add more arguments, repeat [Step 5](#) through [Step 7](#).
- 9 When you finish adding arguments, click *Finish*, or continue with [Step 10](#) to add script entities.
- 10 Click the arrow under *Add Script Entity* to display the list of available entities, then select the type of entity you want.  
A drop-down list of entities is displayed in the *Script Entities* table.
- 11 Select the entity you want from the drop-down list.
- 12 To add more entities, repeat [Step 10](#) and [Step 11](#).
- 13 Click *Finish*.

### 5.6.6 Assigning a Script to a Rule

You can use Perl scripts to provide additional, customized functionality to your rules (see [“Adding a Script” on page 103](#)). To assign a script to a rule, use drag and drop as described in the following procedure.

---

**NOTE:** If you drag a script that has been set to Conditional, the script is added to the rule conditions.

---

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Click the arrow to display the list of rules.
- 4 Click *Scripts* in the navigation pane.
- 5 Select the script you want to assign to the rule.  
To select multiple scripts in the same category, press the Ctrl key and select the required scripts one at a time, or press the Shift key to select a consecutive list of scripts.
- 6 Drag the selected scripts to your rule.
- 7 Configure script arguments and entities for the scripts if necessary. For more information, see [“Configuring Script Arguments and Entities for a Rule” on page 85](#).

### 5.6.7 Removing Script Arguments and Entities

- 1 To remove a script argument, select the argument, then click *Remove*.
- 2 To remove a script entity, select the icon next to the name of the entity, then click *Remove*.

### 5.6.8 Removing a Script from a Rule

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.

- 3 Use the arrow to display the list of rules, then select the rule from which you want to remove a script.
- 4 Select the script you want to remove in the right pane.  
To select multiple scripts, press the Ctrl key and select the required scripts one at a time, or press the Shift key to select a consecutive list of scripts.
- 5 Click *Remove Script* in the task pane.
- 6 Click *Yes* to confirm the removal. The scripts are removed from the rule.

## 5.6.9 Populating LDAP Group Scripts

Command Control policies give you additional options to control the execution of commands. For example, you can use a policy to restrict the rights and roles of a command so that the command works only for one particular directory, file, network address, or system call.

- ♦ [“Configuring a Command Control Policy” on page 87](#)
- ♦ [“Configuring a Path Policy” on page 87](#)

### Configuring a Command Control Policy

A command control policy is defined by using the policy script arguments. A policy script argument specifies the access rights of the applications based on the path, network, and capability.

- 1 Click *Command Control* on the home page of the console.
- 2 From the *Command Control Sample Scripts*, add the *Enhanced Access Control Policy* script.
- 3 Drag the *Enhanced Access Control Policy* script from *Scripts* to *Authorizing Rule*.
- 4 Click the *Authorizing Rule* and access the *Script Arguments*.
- 5 Create a script argument with a name *policy* and add that policy to the *Value* field.

### Configuring a Path Policy

A Path policy is a type of command control policy that restricts an application from accessing a specific directory based on the path.

The syntax of a Path policy is as follows:

```
path [owner] <path> <capability:capability:!capability>
```

*owner* specifies the file or directory ownership that should match with the current user ID.

*path* specifies a particular directory based on the path. Replace *path* with any of the following options:

**Table 5-3** Path Options

Option	Description
<code>/dir/file</code>	Specifies the file that the application can access in the <code>/dir/directory</code> .
<code>/dir/</code>	Specifies the directory that the application can access.

Option	Description
<code>/dir/f*</code>	Specifies a file that begins with <code>f</code> in the <code>/dir/</code> directory that the application can access.
<code>/dir/*</code>	Specifies that the application can access all the files in the <code>/dir/</code> directory.
<code>/dir/**</code>	Specifies that the application can access all the files and the subdirectories within the <code>/dir/</code> directory.
<code>/dir/**/</code>	Specifies that the application can access all subdirectories that are recursively searched for in the <code>/dir/</code> directory.
<code>/dir/**/*</code>	Specifies that the application can access all the files that are recursively searched for in any subdirectory within the <code>/dir/</code> directory.

`capability` specifies the rights of the application. You can use the `!` symbol in the syntax to denote a logical *not*. For example, `all:!write` grants all the rights except the `write` role.

Replace `capability` with any of the following options:

**Table 5-4** *Capability Options*

Option	Description
<code>privperms</code>	Enables the application with the <code>read</code> , <code>write</code> , and <code>ownership</code> permissions for the specified directory or file.
<code>perms</code>	Enables the application to assign the permissions of a specified directory or file.
<code>read</code>	Enables the application to assign the <code>read</code> permission for a specified directory or file.
<code>write</code>	Gives the application the <code>create</code> and <code>write</code> permissions for the specified directory or file.
<code>unlink</code>	Gives the application the <code>deletion</code> rights for the specified directory or file.
<code>mknod</code>	Enables the application to create system files in the specified directory.
<code>exec</code>	Enables the application to execute the shared files and files for which the application does not have <code>read</code> and <code>write</code> permission.
<code>unsafe</code>	Enables the application to execute any file that does not inherit the policy.
<code>link</code>	Enables the application to create a symbolic link or hard link to another file.
<code>log[=&lt;0-9&gt;]</code>	Enables the application to audit system calls, with an optional risk value of 0-9.
<code>all</code>	Enables the application to have all permissions.

You can use wildcards, regular expressions, and strings in the Path policy. For example, using the word `default` in the following example specifies the default policy.

```
path default all:log
path /opt/oracle/private/** !all:log=9
```

## 5.6.10 Finding a Rule

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 To find a rule from the entire list of rules, click *Find Rule* in the task pane.  
or  
To find a rule in a set of rules, select the parent rule, then click *Find Rule*.
- 4 In the *Rule Filter* field, specify the name of the rule you are looking for, then select *Find*.  
You can use wildcard characters \* and ?. For example, rul\* finds the first rule beginning with “rul”. This field is case sensitive.
- 5 If the rule name you are looking for is displayed, double-click it to return to the navigation pane with the rule selected, or click *Close* to return to the navigation pane without a rule selected.

## 5.6.11 Moving a Rule

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Select the rule you want to move.  
To select multiple rules in the same group, make sure the rules are displayed in the right pane of the navigation pane, then press the Ctrl key and select the required rules one at a time, or press the Shift key to select a consecutive list of rules.
- 4 Drag the selected rule to the location you want.

## 5.6.12 Copying a Rule

You can create a copy of an existing rule in your rule hierarchy, so you can use the same rule in more than one place in the hierarchy, or so you can create a new rule based on your existing rule.

---

**NOTE:** If you want to use the same rule in more than one place and you want any changes you make to the rule to be reflected in the other copy or copies, you should link the rule instead. See [“Linking a Rule” on page 90](#) for details.

---

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Select the rule you want to copy.  
To select multiple rules in the same group, make sure the rules are displayed in the right hand pane of the navigation pane, then press the Ctrl key and select the required rules one at a time, or press the Shift key to select a consecutive list of rules.
- 4 To create the copy, press the Ctrl key and drag the selected rule to the desired location
- 5 (Optional) Use the *Modify Rule* option to rename or modify the copy.
- 6 Move the rule to the correct position according to the order in which you want to process your rules. See [“Adding a Rule” on page 83](#) for details.

### 5.6.13 Linking a Rule

If you want a specific rule to be used in different places in your rules hierarchy, you can create a linked rule. Any changes you make to the linked rule are reflected in all the instances of the rule in your hierarchy. If you simply copy the rule, any changes made to the original rule or to one of its copies are not reflected in the other copies.

Changes to sub-rules of a linked rule are not linked. For example if you add or modify a rule under a linked rule, the change is not reflected in other instances of the linked rule.

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Select the rule you want to link.

To select multiple rules in the same group, make sure the rules are displayed in the right pane of the navigation pane, then press the Ctrl key and select the required rules one at a time, or press the Shift key to select a consecutive list of rules.

- 4 To create the links, press the Ctrl key and the Shift key at the same time, then drag the selected rule to the location you want.

A linked rule is displayed with an arrow .

### 5.6.14 Deleting a Rule

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Select the rule you want to delete.

To select multiple rules in the same group, make sure the rules are displayed in the right pane of the navigation pane, then press the Ctrl key and select the required rules one at a time, or press the Shift key to select a consecutive list of rules.

- 4 Click *Delete Rules* in the task pane.
- 5 Click *Finish* to delete the rule and all rule children.

### 5.6.15 Viewing Pseudocode

The pseudocode for a rule provides a simplified representation of the actual code that is processed when the rule is activated. For complex rules, this can assist you with understanding what happens in different situations.

To view the pseudocode for a rule:

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Rules* in the navigation pane.
- 3 Select the rule for which you want to view the pseudocode.

- 4 Click *Pseudocode* in the task pane.

You can copy the pseudocode by using Ctrl+A or Ctrl+C, then paste it into a document for printing.

- 5 Click *Close*.

## 5.7 Command Control Groups

Command Control has three types of groups:

**User Groups:** Contain users with similar responsibilities. This allows you to use the group as a condition for a rule, which either allows or denies the users the rights to run commands.

**Host Groups:** Contains hosts with similar content. This allows you to use the group as a condition for a rule that either allows or denies the rights to run the command on a host.

**Account Groups:** Combine host groups and user groups to be used together in setting rule conditions. Account groups can also contain other account groups. You can also use account groups as script entities.

For example, you could create a Web Account Group, and to this group you could add a user group that contains all the Web server managers and a host group that contains all the host that are Web servers. You could then use the Web Account Group as a condition when creating rules for Web server management.

The following sections explain how to manage these groups:

- ◆ [Section 5.7.1, “User Groups,” on page 91](#)
- ◆ [Section 5.7.2, “Host Groups,” on page 93](#)
- ◆ [Section 5.7.3, “Adding an Account Group,” on page 94](#)
- ◆ [Section 5.7.4, “Modifying an Account Group,” on page 95](#)
- ◆ [Section 5.7.5, “Deleting an Account Group,” on page 95](#)
- ◆ [Section 5.7.6, “Copying a Group,” on page 96](#)
- ◆ [Section 5.7.7, “Moving a Group,” on page 96](#)
- ◆ [Section 5.7.8, “Enhanced Access Control,” on page 96](#)

### 5.7.1 User Groups

User groups contain users who are allowed, or not allowed, to submit or run commands controlled by your rules. You can add user groups to your rule conditions to control whether the rule is processed, depending on the user who is submitting a command or the user who is specified to run a command. You can also use user groups as script entities.

Command Control has two default user groups. Do not modify these groups.

**Everyone:** Use this group to match against any user who has a local account on the hosts where Privileged User Manager is installed.

**Submit User:** Use this group to match against the user that submitted the privileged request. This is useful if you want to ensure that a rule only authorizes access to the account that submitted the request. For example when adding a crush login shell, you should add a clause to the rule that ensures that the run user is in the Submit User group. This ensures that a user cannot use the `-u` option in `usrun` to gain access to other accounts.

You can search for a specific user in a user group by using suitable regular expressions, strings, or wild cards in the command. For example, the wildcards that you can use in the command could be `vi *` or `/usr/bin/vi *`.

To add a regular expression term to the list, prefix the regular expression with `=~`. For example,

```
=~/^vi .*$/
```

```
=~/^user*/
```

The following sections explain how to manage user groups:

- ♦ [“Adding a User Group” on page 92](#)
- ♦ [“Modifying a User Group” on page 92](#)
- ♦ [“Deleting a User Group” on page 93](#)

## Adding a User Group

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Account Groups*, then expand the list.
- 3 Click *User Groups*.
- 4 To add a user group at the top level, click *Add User Group* in the task pane. To add a user group to a category, select the category and click *Add User Group* in the task pane.
- 5 Specify a name for the user group.
- 6 Click *Finish*.  
User groups are represented by the  icon.
- 7 To configure the user group, continue with [“Modifying a User Group” on page 92](#).

## Modifying a User Group

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Account Groups*, then click *User Groups* in the navigation pane.
- 3 Select the user group you want to modify.
- 4 Click *Modify User Group* in the task pane, then configure the following fields:

**Name:** Specify a name for the group.

**Disabled:** Select this check box to disable the group. A disabled user group is dimmed.

**Description:** Describe the purpose of this user group.

**Manager Name, Manager Tel., Manager Email:** Specify the name, telephone number, and e-mail address of the manager of this user group. The manager details can be used in the Compliance Auditor.

If these details have been entered in the manager’s Framework user account details (see [“Modify User: Account Details” on page 52](#)), they can be entered automatically by selecting the manager’s username from the drop-down list. This option is only available if you belong to a Framework user group with the read role defined for the auth module (see [Section 4.2.4, “Configuring Roles,” on page 58](#)).

**Users:** Add or change the users you want to include in this group. You can type the user names, one on each line, or paste them from elsewhere. You can use the *Sort* button to sort the list of users into alphabetical order.

**User Groups:** From the list of groups you have already defined, select the user groups you want to include as subgroups of this user group. You can also add subgroups to a user group by dragging the groups to the target user group in the navigation pane.

**5** Click *Finish*.

You can now use this user group in rule conditions or as a script entity.

### Deleting a User Group

**1** Click *Command Control* on the home page of the console.

**2** Click *Account Groups*, then click *User Groups* in the navigation pane.

**3** Select the required user group.

To select multiple user groups, press the Ctrl key and select the required user groups one at a time, or press the Shift key to select a consecutive list of user groups.

**4** Click *Delete User Group* in the task pane. The selected user groups are listed.

**5** Click *Finish*.

The user groups are deleted, and are also removed from any account group, rule conditions, and script entities where they have been defined.

## 5.7.2 Host Groups

Host groups contain hosts that are allowed, or not allowed, to submit or run commands controlled by your rules. You can add host groups to your rule conditions to control whether the rule is processed, depending on the host that is submitting a command or the host specified to run a command. You can also use host groups as script entities.

Command Control has two default host groups. Do not modify these groups.

**All Hosts:** Use this group to match against any host that have been registered with the Framework. Use the Hosts console to view the hosts that are included has matches for this group.

**Submit Host:** Use this group to match against the host from which the privileged request was made. This is useful if you want to ensure that a rule only authorizes access to the host from which the privileged request was made. This ensures that a user cannot use the `-h` option in `usrun` to gain access to other hosts.

You can search for a specific host in a host group by using suitable regular expressions, strings, or wild cards in the command. For example, the wildcards that you can use in the command could be `vi *` or `/usr/bin/vi *`.

To add a regular expression term to the list, prefix the regular expression with `=~`. For example,

```
=~/^vi .*$/
```

```
=~\w+\.novell\.com
```

The following sections explain how to manage host groups:

- ◆ [“Adding a Host Group” on page 94](#)
- ◆ [“Modifying a Host Group” on page 94](#)
- ◆ [“Deleting a Host Group” on page 94](#)

## Adding a Host Group

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Account Groups*, then click *Host Groups* in the navigation pane.
- 3 To add a host group at the top level, click *Add Host Group* in the task pane. To add a host group to a category, select the category and click *Add Host Group* in the task pane.
- 4 Specify a name for the host group.
- 5 Click *Finish*.  
Host groups are represented by the  icon.
- 6 To configure the host group, continues with [“Modifying a Host Group” on page 94](#).

## Modifying a Host Group

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Account Groups*, then click *Host Groups* in the navigation pane.
- 3 Select the host group you want to modify.
- 4 Click *Modify Host Group* in the task pane, then configure the following fields:
  - Name:** Specify a name for the group.
  - Disabled:** Select this check box to disable the group. A disabled host group is dimmed.
  - Description:** Describe the purpose of this host group.
  - Hosts:** Add or change the hosts you want to include in this group. You can type the host names, one on each line, or paste them from elsewhere. You can use the *Sort* button to sort the list of hosts into alphabetical order.
  - Host Groups:** From the list of groups you have already defined, select the host groups you want to include as subgroups of this host group. You can also add subgroups to a host group by dragging the groups to the host group in the navigation pane.
- 5 Click *Finish*. You can now use this host group in rule conditions or as a script entity.

