OPSWARE INC
Automating IT™

Opsware® System 4.6 Administration Guide

Table of Contents

Chapter 2: Opsware System Administration   33

Preface

Welcome to Opsware System 4.6 — an enterprise-class software solution that enables customers to get all the benefits of Opsware Inc.'s data center automation platform and support services. Opsware System 4.6 provides a core foundation for automating formerly manual tasks associated with the deployment, support, and growth of server and server application infrastructure.

About this Guide

This guide describes how to administer Opsware System 4.6, including how to set up users, user groups, administrators, customers, and facilities. It provides information about permissions that need to be set for different levels of operation and access, and it discusses how to monitor and diagnose the health of the Opsware System components.

This guide is intended for Opsware administrators who will create and modify customer accounts, populate drop-down lists, resolve database conflicts in multiple core environments, monitor logs, and stop and restart components.

Contents of this Guide

This guide contains the following chapters:

Chapter 1: Opsware System Architecture – provides an overview and diagrams of the Opsware System architecture, showing how the Opsware System components and subsystems interact both in single core and multiple core environments. Each of the components and its function is introduced.

Chapter 2: Opsware System Administration – provides information about how to create and delete users, user groups, and administrators and how to assign permissions to each. It describes how to create, update, and delete customers and facilities, and describes best practices for avoiding and resolving database conflicts in multiple core environments. Monitoring user sessions and defining server attributes are also discussed.

Chapter 3: System Health – provides information about possible Opsware System problems, how to contact support, and how to test and diagnose both Opsware System components and managed servers. It describes how to locate component logs, stop and restart Opsware System components, and restart order dependencies. It also discusses how to administer the Opsware Access & Authentication Directory.

Conventions in this Guide

This guide uses the following typographical and formatting conventions.

NOTATION	DESCRIPTION
Bold	Defines terms.
Italics	Identifies guide titles and provides emphasis.
`Courier`	Identifies text of displayed messages and other output from Opsware System programs or tools.
`Courier Bold`	Identifies user-entered text (commands or information).
`Courier Italics`	Identifies variable user-entered text on the command line or within example files.

Icons in this Guide

This guide uses the following icons to indicate important information.

ICON	DESCRIPTION
	This icon is a note. It identifies especially important concepts that warrant added emphasis.
	This icon is a requirement. It identifies a task that must be performed before an action under discussion can be performed.
	This icon is a tip. It identifies information that can help simplify or clarify tasks.
	This icon is a warning. It is used to identify significant information that must be read before proceeding.

Guides in the Documentation Set and Who Should Read Them

The *Opsware System 4.6 Administration Guide* is intended to be used by the Opsware administrator who is responsible for performing the tasks described herein.

The *Opsware System 4.6 User's Guide* is intended to be used by system administrators who perform the day-to-day functions of managing servers, provisioning operating systems, uploading packages, setting up software trees and node hierarchies, attaching software applications and installing them on servers, managing patches, reconciling servers with software, creating and executing scripts, tracking configuration, and deploying and rolling back code and content.

The *Opsware System 4.6 Installation Guide* is intended to be used by system administrators who are responsible for the installation of Opsware System 4.6 in a facility. It documents how to run the Opsware System Installer and how to configure each of the components.

Chapter 1: Opsware System Architecture

Opsware System 4.6 Overview

This section provides an overview of the Opsware System and contains the following topics:

• About Opsware System 4.6

• Supported Operating Systems

• Browsers That Opsware System 4.6 Supports

• Configuring Browsers for Opsware System 4.6

• The Opsware System Environment

• Model-Based Control

About Opsware System 4.6

The Opsware System manages operational activities across disparate, distributed environments, and provides the visibility and control needed to manage a complex operational environment.

The Opsware System is expandable and scalable. As your operational environment grows, you can add additional Opsware Systems for more functionality.

In this guide, Opsware System 4.6 is referred to as the *Opsware System*.

Supported Operating Systems

Table 1-1 shows the operating systems for managed servers that the Opsware System currently supports. The Opsware System also supports a large selection of popular third-party software.

Table 1-1: Supported Operating System Versions for Managed Servers

OPERATING SYSTEM	VERSIONS
AIX	AIX 4.3 AIX 5.1 AIX 5.2
Fujitsu Solaris	Fujitsu 2.8 Fujitsu 2.9
HP-UX	HP-UX 10.20 HP-UX 11.00 HP-UX 11.11/11i
Red Hat Linux	Red Hat Linux 6.2 Red Hat Linux 7.1 Red Hat Linux 7.2 Red Hat Linux 7.3 Red Hat Linux 8.0 Red Hat Linux Advanced Server 2.1 Red Hat Linux Advanced Server 3.0 Red Hat Linux Enterprise Server 2.1 Red Hat Linux Enterprise Server 3.0 Red Hat Linux Workstation 3.0
Sun Solaris	SunOS 5.6 SunOS 5.7 SunOS 5.8 SunOS 5.9
SuSE Linux	SuSE Linux Enterprise Server 8.0 SuSE Linux Standard Server 8.0 SuSE Linux Standard Server 9.0

Table 1-1: Supported Operating System Versions for Managed Servers

OPERATING SYSTEM	VERSIONS
Windows	Windows NT 4.0 Windows 2000 Server Family Windows Server 2003

Opsware Inc. is continually expanding support to include additional software, hardware, servers, and operating systems. By using commercially available packaging software, you can add software support, or Opsware Inc. can provide these services.

See the *Opsware System 4.6 Installation Guide* for information about supported operating systems for the Opsware System core components and the hardware requirements.

Browsers That Opsware System 4.6 Supports

Table 1-2 shows the browsers that the Opsware Command Center, the Web-based interface for the Opsware System, supports.

