



N1 Grid Service Provisioning System 5.0 System Administration Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817-6501
December 2004

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, N1 and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, N1 et Solaris sont des marques de fabrique ou des marques déposées, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



041202@10536



Contents

Preface	11
1 N1 Grid Service Provisioning System Overview	15
Configuring the N1 Grid Service Provisioning System – Process Overview	15
About the N1 Grid Service Provisioning System	16
Starting the N1 Grid Service Provisioning System Applications	16
System User Interfaces	18
▼ How to Launch the Browser Interface	19
▼ How to Launch the Single-Command Mode CLI	19
▼ How to Launch the Interactive Mode CLI	19
2 Managing Hosts, Host Types, Host Sets, and Host Searches	21
Setting Up Hosts in the N1 Grid Service Provisioning System – Process Overview	21
About Hosts	22
Physical Hosts	23
▼ How to Create a Physical Host Record	23
Preparing a Physical Host	25
▼ How to Prepare a Physical Host	26
Virtual Hosts	27
▼ How to Create a Virtual Host Manually	29
▼ How to Create a Host With a Targetable Component	30
Viewing Hosts	31
▼ How to View Hosts	31
Editing Hosts	32
▼ How to Modify Host-Type Attributes of a Host	32

▼ How to Add a Host to or Remove a Host From a Host Set	33
▼ How to Modify a Host Description	34
Hiding Hosts	34
▼ How to Hide a Host	34
▼ How to Display a Hidden Host	35
Deleting a Host	36
▼ How to Delete a Host	37
▼ How to Delete a Host Created by a Targetable Component	38
CLI Commands for Networking Hosts	38
About Host Types	38
Default Host Type	39
Host-Type Attributes	40
Creating Host Types	40
▼ How to Create Host Types	40
▼ How to Create a Host Type Based on an Existing Host Type	41
▼ How to View Host Types	42
Editing Host Types	43
▼ How to Edit Host Types	43
▼ How to Delete Host Types	44
About Host Sets	45
Platform Host Sets	45
Managing Host Sets	46
▼ How to Create Host Sets	46
▼ How to View Host Sets	46
▼ How to Edit or Hide Host Sets	47
▼ How to Delete Host Sets	48
About Host Searches	49
Creating and Editing Host Searches	49
▼ How to Create or Edit Host Searches	49
▼ How to View Host Searches	51
▼ How to Delete Host Searches	52
3 Controlling Access Using Permissions	53
Configuring Permissions	53
System-Wide Permissions	53
Folder-Specific Permissions	54
Create, Edit, and Delete Permission	55
Run Component Procedures Permission	56

	Check In Current and Configure	57
	Allow on Host Set	57
	Equivalent Permissions in the CLI	58
	Non-Configurable Permissions	58
4	Managing Users	59
	Planning User Groups and User Accounts	59
	Managing User Groups	60
	Default User Groups	60
	▼ How to Create User Groups	62
	▼ How to View User Groups	63
	Editing User Groups	63
	▼ How to Edit User Groups	64
	Deleting User Groups	65
	▼ How to Delete User Groups	65
	About User Accounts	66
	Default User Account	66
	Managing User Accounts	66
	▼ How to Create User Accounts	66
	Viewing User Accounts	68
	▼ How to View User Accounts	68
	▼ How to Deactivate or Hide User Accounts	69
	▼ How to Change a User's Group Membership	69
	▼ How to Change Your Password From the Log In Page	70
	▼ How to Change a User's Password from the User Account's Details Page	71
	▼ How to Change a User's Authentication Method	71
5	Plug-In Administration	73
	About Plug-Ins	73
	Importing Plug-Ins	74
	▼ How to Import a Plug-In	74
	▼ How to View a Plug-In	75
	Editing Plug-Ins	75
	▼ How to Edit a Plug-In	76
	Upgrading Plug-Ins	76
	▼ How to Upgrade a Plug-In	77
	Deleting Plug-Ins	78

▼ How to Delete a Plug-In	78
6 Setting Up Folders	81
Introduction to Folders	81
Creating Folders	82
▼ How to Create Folders	82
Viewing Folders	84
▼ How to View Folders	84
Moving Folders	85
▼ How to Move Folders	85
▼ How to Rename Folders	86
▼ How to Change the Description of Folders	86
Granting Folder Permissions to User Groups	87
▼ How to Modify the Folder Permissions of a User Group	87
Changing Folder Owners	88
▼ How to Change Folder Owners	88
Deleting Folders	89
▼ How to Delete Folders	89
7 Managing Objects	91
About Showing and Hiding Objects	91
About Finding Objects	92
Finding an Object From the Object's Table	92
Finding an Object Using Filters	93
Finding a Host Using Host Searches	93
About Categories	93
▼ How to Create a Category	93
▼ How to Associate a Category With an Object	94
▼ How to View Categories	95
▼ How to Filter Objects by Category	95
▼ How to Remove a Category From an Object	96
About Labels	96
▼ How to Add, Edit, or Remove a Label	97
▼ How to View a Label	97
8 Monitoring System Activity	99
Configuring Email Notification	99

▼ How to Set Up Notification Rules	100
▼ How to View Notification Rules	101
▼ How to Edit Notification Rules	102
▼ How to Delete Notification Rules	103
Event Severities	104
Viewing Logs	105
▼ How to View Remote Agent and Local Distributor Log Files	105
Viewing the Run History of a Plan or Comparison	106
▼ How to View Run Histories	106
▼ How to Delete Run Histories	107
9 Backing Up and Restoring	109
Backing Up and Restoring the Master Server	109
▼ How to Back Up a Master Server	109
▼ How to Restore a Master Server	111
Backing Up and Restoring Remote Agents	112
A Authentication Methods	113
Setting Up Additional Authentication Methods	113
▼ How to Modify External Authentication Methods	113
B Commonly Updated Configuration Variables	119
Glossary	121
Index	127

Tables

TABLE 1-1	Start Commands for Solaris OS, Red Hat Linux, and IBM AIX Applications	17
TABLE 1-2	Names of Services to Start for the Windows Master Server, Local Distributor, and Remote Agent	18
TABLE 3-1	HTML and CLI permission names	58
TABLE 8-1	Event Severities	104

Preface

The N1 Grid Service Provisioning System 5.0 System Administration Guide provides procedures for configuring the provisioning system after installation and managing the provisioning system as your environment changes.

Who Should Use This Book

This book is for system administrators who are responsible for configuring and maintaining the N1 Grid Service Provisioning System software.

How This Book Is Organized

[Chapter 1](#) provides an overview of the N1 Grid Service Provisioning System software product and high-level instructions of what needs to be done to configure the provisioning system after installation. A brief discussion of the provisioning system's user interfaces is also included.

[Chapter 2](#) describes how to manage your servers with the provisioning system.

[Chapter 3](#) describes the different types of permissions available in the provisioning system.

[Chapter 4](#) describes how to set up and manage user accounts with the provisioning system.

[Chapter 5](#) describes how to import and perform basic administrative tasks with third-party plug-ins.

[Chapter 6](#) describes how to perform folder administration tasks.

[Chapter 7](#) describes different ways you can manage objects within the provisioning system.

[Chapter 8](#) describes how to set up email notifications using notification rules and how to view host logs.

[Chapter 9](#) describes how to perform backups in the provisioning system.

[Appendix A](#) describes how to set up new authentication methods within the provisioning system.

[Appendix B](#) provides a description of several `config.properties` variables.

Related Books

The following books belong to the N1 Grid Service Provisioning System 5.0 documentation collection. All books listed below are available on docs.sun.com (<http://docs.sun.com>).

TABLE P-1 N1 Grid Service Provisioning System 5.0 Documentation

Application	Title
Software Installation	<i>N1 Grid Service Provisioning System 5.0 Installation Guide</i>
Plan and component creation	<i>N1 Grid Service Provisioning System 5.0 Plan and Component Developer's Guide</i>
Plan and component deployment	<i>N1 Grid Service Provisioning System 5.0 Operation and Provisioning Guide</i>
Command reference	<i>N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual</i>
XML Schema Reference	<i>N1 Grid Service Provisioning System 5.0 XML Schema Reference Guide</i>
Release Notes	<i>N1 Grid Service Provisioning System 5.0 Release Notes</i>
Plug-in Development	<i>N1 Grid Service Provisioning System 5.0 Plug-in Development Guide</i>

Accessing Sun Documentation Online

The docs.sun.comSM Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is <http://docs.sun.com>.

Ordering Sun Documentation

Sun Microsystems offers select product documentation in print. For a list of documents and how to order them, see “Buy printed documentation” at <http://docs.sun.com>.

Typographic Conventions

The following table describes the typographic changes that are used in this book.

TABLE P-2 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .

TABLE P-2 Typographic Conventions (Continued)

Typeface or Symbol	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . Perform a <i>patch analysis</i> . Do <i>not</i> save the file. [Note that some emphasized items appear bold online.]

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-3 Shell Prompts

Shell	Prompt
C shell prompt	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

N1 Grid Service Provisioning System Overview

This chapter describes the N1 Grid Service Provisioning System. It also provides information about how to configure the provisioning system after you have installed the product software.

Configuring the N1 Grid Service Provisioning System – Process Overview

After you have installed the product and set up the communication protocol, you need to set up the provisioning system. Setting up the provisioning system involves the following general actions.

1. Start the provisioning system applications.
See [“Starting the N1 Grid Service Provisioning System Applications”](#) on page 16.
2. Select and set up your user interface.
See [“System User Interfaces”](#) on page 18.
3. Configure your hosts.
See [Chapter 2](#).
4. If necessary, configure external authentication login modules.
See [Appendix A](#).
5. Create user groups.
See [“Planning User Groups and User Accounts”](#) on page 59
6. Create user accounts.
See [“About User Accounts”](#) on page 66.
7. Set up the provisioning system’s work space hierarchy.

For more information on setting up work spaces, see [Chapter 6](#).

8. Import and configure plug-ins.

See [Chapter 5](#).

About the N1 Grid Service Provisioning System

The N1 Grid Service Provisioning System software is an enterprise-class software platform that automates the deployment, configuration, and analysis of applications in data centers.

The provisioning software applies an object-oriented approach to:

- Application components
- Tasks that IT operators perform on application components: configuration, deployment, and analysis.

This object-oriented approach ensures that all the intelligence about an application is automatically taken into account every time that application is acted upon. This consistency makes data center operations more accurate and less prone to error. Through knowledge of what an application requires as a whole, IT operators gain unprecedented control over applications and data center operations.

The provisioning system can help you perform the following tasks.

- Automate and manage software rollouts, patches, and upgrades
- Develop models of your existing deployment processes
- Determine what software is installed on your hosts
- Compare the configurations of hosts
- Monitor and maintain documented and consistent configurations

For more information about the different provisioning system applications, see “Overview of N1 Grid Service Provisioning System 5.0 Applications” in *N1 Grid Service Provisioning System 5.0 Installation Guide*.

Starting the N1 Grid Service Provisioning System Applications

Before you can begin to configure the provisioning system you must start the Master Server, Remote Agents, and Local Distributors.

Starting Applications on Solaris OS, Red Hat Linux, and IBM AIX Systems

The following table lists the commands to start the N1 Grid Service Provisioning System software applications on Solaris OS, Red Hat Linux, and IBM AIX systems. *N1SPS-home* is the home directory of the application.



Caution – Do not use the Bourne shell to start the Master Server or other N1 Grid Service Provisioning System software applications. If you start the Master Server process using the `cr_server start` command in a Bourne shell, and if a `^C` command is issued to any subsequent command in the same shell that started the Master Server, then the database and Master Server processes stop.

In the *N1SPS5.0-home/server/bin/roxdb.out* file, the following messages appear as the most recent entries:

```
DEBUG: fast shutdown request
DEBUG: aborting any active transactions
```

TABLE 1-1 Start Commands for Solaris OS, Red Hat Linux, and IBM AIX Applications

Application	Path to Command	Command to Start
Master Server	<i>N1SPS5.0-home/server/bin/</i>	<code>cr_server start</code>
Local Distributor	<i>N1SPS5.0-home/ld/bin/</i>	<code>cr_ld start</code>
Remote Agent	<i>N1SPS5.0-home/agent/bin/</i>	<code>cr_ra start</code>
CLI Client	<i>N1SPS5.0-home/cli/bin/</i>	<code>cr_cli CLI-command</code>
Jython version of CLI Client	<i>N1SPS5.0-home/cli/bin/</i>	<code>cr_clij CLI-command</code>

Starting Applications on Windows Systems

On Windows systems, you start the Master Server, Local Distributor, and Remote Agent in the Services Panel. You start the CLI Client from a DOS window.

To start the Master Server, Local Distributor, or Remote Agent, click the Start menu, then Programs -> Administrative Tools -> Services. In the Services panel, find the name of the application and start it.

TABLE 1-2 Names of Services to Start for the Windows Master Server, Local Distributor, and Remote Agent

Application	Name of Service to Start
Master Server	N1 Grid Service Provisioning System software Server
	N1 Grid Service Provisioning System software PostgreSQL Server
	N1 Grid Service Provisioning System software IPC Daemon
	N1 Grid Service Provisioning System software Database Preparer
Local Distributor	N1 Grid Service Provisioning System software Distributor
Remote Agent	N1 Grid Service Provisioning System software Agent

System User Interfaces

After you have started the provisioning system applications, you can access the provisioning system through three interfaces.

Browser Interface

This primary interface to the provisioning system interacts with the Master Server through the Master Server's built-in web server. The object-oriented design allows quick navigation and the ability to monitor the system without the need to reference commands. Automation is available to ease component and plan creation. Some activities can only be accomplished by using this interface, such as performing bulk move procedures.

For information about how to access the browser interface, see [“How to Launch the Browser Interface”](#) on page 19.

Command-Line Interface: Single-Command Mode

This CLI mode enables you to issue a single command from a shell to the Master Server. You can also use the CLI to write scripts in languages other than Jython.

For information about how to access and use the non-interactive CLI interface, see *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Command-Line Interface: Interactive mode

This CLI mode allows you to embed CLI commands into sophisticated object-oriented scripts. The Jython Interpreter also allows you to submit multiple CLI commands to the Master Server.

For information about how to access and use the CLI Jython interface, see *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

▼ How to Launch the Browser Interface

Before You Begin The Master Server must be installed and running. For information about how to start the Master Server, see [“Starting the N1 Grid Service Provisioning System Applications” on page 16.](#)

- Steps**
1. **In a web browser, navigate to the home page of the Master Server.**
The URL for the Master Server is configured during the installation process. If you don't know the web address for your master server, consult the person who installed and configured the N1 Grid Service Provisioning System software.
The default location is `http://masterserver-hostname:8080`.

2. **Type your user name and password in the appropriate fields.**

Note – The provisioning system is installed with an administrator account called `admin`. The password for this account is set during installation.

For more information about the `admin` user account, see [“Default User Account” on page 66.](#)

3. **Click Log In.**

Tip – To log in without using your mouse, press the Tab key after you type the password. Then press the Return key to initiate the login process.

▼ How to Launch the Single-Command Mode CLI

For information about how to access and use the single-command CLI interface, see *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

- Step**
- **From a machine where the CLI Client application is installed, type the command you need to run.**

```
N1SPS-home/cli/bin/cr_cli -cmd subsystem.object.command -u user -p password
```

For more information on command syntax and command output formats, see Chapter 1, “Using the Command-Line Interface,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

▼ How to Launch the Interactive Mode CLI

For information about how to access and use the CLI Jython interface, see the *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

- Steps** 1. From a machine where the CLI Client application and Jython are installed, start the Interactive Mode CLI.

```
N1SPS-home/cli/bin/cr_cli.j
```

2. Include the following code at the beginning of your script.

```
from clui import *
app=PyCLUI()
app.execStr(CLI command)
App.close()
```

Note – The assignment `app=PyCLUI()` calls the CLI. The `App.close()` call deletes the instance of this Jython class.

Managing Hosts, Host Types, Host Sets, and Host Searches

You can manage servers through the provisioning system by modeling and storing them as hosts.

By managing and organizing hosts efficiently, the provisioning system enables you to reduce the complexity of deploying, configuring, and analyzing applications in data centers.

This chapter discusses the following topics.

- [“About Hosts” on page 22](#)
- [“About Host Types” on page 38](#)
- [“About Host Sets” on page 45](#)
- [“About Host Searches” on page 49](#)

Setting Up Hosts in the N1 Grid Service Provisioning System – Process Overview

The following outline shows the tasks that you need to perform when setting up the provisioning system.

1. Determine how many hosts you need and the roles that the different hosts will have.
2. If necessary, define the host types that will map to the hosts that you need to create. See [“Creating Host Types” on page 40](#).
3. If necessary, install the Remote Agent or Local Distributor software on your physical servers. See the *N1 Grid Service Provisioning System 5.0 Installation Guide*.
4. Working from the Master Server down the communication chain, create host records for the Local Distributors and Remote Agents in the provisioning system.

See “Physical Hosts” on page 23.

Note – For each host, you need to define the connection between the physical host and the next host closer to the Master Server in the communication chain. Therefore, the host’s parent must be added to the provisioning system before you can add the host itself.

5. Prepare the hosts.
See “Preparing a Physical Host” on page 25.
6. **Optional.** Create virtual hosts.
See “Virtual Hosts” on page 27.

About Hosts

The provisioning system has two types of hosts.

- **Physical hosts** – Represent physical servers that are connected to the network. Within the provisioning system, a physical host can contain the following applications.
 - Remote Agent – A physical host that can be used as a target for deploying or modeling components. You can also compare remote agents to one another.
 - Local Distributor – A physical host that is used as a distributor of components.
- **Virtual hosts** – Represent services that act as a host for other services. For example, a virtual host can represent an application server that acts as a host for web applications.

The provisioning system uses hosts in the following ways.

- **Deploying.** If you need to deploy a file or an application (a component) to a server, the server must be registered as a host.
- **Distributing.** A server used as an intermediary to help distribute deployed components to other hosts must be registered as a host.
- **Capturing.** If you want to capture an application (or add a file to the data repository) so that it can be deployed to other hosts, the server that contains the application must be registered as a host.
- **Comparing.** If you want to compare one server with another, both servers must be registered as hosts.

Physical Hosts

A physical host represents a server in your network that has been registered with the provisioning system.

Physical Host Characteristics

- Each host must have a unique IP address.
- Each host must have a unique host name.
When a server name is used to identify a new host, the server name is resolved to an IP address when the host is saved.
- When a Remote Agent and a Local Distributor are installed on the same host, they must have different, dedicated ports.
- Host locking prevents different users from simultaneously attempting to update the same host.

Before the provisioning system can work with a server, you must perform the following steps.

1. Install Remote Agent or Local Distributor software on the server.
2. Register the server as a physical host by creating a host record for the server.
3. Prepare the physical host.