## Deleting a Host Group

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Account Groups*, then click *Host Groups* in the navigation pane.
- 3 Select the host group you want to delete.  
To select multiple host groups, press the Ctrl key and select the required host groups one at a time, or press the Shift key to select a consecutive list of host groups.
- 4 Click *Delete Host Group* in the task pane. The selected host groups are listed.
- 5 Click *Finish*.  
The host groups are deleted, and are also removed from any account group, rule conditions, and script entities in which they have been defined.

## 5.7.3 Adding an Account Group

To add a new account group:

- 1 Click *Command Control* on the home page of the console.

- 2 Click *Account Groups* in the navigation pane.
- 3 To add an account group at the top level, click *Add Account Group* in the task pane. To add an account group to a category, select the category and click *Add Account Group* in the task pane. For information about categories, see [Section 5.5.6, “Adding a Category,” on page 81](#).
- 4 Specify a name for the account group.
- 5 Click *Finish*.  
Account groups are represented by the  icon.
- 6 To configure the group, continue with [“Modifying an Account Group” on page 95](#).

## 5.7.4 Modifying an Account Group

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Account Groups* in the navigation pane.
- 3 Select the account group you want to modify.
- 4 Click *Modify Account Group* in the task pane, then modify the following fields:
  - Name:** Change the name of the group.
  - Disabled:** To disable the account group, click *Disabled*. A disabled account group is dimmed.
  - Description:** Add or change the description.
  - Manager Name, Manager Tel., Manager Email:** Specify the name, phone number, and e-mail address of the manager of the users in this account group.  
If these details have been entered in the manager’s Framework user account details (see [“Modify User: Account Details” on page 52](#)), they can be entered automatically by selecting the manager’s username from the drop-down list. This option is only available if you belong to a Framework user group with the read role defined for the auth module (see [Section 4.2.4, “Configuring Roles,” on page 58](#)).The manager details can be used in the Compliance Auditor.
  - User Groups, Host Groups, Account Groups:** From the lists of groups you have already defined, select or remove the user groups, host groups, and account groups. You can also add groups to an account group by dragging the groups to the target account group in the navigation pane.
- 5 Click *Finish*. You can now use this account group in rule conditions or as a script entity.

## 5.7.5 Deleting an Account Group

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Account Groups* in the navigation pane.
- 3 Select the account group you want to delete.  
To select multiple account groups, display the groups in the right pane, press the Ctrl key and select the required account groups one at a time, or press the Shift key to select a consecutive list of account groups.
- 4 Click *Delete Account Group* in the task pane. The selected account groups are listed.

**5** Click *Finish*.

The account groups are deleted, and are also removed from any other account groups, rule conditions, and script entities where they have been defined.

## 5.7.6 Copying a Group

- 1 Click *Command Control* on the home page of the console.
- 2 Click the category of the group you are copying such as *Account Groups*, *Host Groups*, or *User Groups*.
- 3 Select the group you want to copy.

To select multiple groups in the same category or group, make sure the groups are displayed in the right pane of the navigation pane, then press the Ctrl key and select the required groups one at a time, or press the Shift key to select a consecutive list of groups.

- 4 To create the copy, press the Ctrl key and drag the selected group to the desired location
- 5 If necessary, use the appropriate *Modify Group* option to rename or modify the copy.

## 5.7.7 Moving a Group

- 1 Click *Command Control* on the home page of the console.
- 2 Click the category of the group you are copying such as *Account Groups*, *Host Groups*, or *User Groups*.
- 3 Select the group you want to move.

To select multiple groups in the same category or group, make sure the groups are displayed in the right pane of the navigation pane, then press the Ctrl key and select the required groups one at a time, or press the Shift key to select a consecutive list of groups.

- 4 Drag the selected group to the desired location.

You can also drag account groups, user groups, and host groups into an account group. This does not delete the groups from their original location.

## 5.7.8 Enhanced Access Control

Command Control policies give you additional options to control the execution of commands. For example, you can use a policy to restrict the rights and roles of a command so that the command works only for one particular directory, file, network address, or system call.

- ♦ [“Configuring a Command Control Policy” on page 96](#)
- ♦ [“Configuring a Path Policy” on page 97](#)

### Configuring a Command Control Policy

A command control policy is defined by using the policy script arguments. A policy script argument specifies the access rights of the applications based on the path, network, and capability.

- 1 Click *Command Control* on the home page of the console.
- 2 From the *Command Control Sample Scripts*, add the *Enhanced Access Control Policy* script.
- 3 Drag the *Enhanced Access Control Policy* script from *Scripts* to *Authorizing Rule*.

- 4 Click the *Authorizing Rule* and access the *Script Arguments*.
- 5 Create a script argument with a name *policy* and add that policy to the *Value* field.

### Configuring a Path Policy

A Path policy is a type of command control policy that restricts an application from accessing a specific directory based on the path.

The syntax of a Path policy is as follows:

```
path [owner] <path> <capability:capability:!capability>
```

*owner* specifies the file or directory ownership that should match with the current user ID.

*path* specifies a particular directory based on the path. Replace *path* with any of the following options:

**Table 5-5** Path Options

Option	Description
<i>/dir/file</i>	Specifies the file that the application can access in the <i>/dir/directory</i> .
<i>/dir/</i>	Specifies the directory that the application can access.
<i>/dir/f*</i>	Specifies a file that begins with <i>f</i> in the <i>/dir/directory</i> that the application can access.
<i>/dir/*</i>	Specifies that the application can access all the files in the <i>/dir/</i> directory.
<i>/dir/**</i>	Specifies that the application can access all the files and the subdirectories within the <i>/dir/directory</i> .
<i>/dir/**/</i>	Specifies that the application can access all subdirectories that are recursively searched for in the <i>/dir/directory</i> .
<i>/dir/**/*</i>	Specifies that the application can access all the files that are recursively searched for in any subdirectory within the <i>/dir/directory</i> .

*capability* specifies the rights of the application. You can use the *!* symbol in the syntax to denote a logical *not*. For example, *all:!write* grants all the rights except the *write* role.

Replace *capability* with any of the following options:

**Table 5-6** Capability Options

Option	Description
<i>privperms</i>	Enables the application with the <i>read</i> , <i>write</i> , and <i>ownership</i> permissions for the specified directory or file.
<i>perms</i>	Enables the application to assign the permissions of a specified directory or file.
<i>read</i>	Enables the application to assign the <i>read</i> permission for a specified directory or file.

Option	Description
write	Gives the application the <code>create</code> and <code>write</code> permissions for the specified directory or file.
unlink	Gives the application the <code>deletion</code> rights for the specified directory or file.
mknod	Enables the application to create system files in the specified directory.
exec	Enables the application to execute the shared files and files for which the application does not have <code>read</code> and <code>write</code> permission.
unsafe	Enables the application to execute any file that does not inherit the policy.
link	Enables the application to create a symbolic link or hard link to another file.
log[=<0-9>]	Enables the application to audit system calls, with an optional risk value of 0-9.
all	Enables the application to have all permissions.

You can use wildcards, regular expressions, and strings in the Path policy. For example, using the word `default` in the following example specifies the default policy.

```
path default all:log
path /opt/oracle/private/** !all:log=9
```

## 5.8 Commands

Command definitions contain the commands you want to control. A command definition can contain a single command, or several commands that you want to control in the same way. You can also specify a command that you want to run in place of a submitted command.

- ◆ [Section 5.8.1, “Adding a Command,” on page 98](#)
- ◆ [Section 5.8.2, “Modifying a Command,” on page 99](#)
- ◆ [Section 5.8.3, “Setting the Command Risk,” on page 101](#)
- ◆ [Section 5.8.4, “Removing a Command Risk,” on page 101](#)
- ◆ [Section 5.8.5, “Copying a Command,” on page 101](#)
- ◆ [Section 5.8.6, “Moving a Command,” on page 102](#)
- ◆ [Section 5.8.7, “Deleting a Command,” on page 102](#)
- ◆ [Section 5.8.8, “Importing Sample Commands,” on page 102](#)

### 5.8.1 Adding a Command

You can add command definitions to your rule conditions to control whether the rule is processed, depending on the command that is submitted by the user. You can also use commands as script entities.

To add a new command:

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Commands* in the navigation pane.

- 3 To add a command at the top level, click *Add Command* in the task pane. To add a command to a category, select the category and click *Add Command* in the task pane.
- 4 Specify a name for the command. This can be different from the name of the actual command you want to control.
- 5 Click *Finish*.
- 6 To configure the command, continue with [Section 5.8.2, “Modifying a Command,” on page 99](#).

## 5.8.2 Modifying a Command

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Commands* in the navigation pane.
- 3 Select the command you want to modify.
- 4 Click *Modify Command* in the task pane.
- 5 Configure the following fields:

**Name:** Specify a different name for the command.

**Disabled:** Select this check box to disable the command. A disabled command is dimmed.

**Description:** Describe the purpose of this command.

**Rewrite:** In the *Rewrite* field, define a command to be used in place of the commands listed in the *Command* field. You can also enter command arguments. Positional parameters can be used, as described in [“Using the Command Rewrite Functionality for Command Arguments” on page 99](#). To use the *Rewrite* field to enable auditing of the command, see [“Configuring Auditing with the Rewrite Functionality” on page 100](#)

**Commands:** Define one or more commands, one on each line. You can also enter command arguments. For example:

```
vi *
/usr/bin/vi *
```

To add a regular expression term to the list, prefix the regular expression with `=~`. For example,

```
=~/^vi .*$/
=#/usr/bin/vi .*#
```

You can copy and paste a list of commands from elsewhere. You can use the *Sort* button to sort the commands into alphabetical order.

**Sub Commands:** From the list of command definitions you have already created, select the subcommands you want to include in this command definition. You can also add subcommands to a command definition by dragging them to the command definition in the navigation pane.

- 6 Click *Finish*.

### Using the Command Rewrite Functionality for Command Arguments

The following table provides examples showing how the command rewrite functionality provided on the Modify Command page can be used with positional parameters to replace the submitted command and parameters. The examples use the `echo` command as the rewritten command to display the selected parameters on the screen.

**Table 5-7** *Command Rewrite Examples*

Function	Rewrite	Submitted Command	Executed Command
Insert all arguments (\$0 is not displayed)	echo \$*	ls passwd shadow fstab	echo passwd shadow fstab
Insert argument 'r;n'	echo \$3	ls passwd shadow fstab	echo fstab
Insert all but argument 'n' (\$0 is not displayed)	echo \${^2}	ls passwd shadow fstab	echo passwd fstab
Insert arguments from 'n' to end	echo \${2-}	ls passwd shadow fstab	echo shadow fstab
Insert arguments from 0 to 'n'	echo \${-2}	ls passwd shadow fstab	echo ls passwd shadow
Insert arguments from 'm' to 'n'	echo \${1-2}	ls passwd shadow fstab	echo passwd shadow
Insert the total number of arguments	echo \$#	ls passwd shadow fstab	echo 3
Insert contents of argument \$#	echo \${\$#}	ls passwd shadow fstab	echo fstab

### Rewrite Example Using ufsdump

In this example, the administrator usually does a backup of the system by using the following command:

```
ufsdump -0f /dev/rmt/0 /usr
```

Assume that new tape drive is installed on the host, and it must be used for the backup. In addition, the administrator must make sure that it is working correctly by using the `-v` flag to verify the tape.

You can ensure that the administrator doesn't need to remember the changes by using the *Rewrite* field to create a command definition for the original command:

```
$0 -v $1 /dev/rmt/1 ${$#}
```

When the administrator enters the original command, the following command runs instead:

```
ufsdump -v -0f /dev/rmt/1 /usr
```

### Configuring Auditing with the Rewrite Functionality

To enable auditing of the command, add the following to the *Rewrite* field:

```
-o audit <n>
```

Replace `<n>` with one of the following values:

- ◆ **0:** Disables auditing. It has the same effect as removing the audit setting from the *Rewrite* field.
- ◆ **1:** Enables auditing of all commands that are not built into the user's shell.

- ♦ **2:** Enables auditing of all commands, including commands that are built into the user's shell. This level of auditing can affect login times.

### 5.8.3 Setting the Command Risk

This option allows you to set a value representing the relative risk of a command when using the rush or crush clients with the session auditing option (see [Section 5.2, “Integrating Command Control into User Environments,” on page 66](#)). When you view a Command Control Keystroke Report, the commands with different risk values are represented in different colors.

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Commands* in the navigation pane.
- 3 Click *Command Risk* in the task pane.
- 4 Click *Add*.
- 5 Set a value for the command risk.
- 6 To use a regular expression instead of a simple pattern match for the command, select the *Regex* check box.
- 7 Specify the command you want to set a risk value for, or the regular expression. You can use wildcard symbols.
- 8 If you want to base the risk level on the directory in which the command is running, define a working directory.
- 9 If you want to base the risk level on who is running the command, define a user.
- 10 If you want to base the risk level on the host where the command is running, define a host.
- 11 If you want to change the order in which the commands are listed, use the arrow buttons.
- 12 Click *Finish*.

### 5.8.4 Removing a Command Risk

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Commands* in the navigation pane.
- 3 Click *Command Risk* in the task pane.
- 4 Select the entry, then click *Remove*.

### 5.8.5 Copying a Command

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Commands* in the navigation pane.
- 3 Select the command you want to copy.  
To select multiple commands in the same category, press the Ctrl key and select the required commands one at a time, or press the Shift key to select a consecutive list of commands.
- 4 To create the copy, press the Ctrl key and drag the selected command to the desired location
- 5 If necessary, use the *Modify Command* option to rename or modify the copy.

## 5.8.6 Moving a Command

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Commands* in the navigation pane.
- 3 Select the command you want to move.

To select multiple commands in the same category, press the Ctrl key and select the required commands one at a time, or press the Shift key to select a consecutive list of commands.

- 4 Drag the selected command to the desired location.

## 5.8.7 Deleting a Command

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Commands* in the navigation pane.
- 3 Select the command you want to delete.

To select multiple commands in the same category, press the Ctrl key and select the required commands one at a time, or press the Shift key to select a consecutive list of commands.

- 4 Click *Delete Command* in the task pane. The selected commands are listed.
- 5 Click *Finish*.

The commands are deleted, and are also removed from any rule conditions and script entities in which they have been defined.

## 5.8.8 Importing Sample Commands

Privileged User Manager ships with the following types of sample commands that you can import and use as is or import and modify to fit your needs:

- ♦ Shell commands (*ksh, sh, csh, bash*)
- ♦ vi commands
- ♦ System commands (*kill, mount, passwd, date, mkdir, useradd, chgrp, chown*)
- ♦ User commands (*env, ls, id, cat uname*)

To import these sample commands, click *Command Control > Import Samples > Sample commands*.

## 5.9 Scripts

You can use Perl scripts to provide additional, customized functionality to your rules. You can also use scripts in rule conditions. Privileged User Manager contains the embedded Perl interpreter version 5.8.9. You can use any of the core Perl modules for your script. It is not recommended that you install any CPAN Perl modules into the embedded Perl interpreter. If you create a script, be aware that any time consuming tasks within the script affect response times.

- ♦ [Section 5.9.1, “Adding a Script,” on page 103](#)
- ♦ [Section 5.9.2, “Modifying a Script,” on page 103](#)
- ♦ [Section 5.9.3, “Copying a Script,” on page 103](#)
- ♦ [Section 5.9.4, “Moving a Script,” on page 104](#)

- ♦ [Section 5.9.5, “Deleting a Script,” on page 104](#)
- ♦ [Section 5.9.6, “Sample Scripts,” on page 104](#)

## 5.9.1 Adding a Script

You can add your own custom attributes for account groups, user groups, host groups, commands, and access times to provide additional parameters for use in your scripts. See [“Defining Custom Attributes” on page 79](#) for details.

To add a new script:

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Scripts* in the navigation pane.
- 3 To add a script at the top level, click *Add Script* in the task pane. To add a script to a category, select the category and click *Add Script* in the task pane.
- 4 Specify a name for the script.
- 5 Click *Finish*.
- 6 To configure the script, continue with [Section 5.9.2, “Modifying a Script,” on page 103](#).