Table 1-2: Supported Browsers

BROWSER	WINDOWS 2000	WINDOWS 2003	WINDOWS XP	LINUX 6.2+	SUNOS 5.6+	APPLE OS X
Microsoft Internet Explorer 5.5	X					
Microsoft Internet Explorer 6.0	X	X	X			
Mozilla 1.6		X	X	X	X	X

Configuring Browsers for Opsware System 4.6

You must configure your browsers to accept cookies and to use Javascript to function correctly with the Opsware System. Also, Opsware Inc. recommends using a Web browser that provides 128-bit encryption.

The Opsware System Environment

In an Opsware System-managed environment, the following two main components are installed in your facility that provide the core Opsware System platform support and the infrastructure used to run your operational environment:

• *Opsware System Core Technology* – The set of back-end processes needed to manage the environment such as the Software Repository, the Model Repository, the Command Engine, the Data Access Engine, and so forth.

• *Managed Environment* – All servers that the Opsware System manages by virtue of the Opsware Agent, which resides on each managed server and performs tasks such as installing or removing software, installing or removing patches, and so forth. The OS Build Agent also resides on each server, and is responsible for registering a bare metal server with the Opsware System and guiding the OS installation process. See Figure 1-1.

Figure 1-1: Opsware System Environment
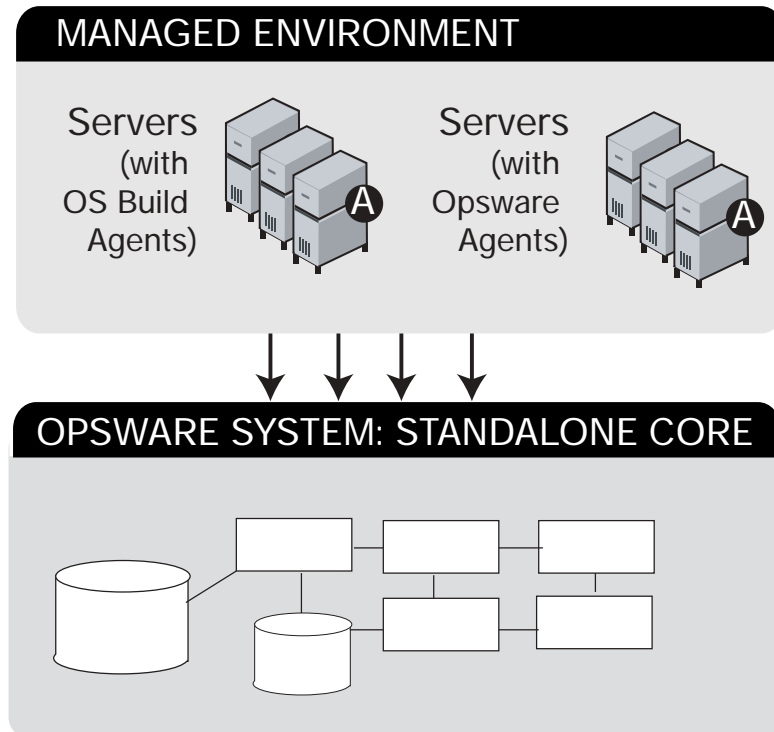
Model-Based Control

The Opsware System utilizes a model-based control approach to accomplish infrastructure management, as Figure 1-2 shows.

Users and administrators interact with the Opsware Command Center, a Web-based front-end application, to accomplish Opsware System tasks such as server management, software distribution, patch management and installation, inventory reporting, system diagnosis, and code and content deployment to the operational environment. The Opsware System tracks the operational environment through a back-end system and data model that has the following key components:

- *Model Repository* - a data repository that stores information about the hardware and software deployed in the operational environment. All Opsware System components work from, or update, a data model of information maintained in the Model Repository for all servers that the Opsware System manages.

- *Software Repository* - a central repository for all software that the Opsware System manages and deploys in the operational environment

- *The Command Engine* - a system for running distributed programs across many servers

• *An Opsware Agent* - on each Opsware System-managed server. Whenever the Opsware System needs to enact change on servers or query servers, it sends requests to the Opsware Agents.

Figure 1-2: The Opsware System Approach to Infrastructure Management
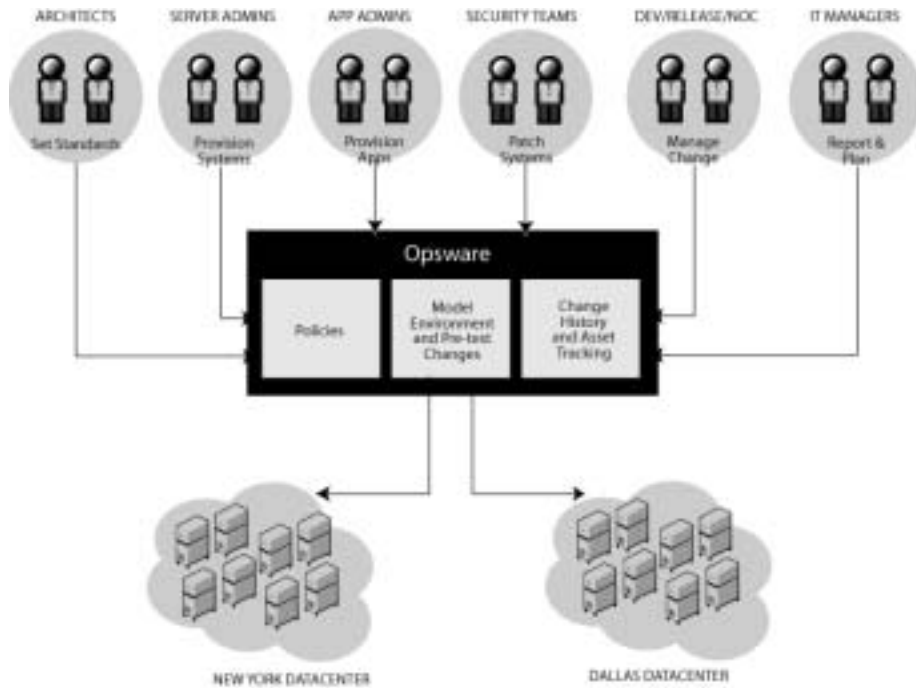
The Opsware System Components

This section explains in detail the Opsware System components and contains the following topics:

• Opsware System Components Overview

• Access & Authentication Directory

• Boot Server

• Build Manager

• Command Engine

• Data Access Engine

• Media Server

- Model Repository

- Model Repository, Multimaster Component

- Opsware Agents

- Dormant Opsware Agents

- Opsware Command Center

- OS Build Agent

- Software Repository

- Software Repository Replicator

- Software Repository, Multimaster Component

- Web Services Data Access Engine

Opsware System Components Overview

The Opsware System utilizes an agent-server architecture. Each server under management has an Opsware Agent that can perform tasks remotely. The *server* portion of the platform consists of multiple, integrated systems, each with a unique purpose.