The provisioning system prepares a host by installing plug-in specific software on the host's Remote Agent in the form of system service components. A host is considered prepared if all system service components that are applicable to the host have been installed on it. After a host has been prepared, you can perform operations on it, such as running comparisons and deploying applications.

Because each physical host represents a specific server, do not attempt to migrate a physical host record to another server. For example, if you need to replace a server that was registered as a Local Distributor, replace the server and create a new Local Distributor. If you attempt to migrate the old Local Distributor record to the new server, the Master Server might contain records indicating that the new server contains files that haven't been installed.

▼ How to Create a Physical Host Record

You need to create a physical host record any time you need to replace a server or add a server to the provisioning system.

This procedure describes how to create new hosts by using the browser interface. You can also create hosts by using the following command.

```
hdb . h . add – Creates a new host.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin

- You must belong to a user group that has *Create*, *Edit*, *Delete* permission on hosts.

- Determine if the Remote Agent or Local Distributor software has been installed on the server.

The host record can be created when Remote Agent and Local Distributor software has not been installed. However, if you plan to use the host for other operations, the Remote Agents and Local Distributors must be running.

To install the Remote Agent and Local Distributor software, see the *N1 Grid Service Provisioning System 5.0 Installation Guide*.

- Determine if the Remote Agent or Local Distributor has been started.

The host record can be created when Remote Agents and Local Distributors are not running. However, if you plan to use the host for other operations, the Remote Agents and Local Distributors must be running.

See the [“Starting the N1 Grid Service Provisioning System Applications”](#) on page 16.

- Ensure the parent host record has been created.

For each host, you need to define the connection between the physical host and the next host closer to the Master Server in the communication chain. Therefore, the host’s parent must be added to the provisioning system before you can add the host itself.

Note – The Master Server is created by the system during installation. You will not need to manually create the Master Server before creating other physical hosts.

Steps

1. **Go to the Hosts page.**

See [“How to View Hosts”](#) on page 31.

2. **In the Host text box, type a name for the new host.**

3. **From the menu, select a host type for the host.**

4. **In the Description text box, type a brief description of the new host.**

5. **Click Create.**

The Edit page for the new host is displayed.

6. **Ensure that the host name, description, and host type display properly.**

7. **(Optional) Update the host-type attribute values.**

If the host type that you selected has attributes associated with it, the host’s Edit page lists the attributes and default values in the attribute table. To assign a

non-default value to an attribute, select the checkbox next to the default value. In the text box that appears, type the new attribute value.

8. **Indicate that you are creating a physical host by selecting Physical Host.**
9. **Indicate the type of provisioning system application that is installed on this host.**
10. **Configure the connection type, network address, and parent information for this host.**

Note – For information on configuring SSL and SSH network connections, see Chapter 6, “Configuring the N1 Grid Service Provisioning System 5.0 to Use Secure Shell,” in *N1 Grid Service Provisioning System 5.0 Installation Guide*, or Chapter 7, “Configuring the N1 Grid Service Provisioning System 5.0 for SSL,” in *N1 Grid Service Provisioning System 5.0 Installation Guide*.

11. **(Optional) To add this host to a host set, click Add to Host Sets in the Host Relationships area of the page.**

In the window that appears, select the host set to which you want to add the host.

For more information on host sets, see “[About Host Sets](#)” on page 45.

12. **(Optional) To hide the host, select Hidden at the bottom of the page.**

13. **Click Save.**

Preparing a Physical Host

The provisioning system prepares a host by installing plug-in specific software on the host’s Remote Agent in the form of system service components. A host is considered prepared if all system service components that are applicable to the host have been installed on it. After a host has been prepared, you can perform operations on it, such as running comparisons and deploying applications.

You might need to re-prepare a physical host in the following scenarios.

- A new system service component is installed in the Master Server that needs to be deployed to the host.

This scenario is likely to occur if a new plug-in is imported.

- The host’s operating system has been upgraded and new system service components are required.

Whenever an operating system is upgraded, prepare the host again. The Master Server might not recognize that the host needs to be prepared again until the next time you attempt to perform an operation on the host.

The browser interface provides several options for when and how to prepare a host. You can prepare a host at the following times.

- When you create the host
- When you view the host's details
- When you view the Hosts table

The browser interface will also prompt you to prepare a host if you attempt to perform an operation on an unprepared host.

▼ How to Prepare a Physical Host

This procedure provides instructions for preparing a host through the Hosts table. This option allows you to prepare several hosts at the same time, which can be useful when you update the entire network after a provisioning system upgrade.

You can also prepare hosts by using the following command.

```
pe.h.prep – Prepares a host.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin Before you can prepare a physical host, you must create the host in the provisioning system. For more information on creating a physical host, see [“How to Create a Physical Host Record” on page 23](#).

To prepare a physical host, you must belong to a user group that has Create, Edit, Delete permission on hosts.

- Steps**
- 1. Go to the Hosts page.**
See [“How to View Hosts” on page 31](#).
 - 2. From the Hosts page, select Indicate Un-prepared Hosts above the Hosts table.**
The Hosts table reloads. Hosts that need to be prepared are highlighted in yellow.
 - 3. Select each host that you need to prepare.**
You can prepare more than one host at a time.
 - 4. In the last row of the Hosts Table, click Prepare Host.**
A window appears and provides the status for the host preparation. When the host preparation completes successfully, the window automatically closes. If host preparation fails, error messages appear in the content frame.

Virtual Hosts

A virtual host represents a service that acts as a host for other services.

Virtual Host Characteristics

- Multiple virtual hosts can exist on a single physical server.
- Every virtual host has a parent host, which can be either a physical host or another virtual host.
- Virtual hosts can be nested arbitrarily, although the outermost parent host must be a physical host.
- You can run comparisons against virtual hosts.
- In most cases, virtual hosts will be created as a by-product of installing plug-in related components.

Note – To determine if a component will create a host, view the component's Details page. If the Creates Host field has a location listed, the component will create a host at that location upon component installation.

Virtual hosts represent logical deployment targets. For example, a WebLogic virtual host can represent a WebLogic application server as a deployment target.

The virtual host provides an interface between the component being deployed (in this example, the web application) and the component being targeted (in this example, the application server).

Rather than deploy web applications directly to WebLogic, which would entail working with the application server's custom administrative console on each host, by using the virtual hosts within the provisioning system, you can deploy a web application to the virtual host and the virtual host will provide WebLogic with specific configuration information. This functionality is particularly useful if you need to deploy a web application to several hosts at once.

The following list includes examples of information supplied to a virtual host and used in the interaction between a virtual host and the service running on a remote machine.

- user name and password for the service
- port numbers
- connection types

By using the provisioning system rather than performing deployments manually, you can receive the following benefits.

- Simultaneous deployments to all hosts in a host set.
- Version control that tracks all deployment and configuration data for reference, reconstruction, and automated rollback to previous states.
- Logging and reporting that creates detailed logs of every action taken by the system across all applications and managed servers. These logs provide a complete audit history of every change made to every host.
- Differential deployments that enable delta-only distribution of large content directories, thereby speeding up deployments significantly and optimizing incremental directory updates.

Virtual hosts are also useful for maintaining multiple instances of an application on a single physical host. These hosts establish a clear distinction between physical servers and the applications running on them. This distinction is especially useful when managing multiple applications that must interact with one another to function.

Because each virtual host represents a specific service, do not attempt to migrate a virtual host definition to another server. For example, if you need to migrate your WebLogic application server to a new physical host, migrate the application and create a new virtual host. If you attempt to migrate the virtual host to a different physical host, the Master Server will contain records with inaccurate information about the migrated application server.

Targetable Components

Hosts can be created manually or they can be created by a special type of component, the targetable component. Targetable components are available with many of the third-party plug-ins.

The targetable component represents a service that components and plans need to use as an installation or call target. When a targetable component is installed, the component creates the virtual or physical host, which acts as the interface between the targetable component and the other components and plans.

When the targetable component is installed, the host is created with a predefined set of information. This information set includes the following values.

- Host name
- Host type with default values
This host type is included with the plug-in.
- Type of host: physical or virtual
- The ID of the targetable component that created the host

After the targetable component creates the host, you can modify the host's description and host-type attributes.

If a new version of the targetable component is installed, the targetable component's host is updated with the new targetable component information.

When the targetable component is uninstalled, the host that it created is also deleted.

▼ How to Create a Virtual Host Manually

For information about virtual hosts and their uses, see [“Virtual Hosts” on page 27](#).

This procedure describes how to create virtual hosts by using the browser interface. You can also create virtual hosts by using the following command.

```
hdb . h . add – Creates a new host.
```

For a detailed description of this command, see Chapter 7, *“hdb: CLI Commands for Managing Hosts,”* in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin Before you add a virtual host, you need to have a physical host setup that will contain the virtual host. You also need to have the deployment target set up.

To create a virtual host, you must belong to a user group that has `Create`, `Edit`, `Delete` permission on hosts.

- Steps**
- 1. Go to the Hosts page.**
See [“How to View Hosts” on page 31](#).
 - 2. In the top row of the Hosts table, perform the following steps.**
 - a. Type a name for the new host in the Host text box.**
 - b. From the menu, select a host type for the host.**
 - c. Type a brief description of the new host in the Description text box.**
 - d. Click Create.**
The Edit page for the new host is displayed.
 - 3. Ensure that the host name, description, and host type display properly.**
 - 4. If necessary, update the host-type attribute values.**
If the host type that you selected has attributes associated with it, the Edit page lists the attributes and default values in the attribute table. To assign a custom value to an attribute, click the checkbox next to the default value. In the text box that appears, type the new attribute value.
 - 5. Indicate that you are creating a virtual host by selecting Virtual Host.**

6. **Type the name of the parent host or click *Select from List* to select the host name from the window that appears.**

Every virtual host has a parent host that is either a physical host or a virtual host. Virtual hosts can be nested arbitrarily, although the outermost parent host must be a physical host.

7. **(Optional) To add this host to a host set, click *Add to Host Sets* in the *Host Relationships* area of the page.**

Select the host set to which you want to add the host.

8. **Click *Save*.**

▼ How to Create a Host With a Targetable Component

This procedure describes how to use a targetable component to create hosts by using the browser interface. You can also create hosts from targetable components by using the following commands.

- `pdb.p.genplan` – Creates the plan from the install component procedure.
- `pe.p.run` – Runs the plan.
- `hdb.h.mod` – Configures the host.

Before You Begin

To create a virtual host with a targetable component, you must have *Run Component Procedures* and *Allow of Host Sets* permissions.

Steps

1. **Install the targetable component.**
2. **Go to the *Hosts* page to see the new host and click *Details*.**
3. **(Optional) Update the host's description.**

4. **If necessary, update the host-type attribute values.**

If the host type that you selected has attributes associated with it, the *Edit* page lists the attributes and default values in the attribute table. To assign a custom value to an attribute, click the checkbox next to the default value. In the text box that appears, type the new attribute value.

5. **Click *Save*.**
6. **If the new host is a physical host, prepare the host.**
See [“How to Prepare a Physical Host”](#) on page 26.

Viewing Hosts

When you view hosts, you see the list of virtual and physical hosts along with text describing the host. If the host is physical, information about the server's operating system is also displayed.

You might need to view hosts in the following situations.

- You need to determine which hosts are currently registered in the provisioning system.
- You need to know which hosts represent a particular host search.
You can filter hosts according to a host search. For more information on host searches, see [“About Host Searches” on page 49](#).
- You need to determine each host's functional role within the provisioning system.
For example, if a host is a Remote Agent, you have control to actively manage applications and components on that host. If the host is only a Local Distributor, you do not manage the host because it is merely a conduit for data to Remote Agents.

▼ How to View Hosts

This procedure describes how to view hosts by using the browser interface. You can also view hosts by using the following commands.

- `hdb.h.1a` – Displays all hosts.
- `hdb.h.1o` – Displays detailed information about a particular host.

For a detailed description of these commands, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

- Steps**
1. **From the navigation menu, choose Host Setup.**
The Hosts Setup page is displayed.
 2. **In the Hosts Setup page, click Hosts.**
The Hosts page lists all the unhidden hosts.
 3. **(Optional) To include the hidden hosts, select Show Hidden Hosts in the upper right corner of the page.**
 4. **(Optional) View more information about a host or host search.**
 - **To view the details of a specific host, click Details in the row listing the host that you want to view.**
The Details page presents the configuration details of the host.

- If you would like to view the list of hosts in a particular host search, you can select the host search from the menu above the hosts table.

The browser interface displays the list of hosts that belong to the selected host search.

Note – If a host was created by a targetable component, the component name displays next to the Created by Component label.

Editing Hosts

You can modify host information after the host has been added to the provisioning system. Hosts are locked when a user saves a modification to prevent different users from simultaneously updating the same host.

You might want to edit a host in the following situations.

- You need to update the host-type attributes of a particular host.
See [“How to Modify Host-Type Attributes of a Host”](#) on page 32.
- You want to add the host to or remove the host from a new or existing host set.
See [“How to Add a Host to or Remove a Host From a Host Set”](#) on page 33.
This task can also be performed from a host set’s Details page. For more information, see [“How to Edit or Hide Host Sets”](#) on page 47.
- You want to modify the host description.
See [“How to Modify a Host Description”](#) on page 34.

▼ How to Modify Host-Type Attributes of a Host

This procedure describes how to modify host-type attributes by using the browser interface. You can also edit hosts by using the following command.

`hdb . h . mod` – Modifies a host.

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

For information about host-type attributes, see [“About Host Types”](#) on page 38.

Before You Begin To modify host-type attributes, you must belong to a user group that has Create, Edit, Delete permission on hosts.

Steps 1. Go to the Details page of the host you plan to edit.

See “How to View Hosts” on page 31.

2. On the Details page, click the Edit button near the bottom of the page.
3. To update the host’s host-type attributes, select the attribute value you want to modify.
An Attribute Value field appears and is highlighted.
4. Type the new value in the Attribute Value field.
5. When you have finished your modifications, click Save.

▼ How to Add a Host to or Remove a Host From a Host Set

This procedure describes how to modify a host’s host set membership by using the browser interface. You can also edit hosts by using the following command.

```
hdb.hs.mod – Modifies a host.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Note – You cannot perform this task on a host created by a targetable component.

Before You Begin To modify a host’s host set membership, you must belong to a user group that has Create, Edit, Delete permission on hosts.

- Steps**
1. Go to the Details page of the host you plan to edit.
See “How to View Hosts” on page 31.
 2. On the Details page, click the Edit button near the bottom of the page.
 3. Update the host’s host set membership.
 - To add the host to a host set, click Add to Host Sets.
The Add to Host Sets window appears.
 - a. Select the host set to which you want to add the host.
You can select multiple host sets at once.
 - b. Click Add Host to Selected Host Sets.
 - To remove a host from a host set, select the host set in the Member of Host Sets Field and click Remove from Selected Host Sets.

4. When you have finished your modifications, click Save.

▼ How to Modify a Host Description

This procedure describes how to modify a host description by using the browser interface. You can also edit hosts by using the following command.

```
hdb.h.mod – Modifies a host.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To edit a host, you must belong to a user group that has Create, Edit, Delete permission on hosts.

- Steps**
1. Go to the Details page of the host you plan to edit.
See “How to View Hosts” on page 31.
 2. On the Details page, click the Edit button near the bottom of the page.
 3. To change the host description, edit the text located in the Description text field.
 4. When you have finished your modifications, click Save.

Hiding Hosts

Because the objects you store in the provisioning system are deeply interconnected, you cannot always delete a host. You can choose to hide objects that you are not currently using to keep your work space uncluttered.

Hiding an object does not make it inactive or unusable. For example, if a host is contained in a host set, hiding the host does not remove the host from the host set. Therefore, if you run a plan on the host set, the host will still be included in the plan run.

Note – When you hide a host, the host is hidden for all users and user groups.

For information on deleting a host, see “How to Delete a Host” on page 37.

▼ How to Hide a Host

This procedure describes how to hide hosts by using the browser interface. You can also hide hosts by using the following command.

```
hdb.h.mod -hide true – Hides a host.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To hide a host, you must belong to a user group that has *Create, Edit, Delete* permission on hosts.

Steps 1. **Go to the Details page of the host you plan to hide.**

See “[How to View Hosts](#)” on page 31.

2. **Click Edit.**

3. **Select Hidden at the bottom of the page.**

4. **Click Save.**

The host no longer appears in the Hosts page.

5. **(Optional) If you hide the host because you will no longer use it, verify that the host is not used in any host sets or found by a host search.**

In the Host Relationships section of the host’s Details page, verify that no host sets are listed in the Member of Host Sets box.

Note – If a host search includes a hidden host and the host search is included in a host set, plans that are run on the host set will be run against the hidden host.

▼ How to Display a Hidden Host

You might need to display a hidden host in the following situations.

- You need to perform a task on a host that is hidden.
- Your use of the host has changed and you now want to actively manage or delete it.

This procedure describes how to view hidden hosts by using the browser interface. You can also view hidden hosts by using the following command.

```
hdb.h.la -sh true – Displays all hidden hosts.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Steps 1. **Go to the Hosts page.**

See [“How to View Hosts”](#) on page 31.

2. **In the upper right corner of the Hosts page, select Show Hidden Hosts.**
The Hosts page now lists hidden hosts as well as displayed hosts. Hidden hosts appear in a row colored dark gray.
3. **(Optional) To remove the hide feature from this host, perform the following steps.**
 - a. **Find the row listing the host you want to view and click Details.**
A Details page for the host appears. The Details page presents the configuration details of the host.
 - b. **Deselect Hidden at the bottom of the page.**
 - c. **Click Save.**

Deleting a Host

By deleting a host, you remove the host from the provisioning system. If you want to keep the host but not view it, see [“How to Hide a Host”](#) on page 34.

When you delete a host, the provisioning system will also delete the following objects.

- All virtual hosts nested in the host.
- All component instances that were installed on the host.

Note – If the component instances installed on the host were created from a resource that was originally collected from the host, the component is updated to remove the reference to the host.

You might want to delete a host in the following situations.

- You need to decommission a server that was registered with the provisioning system.
- You no longer want a record of the host.

▼ How to Delete a Host



Caution – Deleting a host is permanent and cannot be undone. Deleting a host might also be time consuming.