## 5.9.2 Modifying a Script

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Scripts* in the navigation pane.
- 3 Select the script you want to modify.
- 4 Click *Modify Script* in the task pane.
- 5 Configure the following fields:
  - Name:** Specify a different name for the script.
  - Conditional script:** Select the check box to set the script to be conditional. Scripts defined as conditional can be used in rule conditions. The return codes are limited to 1 for true and 0 for false.
  - Disabled:** Select the check box to disable the script. A disabled script is dimmed.
  - Description:** Describe the purpose of the script.
  - Script:** Specify the text of your script in the text box by typing it or by pasting it from elsewhere. The possible return codes you can use in your script for processing by the Command Control software are shown below this field.  
For some sample scripts, see [Section 5.9.6, “Sample Scripts,” on page 104](#).
- 6 Click *Finish*.

You can now assign your script to a rule, or you can specify it in rule conditions if you have set the script to be conditional.

## 5.9.3 Copying a Script

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Scripts* in the navigation pane.

- 3 Select the script you want to copy.

To select multiple scripts in the same category, press the Ctrl key and select the required scripts one at a time, or press the Shift key to select a consecutive list of scripts.

- 4 To create the copy, press the Ctrl key and drag the selected script to the desired location.

- 5 If necessary, use the *Modify Script* option to rename or modify the copy. For details, see [“Modifying a Script” on page 103](#).

## 5.9.4 Moving a Script

- 1 Click *Command Control* on the home page of the console.

- 2 Click *Scripts* in the navigation pane.

- 3 Select the script you want to move.

To select multiple scripts in the same category, press the Ctrl key and select the required scripts one at a time, or press the Shift key to select a consecutive list of scripts.

- 4 Drag the selected script to the desired location.

## 5.9.5 Deleting a Script

- 1 Click *Command Control* on the home page of the console.

- 2 Click *Scripts* in the navigation pane.

- 3 Select the script you want to delete.

To select multiple scripts in the same category, press the Ctrl key and select the required scripts one at a time, or press the Shift key to select a consecutive list of scripts.

- 4 Click *Delete Script* in the task pane. The selected scripts are listed.

- 5 Click *Finish*.

## 5.9.6 Sample Scripts

Privileged User Manager ships with the following sample scripts that you can import and use:

- ♦ Display message scripts
- ♦ Password validation scripts
- ♦ Alternate validation scripts
- ♦ Email scripts
- ♦ Modify environment script
- ♦ Emulate `su` script
- ♦ Secure `vi` script
- ♦ Change User Shell from `Crus` to `Csh`
- ♦ Request Remote Run Host  
novdocx (en) 16 April 2010  
Execute RunUsers Profile  
Truncate stdin/stdout/stderr

Log Originating IP Address in Auditing

Secure vi

Modify Environment

rush Illegal Commands

Enable Command Checksum

Enable rush Command Auditing

Set rush Pre-run Profile

Emulate su

Emulate su-

Enhanced Access Control Policy

- ◆ Populate LDAP Groups

Before creating your own Perl script, check out the sample scripts to see if one is available that meets your needs or one that can be modified to meet your needs. To understand what is available, see the sample scripts in the following sections.

- ◆ [“Modify Environment Script” on page 105](#)
- ◆ [“rush Illegal Commands Script” on page 107](#)

To import a sample script, click `Command Control > Import Samples > Sample Perl Script`.

## Modify Environment Script

This script is used to process environment variables. It has a number of script arguments that can add, delete, clear, and keep environment variables.

Argument	Description
<code>clearenv=1:</code>	Clears all environment variables (unless specifically kept using <code>keepenv</code> )
<code>keepenv=VAR:</code>	Specifically keeps environment variables. As soon as this is set, all other environment variables are deleted.
<code>setenv=VAR=val:</code>	Sets up a specific environment variable.
<code>unsetenv=VAR:</code>	Deletes a specific environment variable.
<code>defaultenv=#:</code>	Sets the default environment:  <b>0:</b> Sets up no default environment variables.  <b>1:</b> Sets up all default environment variables.  <b>2:</b> Sets up default environment variables that do not already exist in the environment.

## Sample Environment Script

```
my $e=$meta->child("Environment");
return(1) if(! $e);

my $n=$e->node_args();
my %env=();

while($n) {
    $env{$1}=$2 if($n->key() ne "items" && $n->value() =~ /^(.*)=(.*)$/);
    $n=$n->next();
}

my %keepenv=();
my $clearenv=0;

for(my $a=$args->node_args();$a;$a=$a->next()) {
    if($a->key() eq "clearenv" && $a->value() > 0) {
        $clearenv=1;
    } elsif($a->key() eq "keepenv" && $a->value() ne "") {
        $keepenv{$a->value()}=1;
    } elsif($a->key() eq "defaultenv" && $a->value >= 0) {
        $meta->child("Job")->arg_int("job_default_env",$a->value());
    }
}

if(scalar %keepenv || $clearenv) {
    while(my ($key,$val) = each %env) {
        delete $env{$key} if(! $keepenv{$key});
    }
}

for(my $a=$args->node_args();$a;$a=$a->next()) {
    if($a->key() eq "unsetenv" && $a->value() ne "") {
        delete $env{$a->value()};
    } elsif($a->key() eq "setenv" && $a->value() =~ /^(.*)\s*=\s*(.*)$/ ) {
        $env{$1}=$2;
    }
}

$meta->del($e);
$e=$meta->add_node("Environment");

my $items=0;

while(my ($key,$val) = each(%env)) {
    $e->arg("arg-$items","$key=$val");
    $items++;
}

$e->arg_int("items","$items");

return(1);
```

## rush Illegal Commands Script

When using the rush shell, Command Control has the ability to restrict the commands being run (even as root). This sample script is named `illegalcmd`, and it restricts the use of the `passwd` command.

This script does not restrict a user that initiates another shell from within a session. When a user does this, Command Control cannot continue a full audit or control the illegal commands, although the session is still captured

```
#to set script argument - name=illegalcmd value= kill *
my $t=$meta->get_params('Ticket');
if(! $t) {
$t=$meta->add_param('Ticket');
}

my $i=$t->get_params('IllegalCmds');
if(! $i) {
$i=$t->add_param('IllegalCmds');
}

my @illegal = $args->arg_values('illegalcmd');

#my @illegal=("echo","ls -l","passwd","/usr/bin/ls -l","ksh","echo date");
foreach my $b (@illegal) {
my $c=$i->add_param('Command');
$c->arg("cmd",$b);
}
return 1;
```

## 5.10 Access Times

You can restrict the times when a rule is valid by defining an access time and adding it to the rule conditions. You can also use access times as script entities.

- ◆ [Section 5.10.1, “Adding an Access Time,” on page 107](#)
- ◆ [Section 5.10.2, “Modifying an Access Time,” on page 108](#)
- ◆ [Section 5.10.3, “Copying an Access Time,” on page 108](#)
- ◆ [Section 5.10.4, “Moving an Access Time,” on page 108](#)
- ◆ [Section 5.10.5, “Deleting an Access Time,” on page 109](#)

### 5.10.1 Adding an Access Time

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Access Times* in the navigation pane.
- 3 To add an access time at the top level, click *Add Access Time* in the task pane. To add an access time to a category, select the category and click *Add Access Time* in the task pane.
- 4 Specify a name for the access time, for example, `Office hours`.
- 5 Click *Finish*.
- 6 To configure the access time, continue with [Section 5.10.2, “Modifying an Access Time,” on page 108](#).

## 5.10.2 Modifying an Access Time

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Access Times* in the navigation pane.
- 3 Select the access time you want to modify.
- 4 Click *Modify Access Time* in the task pane.
- 5 Modify the access time as desired:
  - ♦ Change the name of the access time.
  - ♦ Specify a description of the access time.
  - ♦ Click *Disabled* to disable the access time. A disabled access time is dimmed.
  - ♦ Set the access time as described in [Step 6](#).
- 6 Set the access time in multiples of half-hourly intervals. The default access time is set to *Deny Access* for the whole week, shown in the calendar as blue.
  - ♦ To allow access at specific times, click and drag across the days and times until the hours when you want to grant access are shown in green,
  - ♦ To allow access for the majority of times and deny access for specific times, click the *Grant Access* box below the table to grant access for the whole week, then click and drag across the days and times until the hours when you want to deny access are shown in blue.

For example, to allow access only during the hours from 9:00 to 18:00 from Monday to Friday:

  - 6a Ensure that the whole week is set to Deny Access (blue).
  - 6b Click in the calendar on 9 on Monday morning, then drag to 18 and down to Friday. This creates a green block representing the times when access is allowed.
- 7 Click *Finish*. You can now use this access time in rule conditions or as a script entity.

## 5.10.3 Copying an Access Time

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Access Times* in the navigation pane.
- 3 Select the access time you want to copy.

To select multiple access times in the same category, press the Ctrl key and select the required access times one at a time, or press the Shift key to select a consecutive list of access times.
- 4 To create the copy, press the Ctrl key and drag the selected access time to the desired location.
- 5 If necessary, rename or modify the copy by using the *Modify Access Time* option, as described in [“Modifying an Access Time” on page 108](#).

## 5.10.4 Moving an Access Time

- 1 Click *Command Control* in the navigation pane on the home page of the console.
- 2 Click *Access Times* in the navigation pane.
- 3 Select the access time you want to move.

To move multiple access times in the same category, press the Ctrl key and select the required access times one at a time, or press the Shift key to select a consecutive list of access times.

- 4 Drag the selected access time to the desired location.

### 5.10.5 Deleting an Access Time

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Access Times* in the navigation pane.
- 3 Select the access time you want to delete.

To select multiple access times in the same category, press the Ctrl key and select the required access times one at a time, or press the Shift key to select a consecutive list of access times.

- 4 Click *Delete Access Time* in the task pane.
- 5 Click *Finish*.

The access time is deleted, and is also removed from any rule conditions and script entities in which it has been defined.

## 5.11 Command Control Reports

You can configure customized reports of the contents of the Command Control configuration database, which are dynamically created and e-mailed to the specified person at defined intervals. You can use Perl template scripting to extract the required information and format it into an e-mail for the target person. An option is available for sending your reports to the Compliance Auditor for escalation management.

To use this feature, you must provide details of your e-mail server to the Messaging Component (msgagt) so that reports can be e-mailed. See [“Configuring SMTP Settings for the Messaging Component Package” on page 39](#) for details.

- ♦ [Section 5.11.1, “Adding a Command Control Report,” on page 109](#)
- ♦ [Section 5.11.2, “Modifying a Command Control Report,” on page 110](#)
- ♦ [Section 5.11.3, “Copying a Command Control Report,” on page 110](#)
- ♦ [Section 5.11.4, “Moving a Command Control Report,” on page 111](#)
- ♦ [Section 5.11.5, “Deleting a Command Control Report,” on page 111](#)

### 5.11.1 Adding a Command Control Report

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Reports* in the navigation pane.
- 3 To add a report at the top level, click *Add Report* in the task pane. To add a report to a category, select the category and click *Add Report* in the task pane.
- 4 Specify a name for the report.
- 5 Click *Finish*.
- 6 To configure the report, continue with [Section 5.11.2, “Modifying a Command Control Report,” on page 110](#).

## 5.11.2 Modifying a Command Control Report

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Reports* in the navigation pane.
- 3 Select the report you want to modify.
- 4 Click *Modify Report* in the task pane.
- 5 Modify the report as desired:
  - ♦ Change the name of the report.
  - ♦ Click *Disabled* to disable the report. A disabled report is dimmed.
  - ♦ Set the *Run Report* settings to determine the time of the first report and subsequent frequency of each report. You can set the initial date by using the calendar and type in the time, then set the frequency as required.
- 6 Select the e-mail options you want:
  - 6a In the *Email To* field, specify the e-mail address of the person you want to send the report to.
  - 6b In the *Email From* field, specify the e-mail address of the person you want to send the report from.
  - 6c In the *Email Subject* field, specify a subject for the e-mail.
  - 6d If you want the e-mail to be displayed in HTML, select the *HTML* check box.
  - 6e If you require a receipt, select the *Receipt* check box.
  - 6f Enter a Perl script in the *Report Template* field to control how the e-mail will be formatted and what it will contain.
- 7 If you want the report to be available to be audited by the Compliance Auditor, select the *Audit* check box.
- 8 If you want to send an e-mail while testing this report, select the *Send email* check box.
- 9 (Optional) Click *Test Report* to view the report that will be sent to the defined e-mail address. The report is not shown here in HTML format. If there are errors in the Report Template, these are shown.
- 10 Click *Back* to return to the report configuration page.
- 11 Click *Finish*.

## 5.11.3 Copying a Command Control Report

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Reports* in the navigation pane.
- 3 Select the report you want to copy.

To select multiple reports in the same category, press the Ctrl key and select the required reports one at a time, or press the Shift key to select a consecutive list of reports.
- 4 To create the copy, press the Ctrl key and drag the selected report to the desired location.
- 5 If necessary, use the *Modify Report* option to rename or modify the copy, as explained in [“Modifying a Command Control Report” on page 110](#).

## 5.11.4 Moving a Command Control Report

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Reports* in the navigation pane.
- 3 Select the report you want to move.

To select multiple reports in the same category, press the Ctrl key and select the required reports one at a time, or press the Shift key to select a consecutive list of reports.

- 4 Drag the selected report to the desired location.

## 5.11.5 Deleting a Command Control Report

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Reports* in the navigation pane.
- 3 Select the report you want to delete.

To select multiple reports in the same category, press the Ctrl key and select the required reports one at a time, or press the Shift key to select a consecutive list of reports.

- 4 Click *Delete Report* in the task pane. The selected reports are listed.
- 5 Click *Finish*. The reports are deleted.

## 5.12 Test Suites

Command control test suites allow you to test your rules by running specified commands, submit users and other input values through your rule configuration, and check to make sure the result is as expected. Each test suite can contain a number of test cases where you specify the expected outcome for one or more input values.

- ♦ [Section 5.12.1, “Adding a Test Suite,” on page 111](#)
- ♦ [Section 5.12.2, “Adding or Modifying a Test Case,” on page 112](#)
- ♦ [Section 5.12.3, “Running a Test Suite,” on page 113](#)
- ♦ [Section 5.12.4, “Viewing a Test Suite,” on page 114](#)
- ♦ [Section 5.12.5, “Modifying a Test Suite,” on page 114](#)
- ♦ [Section 5.12.6, “Deleting a Test Case,” on page 114](#)
- ♦ [Section 5.12.7, “Deleting a Test Suite,” on page 115](#)
- ♦ [Section 5.12.8, “Importing a Test Suite,” on page 115](#)
- ♦ [Section 5.12.9, “Exporting a Test Suite,” on page 115](#)

### 5.12.1 Adding a Test Suite

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Test Suites* in the task pane.
- 3 Click *Add Test Suite* in the task pane.
- 4 Specify a name for the test suite.
- 5 Specify a description for the test suite.

- 6 Click *Finish*.
- 7 Continue with [“Adding or Modifying a Test Case” on page 112](#) to add test cases to your test suite.

## 5.12.2 Adding or Modifying a Test Case

A test case allows you to emulate an end user running a command through the Command Control system.

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Test Suites* in the task pane.
- 3 Select the test suite for which you want to add a test case, or modify an existing test case.
- 4 Click *View Test Suite* in the task pane.
- 5 Do one of the following:
  - ♦ To add a new test case, click *Add Test Case* in the task pane.
  - ♦ To modify a test case, select the test case, then click *Modify Test Case*.
- 6 Specify the values and the expected results that you want to run through the rule configuration. (To review the rule configuration you want to test with this case, see [Section 5.6.2, “Modifying a Rule,” on page 83](#).)

Enter a single value in each field. The purpose of the test case is emulate the user performing a `usrun` command from the command line.

- ♦ To create a test case that can be used for general testing and could possible match multiple rules, supply only submit information for the test case.
- ♦ To create a test case that matches only one rule, use the expected fields to specify values that match a single rule.