The Opsware System is made up of the following multiple integrated systems. The sections that follow describe each of these integrated systems.

- *Access & Authentication Directory* - an authentication directory that stores user account information

- *Boot Server* - part of the OS Provisioning subsystem, supports network booting of Sun and x86 systems

- *Build Manager* - facilitates communication between components for OS provisioning

- *Command Engine* - a system for running distributed programs across many servers

- *Data Access Engine* - an XML-RPC interface to the Model Repository

- *Media Server* - provides network access to vendor-supplied media used during OS provisioning

- *Model Repository* - Opsware System's data repository

- *Model Repository, Multimaster Component* - the application that propagates and synchronizes changes from each Model Repository database to all other Model Repository databases

- *Opsware Agent* - an intelligent agent that runs on each server that the Opsware System manages

- *Opsware Command Center* - the primary user interface to the Opsware System

- *OS Build Agent* - responsible for registering a bare metal server with the Opsware System and guiding the OS installation process

- *Software Repository* - the central repository for all software that the Opsware System manages

- *Software Repository Replicator* - serves as backup for Software Repositories in a multimaster mesh, ensuring that packages are available, even if one of the Software Repositories becomes unavailable

- *Software Repository, Multimaster Component* - aids in transferring software from the Software Repository in one facility to the Software Repository in another facility in a multimaster mesh

- *Web Services Data Access Engine* - provides increased performance from the Model Repository to other Opsware System components

Access & Authentication Directory

The Access & Authentication Directory is the central repository used to store user account information. It is implemented by using an LDAP directory that contains account information for specific individuals, the groups they belong to, and the privileges they possess. The Access & Authentication Directory is used for the following purposes:

- Authenticate people when they use the Opsware System services

- Control access to critical Opsware System resources

For example, when an individual accesses the Opsware Command Center, a login screen prompts the person for a username and password. When the user clicks the Sign in button, the Opsware Command Center authenticates the username and password in the Access & Authentication Directory. If the username exists in the Access & Authentication Directory, and the password is correct, the Opsware Command Center allows access. Additionally, the Opsware Command Center retrieves information about the user that describes the user's role and grants read-only or read-write access to specific Opsware System functionality, as appropriate.

See "User Groups and Opsware System Permissions" on page 44 in Chapter 2 for information about how to create users and grant them access to Opsware System functionality.

Boot Server

The Boot Server, part of the OS Provisioning Subsystem, supports network booting of Sun and x86 systems with inetboot and PXE respectively. The processes used to provide this support include the Internet Software Consortium DHCP server, and Sun Solaris TFTP and NFS.

Build Manager

The Build Manager component facilitates communications between OS Boot Agents and the Command Engine. It accepts OS provisioning commands from the Command Engine, and it provides a runtime environment for the platform-specific build scripts to perform the OS provisioning procedures.

Command Engine

The Command Engine is a system for running distributed programs across many servers (usually Opsware Agents). Command Engine scripts are written in Python and run on the Command Engine server. Command Engine scripts can make Remote Procedure Calls (RPC) on Opsware Agents. These calls are delivered in a secure manner and are auditable by using data stored in the Model Repository.

Tools (such as Code Deployment & Rollback) can use Command Engine scripts to implement part of their functionality.

Data Access Engine

The Data Access Engine provides an XML-RPC interface to the Model Repository that simplifies interaction with various clients such as the Opsware Command Center, system data collection, and monitoring agents on servers.

Because interactions with the Model Repository go through the Data Access Engine, clients are less impacted by changes to the Model Repository's schema. The Data Access Engine allows features to be added to the Opsware System without requiring system-wide changes.

9

Media Server

The Media Server is also part of the OS Provisioning Subsystem, and is responsible for providing network access to the vendor-supplied media used during OS provisioning. The processes used to provide this support include the Samba SMB server and Sun Solaris NFS.

Model Repository

The Model Repository is implemented as an Oracle database. All Opsware System components work from, or update, a data model of information maintained for all servers that the Opsware System manages. The Model Repository contains essential information necessary to build, operate, and maintain:

• A list of all servers under management

• The hardware associated with these servers, including memory, CPUs, storage capacity, and so forth

• The configuration of those servers, including IP addresses

• The operating system, system software, and applications installed on servers

• Information on other software available for installation on servers and how it is bundled

• Authentication and security information

When running the Opsware System from multiple facilities, the Model Repository runs in multimaster mode. Running the Opsware System in multimaster mode allows users to maintain multiple Opsware System cores, all write-able, and all kept synchronized, and allows management across facilities, as well.

Model Repository, Multimaster Component

When the Opsware System is running in multiple facilities, the Model Repository, Multimaster Component is installed with the Opsware System core. The Model Repository, Multimaster Component is a Java program that is responsible for keeping multiple Model Repositories synchronized and propagating changes from the originating Model Repository to all other Model Repository databases. Every Model Repository instance has one Model Repository, Multimaster Component instance.

Each Model Repository, Multimaster Component consists of the sender (Outbound Model Repository, Multimaster Component) that polls the Model Repository and sends unpublished transactions, and the receiver (the Inbound Model Repository, Multimaster Component) that receives the transactions and applies them to the local Model Repository instance.