This procedure describes how to delete hosts by using the browser interface. You can also delete hosts by using the following command.

```
hdb . h . del – Deletes a host.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Note – If the host that you want to delete was created by a targetable component, see [“How to Delete a Host Created by a Targetable Component” on page 38](#).

Before You Begin

You cannot delete a host that is referenced by one of the following objects.

- A comparison
- A resource
- A snapshot
- A Remote Agent, Local Distributor, or CLI application
- A targetable component

A host created by a targetable component cannot be deleted while the targetable component is still installed. However, if the parent host that contains the targetable component host is deleted, the targetable component host is also deleted.

To delete a host, you must belong to a user group that has `Create`, `Edit`, `Delete` permission on hosts.

Steps

1. **Go to the Details page of the host that you plan to delete.**

See [“How to View Hosts” on page 31](#).

2. **Click the Delete button to delete the selected host.**

You are asked to verify that you want to delete the host.

3. **Click the Continue to Delete button.**

After the host has been deleted, the Hosts table no longer displays the host.

▼ How to Delete a Host Created by a Targetable Component

This procedure describes how to delete a host that was created by a targetable component by using the browser interface. You can also delete hosts by using the following command.

- `pdb.p.genplan` – Creates the plan from the uninstall component procedure.
- `pe.p.run` – Runs the plan.

Before You Begin To create a virtual host with a targetable component, you must have `Run Component Procedures` and `Allow of Host Sets` permissions.

- Step**
- **Uninstall the targetable component.**
After the targetable component is uninstalled, the `Hosts` table no longer displays the host.

CLI Commands for Networking Hosts

The following `net` commands perform networking tasks related to hosts.

<code>net.genconf</code>	Generates the Transport Configuration file for a provisioning system application. You will need to run this command manually if you updated the network configuration information of a host after the host was created.
<code>net.ping</code>	Checks connectivity to a Remote Agent or Local Distributor by executing the TCP/IP <code>ping</code> command.
<code>net.traceroute</code>	Determines the route to a Remote Agent or Distributor.

About Host Types

A host type is a container for variables that are assigned to a particular host. These variables are called host-type attributes and are based on some combination of host characteristics, such as function, location, or configuration. A host type determines the set of host-type attributes that are assigned to a host.

Creating a host type or using an existing host type provides two benefits.

- Provides the ability to group similar hosts
- Provides a framework for dynamically configured variables, known as host-type attributes

By assigning each host a host type, grouping hosts by their most important characteristics is easier. For example, you can define a host type for the hosts that are used as web servers, another host type for hosts that are used for database servers, and so on.

Host types also enable you to set configuration values dynamically when a data center operation is performed. A web server host type might include attributes such as `Location`, `WebServer_User`, `WebServer_Group`, `HTTP_Port`, and `HTTPS_Port`. A database server host type might include a different set of attributes such as `Location`, `DB_Directory`, `DB_InstanceName`, `DB_Port`, and `DB_SecurePort`. If, in this example, you define a host type for web servers and the host type includes an HTTP port variable, when you perform an installation on hosts of this host type, your installation plan could automatically set the HTTP port variable of each web server based on the HTTP port attribute value set for each host.

Every host that you manage with the provisioning system will be assigned a host type, either the default host type, `crhost`, or a host type that you have defined. Each host can only have one host type.

When hosts are created by targetable components, the host type is automatically assigned. For more information about targetable components, see [“Targetable Components” on page 28](#).

Default Host Type

The N1 Grid Service Provisioning System software includes the default host type, `system#crhost`. By default, all hosts are assigned this host type. The `system#crhost` host type cannot be modified or deleted. The `system#crhost` type includes two attributes:

- `sys.hostName`
- `sys.description`

There are no default values for these attributes. The name value must be specified for each host when you define the host. The description value is optional.

Host-Type Attributes

A host type is distinguished by its name and its attributes. An attribute is a name-value pair.

Host-Type Attribute Characteristics

- Each attribute can have a default value.
- Attribute names must begin with a letter and cannot include spaces.
- Attribute names must be unique within a particular host type.
- Adding new attributes to the host type results in the attributes and their defaults being added to every host of that type.
- Removing attributes from a host type permanently deletes those attributes from every host of that type.
- Attributes cannot be renamed. Attributes can only be deleted and added again with a new name, which results in the attribute value being reset to the default on all hosts of that type.
- Attribute values can be literal values or can represent variables determined by your configuration.

These variables are called substitution variables. One of the most common substitution variables is : [container:installPath].

When you define a host, you can selectively override the host type's default attribute values. Default values that are not overwritten will be applied to the host. If a host-type attribute's default value is changed after hosts are created, all hosts that did not override the default value are updated to reflect the new attribute value.

Creating Host Types

Host types enable you to group hosts by particular characteristics. Host types also enable you to set up dynamic configuration values.

For more information on substitution variables and configuration generation, see Chapter 6, "Configuration Generation," in *N1 Grid Service Provisioning System 5.0 Plan and Component Developer's Guide*.

▼ How to Create Host Types

This procedure describes how to create host types by using the browser interface. You can also create host types by using the following command.

hdb.ht.add – Creates a new host type.

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Note – If you would like to create a new host type from an existing host type, see “[How to Edit Host Types](#)” on page 43.

Before You Begin To create a host type, you must belong to a user group that has `Create`, `Edit`, `Delete` permission on hosts.

- Steps**
1. **Go to the Host Types page.**
See “[How to View Host Types](#)” on page 42.
 2. **In the top row of the table on the Host Types page, type a name and brief description for the new host type and click Create.**
The Host Types Details page for the new host type is displayed.
 3. **Perform the following steps for each attribute that you plan to assign to the new host type.**
 - a. **Type a name for the attribute in the Attribute column.**
 - b. **(Optional) To assign a default value to the attribute , type a value in the Default Value column.**
If you do not assign a value to the attribute, the default value is an empty string.
 - c. **In the action column, click Add.**
 4. **(Optional) To hide the host type, select Hidden.**
 5. **To save the host type, click Save.**

▼ How to Create a Host Type Based on an Existing Host Type

Before You Begin To edit a host type, you must belong to a user group that has `Create`, `Edit`, `Delete` permission on hosts.

- Steps**
1. **Go to the Host Types page.**
See “[How to View Host Types](#)” on page 42.

2. **On the Host Types page, find the row listing the host type that you plan to edit and click Details.**
A Details page appears for the selected host type.
3. **Use the controls on the Edit page to modify the host type.**
4. **To save the modified host type as a new host type, click Save As.**
Another Details page appears.
5. **Type the name of the new host type.**
The new host type name must be unique to existing host types.
6. **Click Continue to Save As.**

▼ How to View Host Types

This procedure describes how to view host types by using the browser interface. You can also view host types by using the following commands.

- `hdb.ht.la` – Displays all host types
- `hdb.ht.lo` – Displays detailed information about a particular host type

For a detailed description of these commands, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

- Steps**
1. **From the navigation menu, choose Administrative.**
The browser interface displays the Administrative page.
 2. **In the Administrative page, click Host Types.**
The browser interface displays the Host types page.
 3. **(Optional) To include hidden host types, select Show Hidden Host Types in the upper right corner of the page.**
 4. **(Optional) To view the details of a specific host type, click Details in the row that lists the host type that you want to view.**
A Details page for the Host Type is displayed.

Editing Host Types

You can make the following modifications to the host type.

- Add, change, or remove attributes
- Change the name or description of the host type

For more information about the impact of updating host-type attributes, see [“Host-Type Attributes” on page 40](#).

▼ How to Edit Host Types

If you need to edit the host-type attributes of a particular host, see [“How to Modify Host-Type Attributes of a Host” on page 32](#).

This procedure describes how to edit host types by using the browser interface. You can also edit host types by using the following command.

```
hdb.ht.mod – Edits a host type.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To edit a host type, you must belong to a user group that has `Create`, `Edit`, `Delete` permission on hosts.

- Steps**
1. **Go to the Details page of the host type that you want to edit.**
See [“How to View Host Types” on page 42](#).
 2. **Use the controls on the Details page to modify the host type.**
 - **Add an attribute.**
 - a. **Type the name of the new attribute in the Attribute field.**
 - b. **(Optional) Type the default value in the Default Value field.**
If you do not assign a value to the attribute, the default value is an empty string.
 - **Remove an attribute by clicking Remove in the Actions column.**
 3. **(Optional) To hide the host type, click Hidden.**
 4. **To save the host type, click Save.**

▼ How to Delete Host Types



Caution – Deleting a host type is permanent and cannot be undone. Deleting a host type might also be time consuming.

This procedure describes how to delete host types by using the browser interface. You can also delete host types by using the following command.

`hdb.ht.del` – Deletes a host type.

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin You cannot delete a host type if any hosts of that type exist. You also cannot delete a host type that is owned by a plug-in.

To delete a host type, you must belong to a user group that has `Create`, `Edit`, `Delete` permission on hosts.

Steps 1. **Go to the Details page of the host type that you plan to delete.**

See “[How to View Host Types](#)” on page 42.

2. **Click the Delete button at the bottom of the page.**

Another page appears that provides the option to cancel the deletion.

Note – You cannot delete host types that are in use. If a host type that you try to delete is being used, a message will list the host type and the reason that it was not deleted.

3. **Verify that you are deleting the correct host type and click the Continue to Delete button.**

The host type that you deleted is no longer listed in table on the Host Types page.

About Host Sets

The provisioning system enables you to group hosts into logical groups called host sets. After hosts are grouped into host sets, you can perform operations on them in the same manner as a single host. Host sets enable you to perform operations on a large number of hosts simultaneously.

A host set can include any of the following objects.

- Hosts
- Host searches, which add all the hosts found by the host searches to the host set
- Host sets

When you define a host set to include a host search, the host search executes and the initial list of hosts is determined. Any time an operation is performed on the host set, the host search is performed again and an updated list of hosts is generated. For example, when the provisioning system uses the host set as a target for an installation or a comparison, the list of hosts that meet the host search criteria is regenerated.

Host sets based on host searches enable the provisioning system to group hosts based on their actual configurations at runtime. Host searches ensure that your host set contains all the hosts that match the search criteria, rather than requiring you to remember which hosts are configured a certain way.

Host sets enable you to perform operations on groups of hosts. A host can appear in a host set only once. If two host searches happened both to retrieve the same hosts, each of these hosts would appear in the host set just once. Duplicate entries are removed from the host set list.

Platform Host Sets

The provisioning system includes a hierarchy of nested host sets that are based on the OS platforms of hosts. The top host set is *any*, which includes hosts running any of the supported OS platforms.

Platform host sets are used to limit the hosts that a particular component or resource may be installed on. Platform host sets typically contain a host search that matches one or more platform-specific properties. The provisioning system includes several predefined platform host sets (and associated host searches).

Each platform host set that is included with the provisioning system will be prefixed with `system#` (for example `system#Windows 2000 Server`). The prefix indicates that they are defined by the system plug-in.

Managing Host Sets

Host sets enable you to perform operations on a large number of remote targets simultaneously.

▼ How to Create Host Sets

This procedure describes how to create host sets by using the browser interface. You can also create host sets by using the following command.

```
hdb.hs.add – Creates a host set.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To create a host set, you must belong to a user group that has *Create*, *Edit*, *Delete* permission on hosts.

- Steps**
- 1. Go to the Host Sets page.**
See “[How to View Host Sets](#)” on page 46.
 - 2. In the top row of the Host Sets table, type a name and a brief description for the host set you plan to create and click Create.**
A Details page for the new host set is displayed. This page includes the list of current member hosts and controls to add or remove hosts from the host set.
 - 3. Use the controls on this page to define the host set.**
Each time you add a host, host set, or host search to the host set, the Host Sets page updates the list of hosts in the host set.
 - 4. Click Save.**

▼ How to View Host Sets

You can survey how hosts are currently grouped within the provisioning system.

This procedure describes how to view host sets by using the browser interface. You can also view host sets by using the following commands.

- `hdb.hs.la` – Displays all host sets.
- `hdb.hs.lo` – Displays detailed information about a particular host set.

For a detailed description of these commands, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

- Steps**
- 1. From the navigation menu, choose Host Setup.**
The Host Setup page is displayed.
 - 2. In the Host Setup page, click Host Sets.**
The Host Sets page is displayed. By default, this page lists all the unhidden host sets.
 - 3. (Optional) To include hidden host sets, select Show Hidden Host Sets in the upper right corner of the page.**
 - 4. (Optional) To view the details of a specific host set, click Details in the row listing the host set that you want to view.**
The Details page for the host set lists the host set’s member criteria along with the hosts that are included in the host set.

▼ How to Edit or Hide Host Sets

This procedure describes how to edit host sets by using the browser interface. You can also edit host sets by using the following command.

```
hdb.hs.mod – Edits a host set.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To edit a host set, you must belong to a user group that has *Create*, *Edit*, *Delete* permission on hosts.

- Steps**
- 1. Go to the Details page for the host set that you need to edit.**
See “[How to View Host Sets](#)” on page 46.
The Details page for the host set is displayed. This page will list the host set’s member criteria and the hosts that currently match this criteria and are included in the host set.
 - 2. Use the controls on this page to define the host set.**
Each time you add a host, host set, or host search to the host set, the Details page updates the list of hosts in the host set.
 - 3. (Optional) To hide the host set, select Hidden.**

4. Click **Save**.

▼ How to Delete Host Sets



Caution – Deleting a host set is permanent and cannot be undone. Deleting a host set might also be time consuming.

You can remove unused host sets from the provisioning system. If you want to keep the host set but not view it, see [“How to Edit or Hide Host Sets”](#) on page 47.

You can also delete host sets by using the following command.

```
hdb.hs.del – Deletes a host set.
```

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin

You cannot delete the host set if a user group has `execute` permission on the host set

You also cannot delete a host set that is referenced by one of the following objects.

- A comparison
- A notification rule
- A component
- A plug-in

You can delete a host set that was created by a plug-in only by deleting the plug-in.

To delete a host set, you must belong to a user group that has `Create`, `Edit`, `Delete` permission on hosts.

Steps

1. **Go to the Details page of the host set you plan to delete.**

See [“How to View Host Sets”](#) on page 46.

2. **At the bottom of the Details page, click Delete.**

You are asked to verify that you want to delete the host search.

3. **Confirm that you want to delete the host set by clicking Continue to Delete.**

After the host set is deleted, the Host Sets page is updated.

About Host Searches

A host search is a query that is run on hosts in the provisioning system. Host searches yield a list of hosts that match criteria that you specify in your host-search query. You can use host searches to create lists of hosts that have the same host type or are configured with the same subnet masks.

Host Searches are objects with names and attributes. Therefore, when you create a host search, you assign it a name, a description, search parameters, and then save it. Because host searches are objects and saved, you can invoke them whenever you need to determine which hosts match the criteria described in the host search.

Invoking host searches provides you with current information about host configurations.

Hosts searches can be defined to search for any attribute on a host including host name, description, host type, and any user defined attribute associated with the host type. Additionally, hosts can be searched for by their operating systems and version.

By defining a host set (a group of hosts) to include the list of hosts that are returned by a specific host search, you can dynamically generate a list of hosts for use as targets based on the host configurations at runtime.

Note – For more information on host sets, see [“About Host Sets” on page 45](#).

Creating and Editing Host Searches

You might want to create or edit a host search in the following situations.

- You want to perform operations on hosts with a particular attribute.
- You need to query the provisioning system about hosts with particular configurations.

▼ How to Create or Edit Host Searches

You can generate a dynamic list of hosts according to specified criteria. You might want to edit a host search when your criteria changes.

This procedure describes how to create or edit host searches by using the browser interface. You can also create or edit host searches by using the following commands.

- `hdb.hr.add` – Creates a host search.
- `hdb.hr.mod` – Edits a host search.

For a detailed description of these commands, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To create or edit a host search, you must belong to a user group that has `Create`, `Edit`, `Delete` permission on hosts.

Steps 1. **Go to the Host Searches page.**

For more information about navigating to the Host Searches page, see “[How to View Host Searches](#)” on page 51.

2. **Create a host search or display the Details page for an existing host search.**

- **If you are creating a host search, in the top row of the Host Searches table, type a name and a brief description for the new host search and click Create.**
- **If you are editing an existing host search, find the row listing the desired host search and click Details.**

A Details page for the host search is displayed. This page displays the search criteria for the host search.

3. **Add an attribute value as a search criterion.**

- a. **Use the menu in the attribute column to select an attribute to be added to the search.**

This menu will list any attributes you have defined as part of host types, as well as standard attributes.

- b. **Use the menu in the Comparison column to select either Equals or Contains.**

- c. **In the Value column of the table, type the value of the attribute that you plan to search against.**

The search value can include the wildcard patterns `.` and `*`.

- d. **In the Action column of the table, click Add to add this attribute-comparison-value criterion to the search.**

You can add as many of these criteria as you like.

4. **Indicate whether the search should look only for physical hosts, only for virtual hosts, or for both.**

5. **Indicate whether the search should be limited to hosts running a specific provisioning system application: the Master Server, a Local Distributor, or a Remote Agent.**
6. **(Optional) To view the list of hosts that match the criteria you have defined so far, click Refresh Listing.**
7. **When you have defined your host search, click Save.**
The Host Searches page lists the new or modified host search.

▼ How to View Host Searches

This procedure describes how to view host searches by using the browser interface. You can also view host searches by using the following commands.

- `hdb.hr.la` – Displays all host searches.
- `hdb.hr.le` – Lists all hosts returned by a host search.
- `hdb.hr.lo` – Displays detailed information about a particular host search.

For a detailed description of these commands, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

- Steps**
1. **From the navigation menu, choose Host Setup.**
The Host Setup page is displayed.
 2. **In the Host Setup page, click Host Searches.**
The browser interface displays the Host Searches page. By default, this page lists all the unhidden host searches.
 3. **(Optional) To include the hidden host searches, select Show Hidden Host Searches in the upper right corner of the page.**
 4. **(Optional) To view the details of a specific host search, click Details in the row listing the host set that you want to view.**
The Details page is displayed for the host search that you selected.

▼ How to Delete Host Searches



Caution – Deleting a host search is permanent and cannot be undone. Deleting a host search might also be time consuming.

This procedure describes how to delete host searches by using the browser interface. You can also delete host searches by using the following command.

`hdb.hr.del` – Deletes a host search.

For a detailed description of this command, see Chapter 7, “hdb: CLI Commands for Managing Hosts,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin

If a host search was created by a plug-in, you must delete the plug-in to delete the host search.

To delete a host search, you must belong to a user group that has `Create`, `Edit`, `Delete` permission on hosts.

Steps

- 1. Go to the Details page of the host search that you plan to delete.**
For more information about navigating to the Host Searches page, see [“How to View Host Searches” on page 51](#).
- 2. At the bottom of the Details page, click Delete.**
You are asked to confirm the deletion of the host search.
- 3. Confirm that you want to delete the host search by clicking Continue to Delete.**
After the host search is deleted, the browser interface displays an updated Host Searches page.

Controlling Access Using Permissions

This chapter describes user permissions.

- [“Configuring Permissions” on page 53](#)
- [“System-Wide Permissions” on page 53](#)
- [“Folder-Specific Permissions” on page 54](#)
- [“Non-Configurable Permissions” on page 58](#)

Configuring Permissions

All provisioning system permissions are assigned to user groups. If you want a user to have a particular set of permissions, assign the user to a user group or a set of user groups that can provide the permissions that you want the user to have. There are two types of permissions in the provisioning system: configurable and non-configurable. Configurable permissions can be categorized by two groups: system-wide or folder-specific.