**Command:** (Required) Specify the command the user would run.

For example, if the user would enter the following on the command line:

```
usrun passwd user1
```

Specify the following as the command:

```
passwd user1
```

**Submit User:** (Required) Specify the name of the user who is entering the privileged command.

**Submit Host:** (Required) Specify the name of the host that the submit user is logged in to.

**Run User:** (Optional) When the submit user is requesting to run the command as a specific user with the `usrun` command, specify the username that is being requested. For example, if the user would enter the following on the command line:

```
usrun -u root ksh
```

Specify the following as the run user:

```
root
```

**Run Host:** (Optional) When the submit user is requesting to run the command on a specific host, specify the hostname that is being requested. For example, if the user would enter the following on the command line:

```
usrun -h hosta ksh
```

Specify the following as the run host:

```
hosta
```

**User Input:** (Optional) Use this field to specify the information that a script, associated with the Command Control policy, expects the user to enter.

**Expected command:** (Optional) Use this field to confirm that the command being executed is the correct command. If the command specified in this field does not match the results, the test case fails.

**Expected authorized:** (Optional) Use this field to confirm that the request was authorised. If value in this field does not match the results, the test case fails.

**Expected capture:** (Optional) This field is compared with the result of the authorization request to confirm the capture mode is correct. If this field does not match the results, the test case fails.

**Expected run user:** (Optional) Use this field to confirm that the user context used to execute the command is correct. If this field does not match the results, the test case fails.

**Expected run host:** (Optional) Use this field to confirm that the host on which the command is being executed is correct. If this field does not match the results, the test case fails.

**Expected risk:** (Optional) This field is compared with the result of the authorization request in order to confirm the risk associated with the command being executed is correct. If this field does not match the results, the test case fails.

**Submit Time:** (Optional) Specify the time that the request should appear to be made. This is useful for testing access time restrictions in the policy.

**Custom Input:** (Optional) Use this field to add attributes within the request object. These XML definitions are inserted into the privileged request. For example, you could use this field to configure the group memberships for a user in order to test policies that perform tests on the user's group membership:

```
<Groups>
  <Group name='grpa' />
  <Group name='grpb' />
</Groups>
```

- 7 Click *Finish*. The input values are shown in the *Test Cases* table.
- 8 Repeat [Step 5](#) through [Step 7](#) for any additional test cases you want to include or modify in this test suite.

You can now run the test suite as explained in [“Running a Test Suite” on page 113](#).

### 5.12.3 Running a Test Suite

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Test Suites* in the task pane.
- 3 Select the test suite you want to run.

To select multiple test suites, press the Ctrl key and select the required test suites one at a time, or press the Shift key to select a consecutive list of test suites. Use Ctrl+A to select all test suites.

- 4 Click *Run Test Suites* in the task pane. The results are displayed for each test case as Success or as Failure, along with the reason for the failure.

- 5 Use the buttons on the left and right of the table to find previous successes and failures, and the next successes and failures.
- 6 To view further details on a specific entry, select the entry and click *Details*.  
The configuration for the test case is shown, and a list of rules that have been tested, with configuration settings for each rule. The *Matched* column shows true if the rule conditions were met, and false if the rule conditions were not met.
- 7 Click *Back* to return to the main Run Test Suite page.
- 8 Click *Cancel* to return to the list of test suites.

To use a command line option to run a test suite or to run a specific test case, see [Section 10.2.3, “Running Test Suites,”](#) on page 153.

### 5.12.4 Viewing a Test Suite

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Test Suites* in the task pane.
- 3 Select the test suite you want to view, then click *View Test Suite*.  
From here you can modify the test suite; add, modify and delete test cases; and run the test suite.

### 5.12.5 Modifying a Test Suite

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Test Suites* in the task pane.
- 3 Select the test suite you want to modify.
- 4 Click *View Test Suite* in the task pane.
- 5 Click *Modify Test Suite* in the task pane.
- 6 Modify the test suite as desired:
  - ♦ Change the name of the test suite.
  - ♦ Add or change the description.
  - ♦ Use the *Up* and *Down* buttons to change the order in which the test cases are run.
- 7 Click *Finish*.

### 5.12.6 Deleting a Test Case

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Test Suites* in the task pane.
- 3 Select the test suite from which you want to delete a test case.
- 4 Click *View Test Suite* in the task pane.
- 5 Select the test case you want to delete.
- 6 Click *Delete Test Case* in the task pane.
- 7 Click *Yes* to confirm the deletion. The test case is deleted.

## 5.12.7 Deleting a Test Suite

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Test Suites* in the task pane.
- 3 Select the test suite you want to delete.

To select multiple test suites, press the Ctrl key and select the required test suites one at a time, or press the Shift key to select a consecutive list of test suites.

- 4 Click *Delete Test Suite* in the task pane.
- 5 Click *Yes* to confirm the deletion. The test suite is deleted.

## 5.12.8 Importing a Test Suite

You use the *Import Test Suites* option to restore a previously backed-up test suite, or to test suites from another Framework. You then use the *Export Test Suites* option to obtain configuration details so you can then paste them into a text document for backup or for use on another Framework.

---

**NOTE:** When you import test suites, they are added to your existing configuration and do not overwrite your existing test suites. However, if you import a Command Control database by using the *Import Settings* option, your existing test suites are overwritten.

---

- 1 Access the test suite data you require and copy it.
- 2 Click *Command Control* on the home page of the console.
- 3 Click *Test Suites* in the task pane.
- 4 Click *Import Test Suites* in the task pane.
- 5 Click in the text area, then paste the copied settings by using Ctrl+V, or right-click in the text area and click *Paste*.
- 6 Click *Finish*.

## 5.12.9 Exporting a Test Suite

You can export your Command Control test suites to a text file for backup purposes, or for use in another Framework. You can then use the *Import Test Suites* option to restore the backed-up test suites, or to import the test suites into another Framework.

- 1 Click *Command Control* on the home page of the console.
- 2 Click *Test Suites* in the task pane.
- 3 Select the test suite you want to export.

To select multiple test suites, press the Ctrl key and select the required test suites one at a time, or press the Shift key to select a consecutive list of test suites. To select all test suites, use Ctrl+A.
- 4 Click *Export Test Suites* in the task pane.
- 5 Select the test suite data by using Ctrl+A, or right-click in the text window and click *Select All*.
- 6 Copy the test suite data by using Ctrl+C, or right-click in the text window and click *Copy*.

7 Paste the text into a text document.

8 Click *Finish*.

## 5.13 Deploying Command Control

To deploy Command Control, you first download several modules to your local Package Manager, then install them:

- ♦ [Section 5.13.1, “Command Control Modules,” on page 116](#)
- ♦ [Section 5.13.2, “Auditing Modules,” on page 116](#)
- ♦ [Section 5.13.3, “Compliance Auditor Modules,” on page 116](#)
- ♦ [Section 5.13.4, “Installing Command Control,” on page 116](#)

### 5.13.1 Command Control Modules

The Command Control feature is made up of the following packages:

- ♦ **Command Control Manager:** Holds the rule configuration and is responsible for validating user command requests.
- ♦ **Command Control Agents:** Installed on machines where user commands are to be controlled or audited.
- ♦ **Command Control Console:** Installed into the Framework Manager console. Required for configuring Command Control rules.

### 5.13.2 Auditing Modules

The auditing modules are made up of the following packages:

- ♦ **Audit Manager:** Acts as the repository for auditing information collected by the Framework.
- ♦ **Reporting Console:** Installed into the Framework Manager console. Required for viewing audit information.
- ♦ **Command Reporting Console:** Installed into the main Reporting Console. Required for viewing Command Control audit information.

### 5.13.3 Compliance Auditor Modules

The Compliance Auditor modules are made up of the following packages:

- ♦ **Compliance Auditor:** Holds the compliance auditor rules and audit information.
- ♦ **Compliance Auditor Console:** Installed into the Framework Manager console. Required for configuring compliance auditor rules and for viewing audit information.

### 5.13.4 Installing Command Control

To deploy Command Control:

- 1 Download the required packages to your local Package Manager. See [Section 5.13, “Deploying Command Control,” on page 116](#) for details.

- 2** Install the Command Control Manager package on the host you want to be the Command Control Manager. This can be on any operating system, including Windows.  
See [Section 3.4.6, “Installing Packages on a Host,”](#) on page 37 for details. Command Control Managers can be deployed to as many hosts as you need in order to build an environment with load balancing and failover.
- 3** Install the Command Control Agent package on all UNIX hosts on which you want to implement Command Control.
- 4** Install the Audit Manager package on the host you want to be the Audit Manager, then install the Compliance Auditor package on the same host.  
This can be on any operating system, including Windows, and can be a different host from your Command Control Manager. The auditing packages can be deployed to as many hosts as you need in order to build an environment with load balancing and failover.
- 5** Install the following consoles:
  - ◆ Command Control Console
  - ◆ Reporting Console
  - ◆ Command Reporting Console
  - ◆ Compliance Auditor ConsoleSee [Section 2.2.2, “Adding a Console to the Framework Manager Console,”](#) on page 21 for details.

Command control is now deployed and ready to use.



# Managing Audit Reports

# 6

Privileged User Manager enables auditing of events at several levels, such as keystroke logging, command authorization, and login success or failure. The Reporting console allows you to view these records and manage them.

- ◆ [Section 6.1, “Audit Settings,” on page 119](#)
- ◆ [Section 6.2, “Encryption Settings,” on page 120](#)
- ◆ [Section 6.3, “Syslog Settings,” on page 120](#)
- ◆ [Section 6.4, “Command Control Reports,” on page 121](#)

## 6.1 Audit Settings

Use this page to control the rollover of the audit database files. The default configuration does not encrypt or roll over the audit databases. If your security model requires you to keep audit records available for years, you need to configure the rollover options and move the rolled-over files to an archive location.

- 1 Click *Reporting* on the home page of the console.
- 2 Click *Audit Settings* in the task pane.
- 3 For each audit database file, set the rollover parameters. Rolled-over databases are kept as SQLite databases.

**Time (hours):** Specify the time interval for rolling over the audit file. If the time interval is always reached before the maximum size is reached, the time interval is used for rollover and the size restriction is ignored.

**Size (MB):** Specify the maximum size the file can reach before the audit file is rolled over. If the file always reaches the maximum size before the time interval is reached, the size restriction is used for rollover and the time interval is ignored.

**Protection:** Select *none* to allow the rollover file to be an unencrypted file or select *encrypted* to encrypt the audit database.

Encrypting the file can impact performance of your audit managers. Also, the encrypted file can be decrypted by the Framework Console, but it cannot be displayed on new systems that do not know the encryption keys.

To configure the encryption keys, click *Reporting > Encryption settings*.

- 4 If you want to zip the rollover files or move them to another location, use the *Rollover Script* option to specify a Perl script that can perform these tasks. The script is called whenever an audit database is rolled over.

For example, the following script uses *gzip* to compress the rolled-over file and enters an error message in the *unifid.log* file.

```
if ($DBGPR eq 'cmdctrl') {
system("gzip $AUDIT_FILE");
$ctx->log_error("Audit rollover $DBGPR $AUDIT_FILE");
}
```

- 5 Click *Finish*.

## 6.2 Encryption Settings

Use this page to configure when the randomly generated encryption key is changed.

- 1 Click *Reporting* on the home page of the console.
- 2 Click *Encryption Settings* in the task pane.
- 3 To specify how frequently the key is changed, specify a *Key Rollover* interval, then select the type of interval (years, months, weeks, or days).
- 4 (Optional) In the *Key* list, disable or enable keys.

Each time a new key is generated, it is added to the list.

If you disable a previous key, Privileged User Manager re-encrypts all database with the old key to the latest key. This can be very time-intensive and can affect performance until it is completed.

If you disable the null cipher key, Privileged User Manager encrypts all unencrypted databases with the latest key. This takes precedence over the encryption setting on the Audit Settings page. This can be very time-intensive and can affect performance until it is completed.

- 5 Click *Finish*.

## 6.3 Syslog Settings

Use this page to configure Privileged User Manager so that it can send syslog messages to a syslog server. This server can be a Sentinel server, a Sentinel Log Manager, or a syslog server that supports TCP with optional TLS or SSL support. Older syslog servers require UDP for the transport protocol.

To configure communication with a syslog server:

- 1 Click *Reporting* on the home page of the console.
- 2 Click *Syslog Settings* in the task pane.
- 3 Configure the following fields:

**Syslog host:** Specify the DNS name or IP address of the syslog server.

**Port:** Specify the port the syslog server is listening on for syslog events. The default port is 514. The default port for a Sentinel server or a Sentinel Log Manager is 1468.

**SSL:** Select the check box to enable SSL communication with a Sentinel server. For a syslog server, do not select this box.

- 4 In the Event table, select the events and the format. All possible events are select:

**Session Failure:** Sends an event when a Privileged User Manager session fails.

**Start Session:** Sends an event when a user starts a Privileged User Manager session on a host.

**Session Terminate:** Sends an event when a user logs out of the Privileged User Manager session.

**Command Audit:** If you have enabled auditing on the user's session or on commands, this option sends all audited events as syslog events.

- ♦ For information on configuring commands for auditing, see [“Configuring Auditing with the Rewrite Functionality” on page 100](#).
- ♦ For information on using a .profile file to enable session auditing, see [Section 5.2.4, “Using rush for Complete Session Control,” on page 71](#).

**Privilege Escalation:** Sends an event when a user starts a privileged session.

**4a** To delete an event, highlight it, then click *Remove*.

**4b** To configure the format, click the format text box and specify a format string.

The `{{}}` string logs the complete string of the audit record in JSON format. For a Sentinel server, format string must be set to `{{}}`.

If you are sending the events to a syslog server, you can specify strings from the Privileged User Manager templates. For example, the format of the Start Session event could use the following string:

```
User ${StartSession.user}$ initiated a Command Control session from  
${StartSession.host}$
```

This format string would produce output similar to the following:

```
Jan 1 01:20:45 localhost npum: User ctaylor initiated a Command  
Control session from citlaptop
```

**5** Click *Finish*.

## Sentinel Notes

For Privilege User Manager to communicate with a Sentinel server, you need to add a Syslog Connector to the Sentinel console. This connector must be configured to listen on port TCP 514 using SSL and the SSL type must be Open. Configure it to listen specifically for the host that has the Syslog Emitter installed. This is usually the Framework Manager console.

## 6.4 Command Control Reports

After you have installed the Framework Manager, all command control requests have records automatically created in the audit database. The default Sample Report displays all of the collected audit records and any associated keystroke captures. In the Command Control Reporting console, you add reports that can be customized by using the *Filters* tab to display records according to your preferences. You can also assign custom roles to the report, which allows you to restrict the read and write access your Framework Manager users have to these reports.

- ◆ [Section 6.4.1, “Adding a Report,” on page 121](#)
- ◆ [Section 6.4.2, “Viewing Report Data,” on page 122](#)
- ◆ [Section 6.4.3, “Filtering the Viewable Records,” on page 122](#)
- ◆ [Section 6.4.4, “Modifying General Report Information,” on page 124](#)
- ◆ [Section 6.4.5, “Selecting Log Files,” on page 124](#)
- ◆ [Section 6.4.6, “Replaying Keystrokes,” on page 124](#)
- ◆ [Section 6.4.7, “Removing a Report,” on page 125](#)
- ◆ [Section 6.4.8, “Generating an Activity Report,” on page 125](#)

### 6.4.1 Adding a Report

- 1** Click *Reporting* on the home page of the console.
- 2** Click *Command Control Reports* in the navigation pane.
- 3** Click *Add Report* in the task pane.

4 Configure the following fields:

**Name:** Specify a name for the report.

**Description:** (Optional) Describe the purpose of the report.

5 Click *Finish*.

6 Continue with one or more of the following:

- ♦ To modify the report so that it displays only a portion of the available records, continue with [Section 6.4.3, “Filtering the Viewable Records,” on page 122](#).
- ♦ To assign roles so that a specific group of Framework Manager users can view the reports, continue with [Section 6.4.4, “Modifying General Report Information,” on page 124](#).