Opsware Agents

Each server that the Opsware System manages has an intelligent agent running on that server. The Opsware Agent is the *agent of change* on a server. Whenever the Opsware System needs to make changes to servers, it does so by sending requests to the Opsware Agent. Depending on the request, the Opsware Agent might use global Opsware System services (such as the Model Repository and Software Repository) in order to fulfill the request. Some functions that the Opsware Agent supports are:

• Software installation and removal

• Configuration of software and hardware

• Periodically reporting server status

• Auditing of the server

An Opsware Agent is idle unless the Opsware System is trying to perform some change on the server. In addition, each Opsware Agent periodically contacts the Model Repository and registers itself, which allows the Model Repository to keep track of machine status, and know when particular servers are disconnected from and reconnected to the network.

Dormant Opsware Agents

The Opsware Agent Installer can install Opsware Agents even when the Opsware System core is not available to a server. If a newly-installed Opsware Agent cannot contact an Opsware System core, the Opsware Agent runs in a dormant mode. While dormant, it periodically attempts to contact the Opsware System core.

When the Opsware System core becomes available, the Opsware Agent performs the initialization tasks, such as hardware and software registration, that usually take place when the Opsware Agent is first installed.

Opsware Command Center

The Opsware Command Center is the primary user interface to the Opsware System. Through the Web-based user interface, an Opsware System user can provision and maintain systems, and deploy code and content to servers. An Opsware administrator adds users and defines access to specific Opsware System resources.

The Opsware Command Center talks primarily to the Data Access Engines (which communicate with the Model Repository), though they also talk directly to other back-end services to implement some operations. Users accessing the Opsware Command Center are authenticated before they gain access.

OS Build Agent

The OS Build Agent, part of the OS provisioning subsystem, is responsible for registering bare metal servers in the Opsware System. In addition, it is the agent of change on the server during the OS installation process (that the Build Manager manages) through to the point at which the actual Opsware Agent is installed.

Software Repository

The Software Repository is the central repository for all software that the Opsware System manages. It contains software packages for operating systems, applications (for example, BEA WebLogic or IBM WebSphere), databases, customer code, and software configuration information.

Working with the Software Repository, an Opsware Agent can install software running on the server where the Opsware Agent is installed. The Model Repository then updates its record of the software installed on the server. This process of updating the actual software configuration of a server with a specified configuration stored in the Model Repository is called *reconciliation*.

You can install new software, code, or configurations in the Software Repository by first packaging the files, and then uploading them into the Software Repository.

See the *Opsware System 4.6 User's Guide* for information about how to upload software packages to the Software Repository.

Software Repository Replicator

The Software Repository Replicator provides backup functionality for Software Repositories running in a multimaster mesh. In most deployments, the Software Repositories do not all have the same content. If one of the Software Repositories becomes unavailable, this might result in some packages not being available until the Software Repository is back online.

Using the Software Repository Replicator provides redundant storage of Software Repositories and thereby helps to ensure that all packages remain available even when a Software Repository goes offline.

Software Repository, Multimaster Component

The Software Repository, Multimaster Component allows software to be distributed across several Software Repositories and to be transferred from one repository to another on-demand. For example, if a Solaris package that resides on Software Repository (A) is needed for installation in a second data center that contains Software Repository (B) that is part of the same multimaster mesh, the Multimaster Component allows B to discover the presence of the package on A. The package is then transferred and cached at B so that it can be used in the second data center.

Web Services Data Access Engine

The Web Services Data Access Engine provides a public object abstraction layer to the Model Repository. This object abstraction can be accessed through a Simple Object Access Protocol (SOAP) API by third-party integration components, or it can be accessed through a binary protocol by Opsware System components like the Opsware Command Center. It provides increased performance to other Opsware System components.

Interaction Among Opsware System Components

To understand the Opsware System architecture, review the following types of Opsware System component interactions:

• General Interaction Among Components

• Opsware System Security

• OS Provisioning Subsystem

• Patch Management Subsystem

• Application Provisioning

• Code Deployment and Rollback Subsystem

• Script Execution Subsystem

• Integration with AIX and HP-UX Installation Technology

• Component Interaction in Multiple Facilities

General Interaction Among Components

The Opsware Command Center, Command Engine, Access & Authentication Directory, Software Repository, and Opsware Agent interact with the Model Repository through the Data Access Engine.

The Data Access Engine issues queries against the Model Repository. It does not cache query results.

The Software Repository contacts the Access & Authentication Directory to authenticate all clients. It maps the client's IP address to the customer name. The Software Repository performs this mapping to enforce access rules on customer-specific files. Figure 1-3 shows an overview of the components.

Figure 1-3: General Overview of All Components
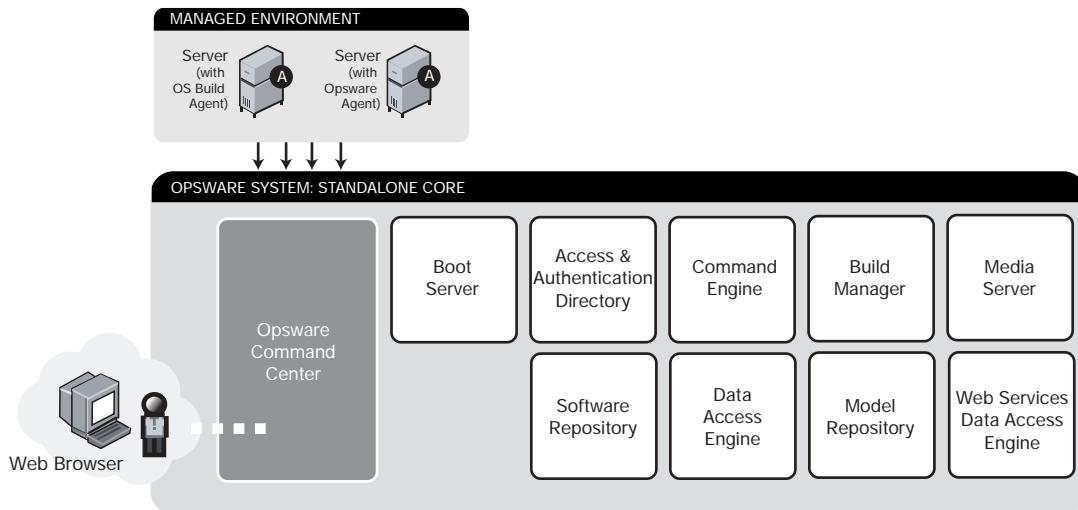
Opsware System Security