System-Wide Permissions

You can set system-wide permissions on the user group’s Details page or by using the `udb.g.mod` command.

For ideas on how you might set up user permissions based on user groups, see [“Planning User Groups and User Accounts” on page 59](#). For instructions on how to implement permissions with your user groups, see [“How to Create User Groups” on page 62](#).

The following list includes all system-wide, user group permissions.

- **Host administration.** Listed in the browser interface as `Hosts`.
You must have host administration permissions if you want to create, edit, or delete hosts, host sets, host searches, or host types.
- **User and user group administration.** Listed as `admin: users & groups`.
You must have permissions on users and user groups to create, edit, or delete users or user groups.
- **Notification rules.**
You must have notification rules permissions to set up email notification for system events.
- **Comparisons.**
You must have comparison permissions to create, edit, delete, or run comparisons. When comparison permissions are granted, the host set to run comparisons on must also be set.

System-wide Permission Characteristics

- All users are given read permissions on all provisioning system objects by default.
- Users are assigned privileges based on the groups they belong to. You can add permissions to a group if you have user and group administration permissions.
- User groups that have write permissions for a particular object type have write permissions on all objects of that type.
For example, if a user is assigned to a group that has write permissions for host administration, that user can make changes to all hosts, host sets, host searches, and host types.
- Users in the `admin` user group have read, write, and execute permissions on all objects by default.
- Permissions for the `admin` group cannot be removed or modified.

Folder-Specific Permissions

You can set folder-specific permissions on the folder's Details page or by using the `fdb.fmp` command. For more information about folders, see [“Introduction to Folders” on page 81](#).

For instructions on how to implement folder-specific permissions, see [“How to Modify the Folder Permissions of a User Group” on page 87](#).

The following list includes all folder-specific permissions.

- [“Create, Edit, and Delete Permission” on page 55](#)
- [“Run Component Procedures Permission” on page 56](#)
This permission is granted implicitly when a user is granted Create, Edit, Delete permission.
- [“Check In Current and Configure” on page 57](#)
Assigns permissions to check in a newer version of a component and to manipulate component variable settings.
- [“Allow on Host Set” on page 57](#)
Assigns the permissions to run plans on a defined host set.

Folder-specific Permission Characteristics

- Objects impacted by folder-specific permissions include folders, plans, and components.
- Folders have owners called owner user groups.
Owner user groups can perform strictly administrative tasks on a folder. For example, owner user groups can change permissions on a folder, including the owner user group permission. Owner user groups do not implicitly have other folder permissions.
- Folder permissions can be set by a member of the owner user group of the folder.
- All users have read permissions on all folders, plans, and components by default.
- Every folder has a set of permissions that are assigned to the user groups. Actions that can be taken on a folder or object contained within the folder are determined by the permissions set on that folder.

Create, Edit, and Delete Permission

When you assign Create, Edit, Delete permission to a user group, members of the user group are able to perform several tasks in the folder. Assigning Create, Edit, Delete permission to a user group automatically assigns Run Component Procedures and Check In Current and Configure permission to the user group.

- Create new subfolders

Note – Subfolders inherit the permissions of the parent folder when the subfolders are created. These permissions can only be modified by the owner user group.

- Create new plans, components, and folders within the current folder
- Check in new versions of plans and components

- Rename subfolders
- Change a folder's description
- Move plans, components, and folders

Note – To move an object, you must have write permission on the parent folder and the destination folder.

- Delete plans and components
- Delete folders

Note – To delete a folder, you must have write permission on the parent folder and all subfolders.

- Perform all tasks allowed by Run Component Procedures and Check In Current permissions.

You need to be particularly careful about assigning `Create`, `Edit`, `Delete` permission to a user group when that user group also has `Allow on Host Set` permission. Users in this type of user group will have free reign over a particular host set since they will be able to write plans and execute them as they choose.

Run Component Procedures Permission

When you assign `Run Component Procedures` permissions to a user group, members of the user group are able to execute plans that are generated by directly running component procedures. These plans are created from component procedures and are stored in the `/system/autogen` folder. `Run Component Procedures` permissions is assigned from the folder that contains the component that generated the plans.

`Run Component Procedures` permission is granted automatically when a user is granted `Create`, `Edit`, `Delete` permission. However, you might want to exclusively grant the `Run Component Procedures` permissions when you want a user group to have permission to run only component procedures or plans that have already been written. This permission is particularly useful if you have one user writing plans and another user running them, and you don't want the users to perform tasks outside of their functional expertise.

Note – Run Component Procedures permission is an extension of the Allow on Host Set permission. You can only run component procedures on the host set defined by the Allow on Host Set permission.

Note – If component procedures contain variable sets that require user input, the user who runs the component procedure will need to belong to a user group that has Check In Current and Configure permission.

There are a few scenarios where you need to be particularly careful about assigning Run Component Procedures permission to user groups.

- Large user groups in folders that contain sensitive materials (for example, system services that perform potentially destructive operations are usually granted execution on all host sets)
- User groups that have permissions to run plans on large host sets

Check In Current and Configure

When you assign Check In Current and Configure permission to a user group, members of the user group are able to check in new versions of components and create, edit or delete component variable settings.

Check In Current and Configure permission is granted automatically when a user is granted Create, Edit, Delete permission. However, you might want to exclusively grant the Check In Current and Configure permission when you want someone to be able to perform deployments but not be able to manipulate the components in a folder.

Operations that are *not* allowed with Check In Current and Configure permission include the following.

- Show and hide objects
- Change an object's category
- Create, edit, delete, or move objects

Allow on Host Set

When you assign Allow on Host Set permission to a user group, members of the user group are able to execute custom plans and generated plans on a single host set. This permission also allows users to create a dependency on a component that is installed on a particular host.

Note – For a user to be able to run component procedures, the user must belong to a user group with `Allow on Host Set` and `Run Component Procedures` permissions assigned. The host set defined by the `Allow on Host Set` permission establishes which hosts the component procedure can be run on.

Equivalent Permissions in the CLI

The CLI and browser interface have a different naming scheme for user permissions. The following table lists the browser interface permission names and their CLI counterparts.

TABLE 3-1 HTML and CLI permission names

HTML User Interface Names	Command-Line Interface Names
Create, Edit, Delete	write
Run Component Procedures	autorun
Check In Current and Configure	checkin-current
Allow on Host Set	execute

Non-Configurable Permissions

The provisioning system also has non-configurable permissions assigned to users and user groups.

- Members of the `admin` user group have all permissions on all objects in the provisioning system and can modify any object regardless of ownership.
See “[admin User Group](#)” on page 60.
- Members of the `admin` user group are also the only users who can import plug-ins and create component types and system services.
- Owners of processes have permission to stop the processes that they started. For example, if you start a comparison, you can stop the comparison.
Users cannot stop a process that they did not initiate unless they belong to the `admin` user group.
- All registered users have read permissions on all objects.

Managing Users

Through the N1 Grid Service Provisioning System, you can create user accounts and define user groups. Because an individual user's permissions are determined by the permissions granted to his or her group, you must consider how you plan to set up accounts and groups.

This chapter covers the following topics.

- [“Planning User Groups and User Accounts” on page 59](#)
- [“About User Accounts” on page 66](#)

Planning User Groups and User Accounts

All permissions are tied to user groups. By adding or removing a user from a user group, you directly impact the tasks that the user can perform. For more information about permissions, see [Chapter 3](#).

When setting up user groups and user accounts, you must also consider how you plan to set up your folders. Folder-based permissions establish permissions on objects that are contained within a folder, which are components, plans, and subfolders. System-wide permissions establish permissions on objects not contained in folders, for example, hosts, comparisons, and users. You can set system-wide permissions in each user group's Details page. You can also set permissions at the folder level.

You should first set up user groups based on the roles you expect users to take. Then you create new user accounts and decide to which user groups they should be added.

When setting up user accounts, consider the following items.

- The type of functions that each user will be performing

- The security level that is required between users

Managing User Groups

A user group is a user-definable object that is used to categorize users and define permissions. By carefully planning out the names for user groups and which permissions to grant each group, you can easily manage individual user permissions by making them part of one or more groups.

User Group Characteristics

- User groups can include one or more individual users.
- User groups can also include one or more user groups.
 - A user group can be a super set of all the user groups included in its member list.
- Nested user groups inherit the permissions of the containing user groups.
 - Since permissions can only be added, the top-level user group should be the least permissive. Nested user groups represent more permissive user groups.
- System-wide permissions are set in a user group's Details page.
 - See ["How to Edit User Groups"](#) on page 64.
- Folder-specific permissions are set in the folder's Details page.
 - See ["How to Modify the Folder Permissions of a User Group"](#) on page 87.

For more information about the different types of user group permissions, see [Chapter 3](#).

Default User Groups

The provisioning system provides three default user groups after installation: `admin`, `registered`, and `universal`. Default user groups cannot be deleted and cannot have their names changed.

admin User Group

The provisioning system provides the `admin` user group after installation to allow initial system configuration. This user group is designed for administrators of the provisioning system.

Members of the `admin` user group have all permissions on all objects in the provisioning system and can modify an object whether or not they own it.

Members of the `admin` user group perform many functions.

- setting up hosts
- adding new user accounts and user groups
- setting permissions of user groups
- creating folders for user groups
- importing plug-ins

The `admin` user group comes with one default user, the `admin` user. However, if you have more than one administrator of the provisioning system, you can add other users to the `admin` user group.



Caution – Since the `admin` user group has complete control over all aspects of the provisioning system, be careful when assigning users to this group.

For more information about the `admin` user, see [“Default User Account” on page 66](#).

For more information about creating new users, see [“How to Create User Accounts” on page 66](#).

registered User Group

The `registered` user group consists of all users that have been created in the provisioning system.

registered User Group Characteristics

- Every user is a member of the `registered` user group, and members cannot be removed.
- Read permissions for all objects are granted to the user group.
Read permissions cannot be revoked.
- New permissions can be assigned to the user group.
The affect of granting a permission to the `registered` group is to allow all users in the system to perform the associated action.
- Although all users are assigned to the `registered` user group, this user group will not display in a user’s user group list.

universal User Group

The `universal` user group includes all users. By default, no permissions are granted to this group. However, new permissions can be granted, which has the affect of allowing anyone to perform the associated operation. The `registered` group may not be removed as a child of this group.

▼ How to Create User Groups

This procedure describes how to create user groups by using the browser interface. You can also create user groups by using the following command.

```
udb.g.add – Creates a new user group.
```

For a detailed description of this command, see “udb.g: Managing User Groups” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin

Before you create user groups, you should determine how you want to organize your users. For information about how to set up user groups, see “[Planning User Groups and User Accounts](#)” on page 59.

To create a user group, you must belong to a user group that has write permissions on users and groups.

Steps 1. Go to the User Groups page.

See “[How to View User Groups](#)” on page 63.

2. In the top row of the table that lists user groups, type a name and a description for the new user group and click Create.

The Details page for the new user group is displayed.

3. Add a user or user group to the group.

Newly created user groups do not contain any members.

- **To add a user, select the user account from the User menu in the Members of Group area and click Add.**

Users added to the user group inherit the permissions given to this user group and all user groups that contain this user group.

- **To add a user group, select the user group from the User Group menu in the Members of Group area and click Add.**

The Details page updates to show the added user or members of the added user group in the Current Group Members field.

Note – The Current Group Members field lists a user only once, even if that user belongs to two or more groups that you have added to the group.

4. In the Permissions of Group Users area of the page, set system-wide permissions for the new user group.

Permissions set in this user group are inherited by members of the user group. These members include individual users as well as other, nested user groups.

Note – If you give the user group comparison permissions, select the host set on which the users in the group can run comparisons.

For more information, see “System-Wide Permissions” on page 53.

5. When you have finished configuring the group, click Save.

The User Groups page lists the new user group.

▼ How to View User Groups

You can view the users and the permissions of a particular user group.

This procedure describes how to view user groups by using the browser interface. You can also view user groups by using the following commands.

- `udb.g.la` – Displays all user groups.
- `udb.g.lo` – Displays detailed information about a particular user group.
- `udb.g.lp` – Displays system-wide permissions granted to a user group.
- `udb.g.lu` – Displays members of a user group.

For a detailed description of these commands, see “udb.g: Managing User Groups” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Steps 1. **From the navigation menu, choose User Setup.**

The User Setup page is displayed.

2. **In the User Setup page, click User Groups.**

This displays the User Groups page, which lists the user groups already defined.

3. **(Optional) To view a list of users within a group or the permissions held by a user group, find the row that lists the group that you want to view and click Details.**

Editing User Groups

Editing user groups allows you perform the following tasks after you have created a user group.

- Add a user to the user group
- Add another user group to the user group
- Remove a user or user group from the user group
- Change system-wide permissions of the user group

▼ How to Edit User Groups

This procedure describes how to edit user groups by using the browser interface. You can also edit user groups by using the following command.

`udb.g.mod` – Edits a user group.

For a detailed description of this command, see “`udb.g: Managing User Groups`” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To edit a user group, you must belong to a user group that has write permissions on users and groups.

Steps 1. **Go to the Details page of the user group that you plan to edit.**

See “[How to View User Groups](#)” on page 63.

2. **(Optional) Add a user or user group to the group.**

Newly created user groups do not contain any members.

- **To add a user, select the user account from the User menu in the Members of Group area and click Add.**

Users added to the user group inherit the permissions given to this user group and all user groups that contain this user group.

- **To add a user group, select the user group from the User Group menu in the Members of Group area and click Add.**

The Details page updates to show the added user or members of the added user group in the Current Group Members field.

Note – The Current Group Members field lists a user only once, even if that user belongs to two or more groups that you have added to the group.

3. **(Optional) In the Permissions of Group Users area of the page, set system-wide permissions for the user group.**

Permissions set in this user group are inherited by members of the user group. These members include individual users as well as other, nested user groups.

Note – If you give the user group comparison permissions, select the host set on which the users in the group can run comparisons.

For more information, see “[System-Wide Permissions](#)” on page 53.

4. (Optional) In the Permissions of Group Users area of the page, select the host set on which users can run comparisons.
5. When you complete your modifications, click Save.

Deleting User Groups

When a user group is deleted, the user group is removed from all user groups to which it belonged. Users and user groups that belonged to the deleted group continue to exist, but they no longer belong to the user group, and therefore, no longer have the permissions granted by the deleted user group.

If a folder granted the user group certain permissions, those permissions are also deleted when the user group is deleted.

▼ How to Delete User Groups

The browser interface provides two options for deleting user groups. You can delete multiple user groups at once from the User Groups page or one at a time from the user group's Details page. This procedure provides instructions on how to delete several user groups at a time.

You can also delete user groups by using the following command.

```
udb.g.del – Deletes a user group.
```

For a detailed description of this command, see “udb.g: Managing User Groups” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin

To delete a user group, the following requirements must be met.

- You must belong to a user group that has write permissions on users and groups.
- The user group must not be the default admin, registered, or universal user groups.
- The user group must not own any folders.

If the user group owns a folder, change the folder's owner user group. Then delete the user group.

Steps

1. Go to the User Groups page.
See “How to View User Groups” on page 63.
2. Select the user groups that you plan to delete.
3. At the bottom of the User Groups table, click Delete.

A verification page lists the user groups that you selected.

4. Click Continue to Delete.

After the user groups have been deleted, the User Groups page updates and the user groups that you deleted no longer appear.

About User Accounts

The N1 Grid Service Provisioning System software restricts access and provides audit trails through the use of user accounts and the plan run history. All users are required to log in to use the application. The Master Server includes a single default account, `admin`. All other accounts are user definable.

Default User Account

The provisioning system's default user account provides initial access to the product and is intended for system administrators. The user name for this account is `admin` and the default password is defined during product installation.

After you have logged in as `admin`, you can set up other user accounts and user groups.

The `admin` user is a member of the `admin` user group. The `admin` user cannot be removed from the `admin` user group. For more information on the role of the `admin` user, see [“admin User Group” on page 60](#).

Managing User Accounts

This section describes how to view and create user accounts. You can also modify user accounts and passwords after an account is created.

▼ How to Create User Accounts

You will need to create new user accounts when users need to access the system.

Note – Any user added to the provisioning system has read permissions on all objects within the system. If you are storing sensitive information in the system, ensure that you only add users that should have access to that sensitive information.

You can also create user accounts by using the following command.

`udb . u . add` – Adds a user account.

For a detailed description of these commands, see “udb.g: Managing User Groups” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin

To create a user account, you must belong to a user group that has write permissions on users and groups.

If you plan to add a user who will be externally authenticated, ensure the following criteria are met.

- The user exists in the external system before the user attempts to login to the provisioning system.
- The user name begins with a letter, does not contain spaces, and does not exceed 32 characters.

Steps

1. **Go to the Users page.**

See “How to View User Accounts” on page 68.

2. **In the top row of the table that lists user accounts, type a name for the new user account and click Create.**

The Details page for the new user account is displayed.

3. **Select an authentication method from the menu.**

If the authentication method that you select has an asterisk, *, you will need to enter a password for the user account in the New Password and Confirm New Password fields.

To add new login configurations to the provisioning system, see [Appendix A](#).

4. **Use the controls in the Member of User Groups area to add this user account to one or more user groups.**

Note – In the provisioning system, permissions are based on user groups rather than on individual user accounts. By adding this user account to a group, you determine the privileges it is assigned. For more information, see “[Folder-Specific Permissions](#)” on page 54.

For more information about adding a user to a user group, see [“How to Edit User Groups” on page 64](#).

5. **(Optional) To hide this user account, select Hidden.**
6. **Click Save.**

Viewing User Accounts

You might want to view user accounts in the following situations.

- You need to audit user accounts to ensure that users who are no longer authorized to access the provisioning system are removed.
- You need to ensure that users belong to the correct groups and have the correct set of permissions.
- You need to ensure that you have permissions to perform your job without any obstacles.

▼ How to View User Accounts

This procedure describes how to view user accounts by using the browser interface. You can also view user accounts by using the following commands.

- `udb.u.la` – Displays all user accounts.
- `udb.u.lo` – Displays detailed information about a particular user account.
- `udb.u.lp` – Displays system-wide permissions granted to a particular user account.

For a detailed description of these commands, see *“udb.g: Managing User Groups”* in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

- Steps**
1. **From the navigation menu, choose User Setup.**
The User Setup page is displayed.
 2. **In the User Setup page, click Users.**
The Users page lists user accounts that are active and not hidden.
 3. **(Optional) To view hidden or deactivated users, select Show Hidden Users.**
The browser interface displays all users. Hidden or deactivated user accounts appear in different colors.
 4. **(Optional) To view the details of a particular user, find the row describing the user and click Details.**
The Details page for the selected user is displayed.