## 6.4.2 Viewing Report Data

1 Click *Reporting* on the home page of the console.

2 Click *Command Control Reports* in the navigation pane.

3 Select the report in the navigation pane.

The navigation pane displays the following information about each instance of the report.

Column	Description
Start Time	Displays the date and time when the report started.
End Time	Displays the date and time when the report ended.
User	Displays the name of the Framework user who issued the command.
Host	Displays the name of the host from which the command was issued.
RunAs	Displays the name of the user who ran the command.
RunHost	Displays the name of the host that the command was run on.
Command	Displays the command that was executed.
Authorized	Displays whether the rule for this command authorized the command.
Capture	Displays whether the rule for this command captured the keystrokes. If a keystroke is present, the <i>Keystroke Replay</i> option is available in the task pane.
Audit Status	If the record has been referenced in the Compliance Auditor, displays the name of the compliance rule and the status.
Audit ID	Displays the unique ID of the audit record.

## 6.4.3 Filtering the Viewable Records

Use the *Filter* tab to build a list of matching conditions that allows you to customize the records that are displayed in the *Report Data* tab. This allows you to build reports that show only the information that your users require.

1 Click *Reporting* on the home page of the console.

2 Click *Command Control Reports* in the navigation pane.

- 3 Select the report in the navigation pane.
- 4 Click the *Filter* tab in the navigation pane.
- 5 Select from the following conditions. You can combine conditions with AND logic, which requires the report to match all conditions that have been joined with an AND. You can also combine conditions with OR logic, which requires the report to match either the conditions before the OR or the conditions after the OR.

**Authorized:** Select this option to use session authorization by the Command Control as a matching criteria. Use the Yes/No drop-down list to specify whether the session matches when the session was authorized or not.

**Session Capture:** Select this option to use session capture as a matching criteria. Use the Yes/No drop-down list to specify whether the report matches when the session capture was authorized or not.

**User:** Select whether you want to match on the submitting user or the run user. For the matching type, select one of the following:

- ♦ Select *Matches* or *Doesn't Match*, then specify an exact value or a value with an asterisk (\*) wildcard such as jo\*.
- ♦ Select *Regex* or *Doesn't Regex*, then specify a regular expression.

**Host:** Select whether you want to match on the submitted host or the run host. For the matching type, select one of the following:

- ♦ Select *Matches* or *Doesn't Match*, then specify an exact value or a value with an asterisk (\*) wildcard such as jo\*.
- ♦ Select *Regex* or *Doesn't Regex*, then specify a regular expression.

**Command:** Select whether you want to match on the submitted command or the audited command. An audited command is a command that has been audited within a session capture. Audited commands are collected when the session used the rush shell with the audit option. For the matching type, select one of the following:

- ♦ Select *Matches* or *Doesn't Match*, then specify an exact value or a value with an asterisk (\*) wildcard such as jo\*.
- ♦ Select *Regex* or *Doesn't Regex*, then specify a regular expression.

**Audit ID:** Select to match the session on the audit ID assigned to the session. For the matching type, select one of the following:

- ♦ Select *Matches* or *Doesn't Match*, then specify an exact value or a value with an asterisk (\*) wildcard such as 4bd\*.
- ♦ Select *Regex* or *Doesn't Regex*, then specify a regular expression.

**Time:** Select to match the session on when it started or when it ended. Select either *Session Start* or *Session End*, select *After* or *Before* for the matching operator, then use the calendar to specify a date and use the time fields to specify the hour and minute.

**():** Select to group conditions so that the record is displayed if it matches the conditions defined by one group in the filter.

- 6 Click *Apply*.
- 7 To view the results, click the *Report Data* tab.

## 6.4.4 Modifying General Report Information

Use the *General* tab to keep the report's name and description in sync with the configured filter and to restrict access to the report by assigning read and update roles.

- 1 Click *Reporting* on the home page of the console.
- 2 Click *Command Control Reports* in the navigation pane.
- 3 Select the report in the navigation pane.
- 4 Click the *General* tab in the navigation pane.
- 5 Modify the values of the following fields:

**Report name:** Specify a new name for the report.

**Description:** Describe the type of records that the report displays.

**Roles:** Specify values if you want to allow users read access to this report and the ability to update specific information such as its name, description, and filters.

- ♦ **Read:** To enable read access, specify a unique name for the read role for this report.
- ♦ **Update:** To enable update rights, specify a unique name for the update role for this report.

If you use the same name for a role on multiple reports, the role grants rights to multiple reports. If you use the same name for both the read role and the update role, the role grants both read and update rights.

To assign these roles to a group, see "[Audit Report Roles](#)" on page 59.

- 6 To save your changes, click *Apply*, or to discard your changes, click *Reset*.

## 6.4.5 Selecting Log Files

Any rolled-over audit database is indexed by the Audit Manager. You use the *Log Files* tab to select which of these rolled-over databases is used to display information in the *Report Data* tab. This allows you to review archived data or current activity.

Only the audit databases currently in the audit directory view are displayed. If an audit database has been taken offline (zipped or moved), it does not appear in the list.

- 1 Click *Reporting* on the home page of the console.
- 2 Click *Command Control Reports* in the navigation pane.
- 3 Select the report in the navigation pane.
- 4 Click the *Log Files* tab in the navigation pane.
- 5 Select the log files that are required for the report.

To include all available log files, select the *All log files* box.

- 6 Click *Apply*.

## 6.4.6 Replaying Keystrokes

Where a rule has been configured to capture session information, you can review the entire session in the report.

- 1 Click *Reporting* on the home page of the console.

- 2 Click *Command Control Reports* in the navigation pane.
- 3 Select the report in the navigation pane.
- 4 In the navigation pane, select the session that you want to review  
Commands for the session data that has been captured are indicated by a Yes in the *Capture* column.
- 5 Click *Keystroke Replay* in the task pane.
- 6 Edit the following fields:
  - Terminal Type:** Change the terminal type if it is set incorrectly.
  - Find:** To find a specify command or string in the report, specify the text in the text box, then click *Find*. If the report contains hundreds of lines, this allows you to find the command you are interested in.
  - Decryption key:** If an encryption password has been defined on the Command Control Audit Settings page to encrypt the sensitive password data in the reports (see “[Defining Audit Settings](#)” on page 77), specify this password in the text box, then click *Refresh* to display the passwords.
  - Show control characters:** Use this option to show or hide control characters on the screen.
  - Show audited commands:** Use this option to show or hide the full list of audited commands. If this option is enabled, the screen shows the actual commands that are being run when a user types a command. You can also view each input command individually by mousing over the command.
  - Show profile commands:** Use this option to show or hide the commands run in the user’s login profile when the user’s rush login shell has auditing configured to level 2.
- 7 From the list of input commands, select a command, then click *Output*.
- 8 Use the *Play*, *Rewind*, and *Pause* buttons to review the data.
- 9 Click *Cancel* to return to the list of reports.

## 6.4.7 Removing a Report

---

**IMPORTANT:** This action can not be undone.

---

- 1 Click *Reporting* on the home page of the console.
- 2 Click *Command Control Reports* in the navigation pane.
- 3 Select the report you want to delete.
- 4 Click *Delete Report* in the task pane.
- 5 Click *Finish*.

## 6.4.8 Generating an Activity Report

The Activity Report option allows you to generate a graphical snapshot of all the audit records currently being displayed in the report. The activity report can then be printed, providing a visual record for managers to see the number of commands each host is processing, the names of users requesting sessions, and the number of session accepted or rejected.

- 1 Click *Reporting* on the home page of the console.

- 2** Click *Command Control Reports* in the navigation pane.
- 3** Select the report you want generate an activity report for.
- 4** Click *Activity Report* in the task pane.  
The navigation pane displays the selected activity report.
- 5** To print the report, click *Print*.
- 6** To return to the list of reports, click *Cancel*.

# Compliance Auditor

# 7

The Compliance Auditor collects, filters, and generates reports of audit data for analysis and sign-off by authorized personnel. The Compliance Auditor can be used in conjunction with Command Control to enable auditors to view security transactions and play back recordings of user activity. Auditors can record notes against each record, creating permanent archives of activity.

Rules can be configured to pull any number of audit events matching a given filter into the Compliance Auditor at specific intervals. Examples of filters include username, host, and command for Command Control. Roles can be assigned to each rule to ensure that an auditor is able to view only extracted records with a matching role defined in his or her user account. In addition, Access Control Levels (ACLs) can be defined to restrict access to individual events, and to prevent users from auditing their own activity.

When an audit event is viewed, auditors can authorize the event, or mark it as unauthorized, escalate it, and assign it to someone else. Each change is recorded in an indelible audit trail within each record, along with any notes made by the auditor. Automatic reports can be generated and e-mailed to the appropriate personnel, and can be used, for example, for daily reporting to managers on audit activity awaiting sign-off, or hourly reporting triggered by an escalation value to notify senior management of activity.

To use the Compliance Auditor:

- ♦ Define roles in user groups to control user access to the Compliance Auditor. See [Section 7.1, “Controlling Access to the Compliance Auditor,”](#) on page 127.
- ♦ Create one or more rules to pull the required events into the Compliance Auditor. See [Section 7.2.1, “Adding or Modifying an Audit Rule,”](#) on page 128.
- ♦ Define ACLs for individual users. See [Section 7.5, “Access Control Levels,”](#) on page 140.
- ♦ View event records and authorize them, or mark them as unauthorized and define further action. See [Section 7.4, “Compliance Auditor Records,”](#) on page 135.
- ♦ Configure auditing reports to be automatically e-mailed to the appropriate personnel. See [Section 7.3.1, “Adding or Modifying an Audit Report,”](#) on page 130.
- ♦ Provide failover and load balancing by installing the Compliance Auditor on multiple hosts. See [Section 7.6, “Deploying the Compliance Auditor,”](#) on page 141.
- ♦ Export and import compliance auditing settings. See [Section 10.7.1, “Exporting and Importing Compliance Auditor Settings,”](#) on page 156.

## 7.1 Controlling Access to the Compliance Auditor

Roles can be used to restrict the Compliance Auditor options available to Framework users. For example, you might want users to be able to audit events, but not administer rules, ACLs, or reports.

To define roles for a user group to control use of the Compliance Auditor:

- 1 Click *Framework User Manager* on the home page of the console.

- 2 (Conditional) To add a new group, click *Groups > Add Group*, specify a name, then click *Finish*.
- 3 To modify an existing group or configure the group you just created, select the group, then click *Modify Group*.
- 4 Select the users you want to be members of this compliance auditing group.
- 5 In the *Roles* option, click *Add*, then add the following roles

Module	Role	Description
secaudit	console	View the Compliance Auditor console.
secaudit	audit	View and edit records.
secaudit	<audit role name>	(Optional) Allows the users to access records generated by the rules configured to use this Audit Role.  If you do not add the <audit role name> role, the users can only access records generated by rules with no Audit Role defined.
audit	read	View a keystroke replay.

Users belonging to this group can access the Compliance Auditor console, view and edit records, and review keystroke logs. If you do not add the <audit role name> role, the users can access all records. If you add the <audit role name> role, the users can access only the records generated by the rules configured to use this Audit Role.

With these roles, the users cannot manage rules, reports, or ACLs. For the roles required for these additional tasks, see [“Compliance Auditor Roles” on page 60](#).

- 6 Click *Finish*.
- 7 To continue setting up the Compliance Auditor, see [Section 7.2.1, “Adding or Modifying an Audit Rule,” on page 128](#).

## 7.2 Compliance Audit Rules

Audit rules specify the events to be pulled in to the Compliance Auditor for viewing and authorization. You can specify:

- ♦ The filters to display the type of event
- ♦ The number of events
- ♦ The time and frequency when the events are pulled in
- ♦ An audit role to restrict access to records of events pulled in by a specific rule

To add or modify a rule, see [Section 7.2.1, “Adding or Modifying an Audit Rule,” on page 128](#).

### 7.2.1 Adding or Modifying an Audit Rule

You can add, modify, and disable audit rules, but you cannot delete them.

- 1 Click *Compliance Auditor* on the home page of the console.

- 2 Click *Audit Rules* in the task pane.
- 3 Select one of the following:
  - ♦ To add a new rule, click *Add Rule* in the task pane
  - ♦ To modify an existing rule, select the rule, then click *Modify Rule*.
  - ♦ To copy an existing rule and modify it, select the rule, then click *Copy Rule*.

- 4 Configure the following fields:

**Rule Name:** Specify a name for your rule.

**Disable:** Select the check box to disable the rule.

By default, disabled rules are not shown in the rule list. You cannot delete a rule.

**Records and All Records:** To collect all records, enable the *All Records* check box, or deselect the *All Records* option and set the number of records to be collected on each audit run.

**Audit Role:** (Optional) Specify the audit role that has been assigned to a group. For configuration information, see [Step 5 in Section 7.1, “Controlling Access to the Compliance Auditor,” on page 127](#).

**Run Filter:** To determine the time and frequency of each audit run, use the calendar to set the initial date, then set the frequency as required.

**Audit Category:** Select the category of events to audit.

**Add Filter:** (Optional) Select one or more filters from the *Add Filter* drop-down list for the type of event you want this rule to pull in, and configure them as required

The filters and configuration options depend on the *Audit Category* selected. For example, you can choose to pull in only those Command Control events that have been submitted by a particular user and that include a session capture.

You can add more than one filter of the same type for filters such as the Command Control Submit User, then select the logic you require from AND or OR. You can also set these filters to be inclusive or exclusive using *matches* or *does not match*.

You can remove a filter by clicking the button to the left of the filter.

- 5 Click *Finish*.

## 7.3 Compliance Audit Reports

You can configure customized reports of events that require compliance auditing. The reports are dynamically created and e-mailed to selected users at defined intervals. You can use filtering and Perl template scripting to extract the appropriate event information and format it into an e-mail for each target user.

Audit reporting uses a tokens object that contains all the user information and other information. You can use keyword anchors in your report configuration, which are replaced by the appropriate values from the tokens object. It is also possible for the Perl code in the report template to set values in the tokens object. Sample report templates are supplied to assist you with creating your own.

- ♦ [Section 7.3.1, “Adding or Modifying an Audit Report,” on page 130](#)
- ♦ [Section 7.3.2, “Sample Command Control Report Template,” on page 131](#)
- ♦ [Section 7.3.3, “Deleting a Report,” on page 135](#)

## 7.3.1 Adding or Modifying an Audit Report

To use this feature, you must provide details of your e-mail server to the Messaging Component (msgagt) so that reports can be e-mailed. See [“Configuring SMTP Settings for the Messaging Component Package” on page 39](#) for details.

To add or modify an audit report:

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Click *Audit Reports* in the task pane.
- 3 Select one of the following:
  - ♦ To add a new report, click *Add Report* in the task pane.
  - ♦ To modify an existing report, select the required report, then click *Modify Report* in the task pane.
  - ♦ To copy an existing report, select the report, then click *Copy Report* in the task pane.
- 4 Configure the following fields:

**Report Name:** Specify a name for the report.

**Disable:** To disable the report, select the *Disable* check box.

By default, disabled reports are not shown on the report list.

**Run Report:** To determine the time and frequency of each audit report, use the calendar to set the initial date, then set the frequency as required

**Report Category:** To limit the report to one category, select the category, or to include all categories, select *All*.

**Report Target:** (Conditional) To send the report to a user or all users in a group, click *User Report* in the *Report Target* section, then select the user or group from the drop-down list.

Ensure that the users' e-mail addresses are defined in the *Account Details* section in the Framework User Account definitions. You must define a keyword anchor in the *Email To* field.

**Report Filter:** Set the *Report Filter* to include the required event records:

- ♦ Select one or more from *New*, *Pending*, *Authorized*, and *Unauthorized*.
- ♦ Select the age of events you want to include in the report. Events older than the number of days you specify are included.
- ♦ Select the escalation level of events you want to include in the report. Events at this escalation level and above are included.