To enable secure communication with the Opsware Agent, the Opsware System automatically issues a unique cryptographic certificate to every server that it manages. The certificate is tied to the server to which it is issued, and cannot be copied and used by a different server. The certificate allows the Opsware Agent to establish a secure https connection to the Opsware System components.

As an additional security measure, the Opsware System performs checks on all requests that an Opsware Agent issues. The Opsware System verifies that the requested operation is appropriate for the particular server and checks the parameters of the request to make sure that they fall within reasonable bounds.

OS Provisioning Subsystem

The OS Provisioning Subsystem supports installation-based provisioning using Red Hat Linux Kickstart, Sun Solaris JumpStart, and Microsoft Windows unattended installation. Image-based provisioning (using Symantec Ghost and Sun Solaris Flash) is not supported out-of-the-box.

Because the OS Provisioning Subsystem supports installation-based provisioning, your organization can keep its OS installations lean. Rather than trying to manage changing software through master images, you can use the OS Provisioning Subsystem to install and remove often changing software, including system patches, system utilities, and third-party agents (such as monitoring, backup, and anti-viral agents). Figure 1-4 shows the OS provisioning subsystem interaction.

See the *Opsware System 4.6 User's Guide* for information about the OS provisioning process.

Figure 1-4: OS Provisioning Subsystem Interaction
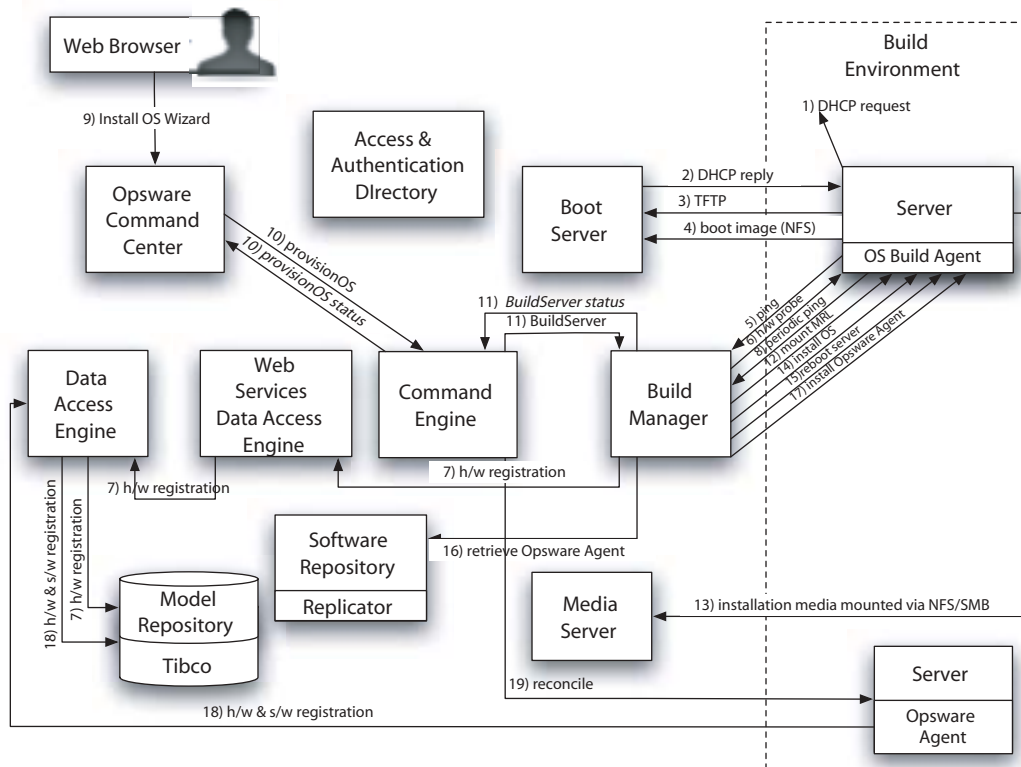
Figure 1-4 shows the following transactions taking place:

1. The DHCP request is a network broadcast.

2. The DHCP reply contains the IP address of the Build Manager for use by Sun Solaris and Red Hat Linux provisioning. For Microsoft Windows provisioning, the DNS configuration in the DHCP reply must resolve the hostname *buildmgr*.

3. TFTP is used to boot the server over the network (using inetboot for Solaris, and PXE for Windows and Linux). Windows and Linux can use a boot floppy instead of PXE.

4. An NFsboot image is used by Solaris and Linux only.

5. The OS Build Agent pings the Build Manager.

6. The Build Manager invokes a Build Script that probes the server's hardware.

7. The hardware detected in Step 6 is then used to register the server with the Opsware System.

8. The OS Build Agent periodically contacts the Build Manager with a ping message. The system remains in this state unless a user begins a provisioning process with the Opsware Command Center, or the server is removed from the network, in which case it will eventually be removed from the Server Pool.

9. A user initiates OS provisioning with the Install OS Wizard in the Opsware Command Center.

10 11 Feedback is provided throughout OS provisioning with status messages passed from the Build Manager to the Command Engine and from the Command Engine to the Opsware Command Center.

12 A Media Resource Locator contains the network location (hostname and path) of an NFS or SMB server from which to retrieve the vendor OS installation media.