▼ How to Deactivate or Hide User Accounts

When you deactivate a user account, the user is no longer allowed to access the provisioning system. Since the provisioning system tracks system activity by user account, you are not able to remove user accounts from the provisioning system.

You can hide a user account as a way to manage the number of users that display in the Users page.

This procedure describes how to deactivate or hide user accounts by using the browser interface. You can also deactivate or hide user accounts by using the following command.

```
udb . u . mod – Edits a user account.
```

For a detailed description of this command, see “udb.g: Managing User Groups” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To deactivate a user account, you must belong to a user group that has write permissions on users and groups.

- Steps**
1. **Go to the Details page of the user account that you plan to edit.**
See “[How to View User Accounts](#)” on page 68.
 2. **In the table that lists users, find the row describing the user you plan to deactivate, and click Details.**
The user’s Details page is displayed.
 3. **To hide or deactivate this user account, select Hidden or Deactivate.**
If you select Deactivate, the user account will automatically be hidden.
 4. **Click Save.**

▼ How to Change a User’s Group Membership

This procedure describes how to change a user’s group membership by using the browser interface. You can also change group membership by using the following command.

```
udb . u . mod – Edits a user account.
```

For a detailed description of this command, see “udb.g: Managing User Groups” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To edit a user account, you must belong to a user group that has write permissions on users and groups.

- Steps**
1. **Go to the Details page of the user account that you plan to edit.**
See “[How to View User Accounts](#)” on page 68.
 2. **In the table that lists users, find the row describing the user account you plan to modify, and click Details.**
The user’s Details page is displayed.
 3. **To change membership in any group, use the controls in the Member of User Groups area to add this user account to or remove it from one or more user groups.**
 4. **Click Save.**

▼ How to Change Your Password From the Log In Page

If you need to need to change your password, you can change your password from the Log In page as long as you are using internal authentication for your password. If you are using external authentication, change your password through that application.

Note – When you use this procedure to change your password, all of your encrypted session variables are re-encrypted using your new password.

- Steps**
1. **In the navigation menu on the Log In page, click Change Password.**
If you are already logged in, click Log Out in the top right corner to navigate to the Log In page.
 2. **Type your user name in the User Name field.**
 3. **Type your current password in the Current Password field.**
 4. **Type the new password you would like to use in the new password field.**
 5. **Confirm that you have typed the new password correctly by typing it in the Confirm New Password field.**
 6. **Click the Change Password button.**

▼ How to Change a User's Password from the User Account's Details Page

You can change anyone's password under the following conditions.

- The password is internally validated.
- You have permission to edit and create users.

If you don't have user and group permissions, see ["How to Change Your Password From the Log In Page"](#) on page 70.

This procedure describes how to change a user's password by using the browser interface. You can also change passwords by using the following command.

```
udb . u . cp – Changes the password of a specific user.
```

For a detailed description of these commands, see *"udb.g: Managing User Groups"* in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Note – When you change the password on a user account and the user account has session variables encrypted by the user account's old password, the user is prompted to clear or restore the session variables the next time the user logs in.

Before You Begin To change the password of someone else's user account, you must belong to a user group that has write permissions on users and groups.

- Steps**
1. **Go to the Details page of the user account.**
See ["How to View User Accounts"](#) on page 68.
 2. **Type the new password in the New Password field.**
 3. **Confirm the new password by typing it in the Confirm New Password field.**
 4. **Click Save.**

▼ How to Change a User's Authentication Method

This procedure describes how to change a user's authentication method by using the browser interface. You can also change authentication methods by using the following command.

```
udb . u . mod – Changes the user's authentication method.
```

Note – When you change a user’s authentication method and the user account has session variables encrypted by the user account’s old password, the user is prompted to clear or restore the session variables the next time the user logs in.

Before You Begin

Before you can change the authentication method of a user account, the new authentication method must be configured with the provisioning system. See [Appendix A](#).

To change a user’s authentication method, you must belong to a user group that has write permissions on users and groups.

Note – You cannot change the authentication method of the admin user. The admin user must always be set to internal authentication.

Steps

1. **Go to the Details page of the user account.**
See [“How to View User Accounts”](#) on page 68.
2. **From the Authentication Method menu, select the new authentication method.**
3. **If you change the user from an external to an internal authentication method, you must supply the user a new password.**
4. **Click Save.**

Plug-In Administration

This chapter describes basic administrative tasks involving plug-ins.

“About Plug-Ins” on page 73

About Plug-Ins

A plug-in is a group of objects that have been created for a particular application and packaged to work with the provisioning system. Plug-ins allow you to work with different applications in the provisioning system without modeling the applications first. Plug-ins can provide a number of different objects that are ready for use. This saves you the time of creating the components and plans from scratch. Rather than require that each customer recreate all the objects that an application requires, plug-ins provide common objects that are useful in any environment.

A plug-in can contain the following objects.

- Components
- Component types
- Plans
- Host sets
- Host searches
- Host types
- Folders

Plug-In Characteristics

- A plug-in owns all objects that it creates. No object that is created by a plug-in can be deleted directly. However, deleting the plug-in itself will result in the deletion of all the objects it owns.
- A plug-in can create folders to contain its components, plans, and subfolders. If the folder is supposed to be nested within other folders, the plug-in also creates the parent folders of the plug-in.
- The `admin` user group is the owner user group of all plug-in folders.
- Only members of the `admin` user group can perform administrative tasks on plug-ins or plug-in-owned objects.
- All plug-in objects, except folders, components, and plans, will use the *plug-in-name#object-name* naming convention.
For example, `system#file` represents an object from the `system` plug-in named `file`.

Importing Plug-Ins

You can install a plug-in by importing the plug-in's JAR file. When you import a plug-in, you lock down the provisioning system until the import is complete.

The following tasks also require a lock-down of the provisioning system. This means that you cannot begin importing a plug-in if one of these operations is being executed. You also cannot start one of these operations when you are importing a plug-in.

- Saving a component type
- Saving a system service
- Moving or renaming a component
- Moving or renaming a plan
- Moving or renaming a folder
- Deleting any object

▼ How to Import a Plug-In

This procedure describes how to import a plug-in by using the browser interface. You can also import a plug-in by using the following commands.

`plg.p.add` – Imports a new or existing version of a plug-in.

Before You Begin To import a plug-in, you must be a member of the `admin` user group.

- Steps**
1. In the navigation menu, expand the **Administrative** category and click **Plug-ins**.
 2. Click **Import**.

The Import page is displayed.

3. Select the plug-in JAR that you plan to import.

- **Type the name of the plug-in JAR file.**
- **Click Browse to select the JAR file from a file browser.**

4. Click Continue to Import.

A status bar displays the progress of importing the plug-in.

After the plug-in finishes importing, the plug-in Details page displays and lists the contents of the JAR file.

5. Click Done.

See Also After installing the plug-in, check to see if there is a plug-in ReadMe file. This file might or might not exist. If the file exists, it should contain information about administering the plug-in, including information on how to set up the plug-in's folder permissions.

▼ How to View a Plug-In

This procedure describes how to view imported plug-ins by using the browser interface. You can also view plug-ins by using the following commands.

- `plg.p.la` – Displays all imported plug-ins.
- `plg.p.lo` – Displays detailed information about a particular plug-in.

Steps **1. In the navigation menu, expand the Administrative category and click Plug-ins.**

2. Select the plug-in you want to view and click Details.

Objects contained by the plug-in and general plug-in information are displayed.

3. (Optional) To view objects contained in the plug-in, select the object you want to view and click Details.

The object's Details page displays.

4. (Optional) If the plug-in contains a ReadMe file, click Download ReadMe.txt file to view the file.

Editing Plug-Ins

You can change the order that plug-ins are displayed in the navigation menu. If you do not configure the order of the plug-ins, the plug-ins display in alphabetical order.

▼ How to Edit a Plug-In

This procedure describes how you can change the menu order of your plug-ins by using the browser interface. You can also perform this task using the following command.

```
plg.p.mod – Modifies the menu order of a plug-in.
```

Before You Begin To change the menu order of your plug-in, you must be a member of the admin user group.

- Steps**
1. **In the navigation menu, expand the Administrative category and click Plug-ins.**
 2. **Select the plug-in to which you'd like to assign a menu order.**
The plug-in's Details page displays.
 3. **Click Edit.**
 4. **In the Menu Order field, type where you want this plug-in displayed in the navigation menu.**

The criteria for menu order is as follows.

- No specified menu order takes precedence over a specified order.
This means that if you want to specify that one plug-in should appear before all other plug-ins, you will need to provide a menu order for all imported plug-ins.
- Low numbers take precedence over high numbers.
A menu order of "1" takes precedence over a menu order of "2".
- Numbers take precedence over letters.
- Capital letters take precedence over lowercase letters.
- If the first numbers or letters are the identical, the first nonidentical character will determine the menu order.

Upgrading Plug-Ins

You can install a plug-in by importing the plug-in's JAR file. You can update a plug-in by importing a JAR file that contains just the contents needed to update the plug-in. This type of JAR file is called a patch and should be identified as a patch on the download site.

Patch Characteristics

- Patches can be applied in series to upgrade multiple versions.
For example, to upgrade from the 1.0 version to the 1.2 version of a plug-in, first the 1.0 -> 1.1 patch is applied, followed by the 1.1 -> 1.2 patch.
- Patches can be created to upgrade from a specific plug-in version to a higher specific version.
It is not possible to create a patch to upgrade from any arbitrary version to a higher version.
- Patches can only add content to a plug-in.
Objects created by any version of an plug-in cannot be deleted.
- Patches cannot be uninstalled.
To remove patch content, you can delete the plug-in and reinstall the older plug-in version. The provisioning system does not store older versions of a plug-in.

When you upgrade a plug-in, the patch adds new and updated objects to the provisioning system. If a plug-in patch attempts to add an object that already exists by the same name, two outcomes are possible.

- The object is updated with a new version number.
If the object is a versioned object, the new object's minor version number increments unless the plug-in definition requires the object to increment by major versions.
- The object overwrites the existing object.
If the object is a non-versioned object, the new object replaces the previous definition of the object.

In both the versioned and unversioned cases, the existing object must have been created by a prior version of the plug-in that is attempting to create the new version of the object.

▼ How to Upgrade a Plug-In

This procedure describes how to upgrade a plug-in by using the browser interface. You can also upgrade a plug-in by using the following commands.

```
plg.p.add – Imports a new or existing version of a plug-in.
```

Before You Begin To upgrade a plug-in, you must be a member of the admin user group.

- Steps**
1. **In the navigation menu, expand the Administrative category and click Plug-ins.**
 2. **Click Import.**
The Import page is displayed.

3. **Select the plug-in JAR that you plan to import.**

- **Type the name of the plug-in JAR file.**
- **Click Browse to select the JAR file from a file browser.**

4. **Click Continue to Import.**

A status bar displays the progress of importing the plug-in.

After the plug-in finishes importing, the plug-in Details page displays and lists the contents of the JAR file.

5. **Click Done.**

See Also After installing the plug-in, check to see if there is a plug-in ReadMe file. This file might or might not exist. If the file exists, it should contain information about administering the plug-in, including information on how to set up the plug-in's folder permissions.

Deleting Plug-Ins

When you delete a plug-in, the provisioning system also deletes all object created by all versions of that plug-in. If any object is being used, the plug-in cannot be deleted.

By deleting a plug-in, you lock down the provisioning system until the deletion is complete.

The following tasks also require a lock-down of the provisioning system. This means that you cannot begin deleting a plug-in if one of these operations is being executed. You also cannot start one of these operations when you are deleting a plug-in.

- Saving a component type
- Saving a system service
- Importing a plug-in
- Moving or renaming a component
- Moving or renaming a plan
- Moving or renaming a folder
- Deleting any object

▼ How to Delete a Plug-In

This procedure describes how to delete a plug-in by using the browser interface. You can also delete plug-ins by using the following command.

```
plg.p.del – Deletes a plug-in.
```

Before You Begin You cannot delete a plug-in if any object owned by the plug-in is being used by the provisioning system.

To delete a plug-in, you must be a member of the admin user group.

- Steps**
- 1. Go to the Details page of the plug-in you plan to delete.**
 - 2. At the bottom of the Details page, click Delete.**
You are asked to verify that you want to delete the plug-in.
 - 3. Click Continue to Delete.**
After the plug-in is deleted, the Plug-in page is updated.

Setting Up Folders

You can organize objects within the N1 Grid Service Provisioning System software in several ways using folders.

- [“Introduction to Folders” on page 81](#)

Introduction to Folders

Folders are directory-like objects that enable you to organize the provisioning system in a number of ways. Folders allow you to apply permissions to objects and to group those objects as needed. By grouping objects and manipulating permissions, you can organize components and plans to reflect how you use them.

- **Organize by user group.**

Every user group can be assigned a work space where they have complete access and control over all objects that exist within their folder. For example, a particular department in your organization might want control over the plans and components that affect their group. The department can set up their folder so that only people within their user group can run the plans or modify the components within the folder. This setup prevents people outside the department from having access to write or execute plans and components.

- **Organize by project or application.**

By creating different folders for each project or application, you can control who is working on which projects. If you want to use objects in several different projects, you can create a shared folder for objects and assign permissions to the user groups who need to access the folder.

Folder Characteristics

- Folders give all users read access to all folders, components, and plans.
- Each folder has an owner and a set of write and execute permissions that limits who can perform tasks with objects within the folder.
- The permissions and owner user group of a folder do not extend to other folders.
- A new folder inherits the permissions of its parent folder, but after the initial folder creation the subfolder's permissions are handled directly.
- Each folder acts as its own namespace.
 - Different types of objects can share the same name within a folder.
 - Two objects of the same type cannot share the same name unless they reside in different folders.
- The folder path persists whether viewing components, plans, or folders. When you are in a particular folder, you can navigate between different views and remain in the same folder.
- Users cannot add components, plans, or subfolders to folders that were created by a plug-in.
- The root folder cannot be renamed, moved or deleted.

Creating Folders

Creating new folders allows you to further organize your components and plans.

When you create a new folder, the path of the folder is determined by the folder that is selected at the time the new folder is created. For information about how to change the folder's path after the folder has been created, see ["How to Move Folders"](#) on page 85.

The owner user group and permissions of the new folder default to the permissions of the parent folder. For information about how to change a folder's permissions, see ["How to Modify the Folder Permissions of a User Group"](#) on page 87.

Note – Folders are no longer automatically created when components or plans are created. If you have scripts that rely on the old methods of folder creation, update your scripts to explicitly create folders before creating plans and components that are intended to reside in these folders.

▼ How to Create Folders

This procedure describes how to create folders by using the browser interface. You can also create folders by using the following command.

fdb.f.add – Creates new folder.

Before You Begin

To create a folder, you must have write permissions on the parent folder.

Note – You do not need to be in the folder’s owner user group to create a subfolder. However, if you are not a member of the owner user group, then you cannot modify the new folder’s owner user group or default permissions.

Steps

1. View the folder to which you plan to add a subfolder.

For the procedure about how to navigate to the desired Folders page, see [“How to View Folders” on page 84](#).

2. In the top row of the table on the Folders page, type a folder name, select an owner user group, and provide a brief description for the new folder and click Create.

The Edit page for the new folder is displayed.

3. Ensure that the correct folder name, owner user group, and description are displayed.

4. (Optional) If you are a member of the owner user group, you can change the folder’s owner user group and set the permissions for each user group that has access to the new folder.

Note – After you configure a user group’s folder permissions, click Add. If you do not click Add, the user group’s permissions will be lost when you save the folder.

For more information on changing folder permissions, see [“How to Modify the Folder Permissions of a User Group” on page 87](#).

5. Click Save.

Troubleshooting If your attempt to create a new folder failed, verify the following conditions.

- The folder’s creation would not result in the folder path exceeding 512 characters (including forward slashes,/).
- The folder’s name contains only letters, numbers, dashes (-), underscores (_), periods (.), and spaces (' ').
- The new folder’s name does not conflict with an existing folder in the parent folder.
- The parent folder was not created by a plug-in.

Only a revision of the plug-in or a plug-in that depends on that plug-in can create new subfolders.

Viewing Folders

You can view folders in two ways: a hierarchical view or a flat view. In each view, you see the folder path, a list of folders that are contained in the folder, and the owners of the folders listed.

The default view is the hierarchical view, which shows the permissions you have for the current folder and the first level of subfolders contained in the folder. The flat view shows the current folder information and displays the entire folder hierarchy below the current folder.

Tip – To view all folders in the system, go to the top level folder (/) and select In Flat View.

▼ How to View Folders

This procedure describes how to view folders by using the browser interface. You can also view folders by using the following commands.

- `fdb.f.la` – Displays all folders.
- `fdb.f.lo` – Displays detailed information about a particular folder.

Steps 1. **From the navigation menu, choose Folders.**

The folder list page presents a list of folders that are contained in the folder, and the owners of the folders listed.

2. **On the Folders page, select the folder that you need to view and click Details.**

Tip – If the folder you need to view is not displayed in the Folder table, click Change Folder. Select the folder from the Folder tree and click Change to Selected Path.

The Details page displays your permissions, the user groups with access to the folder, and links to other tasks that you can perform on the folder.

3. **(Optional) To see the folder's entire hierarchy of subfolders, select Show Flat View.**

The flat view shows all folders nested under the selected folder in a single table.

4. **(Optional) To view the contents of folders, use the Show menu to toggle between components, plans, and subfolders.**

Moving Folders

Moving folders allows you to have the flexibility you need as your work space evolves. You might need to move folders if projects change or your team's work flow changes.

Moving a folder also moves all objects that the folder contains. In addition to moving all components, plans, and subfolders, references to the components and plans will automatically update to reflect the new hierarchy.

When a folder is moved, the folder retains its permissions and owner user group, even if they are different from the destination folder.

If you want to move multiple folders at a time, you can move them from the folder's list page. When you move multiple folders at a time, all folders are moved to the same destination folder. However, if one folder fails to move, none of the folders are moved.

▼ How to Move Folders

This procedure describes how to move by using the browser interface. You can also move folders by using the following command.

```
fdb.f.mv – Moves or renames one folder.
```

Before You Begin To move a folder, you must have write permissions on the parent folder of the folder that you plan to move and the destination folder. You do not need write permissions on other objects that will be affected by moving the folder.

Steps 1. **View the Details page of the folder that you plan to move.**

For more information about how to navigate to a folder's Details page, see [“How to View Folders” on page 84](#).

2. **Click Move.**

The Move Folder window appears.

3. **Select the destination folder to which you want to move the folder.**

You can select the new parent folder in one of two ways.

- Navigate through the folder tree and select the parent folder.
- Type the parent folder's name into the Move to Folder field.