**Email To:** Specify the e-mail address of the user who is to receive the report:

- ♦ If you want the report to be sent to a user who is not defined as a Framework user, specify the user's e-mail address in the *Email To* field.
- ♦ If you want the report to be sent to a user or group defined as the *Report Target* above, specify the following keyword anchor in the *Email To* field:

```
$User.ACT_EMAIL.value$
```

You can view the format in XML of the object tokens passed into the audit report by entering `$<>$` in the *Report Template* field, deselecting the HTML check box, then clicking *Test Report* (ensure that you have defined a *Report Target*). To view just the user subtree, use `$<User>$`.

The tokens that appear are dependent upon what has been configured for the users. If the `ACT_EMAIL.value` token is not present for the target, an email address has not been defined for the user. For user configuration information, see [Section 4.1.3, “Modifying a Framework User,” on page 49](#).

**Email From:** Specify the email address of the user sending the report.

You can also use a keyword anchor in the *Email From* field.

**Receipt:** Select if you want to enable notification when the receiver has read the message. The message is sent to the email address specified in the *Email From* field.

**Email Subject:** Specify a subject for the email message.

This can be a text string or you can use a keyword anchor in the *Email Subject* field. For example, if you wanted to display the target user’s name in the e-mail subject, you could enter the following in the *Email Subject* field.

```
Report for $User.ACT_FULL_NAME.value$
```

**Report Template:** Specify a Perl script in the *Report Template* field to control how the e-mail messages are formatted and what they contain. If you want the messages to be displayed in HTML, select the *HTML* check box.

For an example report template, see [“Sample Command Control Report Template” on page 131](#).

**5** Click *Test Report* to view the report that is sent to each e-mail target.

Use the arrow buttons with the mouse to page through the reports. In the test, the reports are not shown in HTML format. If there are errors in the *Report Template*, these are shown.

**6** Click *Back* to return to the report configuration screen.

**7** Click *Finish*.

## 7.3.2 Sample Command Control Report Template

If you are using this sample as a base for your own report templates, select *HTML* to correctly display the messages. The sample displays a message to the recipients of the e-mail messages, requesting them to log in to the Compliance Auditor and review activity. It extracts selected events and lists them in tables according to the age of the events, and provides information about the events.

As shown in the sample, you can use the user name keyword anchor `$User.ACT_FULL_NAME.value$` to display a user’s name in the e-mail, if you are using the Report Target option. You must ensure that a *Display name* is entered for the user in the *Account Details* section in the Framework User Account definitions.

```
<%!  
my @lv10;  
my @lv11;  
my @lv12;  
my @lv13;  
my @gt0;  
my @gt5;  
my @gt10;  
my @gt20;  
%>  
<%  
my @audit_records = @{$tokens->{'AuditRecords'}->{'AuditRecord'}} if
```

```

(defined($tokens->{'AuditRecords'}) && defined($tokens->{'AuditRecords'}-
>{'AuditRecord'}));
foreach my $ar (@audit_records) {
    my $age = $ar->{'age'};
    my $lvl = $ar->{'level'};

    if ($age > 5 && $age < 10) {
        push(@gt5,$ar);
    } elsif ($age >= 10 && $age < 20) {
        push(@gt10,$ar);
    } elsif ($age >= 20) {
        push(@gt20,$ar);
    } else {
        push(@gt0,$ar);
    }
    if ($lvl == 1) {
        push(@lvl1,$ar);
    } elsif ($lvl == 2) {
        push(@lvl2,$ar);
    } elsif ($lvl >= 3) {
        push(@lvl3,$ar);
    } else {
        push(@lvl0,$ar);
    }
}
%>
<%
my $total = @audit_records;
if ($total > 0) {
%>
<style type="text/css">
<!--
.style1 {
color: #000000;
font-family: Arial, Helvetica, sans-serif;
font-size: 12px;
}
.style2 {
color: #000000;
font-family: Arial, Helvetica, sans-serif;
font-size: 12px;
font-weight:bold;
}
.style4 {
color: #000000
}
-->
</style>
<p class="style1"> Hello $User.ACT_FULL_NAME.value$,<br/>
<br/>
    This is an automated event notification email from the Compliance Auditor.
<br/>
<br/>

    It is the responsibility of management to log into the Compliance Auditor
each day and review their team's keystroke logs. <br/> <br/>

    Please log on to the Compliance Auditor at your earliest convenience using
this link: <a href="https://admin.company.com">https://admin.company.com</a></

```

```

p>

<%
my $gt0 = @gt0;
%>
<span class="style2">Events &lt; 5 days old (<%= "$gt0" %>)</span>
<table border="1">
  <tr class="style1">
    <td>Time</td>
    <td>User</td>
    <td>Run As</td>
    <td>Host</td>
    <td>Command</td>
  </tr>
  <%
foreach my $ar (@gt0) {
  my $cmd = $ar->{'cmdctrl'}->{'cmd'};
  my $usr = $ar->{'cmdctrl'}->{'user'};
  my $ras = $ar->{'cmdctrl'}->{'runAs'};
  my $hst = $ar->{'cmdctrl'}->{'host'};
  my $tme = $ar->{'cmdctrl'}->{'time'};
  $tme = localtime($tme);
%>
  <tr class="style1">
    <td><%= "$tme" %></td>
    <td><%= "$usr" %></td>
    <td><%= "$ras" %></td>
    <td><%= "$hst" %></td>
    <td><%= "$cmd" %></td>
  </tr>
  <%
}
%>
</table>
<br/>

<%
my $gt5 = @gt5;
%>
<span class="style2">Events &gt; 5 days old (<%= "$gt5" %>)</span>
<table border="1">
  <tr class="style1">
    <td>Time</td>
    <td>User</td>
    <td>Run As</td>
    <td>Host</td>
    <td>Command</td>
  </tr>
  <%
foreach my $ar (@gt5) {
  my $cmd = $ar->{'cmdctrl'}->{'cmd'};
  my $usr = $ar->{'cmdctrl'}->{'user'};
  my $ras = $ar->{'cmdctrl'}->{'runAs'};
  my $hst = $ar->{'cmdctrl'}->{'host'};
  my $tme = $ar->{'cmdctrl'}->{'time'};
  $tme = localtime($tme);
%>
  <tr class="style1">
    <td><%= "$tme" %></td>

```

```

        <td><%= "$usr" %></td>
        <td><%= "$ras" %></td>
        <td><%= "$hst" %></td>
        <td><%= "$cmd" %></td>
    </tr>
<%
}
%>
</table>
<br/>

<%
my $gt10 = @gt10;
%>
<span class="style2">Events &gt; 10 days old (<%= "$gt10" %>)</span>
<table border="1">
    <tr class="style1">
        <td>Time</td>
        <td>User</td>
        <td>Run As</td>
        <td>Host</td>
        <td>Command</td>
    </tr>
<%
foreach my $ar (@gt10) {
    my $cmd = $ar->{'cmdctrl'}->{'cmd'};
    my $usr = $ar->{'cmdctrl'}->{'user'};
    my $ras = $ar->{'cmdctrl'}->{'runAs'};
    my $hst = $ar->{'cmdctrl'}->{'host'};
    my $tme = $ar->{'cmdctrl'}->{'time'};
    $tme = localtime($tme);
%>
    <tr class="style1">
        <td><%= "$tme" %></td>
        <td><%= "$usr" %></td>
        <td><%= "$ras" %></td>
        <td><%= "$hst" %></td>
        <td><%= "$cmd" %></td>
    </tr>
<%
}
%>
</table>
<br/>

<%
my $gt20 = @gt20;
%>
<span class="style2">Events &gt; 20 days old (<%= "$gt20" %>)</span>
<table border="1">
    <tr class="style1">
        <td>Time</td>
        <td>User</td>
        <td>Run As</td>
        <td>Host</td>
        <td>Command</td>
    </tr>
<%
foreach my $ar (@gt20) {

```

```

my $cmd = $ar->{'cmdctrl'}->{'cmd'};
my $usr = $ar->{'cmdctrl'}->{'user'};
my $ras = $ar->{'cmdctrl'}->{'runAs'};
my $hst = $ar->{'cmdctrl'}->{'host'};
my $tme = $ar->{'cmdctrl'}->{'time'};
$tme = localtime($tme);
%>
<tr class="style1">
  <td><%= "$tme" %></td>
  <td><%= "$usr" %></td>
  <td><%= "$ras" %></td>
  <td><%= "$hst" %></td>
  <td><%= "$cmd" %></td>
</tr>
<%
}
%>
</table>
<br/>

<p class="style2">Total Events = <%= $total %></p>

<%
}
%>

```

### 7.3.3 Deleting a Report

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Click *Audit Reports* in the task pane.
- 3 Select the report you want to delete.
- 4 Click *Delete Report* in the task pane.
- 5 Click *Finish* to confirm the deletion.

## 7.4 Compliance Auditor Records

The Compliance Auditor main page lists the records (events) collected according to defined audit rules.

By default, all new and pending events are displayed, as indicated in the *Status* column. To view authorized and unauthorized events, select the appropriate check boxes and click *Refresh*. Pending events are events that have been viewed and their records edited, but they have not been classified as authorized or unauthorized. You can click any of the column headings to sort by that column.

To view events for a specific time period, select the *From* and *To* check boxes, select the required dates, specify the required times, and click *Refresh*.

The table displays the following information about each event:

Column	Description
First	The color-coded indicators for Command Control command risk level and rule risk level, ranging from green (low) to red (high). For more information, see <a href="#">“Setting the Command Risk” on page 101</a> .
Level	The escalation level set by the auditor editing the event record.
Status	The status of the event, indicating whether an auditor has classified the event as authorized or unauthorized. New events have not been viewed. Pending events have been viewed and edited, but have not been marked as authorized or unauthorized.
Time	The date and time the event occurred.
Event	A description of what the record contains.
Note	Any notes made by the auditor when editing the event record.
Assigned	The user the event has been assigned to by the auditor of the event record.
Rule	The audit rule that pulled in the event.
Type	The type of event.
Size	The size of the keystroke capture with the total time of the session displayed between parentheses.
Event ID	The unique event ID.

From this page, you can perform the following tasks:

- ◆ [Section 7.4.1, “Viewing a Compliance Audit Record,” on page 136](#)
- ◆ [Section 7.4.2, “Viewing and Editing a Command Control Keystroke Report,” on page 137](#)
- ◆ [Section 7.4.3, “Viewing a Change Management Audit Record,” on page 138](#)
- ◆ [Section 7.4.4, “Viewing a Report Audit Record,” on page 138](#)
- ◆ [Section 7.4.5, “Editing an Audit Record,” on page 138](#)
- ◆ [Section 7.4.6, “Archiving Records,” on page 139](#)
- ◆ [Section 7.4.7, “Managing Archived Records,” on page 139](#)

## 7.4.1 Viewing a Compliance Audit Record

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Select the record you want to view.
- 3 Click *View Record* in the task pane.

Record data for this event is shown, including the submit user and host, the run user and host, the command, whether it was authorized by Command Control, and whether the session was captured.

From here you can view a Command Control keystroke report, if it exists, or edit the record. If a keystroke report exists, you must review it before you can edit the record. See [Section 7.4.2, “Viewing and Editing a Command Control Keystroke Report,” on page 137](#) for more information.

## 7.4.2 Viewing and Editing a Command Control Keystroke Report

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Select the record for which you want to view a keystroke report.
- 3 Click *View Record* in the task pane.
- 4 Click *View Keystroke Report* in the task pane, or click the *Keystroke* button.

The text that the user entered during the session is shown on the Input page. The first column displays color-coded indicators for command risk level and rule risk level, ranging from green (low) to red (high). For more information, see [“Setting the Command Risk” on page 101](#) and [“Modifying a Rule” on page 83](#).

- 5 On the Command Control Keystroke Report page, edit the following fields:

**Terminal Type:** Change the terminal type if it is set incorrectly.

**Find:** To find a specify command or string in the report, specify the text in the text box, then click *Find*.

**Decryption key:** If an encryption password has been defined on the Command Control Audit Settings page to encrypt the sensitive password data in the reports (see [“Defining Audit Settings” on page 77](#)), specify this password in the text box, then click *Refresh* to display the passwords.

**Show control characters:** Use the *Show control characters* check box to show or hide control characters on the screen.

**Show audited commands:** Use the check box to show or hide the full list of audited commands. If this option is enabled, the screen shows the actual commands that are being run when a user types a command. You can also view each input command individually by mousing over the command.

**Show profile commands:** Use the check box to show or hide the commands run in the user’s login profile when the user’s rush login shell has auditing configured to level 2.

- 6 (Optional) To see the keystroke text being played back with the screen output, click *Output*.  
You can start the playback from a specific line in the input by selecting that line before clicking *Output*.
  - ♦ Click *Play* to play the keystroke entries and view the output.
  - ♦ Click *Rewind* to go back to the beginning.
  - ♦ Click *Pause* to pause the playback.
  - ♦ Click *Forward* to skip any pauses in the playback where the user might have taken a break from typing.
  - ♦ Set the *Playback Speed* to *Real Time*, *Double Speed*, or *Full Speed*.
  - ♦ Set the *Scrollbar* field to the amount of text you want to be able to scroll back through, in kilobytes.
  - ♦ Change the *Terminal Type* to the one you want.
- 7 Click *Cancel* to return to the record list.

### 7.4.3 Viewing a Change Management Audit Record

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Select the Command Control Change Management record you want to view.  
The record type is shown in the *Type* column. You might need to scroll to the right to see this column.
- 3 Click *View Record* in the task pane.  
Information about the Change Management action is displayed, including the name of the user who made changes to the database, and any entries the user made when committing the Command Control transaction.
- 4 To edit the record, see [Section 7.4.5, “Editing an Audit Record,” on page 138](#).

### 7.4.4 Viewing a Report Audit Record

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Click the record you want to view.  
The record type is shown in the *Type* column. You might need to scroll to the right to see this column.
- 3 Click *View Record* in the task pane.  
Record data for this report is shown, including the contents of the report sent.
- 4 To edit the record, see [Section 7.4.5, “Editing an Audit Record,” on page 138](#).

### 7.4.5 Editing an Audit Record

For each event listed in the Compliance Auditor, you can edit the audit record to authorize the event, or mark it as unauthorized, escalate it, and assign it to another user. You can also add notes for display in the event record, and comments that are permanently recorded in the event history.

---

**NOTE:** For Command Control events for which a keystroke report exists, you must view the keystroke report before editing the audit record. See [Section 7.4.2, “Viewing and Editing a Command Control Keystroke Report,” on page 137](#) for more information.

---

To edit an audit record:

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Select the record you want to edit.
- 3 Click *View Record* in the task pane.
- 4 Click *Edit Record*.
- 5 (Optional) Authorize the event:
  - 5a Select the *Authorized* check box.
  - 5b In the *Note* field, specify a note to be displayed on the event list and event record.
  - 5c In the *Comment* field, specify a comment to be permanently displayed in the *History* on the View Record page.

- 6 (Optional) Mark the event as unauthorized:
  - 6a Select the *Unauthorized* check box.
  - 6b If necessary, set an *Escalation Level* to be displayed on the event list.

This can be used as a report filter when setting up reports. See [Section 7.3.1, “Adding or Modifying an Audit Report,”](#) on page 130.
  - 6c If necessary, use the *Assigned to* field to assign the record to a different user.
  - 6d Specify a *Note* or a *Comment* to explain why the event is unauthorized.
- 7 Click *Finish*.

## 7.4.6 Archiving Records

Audit records can be archived from the console or from the command line. For information about the command line options, see [Section 10.7.2, “Managing Compliance Auditor Records,”](#) on page 156.

To archive records from the console:

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Select the records you want to archive.

To select multiple records, press the Ctrl key and select the records one at a time, or press the Shift key to select a consecutive list of records.
- 3 Click *Archive Records* in the task pane.

A list of the selected records is displayed.
- 4 Configure the following fields:

**Comment:** (Required) Specify the reason for the archive.

**Keep Online:** (Optional) Select if you want the archived records to continue to be displayed in the list of records.
- 5 Configure the types of records to archive.

By default, authorized and unauthorized records are selected. New and pending records are not displayed. If you want to archive these records, select the *New* and *Pending* options.