13 Solaris and Linux provisioning use NFS. Windows provisioning uses SMB.

14 The vendor installation program is used to install the OS (Sun Solaris Jumpstart, Red Hat Linux Kickstart, or Windows unattended.txt).

15 The server is rebooted after OS installation.

16 17 The OS Build Agent is used to install the Opsware Agent.

18 Hardware and software registration is performed as part of the Opsware Agent installation.

19 The Reconcile function installs additional software that the vendor installation program did not install.

Steps 12 through 18 are managed by a build script that runs inside the Build Manager. The build script is invoked by provisionOS and manages the OS installation at a micro level. ProvisionOS is a Command Engine script that is responsible for managing the installation process at a macro level.

Patch Management Subsystem

The Opsware System automates the key aspects of patch management, while offering a fine degree of control over how and under what conditions patches are installed.

Because patches are often released to address grave security threats, an organization needs to be able to roll out patches quickly, before systems become compromised. At the same time, however, patches can cause serious problems, from performance degradation to general system failure.

The Patch Management Subsystem allows you to react quickly to newly discovered threats, but it also provides support for strict testing and standardization of patch installation. And, if patches later cause problems even after being tested and approved, the Patch Management Subsystem also allows you to uninstall the patches in a safe and standardized way. See Figure 1-5 through Figure 1-8.

See the *Opsware System 4.6 User's Guide* for information about the patch management process.

Figure 1-5: Patch Management Subsystem: Upload Patch
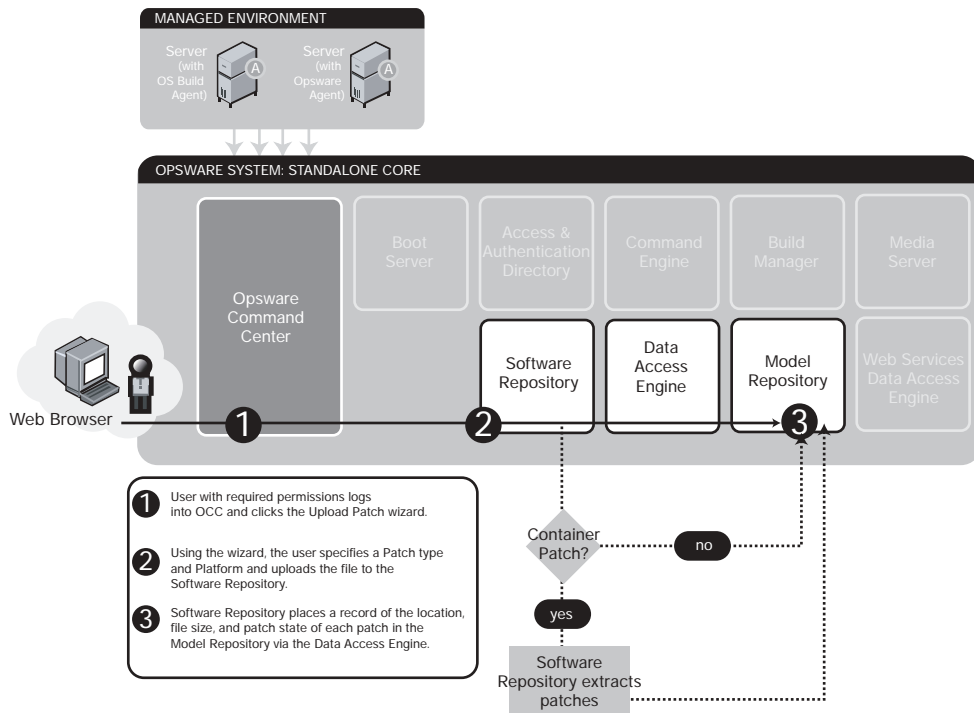
Figure 1-6: Patch Management Subsystem: Install Patch
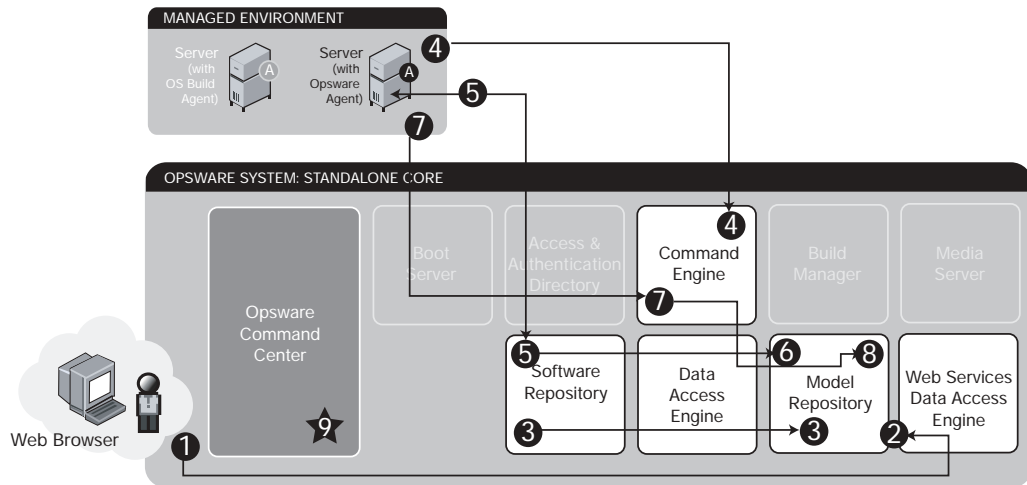
1 User with required permissions logs into OCC and clicks the Install Patch wizard.

2 Using the wizard, the user specifies patches and servers and starts the install process, retrieving patch information from the Model Repository via Web Services Data Access Engine.

3 Software Repository places a record of the location, file size, and patch state of each patch in the Model Repository via the Data Access Engine.

4 Command Engine gets list of installed software from Opsware Agent on managed servers. It compares it to the user-specified list of patches to determine what needs to be installed.

5 Opsware Agent on each managed server downloads patches from the Software Repository and installs them, performing all required install operations and reboots.