4. **Click Move.**

After the move is complete, the folder's Details page shows the folder's new path.

Troubleshooting If your attempt to move a folder failed, verify the following items.

- The moved folder name does not conflict with an existing folder in the destination folder.
- Moving the folder into the new path would not result in a folder path exceeding 512 characters.
- The folder, all of its subfolders, and the destination folder were not created by a plug-in.

▼ How to Rename Folders

This procedure describes how to rename folders and change their descriptions by using the browser interface. You can also rename folders by using the following command.

```
fdisk . f . mv – Moves or renames the folder.
```

Before You Begin When you rename a folder, references to the old folder are automatically updated.

To rename a folder you must have write permissions on the parent folder. You do not need write permissions on other objects that will be affected by renaming the folder.

Steps 1. **View the Details page of the folder that you plan to rename.**

For more information about how to navigate to a folder's Details page, see ["How to View Folders"](#) on page 84.

2. **Click Rename.**

The Rename Folder window appears.

3. **Type the new folder name in the text field and click Rename.**

The folder's Details page shows the folder's new name.

Troubleshooting If your attempt to rename a folder failed, verify the following items.

- The new folder name does not contain special characters.
- Renaming the folder does not result in a path that is greater than 512 characters.
- The new folder name does not conflict with an existing folder in the parent folder.
- The folder and all of its subfolders were not created by a plug-in.

▼ How to Change the Description of Folders

This procedure describes how to change folder descriptions by using the browser interface. You can also change folder descriptions by using the following command.

`fdb.f.mod` – Modifies a folders description.

Before You Begin To change the description of a folder, you must have write permissions on the folder you plan to change.

- Steps**
1. **Go to the Details page of the folder that you plan to edit.**
For more information about how to navigate to a folder's Details page, see ["How to View Folders"](#) on page 84.
 2. **Click Edit.**
 3. **Type the new description in the Description field.**
 4. **Click Save.**

Granting Folder Permissions to User Groups

To change a folder's permissions, you can add or remove user groups from the folders or you can alter the type of permissions granted to the current user groups.

▼ How to Modify the Folder Permissions of a User Group

This procedure describes how to change a folder's permissions by using the browser interface. You can also add or revoke permissions from a particular group by using the following command.

`fdb.f.mp` – Modifies folder permissions and adds user groups to folders.

Before You Begin If you plan to add permissions to a folder while you are creating the folder, you must be a member of the parent folder's owner user group. However, after the folder is created, you must be a member of the folder's owner user group.

- Steps**
1. **View the Details page of the folder you plan to edit.**
For more information about how to navigate to a folder's Details page, see ["How to View Folders"](#) on page 84.
 2. **On the Details page, click Edit.**
 3. **Add or remove user groups from the Permissions table.**
 - **Add a user group.**
 - a. **In the Permissions table, select the user group to which you want to assign folder permissions.**

- b. **Select the permissions that you want to assign to the user group.**

For a description of each of the permissions, see [“Folder-Specific Permissions”](#) on page 54.

- c. **Click Add to register the new user group’s permissions.**

- **Remove a user group by selecting Remove from the Action column.**

4. **Click Save.**

Changing Folder Owners

You can assign administrative permissions of a folder to different groups by changing the folder’s owner user group. You might want to change a folder’s owner user groups if you or another user creates a new subfolder for the purpose of allowing a different user group to have absolute control over all objects contained within the subfolder.

▼ How to Change Folder Owners

This procedure describes how to change a folder’s owner by using the browser interface. You can also change owners by using the following command.

`fdb.f.co` – Changes a folder’s owner.

Before You Begin

To change the owner user group of a folder, you need to be a member of the current owner user group. However, you do not need to be a member of the new owner user group.

Note – You can remove ownership permissions from yourself and everyone else in the current owner group.

- Steps**
1. **View the Details page of the folder you plan to edit.**

For more information about how to navigate to a folder’s Details page, see [“How to View Folders”](#) on page 84.

2. **On the Details page, click Edit.**
3. **Select the new owner user group from the Owner User Group menu.**
4. **To save your changes, click Save.**

Deleting Folders

When you delete a folder, you delete everything contained in the folder, including components, plans, and other folders. If any object fails to delete, none of the contained objects are deleted and the folder remains.

If you want to delete multiple folders at a time, you can delete them from the folder's list page. When deleting multiple folders, you might succeed in deleting some folders even if other folders fail to delete.

▼ How to Delete Folders

This procedure describes how to delete a folder by using the browser interface. You can also delete folders by using the following command.

```
fdb.f.del – Deletes a folder.
```

Before You Begin To delete a folder you must have write permissions on the folder, the parent folder, and all folders contained by the folder that you need to delete.

- Steps**
- 1. View the Details page of the folder that you plan to delete.**
For more information about how to navigate to a folder's Details page, see [“How to View Folders”](#) on page 84.
 - 2. Click the Delete button at the bottom of the page.**
You are asked to verify or cancel the deletion.
 - 3. Verify that you are deleting the correct folder and click the Continue to Delete button.**
The folder that you deleted is no longer listed in the parent folder's table.

Troubleshooting If your attempt to delete a folder failed, verify the following items.

- The folder and all objects contained in the folder (such as plans and components) are not in use.
For more information on deleting components and plans, see *N1 Grid Service Provisioning System 5.0 Plan and Component Developer's Guide*.
- The folder you attempted to delete was not created by a plug-in.
You can only delete folders created by a plug-in by deleting the plug-in.
- The folder you attempted to delete did not contain a subfolder that was created by a plug-in.

Managing Objects

You can organize objects within the N1 Grid Service Provisioning System software in several ways. Folders, categories, and labels are different mechanisms that allow you to manage your components, plans, and comparisons.

In addition to providing a structure to manage components, plans, and comparisons, the provisioning system provides simple tools to help you search for objects throughout the system. This chapter discusses the following topics.

- [“Introduction to Folders” on page 81](#)
- [“About Categories” on page 93](#)
- [“About Labels” on page 96](#)
- [“About Finding Objects” on page 92](#)

About Showing and Hiding Objects

Over time, objects may become obsolete or no longer require active management. To reduce the visual clutter, you can hide objects so that they are not shown in tables or displayed by CLI commands. This feature allows you to keep objects active in the system without needing to view them.

Objects that can be hidden include the following.

- Components
- Plans
- Hosts
- Host sets
- Host searches
- Host types
- Users

When you hide objects, you can still view them on the object's main page by selecting the show hidden option. For example, if you want to view hidden hosts, you can go to the Hosts page and click Show Hidden Hosts. Hidden hosts will display in a different row color.

About Finding Objects

You can find objects within the provisioning system in a number of ways.

- Tables
- Filters—menus that exist above a list page table.
- Host searches

Finding an Object From the Object's Table

You can search for an object from the top row of the appropriate object table. Searches performed are case-insensitive. A search to find all objects that start with a 'c' may return both "component" and "Component".

EXAMPLE 7-1 Finding a Component with the Object Table

If you are looking for a component that is a `system#file` component type, but you don't know what folder the component might be in, you can perform the following steps.

1. Go to the Components page.
2. Select In Flat View above the Components table.

By selecting In Flat View, you see all components that match your criteria in the current folder and in all nested folders.

Note – If you don't have any idea which folder the component might be in, perform this task from the top level folder (/).

3. From the Type column, select `system#file`.

Note – If you know any other details about the component, you can enter the information in the appropriate field to narrow the search results.

4. Click Find.

Finding an Object Using Filters

You can use filters to simplify your view of objects. You can use filters on the following objects.

- **Hosts.** Using the menu above the table on the Hosts page allows you to filter hosts by host search.
- **Comparisons.** Using the Show Comparisons in Category menu allows you to filter comparisons by category.
- **Plans.** Using the Show Category menu allows you to filter plans by category.
- **Components.** Using the Show Category menu allows you to filter components by category.

See [“About Categories” on page 93](#).

Finding a Host Using Host Searches

While host searches are typically used for generating a list of hosts for use in a host set, you can also use host searches to find a particular host.

For information on how to create a host search, [“How to Create or Edit Host Searches” on page 49](#).

About Categories

You can use categories to group objects that span multiple folders. As you organize your work space, you might notice that several applications rely on a single component. Rather than create multiple copies of the component or group multiple applications in one folder, you can create an easy-to-manage solution with categories.

Category Characteristics

- Components, plans, and comparisons can be assigned to categories.
- An object can belong to more than one category.
- Write permission on the object’s folder is necessary to manipulate categories.
- A component or plan’s version number is not affected by adding or removing a category.

▼ How to Create a Category

This procedure describes how to create a category by using the browser interface. You can also create categories by using the following command.

`cat . add` – Adds a new category.

- Steps**
1. In the navigation menu under the Administrative section, click Categories.
 2. In the top row of the categories table, type the name of the new category and a description.
 3. Click Create.

▼ How to Associate a Category With an Object

You can associate categories with components, plans, and comparisons in several ways.

- Add one or more categories to one or more components or plans.
- Replace an existing set of categories with a new set.

This procedure describes how to associate a category with a component, plan, or comparison by using the browser interface. You can also associate categories by using the following commands.

- `cdb . c . sc` – Associates a component with a category.
- `pdb . p . sc` – Associates a plan with a category.
- `cmp . ds . sc` – Associates comparison settings with a category.

Before You Begin To apply a category to a component or plan, you must have write permissions on the folder that contains the component or plan. For example, if you plan to apply a category to a component, you must have write permissions on the folder that contains the component.

To apply a category to a comparison, you must have write permissions on the comparison.

- Steps**
1. View the Details page of the component, plan, or comparison that you plan to associate with a category.
 2. In the Category row, click Apply Categories.
A window is displayed listing the existing categories. If the component or plan belongs to a category, that category's checkbox is selected.
 3. Associate the object with the category.
 - If the category exists, check the category and click Apply Categories.
 - If you need to associate the object with a new category, perform the following steps.
 - a. Type the new category name and description and click Create.

- b. When the category's name displays in the table, click Save.

Note – If you do not click Save before you apply categories or enter another category name, the first category name is deleted.

- c. If the object is associated with other categories, indicate whether you want the new category to be added to the list of categories or to replace the existing list of associated categories.

4. Click Apply Categories.

The window closes. You can now select the desired category from the pull-down category menu and view all objects that are listed in the category.

▼ How to View Categories

- Step** ● In the navigation menu under the Administrative section, click Categories.
The browser interface displays the Categories page.

▼ How to Filter Objects by Category

After you have categories set up and objects assigned to different categories, you can easily filter across folders.

This procedure describes how to sort objects by category by using the browser interface. You can also sort objects by using the following commands.

- `pdb.p.la -cat CategoryID` – Lists plans from specified category.
- `cdb.c.la -cat CategoryID` – Lists components from specified category.
- `cmp.ds.la -cat CategoryID` – Lists comparisons from specified category.

- Step** ● From the Components, Plans, or Comparisons view, select the category you need to view from the Show Category menu.

The browser interface will display the list of components or plans that are in the category that you selected.

Note – The list of objects is not heterogeneous. You can toggle between Plan and Component views by selecting one or the other from the Show menu.

▼ How to Remove a Category From an Object

This procedure describes how to remove a category from a component, plan, or comparison by using the browser interface. You can also remove categories by using the following commands.

- `pdb.p.la plan-ID -cat ""` – Removes all categories from a plan.
- `cdb.c.la component-ID -cat ""` – Removes all categories from a component.
- `cmp.ds.la comparison-ID -cat ""` – Removes all categories from a comparison.

Before You Begin To remove a category from an object, you must have write permissions on the object or the folder that contains the object. For example, if you plan to remove a category from a component, you must have write permissions on the folder that contains the component.

- Steps**
1. **View the Details page of the object from which you plan to remove a category.**
 2. **In the Category row, click Apply Categories.**
An Apply Categories to Selected Objects window is displayed listing the categories. If the component or plan belongs to a category, that category's checkbox is selected.
 3. **Deselect the category that you want to remove.**
 4. **Click Apply Categories.**
The window closes. The object is no longer associated with the category.

About Labels

Labels allow you to mark components in a way that is meaningful to you. For example, rather than rely on a component's version number, which does not correlate to the product release for which the component is used, you can label a component to indicate which product release a component correlates to.

Label Characteristics

- You can mark only components with labels.
- Making a change to a component's label does not cause the component's version number to increase.

▼ How to Add, Edit, or Remove a Label

This procedure describes how to modify a component's label by using the browser interface. You can also modify labels by using the following command.

```
cdb.c.mod -comp ComponentID -label LabelName – Modifies a component's label.
```

Note – When you add, edit, or remove a component's label, the version number of the component increases.

Before You Begin To add, edit, or remove a label, you must have write permissions on the folder that contains the component you plan to modify.

- Steps**
1. Go to the **Details** page of the component to which you plan to add a label.
 2. Click **Edit**.
 3. In the **Label** text field, enter the new or modified label.

Note – If you plan to remove the component's label, erase the current label and leave the text field empty.

4. Click **Check In**.
The browser interface displays a confirmation page.
5. Select **component's new version number**.
You can choose to check the component in as the next major or minor version.
6. Click **Continue to Check In**.
The browser interface displays the Components page with the updated component label.

▼ How to View a Label

- Step**
- Go to the **Details** page of the component with the label you want to view.
The label name will be display next to the field labeled **Label**.

Monitoring System Activity

You can monitor activity on the N1 Grid Service Provisioning System in several ways. This chapter discusses the following topics.

- “Configuring Email Notification” on page 99
- “Viewing Logs” on page 105
- “Viewing the Run History of a Plan or Comparison” on page 106

Configuring Email Notification

You can specify that the provisioning system send email notification when any system, administrative, or custom event occurs. The criteria that you specify is saved in the provisioning system as a notification rule. You can define a notification rule to apply to a single host or a host set, to an event that generates a particular message, or to an event of a specified severity level.

When you create a notification rule, you tell the provisioning system which events you are interested in. When an event occurs that matches your criteria, the provisioning system sends a message to the email address that is specified in the notification rule.

You might want to create notification rules in the following scenarios.

- You need to know when an error has occurred or a plan has failed to run.
- You want real-time feedback about who is logged in to the provisioning system.
- You want to know about anything that happens on a particular host set.

When you set up a notification rule, you specify the email addresses to which the notification will be sent.

When you installed the provisioning software, you specified the following information.

- The email address from which the provisioning software sends notifications

- The subject line that should appear in the header of notification email messages
The default subject line text is `N1 Service Provisioning System Notification`. To change the notification message's subject line, go to the Master Server's `config.properties` file and update the `note.notificationSubject` parameter with the desired subject line message. Once you have updated the file, restart the Master Server for the change to take affect.

▼ How to Set Up Notification Rules

If you plan to create a new notification rule by using an existing rule as a template, see [“How to Edit Notification Rules” on page 102](#).

This procedure describes how to create notification rules by using the browser interface. You can also perform this task by using the following command.

```
rule.add – Creates a new notification rule.
```

Before You Begin To create a notification rule, you must belong to a user group that has write permissions on notification rules.

Steps 1. **In the navigation menu, click Notification Rules.**

The Notification Rules page is displayed. It lists the notifications that are already saved in the database.

2. **In the top row of the Notification Rules table, type a name and a brief description for the rule you plan to create and click Create.**

The Details page for the notification rule is displayed.

The notification rules Details page allows you to perform the following tasks.

- Specify the event or events that will trigger the notification
- Specify the email address to which the notification should be sent

3. **Specify the criteria for the new notification rule.**

There are four ways to specify which events will trigger an email notification.

- **Event Type.** If you know the type of event that you want to monitor, select the event type from the drop-down list.
- **Severity Level.** If you want to monitor events of a particular severity level, select the severity level from the drop-down list.

When you specify the event's severity level, you can specify whether you want a notification email to be triggered when an event of that severity or higher occurs. To determine what the severity of a particular type of event is, refer to [Table 8-1](#).

- **Event Message.** If you want to select events based on the messages they automatically generate, you can specify part or all of an event text message in the event text contains field.

To determine the generated message of the event you want to monitor, set up a notification rule that will trigger a notification email. When you receive a notification email that was triggered by the event you want to monitor, place the message contents in the Event Text Contains field. You can also use an asterisk (*) to match any number of characters.

- **Host or Host Type.** If you want to monitor a specific host or host set, type the name of the hosts in the host field or select a host set from the drop-down list.

To specify multiple hosts by name, lists the names in the On Host field and separate them with spaces, semicolons, or commas.

4. **In the Email a Notification To field, type an email address for the recipient of the notification messages.**

Separate multiple addresses with spaces, semicolons, or commas.

5. **When you have finished setting up the notification rule, click Save.**

The notification rule is added to the Notification Rules page.

▼ How to View Notification Rules

Viewing notification rules allows you to evaluate if the correct people are receiving the correct amount of system information.

You might want to view notification rules in the following situations.

- You need to determine if the correct people are receiving system messages.
- You need to determine why you are receiving email notifications.
- You need to determine why you are not receiving expected email notifications.
- You want to remove yourself from a notification rule's mailing list.

This procedure describes how to view notification rules by using the browser interface. You can also perform this task by using the following commands.

- `rule.la` – Displays all notification rules.
- `rule.lo` – Displays detailed information about a specific rule.

- Steps**
1. **From the navigation menu, choose Notification Rules.**

The Notification Rules page is displayed.

2. **(Optional) To view details for a specific notification rule, find the row that lists the rule that you want to view and click Details.**

▼ How to Edit Notification Rules

Editing notification rules allows you to adjust existing rules to further meet your needs. You also have the option of creating new notification rules based on existing rules.

This procedure describes how to edit notification rules by using the browser interface. You can also perform this task by using the following command.

`rule.mod`— Modifies a notification rule.

Before You Begin To edit a notification rule, you must belong to a user group that has write permissions on notification rules.

- Steps**
1. **From the navigation menu, choose Notification Rules.**
The browser interface displays the notifications page, which lists the notifications that are already saved in the database.
 2. **Select the notification rule that you plan to edit and click Details.**

Note – If you plan to create a new notification rule based on an existing rule, select an existing rule that is similar to the new rule you plan to create.

The browser interface displays the notification rule's Details page.

3. **Specify the criteria for the new notification rule.**
There are four ways to specify which events will trigger an email notification.
 - **Event Type.** If you know the type of event that you want to monitor, select the event type from the drop-down list.
 - **Severity Level.** If you want to monitor events of a particular severity level, select the severity level from the drop-down list.
When you specify the event's severity level, you can specify whether you want a notification email to be triggered when an event of that severity or higher occurs. To determine what the severity of a particular type of event is, refer to [Table 8-1](#).
 - **Event Message.** If you want to select events based on the messages they automatically generate, you can specify part or all of an event text message in the event text contains field.
To determine the generated message of the event you want to monitor, set up a notification rule that will trigger a notification email. When you receive a notification email that was triggered by the event you want to monitor, place the message contents in the Event Text Contains field. You can also use an asterisk (*) to match any number of characters.