---

**IMPORTANT:** After a record is archived, it cannot be modified. If you archive new or pending records, their status can never change.

---
- 6 Click *Finish*.

## 7.4.7 Managing Archived Records

From the Framework Manager console, you can restore an archive and move archives from an online state (viewable in the console) and to an offline state (not viewable in the console) and from an offline state to an online state. You must use the command line options to purge an archive. See [Section 10.7.2, “Managing Compliance Auditor Records,”](#) on page 156.

To manage archived records from the console:

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Click *Manage Archives* in the task pane.

- 3 To restore an archive to an online status, select the archive, then click *Restore*.
- 4 To move an archive from an online status to an offline status, select the archive, then click *Remove*.
- 5 Click *Close*.

## 7.5 Access Control Levels

You can define an Access Control Level (ACL) for your auditors that specifies which events they are allowed to view and restricts auditors from authorizing their own activity.

- ♦ [Section 7.5.1, “Adding or Modifying a User ACL,” on page 140](#)
- ♦ [Section 7.5.2, “Deleting a User ACL,” on page 140](#)

### 7.5.1 Adding or Modifying a User ACL

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Click *Access Control* in the task pane.
- 3 To add a new ACL, click *Add User ACL* in the task pane. To modify an existing ACL, select the required *User* and click *Modify ACL* in the task pane.  
When creating a new user ACL, select the user from the *Username* drop-down list.
- 4 Click *Add*.
- 5 At the bottom of the table, select the attribute from the drop-down list that describes the entity to which you want to control access for the selected user.  
For example, if you do not want this user to be able to audit Command Control events involving a particular command, click *Command*.
- 6 In the *Matches* field, specify the value of the attribute you want to control access to.  
For example, if you do not want this user to be able to audit any Command Control events that involve the `cat /etc/passwd` command, specify this command in this field. You can use wildcard characters in this field.
- 7 Set the *Action* to allow or deny.
- 8 (Optional) Use the arrow buttons to move entries up and down the list.  
You might want to do this if, for example, you are allowing the user to access a restricted list of commands, and using the wildcard `*` to deny access to all other commands. The `allowed commands` entries must be above the `deny all` entry. By default, all commands are allowed.
- 9 (Optional) Remove an attribute by selecting it and then clicking the *Remove* button.
- 10 (Optional) Modify an entry by selecting it, then specifying the changes.
- 11 Click *Finish*.

### 7.5.2 Deleting a User ACL

- 1 Click *Compliance Auditor* on the home page of the console.
- 2 Click *Access Control* in the task pane.
- 3 Select the user for whom you want to delete an ACL.

- 4 Click *Delete User ACL* in the task pane.
- 5 Click *Finish* to delete the ACL for the user.

## 7.6 Deploying the Compliance Auditor

You can provide failover and load balancing by installing the Compliance Auditor on multiple hosts. The Compliance Auditor consists of the following packages:

- ♦ **Compliance Auditor (secaudit):** Holds the compliance auditor rules and audit information.
- ♦ **Compliance Auditor Console:** Installed into the Framework Manager console. Required for configuring Compliance Auditor rules and for viewing audit information.

The Compliance Auditor has the following dependencies:

- ♦ The Compliance Auditor package is shown as an available package only on hosts that have the Audit Manager (audit) deployed.
- ♦ If you want to use the Compliance Auditor reporting facilities, you need to install the Access Manager (auth) on the host with the Compliance Auditor.

To deploy the Compliance Auditor:

- 1 Download the required packages to your local Package Manager. See [Chapter 2, “Managing Package Distribution,” on page 19](#) for details.
- 2 Install the Audit Manager package on the host you want to be the Audit Manager, then install the Compliance Auditor package on the same host.  
This can be on any operating system, including Windows. See [Section 3.4.6, “Installing Packages on a Host,” on page 37](#) for details. The auditing packages can be deployed to as many hosts as you need in order to build an environment with load balancing and failover.
- 3 If you need reporting facilities, install the Access Manager package on the same host as the Compliance Auditor package.
- 4 Install the Compliance Auditor console. See [Section 2.2.2, “Adding a Console to the Framework Manager Console,” on page 21](#) for details.

The Compliance Auditor is now deployed and ready to use.



# Load Balancing and Failover

# 8

The load balancing and failover features work by using a hierarchical view of the hosts associated with the Framework.

The hierarchy of hosts is created by using the Hosts console to group hosts into domains and subdomains, which are representative of your enterprise network structure. This effectively gives them a chain of command, where they always address requests to managers in their immediate subdomain before moving along a branch to another subdomain or parent domain.

To achieve an effective load balancing and failover environment, multiple Framework Manager packages must be deployed across the Framework. The licensing model is not based on how many managers or agents are deployed, but how many hosts the Framework is deployed on. This means that there are no restrictions on how many Framework Manager packages you can deploy.

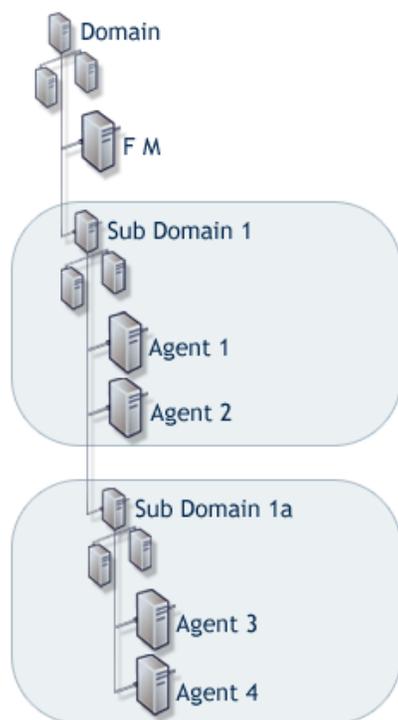
The Registry Manager controls a database that records the location and status of each package deployed on each of the hosts within the Framework. A copy of this information is held at each host by the Registry Agent package that is included as part of the agent installation. The distributed information is used to calculate the route to the appropriate manager for requests from any agent registered on the Framework. The structure of the registry data enables each host to determine which Framework Manager on the Framework should be the target of requests, and which Framework Manager to use if there is a failure or withdrawal of the initially selected Framework Manager.

- ♦ [Section 8.1, “Failover,” on page 143](#)
- ♦ [Section 8.2, “Load Balancing,” on page 144](#)

## 8.1 Failover

The failover feature automatically and transparently redirects requests from a failed or withdrawn Framework Manager to the next available manager of the same type. The agent automatically connects to a manager that is next in line in accordance with your defined hierarchy.

**Table 8-1** *Creating a Failover Environment*



This diagram shows an example of a typical way to create an effective failover environment.

The Framework Manager (FM) is a Windows host. All agents are UNIX hosts.

**Deployment:** Deploy the Command Control Manager package on the Framework Manager, Agent 1, and Agent 3 hosts.

**Who authenticates to whom:** By default, each agent contacts the following host for Command Control authentication:

Agent 1 and 2 contact Agent 1.

Agent 3 and 4 contact Agent 3.

**IMPORTANT:** Windows supports only the Command Control Manager package.

Examples:

1. Agent 3 is downed for maintenance. Agent 4 seeks authentication from Agent 1.
2. Agent 1 is downed because of a broken network card. Agents 2, 3, and 4 seek authentication from the Framework Manager.
3. The Command Control Manager package is removed from the Framework Manager and the Agent is still broken. Agents 2, 3, and 4 seek authentication from Agent 3.

**IMPORTANT:** If an additional subdomain is added, agents under Subdomain 1 and 1a then seek authentication from the new Subdomain if no other Command Control Manager is available.

## 8.2 Load Balancing

Load balancing means the ability to evenly distribute processing and communications activity across the Framework so that no single Framework Manager is overwhelmed by agent requests.

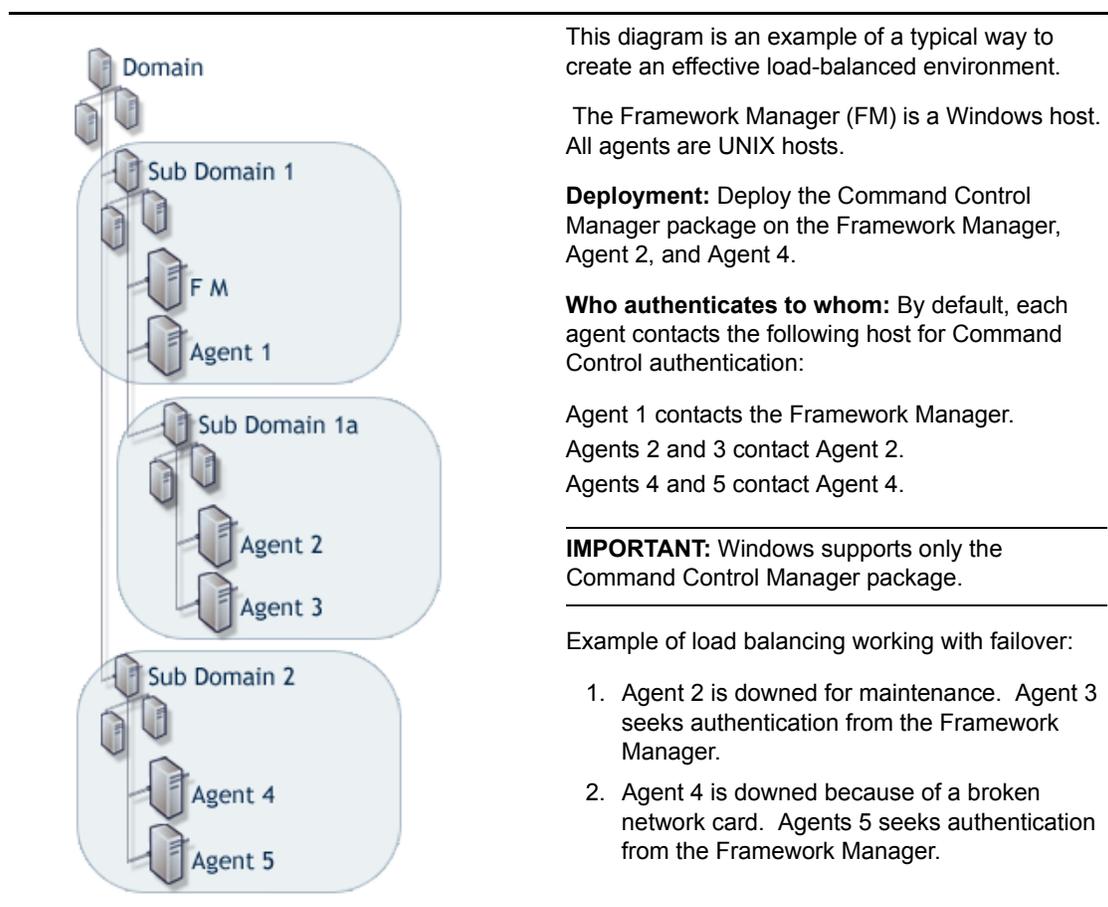
Load balancing is particularly important in situations where it is difficult to predict the number of requests that are directed to a specific category of manager.

The Framework automatically replicates data from the defined primary manager to each additional manager that is deployed in the Framework. Replication takes place automatically when the manager is initially deployed and then again at any stage when the data on the primary manager is modified.

The following packages can be load balanced:

- ♦ **Registry Manager:** Maintains a database of all hosts and modules and provides certificate-based registration features for the hosts.
- ♦ **Package Manager:** Manages a repository for packages.
- ♦ **Administration Agent:** Provides the functionality for the Web-based user interface. Consoles can be installed on the Administration Agent and used to control product features.
- ♦ **Access Manager:** Maintains a list of Framework user accounts and provides authentication services for the Framework. This package must be installed with a local Registry Manager in order to create a secure user authentication token.
- ♦ **Command Control Manager:** Maintains a database of all defined command control rules, commands, and scripts.

**Table 8-2** *Creating a Load Balancing Environment*





# Command Control Components

# 9

The components of the Command Control module interact with other Privileged User Manager modules as indicated in the following table.

**Table 9-1** *Command Control Components*

Component	Description
Command Control Console	<p>Provides the user interface with the Command Control database to allow creation and management of Command Control rules.</p> <p>Available on all supported platforms.</p> <p>Interacts with:</p> <p><b>Framework Manager Console:</b> Provides the main user interface.</p> <p><b>Command Control Manager:</b> Contains the Command Control information.</p>
Command Control Manager	<p>Contains the database of Command Control rules.</p> <p>Available on all supported platforms.</p> <p>Interacts with:</p> <p><b>Command Control Console:</b> Provides a user interface to the Command Control module.</p> <p><b>Command Control Agents:</b> Provides client functionality for the Command Control module.</p>
Command Control Agent	<p>Provides the client for Command Control, including shells and remote execution binaries.</p> <p>Available on all supported UNIX and Linux platforms.</p> <p>Interacts with:</p> <p><b>Command Control Manager:</b> Obtains command authentication information.</p> <p><b>System Information Agent:</b> Obtains remote host information from alternative hosts.</p> <p><b>Audit Manager:</b> Sends collected audit information to the Audit Manager.</p>
System Information Agent	<p>A background agent that gathers host and user information on the agents.</p> <p>Available on all supported UNIX and Linux platforms.</p> <p>Interacts with:</p> <p><b>Command Control Agents:</b> Supply requested information regarding hosts and users.</p>

Component	Description
Audit Manager	<p>Contains the database of all Command Control activity and user-logged sessions.</p> <p>Available on all supported platforms.</p> <p>Interacts with:</p> <p><b>Command Control Reports:</b> Provide a user interface for Command Control audited events.</p> <p><b>Command Control Agent:</b> Receives the auditing information.</p>
Reporting	<p>Provides the primary user interface for auditing. You add plug-ins to view appropriate package information.</p> <p>Available on all supported platforms.</p> <p>Interacts with:</p> <p><b>Framework Manager Console:</b> Provides the main user interface.</p> <p><b>Command Control Reports:</b> Provide the Command Control auditing interface.</p>
Command Control Reports Console	<p>Provides the Command Control auditing interface that reports on all successful/unsuccessful Command Control events, as well as captured sessions and replay.</p> <p>Available on all supported platforms.</p> <p>Interacts with:</p> <p><b>Audit Manager:</b> Contains the audited information.</p> <p><b>Reporting:</b> Parent console of the Command Control Reports console.</p>

# Command Line Options

# 10

The command line options can be run from either a Linux/UNIX Framework Manager or a Windows Framework Manager. Some can be run from an agent machine.

To use the command line options, change to the `unifi` directory:

**Linux/UNIX:** `/opt/novell/npum/sbin/unifi`

**Windows:** `<Drive>\Program Files\Novell\npum\sbin\unifi`

The following sections assume that you are running the commands from the `unifi` directory. If you are using the new Windows power shell, replace the `./` of the syntax with `.\`. If you are not using this new Windows shell, remove the `./` from the command. For example:

**Linux/UNIX:** `./unifi -v`

**Windows Power Shell:** `.\unifi -v`

**Windows:** `unifi -v`

This command displays version information for the Command Control module.

Privileged User Manager supports the following command line options:

- ♦ [Section 10.1, “The unifi Options,” on page 149](#)
- ♦ [Section 10.2, “Command Control Options,” on page 150](#)
- ♦ [Section 10.3, “Package Distribution Options,” on page 153](#)
- ♦ [Section 10.4, “Package Manager Options,” on page 154](#)
- ♦ [Section 10.5, “Registry Agent Options,” on page 154](#)
- ♦ [Section 10.6, “Registry Manager Options,” on page 155](#)
- ♦ [Section 10.7, “Compliance Auditor Options,” on page 156](#)

## 10.1 The unifi Options

The `unifi` binary is located in the `/opt/novell/npum/sbin/unifi` directory for Linux and Unix platforms, in the `./` and in the `\Program Files\Novell\npum\sbin\unifi` directory for the Windows platform. The command has the following syntax:

**Syntax:** `./unifi [options]`

Replace `[options]` with one or more of the following:

Option	Description
<code>-v</code>	Displays the version of the Framework patch.
<code>-s</code>	Displays the service name if you have changed it by modifying the <code>unifi.xml</code> file.