6 When installation is complete, a record of all currently-installed software is stored in the Model Repository via the Data Access Engine.

7 The Opsware Agent on each managed server reports installation status to the Command Engine.

8 Command Engine stores installation status in the Model Repository, via the Data Access Engine.

9 Operation complete status displays on the OCC.

Figure 1-7: Patch Management Subsystem: Uninstall Patch

1 User with required permissions logs into OCC and clicks the Uninstall Patch wizard.

2 Using the wizard, the user specifies servers and a patch to be uninstalled and starts the uninstall process, retrieving server and patch information from the Model Repository via Web Services Data Access Engine.

3 OCC passes Uninstall operation details to the Command Engine.

4 Command Engine gets list of installed software from Opsware Agent on managed servers. It compares it to the user-specified patch to be uninstalled and determines if it does need to be uninstalled.

5 Opsware Agent on each managed server removes the patch from the managed servers and performs all required uninstall operations and reboots.

6 When uninstallation is complete, a record of all currently-installed software is stored in the Model Repository via the Data Access Engine.

7 The Opsware Agent on each managed server reports uninstallation status to the Command Engine.

8 Command Engine stores uninstallation status in the Model Repository, via the Data Access Engine.

9 Operation complete status displays on the OCC.

Figure 1-8: Patch Management Subsystem: Microsoft Patch Update

MANAGED ENVIRONMENT

Server (with OS Build Agent) — Server (with Opsware Agent) — ❶ — Microsoft Patch Database — Recommended Patch List

OPSWARE SYSTEM: STANDALONE CORE

Opsware Command Center · Boot Server · Access & Authentication Directory · Command Engine · Build Manager · Media Server · Software Repository · Data Access Engine · Model Repository · Web Services Data Access Engine

Web Browser

❶ Every 24 hours, the Opsware Agent builds an inventory of software installed on the server. It uses that inventory and the Microsoft Patch Database to determine what Hot fixes and Service Packs are needed to bring the server up to current patch level. This is the Recommended Patch List.

❷ The Recommended Patch List and a full inventory of installed software is stored in the Model Repository via the Data Access Engine.

❸ User with required permissions logs into OCC and clicks the Microsoft Patch Wizard.

❹ Using the Wizard, the user selects the Windows servers, selects some or all of the patches on the Recommended Patch List, then clicks the Install button.

❺ The installation details are passed from the OCC to the Command Engine, which obtains a list of installed software from the Opsware Agent. It compares this list to the user-selected list and determines what actually needs to be installed.

❻ Opsware Agent on the managed server downloads patches from the Software Repository and installs them, performing all required install operations and reboots.

❼ When installation is complete, a record of all currently installed software is stored in the Model Repository via the Data Access Engine.

❽ Install operation status is reported to the Command Engine, which places it in the Model Repository via the Data Access Engine.

★ Operation complete status displays on the OCC.

Application Provisioning

In the Opsware System, software packages reside in a central Software Repository. Opsware administrators upload the software and also specify options that help ensure that the software is installed in a safe and consistent way. Administrators can add pre- and post-install and uninstall scripts to a software package that help control the way that the software is installed. The administrator can use Software Tree nodes to specify dependencies for package installation order. If a user tries to install software on a server that does not have the required prerequisite software, the user is alerted about the need to install the software and the installation operation is stopped on that server. (The installation process continues, however, if other selected servers do have the required software already installed on them.)

The Opsware System maintains detailed information about the state of every server under management in a central database called the Model Repository. This information includes details about software that is installed. You can use the information to check the rollout of software and also to help diagnose common server problems.

Information about the applications is consolidated into the centralized Model Repository. Provisioning a server with applications entails using automation to deliver software and configuration information about applications to any Opsware System-enabled server.

See Figure 1-9 and Figure 1-10. See the *Opsware System 4.6 User's Guide* for information about the application provisioning process.

Figure 1-9: Application Provisioning Step 1: Preview Reconcile

MANAGED ENVIRONMENT

Server (with OS Build Agent) — A

Server (with Opsware Agent) — A

OPSWARE SYSTEM: STANDALONE CORE

Boot Server

Access & Authentication Directory

Command Engine ❷

Build Manager

Media Server

Opsware Command Center

Software Repository

Data Access Engine ❹

Model Repository

Web Services Data Access Engine

Web Browser ❶

❸
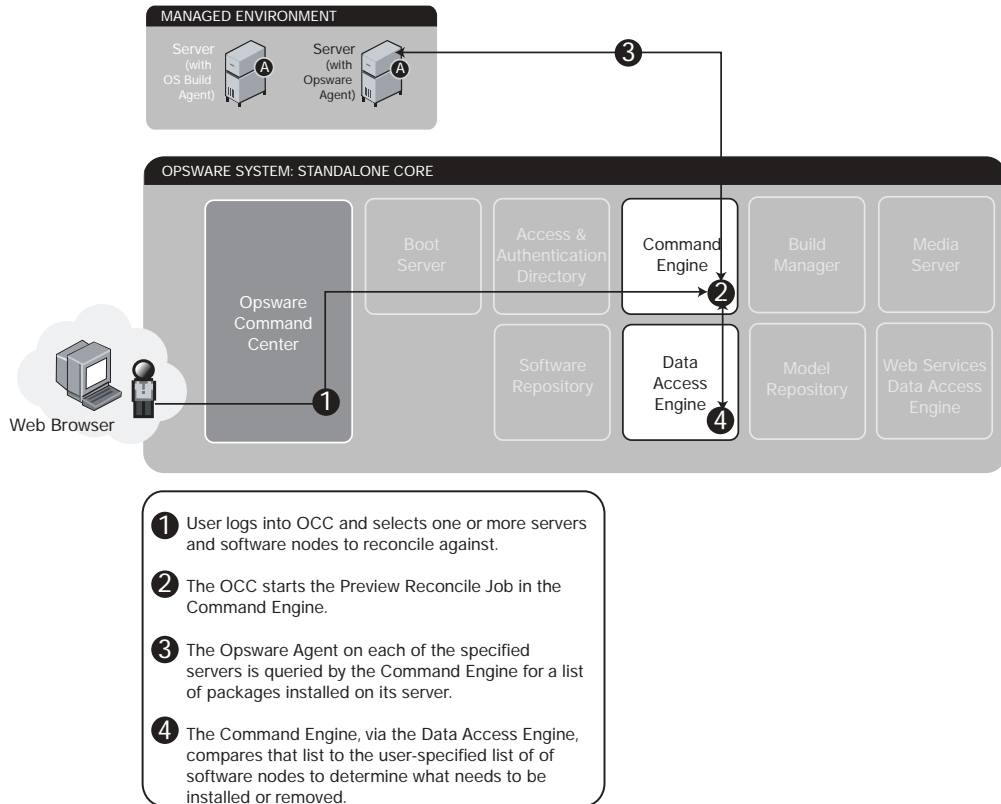
❶ User logs into OCC and selects one or more servers and software nodes to reconcile against.