- **Host or Host Type.** If you want to monitor a specific host or host set, type the name of the hosts in the host field or select a host set from the drop-down list.
To specify multiple hosts by name, lists the names in the On Host field and separate them with spaces, semicolons, or commas.
- 4. **Save the notification rule.**
 - **To save the modified notification rule, click Save.**
 - **To use an existing notification rule as a template for a new rule, click Save As.**
 - a. **Type a new, unique name for the notification rule.**
 - b. **Click Continue to Save.**
The new rule displays in the Notification Rules table.

▼ How to Delete Notification Rules

You might want to delete a notification rule when it is no longer being used or you no longer want to receive email about a particular event.

You can delete notification rules at the following times.

- When you view the Notification Rules table
- When you view the notification rule's Details page

This procedure describes how to delete notification rules by using the browser interface. You can also perform this task by using the following command.

`rule.del`—Deletes a notification rule.

Before You Begin To delete a notification rule, you must belong to a user group that has write permissions on notification rules.

- Steps**
1. **Go to the Notification Rules page.**
For the procedure about how to navigate to the Notification Rules page, see "[How to View Notification Rules](#)" on page 101.
 2. **Select each notification rule that you plan to delete.**
You can delete more than one notification rule at a time.
 3. **In the last row of the Notification Rules table, click Delete.**
Another page is displayed that provides the option to cancel the deletion.
 4. **Verify that you are deleting the correct notification rule and click Continue to Delete.**
The notification rule that you deleted is no longer listed in the table on the Notification Rules page.

Event Severities

The following table defines the severity of each event description.

TABLE 8-1 Event Severities

Severity	Event Description
ERROR	System Error
ERROR	Comparison Failed
ERROR	Plan/Preflight Failed
ERROR	Step Failed
WARNING	System Warning
WARNING	Difference Found
INFO	System Information
INFO	Component Installed
INFO	Component Uninstalled
INFO	Component Control Service Executed
INFO	Resource Installed
INFO	Resource Uninstalled
INFO	Comparison Canceled
INFO	Comparison Cancellation Requested
INFO	Comparison Completed
INFO	Comparison Started
INFO	Node Started
INFO	Node Stopped
INFO	Preflight Completed
INFO	Plan/Preflight Canceled
INFO	Plan/Preflight Cancellation Requested
INFO	Plan/Preflight Completed
INFO	Plan/Preflight Started
INFO	Custom Event Sent
INFO	Step Completed

TABLE 8-1 Event Severities (Continued)

Severity	Event Description
INFO	Step Started
INFO	User Logged In
INFO	User Failed to Log In
INFO	User Logged Out

Viewing Logs

The provisioning system uses four different log files to capture system information.

- Master Server log file at *N1SPS-MasterServer-home/server/bin/cr_server.out*
- Remote Agent log file at *N1SPS-RemoteAgent-home/agent/bin/cr_agent.out*
- Local Distributor log file at *N1SPS-LocalDistributor-home/ld/bin/cr_ld.out*

Remote Agent and Local Distributor logs can be viewed by logging into the physical host or through the browser interface. The Master Server log file is only accessible by logging into the physical host.

Some of the information in the log files might repeat errors that appear on the screen, such as Master Server deployment errors. The log files typically contain more detailed information that is not available through the browser interface's error messages. This more detailed information can help you to troubleshoot errors more effectively.

▼ How to View Remote Agent and Local Distributor Log Files

Viewing the Remote Agent or Local Distributor log files through the browser interface is convenient. You only need to log in to the browser interface and not to each of the Remote Agents or Local Distributors that you want to view.

- Steps**
1. **From the host's Details page in the browser interface, click the View Connection & Configuration & Log button.**
A window appears and indicates whether the Master Server could connect to the host.
 2. **If the host responded normally, view the contents of the host's log.**
 - **To view the most recent lines of the host's log, click the Log button.**

The window displays the most recent lines of the host's log.

- **To view the complete log, click Download Log.**
Another window displays the complete contents of the host's log.

Viewing the Run History of a Plan or Comparison

The provisioning system keeps a record of each plan and comparison that is run. The record is started when a plan or comparison is started, and it allows you to view the progress of plans as they are run.

You might want to use the Run History option in the following scenarios.

- To view the progress of a running plan
- To view the progress of a running comparison
- To audit what is happening within the provisioning system

▼ How to View Run Histories

This procedure describes how to view the run history of plans and comparisons in the provisioning system by using the browser interface. You can also perform this procedure by using the following commands.

- `pe.p.la` – Displays running and completed plans.
- `pe.p.lo` – Displays the progress of a running or completed plan.
- `pe.pi.lo` – Lists the parameters used to run a plan.
- `cmp.dj.la` – Lists running and completed comparisons.
- `cmp.dj.lo` – Displays status of a particular comparison.

For a detailed description of these commands, see Chapter 5, “*cmp: CLI Commands for Running Comparisons*,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual* or Chapter 10, “*pe: CLI Commands for Running Plans*,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

- Steps**
1. **In the navigation menu, expand the Application Deployment category and click Run History.**
A list of recently run plans and comparisons displays.
 2. **(Optional) View more information about a plan or comparison.**
 - **To view information about how many times or by whom a plan or comparison was run, click History.**

- To view the details of a plan or comparison that was run, click **Results**.

▼ How to Delete Run Histories

This procedure describes how to view the run history of plans and comparisons in the provisioning system by using the browser interface. You can also perform this procedure by using the following commands.

- `pe.p.del` – Deletes history of completed plan run.
- `cmp.ds.del` – Deletes history of completed comparisons.

For a detailed description of these commands, see Chapter 5, “*cmp: CLI Commands for Running Comparisons*,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual* or Chapter 10, “*pe: CLI Commands for Running Plans*,” in *N1 Grid Service Provisioning System 5.0 Command-Line Interface Reference Manual*.

Before You Begin To delete a run history, you must belong to a user group that has write permissions on comparisons.

- Steps**
1. Go to the **Run History** page.
 2. Select the plan or comparison history that you want to delete and click **Results**.
The run history’s Details page displays.
 3. Click **Delete**.

Backing Up and Restoring

The provisioning system includes utilities for completely backing up and restoring a Master Server. These utilities are found in the *NISPS-home/server/bin* directory.

By default, the resource store, database contents, data files, configuration data, and keystores are backed up. You can skip backing up these components by using the appropriate command-line arguments.

Backing Up and Restoring the Master Server

▼ How to Back Up a Master Server

Before You Begin

Before you back up the Master Server, any plans or preflights must be stopped, as well as any other tasks, such as comparisons.



Caution – Be sure to specify an output directory when you run the backup script. If you do not specify an output directory, the resultant backup file is saved in the *NISPS-home/server/bin* directory. If you are uninstalling and then reinstalling the Master Server and the backup file is in this directory, the backup file is deleted. Consequently, you will not be able to restore the Master Server.

- Steps**
1. Log into the machine as the user who installed the provisioning system.



Caution – Do not run the backup script as root. Ensure that you are logged in to the provisioning system’s user account.

2. Change to the directory where the backup script is located.

```
% cd NISPS-home/server/bin
```

NISPS-home is the home directory of the application.

3. Start the backup.

```
% ./cr_backup.sh options
```

The following options are available for use with the `cr_backup.sh` command.

Note – By default, the resource store, database contents, data files, configuration data, and keystores are backed up. You can skip backing up these components by using the appropriate command-line arguments.

- | | |
|---------------------|---|
| -b | Master Server base directory. |
| -q | Quiet. No informative messages are printed. |
| -nors | Skip the Resource Store back up. |
| -nodb | Skip the Database, Plug-in, Custom Task backup |
| -noconfig | Skip the Configuration and Keystore backup |
| -nokeystore | Skip the Keystore backup |
| -o <i>directory</i> | Save the backup file in this directory |
| | The backup script verifies that you have write permissions for the directory you specified. If you do not have write permissions for this directory, the script generates an error. |
| | If you do not specify a directory, the files are saved in the <i>NISPS-home</i> /server/bin directory. |
| -z | Compress the resulting backup file |
| -l <i>logfile</i> | Write log output to the <i>logfile</i> file instead of the default <code>logfile</code> file. |
| -gz | GZip the resulting backup file if <code>gzip</code> is in PATH. |
| -shut down yes | Shut down the Master Server without prompting the user. |
| -u | Print this information. |
| -h | Print this information. |

The script warns that if you proceed, the Master Server process will stop. Any tasks, such as searches, plans, and comparisons, that are in progress will also stop.

4. Type `y` to proceed with the backup.

The script shows the progress of the backup and displays the location of the backup file.

The script restarts the Master Server.

▼ How to Restore a Master Server

Before You Begin Before performing a restore, you must have an installation of the Master Server that contains no data.

Steps **1. Log into the machine as the user who installed the provisioning system.**

2. Stop the Master Server.

```
%N1SPS-home/server/bin/cr_server stop
```

3. Change to the directory where the backup script is located.

```
% cd N1SPS-home/server/bin
```

N1SPS-home is the home directory of the application.

4. Start the restoration.

```
% ./cr_restore.sh options
```

The following options are available for use with the `cr_restore.sh` command.

Note – By default, the resource store, database contents, data files, configuration data, and keystores are restored. You can skip restoring these components by using the appropriate command-line arguments.

`-b` Master Server base directory

If you do not use the `-b` option to specify a directory to in which to restore the backup files, the files are put in the current directory, *N1SPS4.1-home/server/bin*. If you do not have write permissions for this directory, the software generates an error.

`-q` Quiet. No informative messages are printed

`-nors` Skip the Resource Store restore

-nodb	Skip the Database, Plug-in, Custom Task restore
-noconfig	Skip the Configuration and Keystore restore
-nokeystore	Skip the Keystore restore
-f <i>backupfile</i>	Restore the contents of the <i>backupfile</i> file
-l <i>logfile</i>	Write log output to the <i>logfile</i> file instead of the default <i>logfile</i> file
-t <i>temp_directory</i>	Use the <i>temp_directory</i> directory to save temporary files
-overwrite yes	Overwrite the existing data during restore
-shutdown yes	Shutdown the Master Server without prompting the user
-u	Print this information
-h	Print this information

The script verifies that the backup files do not contain any errors. The script warns that if you proceed the Master Server process will stop, if it is not already stopped. Any tasks, such as searches, plans, and comparisons, that are in progress will also stop.

5. Type y to proceed.

The script warns that any data currently in the databases will be overwritten with the data in the backup files.

6. Type y to proceed.

The restoration proceeds. The script starts the Master Server.

Backing Up and Restoring Remote Agents

To manually back up the Remote Agent, stop the agent and copy the contents of the *NISPS-home/data* directory to a safe location. To restore the Remote Agent, stop the agent, and copy the contents of the directory that you saved.

Authentication Methods

This appendix provides the requirements for adding authentication methods to the provisioning system, the steps to modify an authentication method, and an example of the configuration files that you need to alter when configuring authentication methods.

Setting Up Additional Authentication Methods

The provisioning system has three requirements with regard to adding authentication methods.

- The provisioning system only allows authentication methods that require a single user name and password to authenticate.

Adding an authentication method that allows multiple challenge and responses results in the user being unable to log in.

- Users are not able to add authentication methods with the name “internal” to the `jaas.config` file.

The `internal` authentication method is automatically available in the `jaas.config` file as a hidden authentication method.

- Changes to the authentication methods require the Master Server to be restarted.

▼ How to Modify External Authentication Methods

This procedure describes how to modify the authentication methods that are used within the provisioning system. After you have added a new authentication method, see [“How to Change a User’s Authentication Method”](#) on page 71.

Before You Begin

For more information about how to add authentication methods to the provisioning system, refer to the following documentation.

- For information on editing the Java Authentication and Authorization Service (JAAS) configuration files, see the Java 2 Platform `javax.security.auth.login.Configuration` class documentation (<http://java.sun.com/j2se/1.4.2/docs/api/>).
- For information about each `LoginModule` bundled with JAAS, see the `com.sun.security.auth.module` Java package documentation (<http://java.sun.com/j2se/1.4.2/docs/guide/security/jaas/spec/>).

Steps

1. **Log into the machine as the user who installed the provisioning system.**
2. **Stop the Master Server.**

```
%NISPS-MasterServer-home/server/bin/cr_server stop
```

3. **Edit the `jaas.config` file.**

This file is located at `NISPS-MasterServer-home/server/lib/security/jaas.config`. This file contains documentation on how to set up authentication using LDAP, Sun Directory Server, and Microsoft Windows 2000 Active Directory Server.

4. **Edit the `config.properties` file.**

This file is located at `NISPS-MasterServer-home/server/config/config.properties`. Set the `userdb.authModules` property. The value of this property should be a comma-separated list of the login configurations specified in the `jaas.config` file. The order that the login configurations are listed in represents the order they display in the Authentication Method menu.

Note – Although `internal` cannot be specified in the `jaas.config` file, you can include the default, internal authentication method in the `userdb.authModules` property's comma-separated list.

5. **Restart the Master Server.**

The new or modified login configurations are used by the provisioning system after restart.

Note – If a user's password changes as the result of a change in their login configuration, they will need to re-encrypt or flush their session variables.

Example A-1 jaas.config File Configured for LDAP and Sun Directory Server

```
// This file shall contain login configuration definitions
// for the system. The login configurations enabled for
// authentication are listed in a comma separated list
// value of the property userdb.authModules in the
// config.properties file.
//
// Please note that login configuration name 'internal'
// is reserved and cannot be specified in this file.
// 'internal' login configuration is always available
// by default and includes only internal login module
// for authentication.
//
// Examples
//
// Using internal login configuration only
//
// internal-auth {
//   com.sun.n1.sps.userdb.UserDBLoginModule Required;
// };
//
//
// Using LDAP authentication
//
// ldap-auth {
//   com.sun.n1.sps.userdb.LdapLoginModule Required
//   userProvider="ldap://LDAPServerHostName/userDN"
//   userFilter="ldapFilter"
//   authIdentity="userDN"
//   useSSL="sslFlag"
//   debug="debugFlag";
// }
corporate-ds {
  com.sun.n1.sps.userdb.LdapLoginModule Required
  userProvider="ldap://sol01.cr.nllab.sfbay.sun.com"
  authIdentity="uid={USERNAME},ou=People,dc=cr,dc=nllab,dc=sfbay,dc=sun,dc=com"
};
//
// where:
//   userProvider consists of the LDAPServerHostName part,
//   which specifies the server host name of the LDAP
//   server, which may include a port number (for more
//   information check: (http://www.ietf.org/rfc/rfc2255.txt)
//   (for example, ldap.sun.com:389) and the position
//   in its directory tree where user entries (userDN)
//   are located.This userDN is optional and if specified
//   is used by the userFilter.
//   userFilter specifies the search filter to use to locate
//   a user's entry in the LDAP directory.The ldapFilter
//   is an LDAP filter string(for more information check:
//   (http://www.ietf.org/rfc/rfc2254.txt).If it contains
//   the special token {USERNAME}, then that token will be
//   replaced with the supplied username value before the
```

```

//      filter is used to search the directory.The userFilter
//      search is based on the userDN specified in the
//      userProvider.
//      authIdentity specifies the distinguished name to use to
//      authenticate a user. It is an LDAP distinguished
//      name string (for more information check:
//      http://www.ietf.org/rfc/rfc2253.txt)
//      It must contain the special token "{USERNAME}"
//      which will be replaced with the supplied username
//      value before the distinguished name is used for
//      authentication.Note that if this option does not
//      contain a distinguished name then the userFilter
//      option must also be specified.
//      sslFlag specifies if a ssl connection should be used
//      to connect to the ldap server. Valid values are
//      true or false. The default value is true.
//      debugFlag specifies if debugging information may be
//      generated when attempting LDAP authentication.
//      Valid values are true or false.The default value is
//      false.
//
//
// Here are some example entries for authenticating with a
// Sun Directory Server.
//
// sun-ldap1 {
//   com.sun.n1.sps.userdb.LdapLoginModule Required
//   userProvider="ldap://ldaphost.example.com/ou=People,dc=example,dc=com"
//   userFilter="( &(uid={USERNAME})) (objectClass=inetOrgPerson) )"
//   };
//
// sun-ldap2 {
//   com.sun.n1.sps.userdb.LdapLoginModule Required
//   userProvider="ldap://ldaphost.example.com"
//   authIdentity="uid={USERNAME},ou=People,dc=example,dc=com"
//   };
//
sun-ldap2 {
  com.sun.n1.sps.userdb.LdapLoginModule Required
  userProvider="ldap://falkland.cr.nllab.sfbay.sun.com:389"
  authIdentity="uid={USERNAME},ou=People,dc=example,dc=com";
};
//
// sun-ldap3 {
//   com.sun.n1.sps.userdb.LdapLoginModule Required
//   userProvider="ldap://ldaphost.example.com/ou=People,dc=example,dc=com"
//   userFilter="( &(uid={USERNAME})) (objectClass=inetOrgPerson) )"
//   authIdentity="uid={USERNAME},ou=People,dc=example,dc=com";
// };
//
//
// Here are some example entries for authenticating with a
// Microsoft Windows 2000 Active Directory server.

```

```

// ads-ldap1 {
//   com.sun.n1.sps.userdb.LdapLoginModule Required
//   userProvider="ldap://ldaphost.example.com/CN=user,DC=example,DC=com"
//   userFilter="(&(userPrincipalName={USERNAME}@example.com)(objectClass=user))"
//   authIdentity="{USERNAME}@example.com"
//   useSSL=false
//   debug=true;
// };
//
// ads-ldap2 {
//   com.sun.n1.sps.userdb.LdapLoginModule Required
//   userProvider="ldap://ldaphost.example.com/CN=user,DC=example,DC=com"
//   userFilter="(&(samAccountName={USERNAME}))(objectClass=user)"
//   authIdentity="{USERNAME}@example.com"
//   useSSL=false
//   debug=true;
// };

```

Example A-2 Section of `config.properties` File Configured for LDAP and Sun Directory Server

```

...
# A comma separated list of login configurations for authenticating users.
userdb.authModules=internal,corporate-ds,sun-ldap2
...

```


Commonly Updated Configuration Variables

The following list describes configuration variables that are commonly updated in the `config.properties` file.

`rsrc.maxMSRepoSize`

The maximum limit on the amount of disk space the Master Server repository is allowed to occupy. The default value is 100 Gigabytes.

`rsrc.maxTransientRepoSize`

The maximum limit on the amount of disk space the Remote Agent or Local Distributor repository is allowed to occupy. The default value is 1 Gigabyte. This value must reflect the largest possible file transfer the Remote Agent or Local Distributor could receive.

`userdb.sessionTimeout`

The duration of time that a session may remain inactive before being removed, in milliseconds. The default value is 30 minutes.