Option	Description
-u <username>	Specifies the username of the user requesting the command. This is used to verify that the user has sufficient rights to execute the command. Most, but not all commands, require authorization.
-p <pwd>	Specifies the password of the user.
-n	Indicates that the user's native account can be used for credentials. This option replaces the -u <username> -p <pwd> options. For information on how to set up native maps, see <a href="#">“Modify User: Native Maps” on page 54</a> .  Most of the module commands require authentication credentials to verify the user's rights to issue the command.

For example:

```
/opt/novell/npum/sbin/unifi -v
```

This command displays the version of the Framework patch.

## 10.2 Command Control Options

The command line options for the Command Control module allow you to perform the following tasks:

- [Section 10.2.1, “Importing and Exporting Command Control Settings,” on page 150](#)
- [Section 10.2.2, “Backing Up and Restoring a Command Control Configuration,” on page 151](#)
- [Section 10.2.3, “Running Test Suites,” on page 153](#)

### 10.2.1 Importing and Exporting Command Control Settings

The following commands allow you to export a current command control configuration and import one. Importing a configuration overwrites any existing rule set; therefore before importing a configuration, you should back up the current configuration (see [Section 10.2.2, “Backing Up and Restoring a Command Control Configuration,” on page 151](#)).

The export command has the following syntax:

**Syntax:** `./unifi -n cmdctrl export [options]`

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the -n option with -u <username> -p <password> options and specify the name and password of a Framework Manager user who has the rights to perform this task.

Replace [options] with one or more of the following:

Options	Description
-f <arg>	Specifies where to export the configuration to. Replace <arg> with a filename or a path and filename.
-c	Specifies that the configuration should be exported in clear text. This option cannot be used with the -p option.

Options	Description
-p <pwd>	Specifies an encryption password for the file. If a password is specified, the password must be entered when importing the file. This option cannot be used with the -c option.

The import command has the following syntax:

**Linux Syntax:** `./unifi -n cmdctrl import [options]`

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the -n option with -u <username> -p <password> options and specify the name and password of a Framework Manager user who has the rights to perform this task.

Replace [options] with one or more of the following:

Options	Description
-f <arg>	Specifies the file to import. Replace <arg> with a filename or a path and filename.
-p <pwd>	Specifies the password that was used to encrypt the configuration when it was exported.

## 10.2.2 Backing Up and Restoring a Command Control Configuration

The following commands can be executed on the primary console or on backup hosts. When they are executed on a backup host, the commands actually execute on the primary console.

**Syntax:** `./unifi -n cmdctrl [option]`

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the -n option with -u <username> -p <password> options and specify the name and password of a Framework Manager user who has the rights to perform this task.

Replace [option] with one of the following:

Option	Description
backup -t <"reason">	Backs up the current command control database. The -t <"reason"> parameter allows you to supply a reason for the backup, and is optional but recommended. Enclose the reason text in double quotes.

Option	Description
<code>listcfg &lt;format&gt;</code>	<p>Lists the backups that are available for restoration. To specify a format, use one of the following:</p> <p><b>-X:</b> For XML output. For example: <code>&lt;a.Item I.version="0" who="admin" reason="Backup 1" I.timestamp="1247146780" I.id="1"/&gt;</code></p> <p><b>-D &lt;date&gt;:</b> For modifying the date format. For example, if you replace <code>&lt;date&gt;</code> with <code>%D</code> for the format, the time stamp is displayed as <code>07/14/09</code> rather than <code>2009-07-14_11-52-56</code>. For possible options, see <code>strftime(3C)</code>.</p> <p><b>-F &lt;fmt&gt;:</b> For specifying what template information is displayed. By default, the following information is displayed.</p> <ul style="list-style-type: none"> <li>◆ <b>id:</b> The unique ID of the backup.</li> <li>◆ <b>who:</b> The ID of the user who created the backup.</li> <li>◆ <b>reason:</b> The reason for the backup, if provided by the user.</li> <li>◆ <b>timestamp:</b> The date and time when the backup occurred.</li> </ul> <p>Replace <code>&lt;fmt&gt;</code> with one or more of these options. Individual options are enclosed with <code>{ }</code> and separated from other options with a comma. The entire string is enclosed in single quotes. For example:</p> <pre>-F '{id}\$,{reason}\$'</pre> <p>This string would print out the following:</p> <pre>1,Basic test rules for session closure</pre>
<code>restore -n &lt;id&gt;</code>	<p>Restores the command control database to the select version. Replace <code>&lt;id&gt;</code> with the version number you want to restore. The current configuration is overwritten.</p> <p>You cannot restore when transactions are enabled (see <a href="#">Section 5.4, "Command Control Transactions," on page 75</a>).</p>
<code>delcfg -n &lt;id&gt;</code>	<p>Deletes the selected backup from the list. Replace <code>&lt;id&gt;</code> with the version number you want to delete.</p> <p>Deleting a backup is permanent and cannot be undone.</p>
<code>backup --?</code>	Displays the usage help for the <code>backup</code> command.
<code>listcfg --?</code>	Displays the usage help for the <code>list</code> command.
<code>restore --?</code>	Displays the usage help for the <code>restore</code> command.
<code>delcfg --?</code>	Displays the usage help for the <code>delete</code> command.

## Sample Commands

To back up the database:

```
./unifi -n cmdctrl backup -t "Added the ls command."
```

To restore the second backup in the list:

```
./unifi -n cmdctrl restore -n 2
```

## 10.2.3 Running Test Suites

The test suite options allow you to run part or all of the Command Control test suites.

**Syntax:** `./unifi -n cmdctrl runTest [option]`

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the `-n` option with `-u <username> -p <password>` options and specify the name and password of a Framework Manager user who has the rights to perform this task.

Replace `[option]` with one or more of the following:

Option	Description
<code>-t &lt;'arg'&gt;</code>	Specifies a specific test suite to run. Replace <code>&lt;arg&gt;</code> with the name of the test suite. This option cannot be used with the <code>-A</code> option
<code>-A</code>	Runs all the test suites. This option cannot be used with the <code>-t</code> option.
<code>-v</code>	Outputs the full debug information to the screen.
<code>-V &lt;file&gt;</code>	Outputs the full debug information to the specified file or the specified path and file.
<code>-o &lt;file&gt;</code>	Outputs the test results to the specified file or the specified path and file.

For example:

```
./unifi -u admin cmdctrl runTest -A -o /tmp/test.log
```

This command writes the results of the test suites to the `test.log` file in the `/tmp` directory.

## 10.3 Package Distribution Options

The following command allows you to import packages into the Package Manager

**Syntax:** `./unifi -n distrib publish [option]`

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the `-n` option with `-u <username> -p <password>` options and specify the name and password of a Framework Manager user who has the rights to perform this task.

Replace `[option]` with one of the following:

Option	Description
<code>-d &lt;directory&gt;</code>	Imports the packages from the specified directory, for example: <code>-d /tmp/framework</code>
<code>-f &lt;package&gt;</code>	Imports the specified package, for example: <code>-f /tmp/framework/xxx.pak</code>

## 10.4 Package Manager Options

The following commands allow you to install and uninstall the packages on the agents.

**Syntax:** `./unifi -n pkgman install <agent> <package>`

**Syntax:** `./unifi -n pkgman uninstall <agent> <package>`

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the `-n` option with `-u <username> -p <password>` options and specify the name and password of a Framework Manager user who has the rights to perform this task.

Replace `<agent>` with the agent name for the host. To view a list of these names, click *Hosts* on the home page of the Framework Console.

Replace `<package>` with the name of the package to install or uninstall. To view a list of package names in the Framework Console, click *Hosts*, select a host, select to display the packages. The name field contains the package name that is used in this command.

---

**NOTE:** You cannot use this command to install or uninstall consoles. It can only be used to install and uninstall modules.

---

## 10.5 Registry Agent Options

The Registry Agent module allows you to perform the following tasks from the command line:

- ♦ [Section 10.5.1, “Registering an Agent,” on page 154](#)
- ♦ [Section 10.5.2, “Finding a Primary Manager Package,” on page 155](#)
- ♦ [Section 10.5.3, “Agent Status,” on page 155](#)

### 10.5.1 Registering an Agent

The following command registers an agent with the Framework Manager. It must be run from the agent machine.

To run the command and be prompted to supply information:

```
./unifi regclnt register
```

You are prompted to supply the IP address of the Framework Manager, the registration port (the default port is 29120), the DNS name or IP address of the agent machine, the agent name, then a Framework Manager username and password.

To run the command with all the parameters on the command line:

```
./unifi regclnt register <manager> 29120 <hostname> <agent name> <admin>  
<password>
```

Replace `<manager>` with the IP address of the Framework Manager, `<hostname>` with the DNS name or IP address of the agent machine, `<agent name>` with the name of the agent, `<admin>` with a Framework Manager username, and `<password>` with the user’s password. For example:

```
./unifi regclnt register manager1 29120 agent1.domain.com agent1 admin novell
```

## 10.5.2 Finding a Primary Manager Package

The following command displays details about primary manager packages. It can be run from any host machine, and displays the primary manager information contained in the local machine's databases.

**Syntax:** `./unifi -n regclnt getManager <package>`

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the `-n` option with `-u <username> -p <password>` options and specify the name and password of a Framework Manager user who has the rights to perform this task.

Replace `<package>` with one of the following: `admin`, `audit`, `auth`, `cmdctrl`, `msgagnt`, `pkgman`, `registry`, `secaudit`, `syslogemit`.

For example, the following command returns details about the primary Command Control Manager:

```
./unifi regclnt getManager cmdctrl
```

## 10.5.3 Agent Status

The following command displays the status of agents within the framework. It can be run from any host machine.

**Syntax:** `./unifi -n regclnt status [option]`

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the `-n` option with `-u <username> -p <password>` options and specify the name and password of a Framework Manager user who has the rights to perform this task.

Replace `[option]` with one or more of the following:

Option	Description
<code>-s &lt;server&gt;</code>	Displays status of the specified host. This option can be repeated on the command line for more than one host.
<code>-o &lt;domain&gt;</code>	Allows you to request status for all agents in a domain. This option can be repeated on the command line for more than one domain.
<code>-S &lt;server&gt;</code>	Confirms whether the host can communicate with the specified agent.
<code>-M &lt;module&gt;</code>	Confirms whether the agent can communicate with the specified module.
<code>-a</code>	Displays the status for all defined hosts.
<code>-c</code>	Provides output in CSV format.
<code>-h</code>	Prevents the display of the CVS header.
<code>-?</code>	Displays the usage message.

## 10.6 Registry Manager Options

The registry command allows you to promote the registry on a backup host to primary status from the command line. After you have promoted the registry to primary, you can log in to the backup console and promote the remaining packages to primary.

**Syntax:** `./unifi -n registry promote`

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the `-n` option with `-u <username> -p <password>` options and specify the name and password of a Framework Manager user who has the rights to perform this task.

## 10.7 Compliance Auditor Options

The command line options for the Compliance Auditor allow you to perform the following tasks:

- ♦ [Section 10.7.1, “Exporting and Importing Compliance Auditor Settings,” on page 156](#)
- ♦ [Section 10.7.2, “Managing Compliance Auditor Records,” on page 156](#)

### 10.7.1 Exporting and Importing Compliance Auditor Settings

The Compliance Auditor now supports the ability to export and import its settings from the command line. You can export and import the following settings:

- ♦ Audit Rules
- ♦ Audit Reports
- ♦ Access Control Levels

**Exporting:** The export command exports only configuration settings; the audit records are not exported. The export includes all rules and reports, even those that have been disabled. The Compliance Auditor does not allow rules or reports to be deleted, because they might be associated with audit records. The exported file is in XML format.

**Importing:** You should import the settings only on a system that hasn’t been configured or on a system where the current configuration is not needed. Every rule and report contains a unique ID, but if that ID already exists on the current system, the rule or report is overwritten by the imported configuration.

**Commands:** The commands use the following syntax:

```
./unifi -n secaudit import -f <file>  
./unifi -n secaudit export -f <file>
```

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the `-n` option with `-u <username> -p <password>` options and specify the name and password of a Framework Manager user who has the rights to perform this task.

Replace `<file>` with the name of the file to import or to create for the export.

### 10.7.2 Managing Compliance Auditor Records

The compliance auditor now supports the ability to archive, restore, and purge the audit records from the command line. These commands can be performed on the Framemaker Manager console machine or from a backup host. When executed from a backup host, a command is actually execute on the primary host.

If a backup host is promoted to be a primary host, the archived database can be placed on the promoted manager and restored.

The `secaudit` command has the following syntax:

```
./unifi -n secaudit [list] [listarchive] [archive] [restore] [purge]
```

If you have not mapped your local account to a Framework Manager user (see [“Modify User: Native Maps” on page 54](#)), replace the `-n` option with `-u <username> -p <password>` options and specify the name and password of a Framework Manager user who has the rights to perform this task.

The `secaudit` command supports the following options:

Option	Description
<code>list &lt;format&gt;</code>	<p>Displays all of the audit records currently stored, including any records already archived, unless archived records have been purged. To view a format other than the default, specify one of the following:</p> <ul style="list-style-type: none"><li><b>-x:</b> For XML output.</li><li><b>-D &lt;date&gt;:</b> For modifying the date format. For example, if you replace <code>&lt;date&gt;</code> with <code>%D</code> for the format, the time stamp is displayed as <code>07/14/09</code> rather than <code>2009-07-14_11-52-56</code>. For possible options, see <code>strftime(3C)</code>.</li><li><b>-F &lt;fmt&gt;:</b> For specifying what template information is displayed. By default, the following information is displayed.<ul style="list-style-type: none"><li>♦ <b>id:</b> The unique ID of the archive.</li><li>♦ <b>who:</b> The ID of the user who created the archive.</li><li>♦ <b>reason:</b> The reason for the archive, if provided by the user.</li><li>♦ <b>timestmp:</b> The date and time when the archive occurred.</li></ul></li></ul> <p>Replace <code>&lt;fmt&gt;</code> with one or more of these options. Individual options are enclosed with <code>{ }</code> and separated from other options with a comma. The entire string is enclosed in single quotes. For example:</p> <pre>-F '\${id}\$,\${reason}\$'</pre>
<code>archive -n &lt;from:to&gt; -p &lt;pwd&gt; -r "&lt;reason&gt;"</code>	<p>Creates a database in the <code>/opt/novell/npum/service/local/secaudit</code> directory with the following format:</p> <pre>sa-2009-06-05_11-38-43.db</pre> <p>Each archived database can then be taken offline (moved to another storage area) and put back in place at any point.</p> <p>Specify values for the following parameters:</p> <ul style="list-style-type: none"><li><b>-n &lt;from:to&gt;:</b> Specifies the records to archive. To archive one record, specify its ID. To archive a range of records, replace <code>&lt;from:to&gt;</code> with the range. For example to archive records 20 to 40, specify <code>20:40</code>. Use the <code>list</code> option to view the IDs of the records.</li><li><b>p &lt;pwd&gt;:</b> (Optional) Specifies a password. If a password is specified for an archive, the same password must be used to restore the archive.</li><li><b>-r "&lt;reason&gt;":</b> (Optional) Specifies a reason for the archive. The text must be included in double quotes.</li></ul>

Option	Description
listarchive <format>	Displays each of the archives that have been created. To view a format other than the default, replace <format> with a supported format. See the list option for valid values.
restore -n <archid> -p <pwd>	<p>Restores an archive set of audit records so that they are displayed in the Compliance Auditor console.</p> <p><b>-n &lt;archid&gt;</b>: Specifies the archive to restore. Use the listarchive option to view the IDs of the archives.</p> <p><b>p &lt;pwd&gt;</b>: (Conditional) Specifies a password. If a password is specified for an archive, the same password must be used to restore the archive.</p>
purge	<p>Purges audit records that have been archived.</p> <p>Records that have been purged no longer appear in the Compliance Auditor console. A restore of the archive makes these records viewable again.</p>