`userdb.reaperInterval`

The duration of time to wait before removing inactive sessions from the session table, in milliseconds. This setting is used for performance tuning and should be roughly correlated with the session time out value. The default value is 60 minutes.

`userdb.enableSessionVariablePersist`

The boolean value indicating whether session variable variables should be saved to the database when a user edits them. The default value is `true`.

`userdb.authModules`

A comma-separated list of valid authentication method names to use for authenticating users. Each method, except `internal`, should be listed the `jaas.config` file as well. The default value is `internal`.

Glossary

abstract component	A component that serves only as a base component for other components to extend. An abstract component cannot be installed and only an abstract component is permitted to declare abstract child elements.
call compatibility	A compatibility type for system service components. This compatibility is also called API compatibility or interface compatibility.
category	A general class in which you can group objects that are stored in multiple folders.
child component	A component that is referenced by a container component. Also called <i>contained component</i> . See also <i>container component</i> .
comparison	A feature that searches for and identifies differences between hosts and component models. The N1 Grid Service Provisioning System supports these three types of comparisons: <ul style="list-style-type: none">■ Model to model – Examines the deployment repository and history that is stored on the master server for two hosts and reports any differences■ Model to install – Compares what the master server reports is installed on a host to what is actually on the host and reports any differences■ Install to install – Examines the contents of two hosts' file systems and reports any differences
component	A logical grouping of source information that defines an application. A component also includes a set of instructions that specifies how to manage the source information. The XML representation of a component includes the following:

	<ul style="list-style-type: none"> ■ List of resources used by the application ■ Installation steps ■ Uninstallation steps ■ Dependencies
component compatibility	A situation where a component can be safely replaced by another. The N1 Grid Service Provisioning System supports two kinds of component compatibility: call compatibility and install compatibility.
component inheritance	The means by which a component obtains attributes and behavior from another component. When you create a component, it inherits any variables, snapshots, and procedures from the associated component type.
component procedure	A program in a component that controls deployment of the component, such as installation, uninstallation, management, and capturing snapshots. Management procedures are defined in the control block.
component repository	A location on the master server where components and their resources are checked in.
component type	A special kind of component that encapsulates behavior that can be reused by other components. A component can inherit the behavior of a component type by extending from it.
component variable	A user-definable name-value pair that is used to make parts of a component accessible and configurable by objects that are external to the component.
composite component	A component that contains only references to other components, both simple and composite. A composite component cannot contain any resources.
composite plan	A plan that is composed solely of subplans, which can be simple or composite subplans. A composite plan is not directly targeted, as each subplan can run on a different set of targets.
configuration generation engine	A software engine on the master server that replaces substitution variable references with the appropriate variable setting values. The engine interacts with the host repository and component repository to resolve values any time that you run a plan to deploy a component.
contained component	A component that is referenced by other components.
container component	A component that contains references to other components.
control	A procedure defined by a component that can be used to manage the deployed applications. For example, a control might be used to start or stop an application. Also called <i>control service</i> .

deployment	Using a plan or component procedure to act on a component. The component's lifecycle includes installation, uninstallation, and application management.
direct-run procedure	A component procedure that can be run directly from the component by using the browser interface.
downstream	In the N1 Grid Service Provisioning System software network hierarchy, the server that is further from the master server. For example, the master server connects downstream to a local distributor. Any remote agents connected to the local distributor are downstream from the local distributor.
execNative call	An optional call out to custom scripts from the XML of a plan or component.
execution plan	See <i>plan</i> .
extend	To base a component on a component type so that the component inherits variables and procedures that are defined by the component type. The component can override variable values and procedure definitions defined by its associated component type.
final component	A component that cannot be extended by another component.
folder	Directory-like containers that enable you to apply permissions to and organize components, plans, and subfolders.
gold server	A reference server that contains files, directories, and other resources that make up an application and that checks in these resources to the master server.
host	A server that is managed by the N1 Grid Service Provisioning System.
host set	A user-defined, logical grouping of hosts that share one or more common attributes, such as physical location or functional group. Use a host set to quickly and easily update applications on all hosts in the set. You can also use a host set to perform model-to-install comparisons between two hosts.
host type	A base class of servers that is bound by a set of common attributes, all of which are user-defined. You can use host types to categorize hosts into logical groupings and to facilitate host searches.
host search	A query run on the host repository that yields a list of hosts whose attributes match those specified by the query. For example, you can use host searches to create a list of hosts that have the same host type, that run the same applications, and that are configured with the same subnet masks.
install compatibility	A compatibility type for component types. This compatibility is also called structural compatibility.

Java Runtime Environment (JRE)	A subset of the Java Development Kit (JDK [®]) for users and developers who want to redistribute the runtime environment. The Java runtime environment consists of the Java virtual machine (JVM), the Java core classes, and supporting files.
Java Virtual Machine (JVM)	The part of the Java runtime environment (JRE) responsible for interpreting bytecodes.
Jython	An implementation of the high-level, dynamic, object-oriented language, Python, seamlessly integrated with the Java platform. The predecessor to Jython, JPython, is certified as 100% pure Java.
label	A means of marking a component version beyond the N1 Grid Service Provisioning System version number. For example, a component version number describes the version of the component. A label can describe the version of the application that the component represents.
local distributor	<p>The application that is installed on a server. The Local Distributor application acts as a link between other servers in the N1 Grid Service Provisioning System in the following ways:</p> <ul style="list-style-type: none"> ■ master server to remote agents ■ master server to other local distributors ■ local distributor to remote agents <p>Local distributors maximize bandwidth efficiency and speed, and can also provide secure network connections for navigating restricted environments.</p>
master server	The application that is installed on a server that manages the N1 Grid Service Provisioning System. The Master Server application can connect to any of the data center environments managed by the N1 Grid Service Provisioning System. The Master Server provides centralized data storage, data processing, and user interfaces.
modeling	To create components and plans that represent an application that you want to deploy with the N1 Grid Service Provisioning System.
nested component	A contained component that, when installed, can provide its services only to its container component. A nested contained component defines a finer-grained unit of functionality required by the container component, but is not otherwise useful to other components.
network protocol	A way to transmit data between devices on a network. The N1 Grid Service Provisioning System uses TCP/IP, SSH, and SSL.
notification email	An email sent by the N1 Grid Service Provisioning System to advise that a system, administrative, or custom event has occurred. The system administrator specifies the rules used to determine when notification emails are sent and the email addresses to which the email is sent.

notification rule	The criteria used by the N1 Grid Service Provisioning System to determine whether an email notification is sent. The system administrator defines the criteria that is used to determine when an email notification is sent.
parent component	A component that contains references to other components. Also called <i>container component</i> . See also <i>contained component</i> .
physical host	A physical server that is connected to the network. Within the provisioning system, a physical host can act as a remote agent or a local distributor.
plan	A sequence of instructions that is used to manipulate one or more components. A plan can also be a sequence of other plans, which enables common instruction sequences to be shared between one or more plans.
plan executor	The software engine on the master server that runs preflights and deployments.
preflight	The simulated execution of a plan to a simulated UNIX environment that finds and reports any errors or potential errors that might affect the deployment. A preflight always precedes a deployment, but you can run a preflight as a standalone operation.
procedure	See <i>component procedure</i> .
provisioning system	The software applications that, when installed on servers, form the N1 Grid Service Provisioning System.
remote agent	The application that is installed on any server in the N1 Grid Service Provisioning System to which components are deployed. The Remote Agent application manages tasks, such as installing software, controlling services, and collecting information to deliver to the master server.
resource	A file that is deployed to a host when a plan is executed. The file might be a directory, a symbolic link, or another kind of file.
server	A computer that manages resources and supplies services to a client. In the N1 Grid Service Provisioning System, a server is a computer on which one of the N1 Grid Service Provisioning System applications has been installed.
session	A period of time that is initiated when you log in. A session persists until you log out or inactivity causes the session to expire. Logically, a session represents the authenticated credentials of a particular user. A session is used to identify the user throughout a series of related requests without reauthentication.

session variable	A variable that is associated with a user session. The user can change session variable values for each login session. Session variable values can also be securely saved for reuse in subsequent sessions.
simple component	A component that contains a single resource. A simple component cannot contain references to other components.
simple plan	A sequential list of steps that are executed on a particular set of target servers. A simple plan does not contain or call other plans.
snapshot	A capture of the resources that are stored on a host during a deployment. The snapshot is used when performing comparisons between a host and its model on the master server (model-to-install).
step	An instruction that can be part of a plan or a component.
substitution variable	A variable that appears in plans, components, or configuration files that is substituted by the configuration generation engine during deployment.
system service	A component that is automatically deployed to all applicable hosts when the hosts are prepared. System services define utility controls and resources that can be used by other components.
targetable component	A component that creates a host that serves as a deployment target for other components when it is installed. When a targetable component is uninstalled, the host it created is automatically deleted.
top-level component	A contained component that, when installed, can be used by any component just as if it had been directly installed by a plan. A top-level contained component defines services that will be used by the container component as well as by other components.
upstream	In the N1 Grid Service Provisioning System network hierarchy, the server that is closer to the master server. For example, the master server is upstream from the local distributor. The local distributor is upstream from any remote agents that are connected to that local distributor.
variable	See <i>component variable</i> .
variable settings	A collection of variable values that can be used to override the default values of one or more component variables. Based on the variable settings that you use, you can specify different values for component variables. You specify the variable settings to use when you run a plan.
virtual host	Services that act as a host for other services. For example, a virtual host can represent an application server that acts as a host for web applications.
XML schema	The language used by the N1 Grid Service Provisioning System to create plans and components.

Index

A

adding
 See also creating
 authentication methods, 113-117
 categories, 94-95
 labels, 97
 login configurations, 113-117
 LoginModules, 113-117
 user groups to folders, 87-88
adding hosts, host sets, 33-34
admin User Account, description, 66
admin user group, description, 60-61
alert system, *See* notification rules
assigning, categories, 94-95
associating, *See* assigning
authentication method, changing, 71-72
authentication methods, setting up, 113-117

B

backing up, 109-112
 master servers, 109-111
 remote agents, 112
browser interface, starting, 19

C

categories
 adding, 94-95
 assigning, 94-95
 creating, 93-94

categories (Continued)
 description, 93-96
 removing, 96
 sorting, 95-96
 viewing, 95
changing
 See also editing
 authentication method, 71-72
 folder descriptions, 86-87
 folder owners, 88
 description, 88
 folder permissions, 87-88
 passwords, 70, 71
command-line interface
 See also Jython command-line interface
 starting, 19
comparisons, finding, 93
components
 See also targetable components
 finding, 93
config.properties file, updating, 100
creating
 categories, 93-94
 folders, 82-83
 description, 82
 host searches, 49-51
 host sets, 46
 host types, 40-41, 41-42
 description, 40
 hosts
 targetable components, 30
 Local Distributors, 23-25
 notifications rules, 100-101

- creating (Continued)
 - physical hosts, 23-25
 - Remote Agents, 23-25
 - user accounts, 66-68
 - user groups, 62-63
 - virtual hosts, 29-30
- crhost, 39

D

- database logs, viewing, 105-106
- deactivating, user accounts, 69
- default, host types, 39
- default user account, *See* admin User Account
- default user groups, description, 60-61
- deleting
 - folders, 89
 - description, 89
 - host searches, 52
 - host sets, 48
 - host types, 44
 - hosts, 37
 - description, 36
 - targetable components, 38
 - notification rules, 103
 - plug-ins, 78-79
 - user groups, 65-66
 - description, 65
- description
 - admin User Account, 66
 - admin user group, 60-61
 - categories, 93-96
 - changing folder owners, 88
 - creating folders, 82
 - creating host types, 40
 - default user groups, 60-61
 - deleting folders, 89
 - deleting hosts, 36
 - deleting user groups, 65
 - folders, 81-89
 - hiding hosts, 34
 - hiding objects, 91-92
 - host searches, 49-52
 - host sets, 45-48
 - host-type attributes, 40
 - host types, 38-44
 - hosts, 22-38

- description (Continued)
 - interfaces, 18-20
 - labels, 96-97
 - moving folders, 85-86
 - notification rules, 99-105
 - platform host sets, 45
 - preparing physical hosts, 25-26
 - Registered User Group, 61
 - run history, 106-107
 - showing objects, 91-92
 - Universal User Group, 61
 - upgrading plug-ins, 76-77
 - user accounts, 66
 - user groups, 60-66
 - viewing folders, 84
 - viewing hosts, 31
- display, hidden hosts, 35-36

E

- editing
 - host descriptions, 34
 - host searches, 49-51
 - host sets, 47-48
 - host-type attributes, 32-33
 - host types, 43
 - hosts, 32
 - labels, 97
 - notifications rules, 102-103
 - user groups, 64-65
- email, *See* notification rules
- email notification, *See* notification rules
- email notifications, *See* notification rules
- event severities, 104-105

F

- filters, 92-93
 - using, 93
- finding
 - comparisons, 93
 - components, 93
 - hosts, 93
 - objects, 92-93
 - plans, 93
 - using filters, 93

- finding objects in, tables, 92
- folder descriptions, changing, 86-87
- folder permissions, changing, 87-88
- folders
 - changing owners, 88
 - description, 88
 - creating, 82-83
 - description, 82
 - deleting, 89
 - description, 89
 - description, 81-89
 - moving, 85-86
 - description, 85-86
 - renaming, 86
 - viewing, 84
 - description, 84

G

- granting permissions, 87-88
- group membership, changing, 69-70

H

- hiding
 - hosts, 34-35
 - description, 34
 - objects
 - description, 91-92
 - user accounts, 69
- host-creating components, *See* targetable components
- host descriptions, editing, 34
- host lists, *See* host searches
- host searches
 - creating, 49-51
 - deleting, 52
 - description, 49-52
 - editing, 49-51
 - viewing, 51
- host sets
 - adding hosts, 33-34
 - creating, 46
 - deleting, 48
 - description, 45-48
 - editing, 47-48

- host sets (Continued)
 - removing hosts, 33-34
 - viewing, 46-47
- host-type attributes
 - description, 40
 - editing, 32-33
 - example, 40
- host types
 - creating, 40-41, 41-42
 - description, 40
 - default, 39
 - deleting, 44
 - description, 38-44
 - editing, 43
 - viewing, 42
- hosts
 - architecture, 21-22
 - creating
 - targetable components, 30
 - creating physical, 23-25
 - creating virtual, 29-30
 - deleting, 37
 - description, 36
 - targetable components, 38
 - description, 22-38
 - display hidden, 35-36
 - editing, 32
 - finding, 93
 - finding with host searches, 93
 - hiding, 34-35
 - description, 34
 - modifying, 32
 - networking, 38
 - overview, 21-22
 - physical, 23
 - preparing physical, 26
 - setting up, 21-22
 - viewing, 31-32
 - description, 31
 - virtual, 27-29

I

- importing, plug-ins, 74-75
- interactive CLI, launching, 19-20
- interfaces
 - See also* command-line interface

interfaces (Continued)
 browser interface, 18-20
 description, 18-20
 Jython command-line interface, 18-20

J

Jython command-line interface, *See* interactive CLI

L

labels
 adding, 97
 description, 96-97
 editing, 97
 removing, 97
 viewing, 97
launching
 See starting
 interactive CLI, 19-20
Local Distributors
 creating, 23-25
local distributors
 starting, 17
 Windows, 17-18
Local Distributors
 viewing log files, 105-106
login configurations, adding, 113-117
LoginModules, adding, 113-117
logs
 description of, 105-106
 viewing, 105-106

M

master servers
 backing up, 109-111
 restoring, 111-112
 starting, 17
 Windows, 17-18
modifying
 See also changing
 See edit
monitoring, *See* email notification rules

moving
 folders, 85-86
 description, 85-86

N

networking, host commands, 38
notification rules
 creating, 100-101
 deleting, 103
 description, 99-105
 editing, 102-103
 email subject lines, 100
 event severities table, 104-105
 viewing, 101
notifications, *See* notification rules

O

objects, finding, 92-93
overview
 hosts, 21-22
 setting up, 15-16
 user groups, 59-60
 users, 59-60

P

passwords
 changing, 70, 71
patches, 76-77
patching, plug-ins, 76-77
permissions, folders, 87-88
physical hosts, 23
 See also hosts
 preparing, 26
 description, 25-26
plans, finding, 93
platform host sets, description, 45
plug-ins
 deleting, 78-79
 importing, 74-75
 patching, 76-77
 upgrading, 77-78
 description, 76-77

- plug-ins (Continued)
 - viewing, 75
- preparing
 - physical hosts, 26
 - description, 25-26

R

- Registered User Group, description, 61
- relocating, *See* moving
- remote agents
 - backing up, 112
- Remote Agents
 - creating, 23-25
- remote agents
 - restoring, 112
 - starting, 17
 - Windows, 17-18
- Remote Agents
 - viewing log files, 105-106
- removing
 - categories, 96
 - labels, 97
- removing hosts, host sets, 33-34
- renaming, folders, 86
- restoring, 109-112
 - master servers, 111-112
 - remote agents, 112
- run history, description, 106-107

S

- saving as, new host type, 41-42
- setting up
 - authentication methods, 113-117
 - overview, 15-16
- severity, event descriptions, 104-105
- show, *See* display
- showing
 - objects
 - description, 91-92
- sorting, categories, 95-96
- starting
 - browser interface, 19
 - command-line interface, 19
 - Jython command-line interface, 19-20

- starting (Continued)
 - local distributors, 17
 - Windows, 17-18
 - master servers, 17
 - Windows, 17-18
 - remote agents, 17
 - Windows, 17-18

T

- tables, finding objects from, 92
- targetable components
 - creating hosts, 30
 - deleting hosts, 38

U

- Universal User Group, description, 61
- updating
 - See also* editing
 - See also* upgrading
 - `config.properties` file, 100
 - email subject lines, 100
 - passwords, 70, 71
- upgrading
 - plug-ins, 77-78
 - description, 76-77
- user accounts
 - changing membership, 69-70
 - creating, 66-68
 - deactivating, 69
 - description, 66
 - hiding, 69
 - viewing, 68
- user groups
 - adding to folders, 87-88
 - creating, 62-63
 - default
 - description, 60-61
 - deleting, 65-66
 - description, 65
 - description, 60-66
 - editing, 64-65
 - overview, 59-60
 - viewing, 63
- users, overview, 59-60

using, host searches, 93
using filters, finding objects, 93

V

viewing

- categories, 95
- database logs, 105-106
- folders, 84
 - description, 84
- host searches, 51
- host sets, 46-47
- host types, 42
- hosts, 31-32
 - description, 31
- labels, 97
- logs, 105-106
- notification rules, 101
- plug-ins, 75
- user accounts, 68
- user groups, 63

virtual hosts, 27-29

- See also* hosts
- creating, 29-30