❷ The OCC starts the Preview Reconcile Job in the Command Engine.

❸ The Opsware Agent on each of the specified servers is queried by the Command Engine for a list of packages installed on its server.

❹ The Command Engine, via the Data Access Engine, compares that list to the user-specified list of of software nodes to determine what needs to be installed or removed.

Figure 1-10: Application Provisioning Step 2: Software Installation and/or Removal Through Reconcile

5 At the end of Reconcile Preview, the OCC displays a list of the software to be installed and/or removed. The user confirms proceeding with the reconcile.

6 The OCC starts the Reconcile job in the Command Engine.

7 The Opsware Agent on each of the specified servers is queried by the Command Engine for a list of packages installed on its server.

8 The Command Engine, via the Data Access Engine, compares that list to the user-specified list of of software nodes to determine what needs to be installed or removed.

9 The Command Engine tells the Opsware Agent to install and/or remove packages.

10 The Opsware Agent downloads packages from the Software Repository, removes any packages that need to be removed, and installs the new packages, performing all necessary install, uninstall, and reboots, if required.

11 The Opsware Agent reports installation status to the Command Engine.

12 The Opsware Agent on each of the specified servers is queried by the Command Engine for a list of packages installed on its server to confirm what was installed and/or removed.

13 The Command Engine stores installation status in the Model Repository via the Data Access Engine.

14 Status of completed installation and removal of packages displays in the OCC via the Command Engine.

Code Deployment and Rollback Subsystem

Before you use Code Deployment and Rollback (CDR) to push code and content, you must upload new or updated files to your Opsware System staging environment. You can use Opsware System-supported content management tools, such as OpenDeploy, scp, or rsync over SSH, to do that.

After you upload the files and test your changes, you can synchronize updates to the production hosts that run your operational environment. You can run specific synchronizations and perform other service deployment operations by selecting CDR menu options available from the Opsware Command Center navigation panel. Figure 1-11 shows the code deployment and rollback subsystem process.

See the *Opsware System 4.6 User's Guide* for information about the process to deploy code and content to servers in the managed environment.

Figure 1-11: Code Deployment and Rollback Subsystem
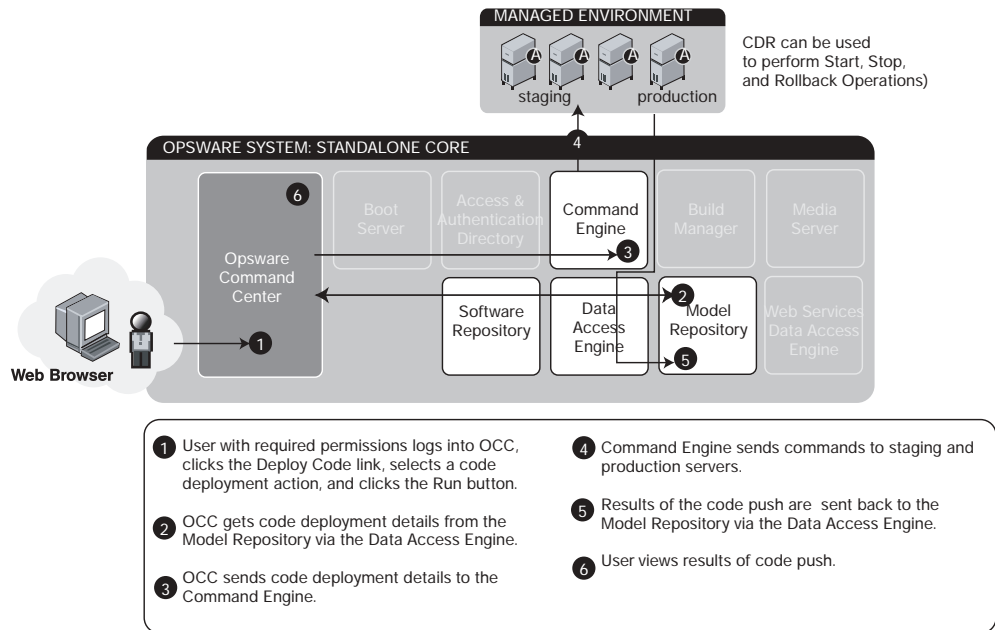
Script Execution Subsystem

The Script Execution Subsystem provides features and tools for automating the management and execution of server scripts. Previously, a user created a script and then manually executed the script at individual servers, one server after another. With the Script Execution Subsystem, a user performs all script tasks at one location-the Opsware Command Center.

From the Opsware Command Center, you can create or upload a script, set it up to run simultaneously across multiple Unix or Windows servers, and monitor it as it executes on each server. After a script runs, job- and server-specific execution results are available for review. You can modify, delete, or rerun a script at a later date. See Figure 1-12 and Figure 1-13.

See the *Opsware System 4.6 User's Guide* for information about the process to create and execute scripts in the managed environment.

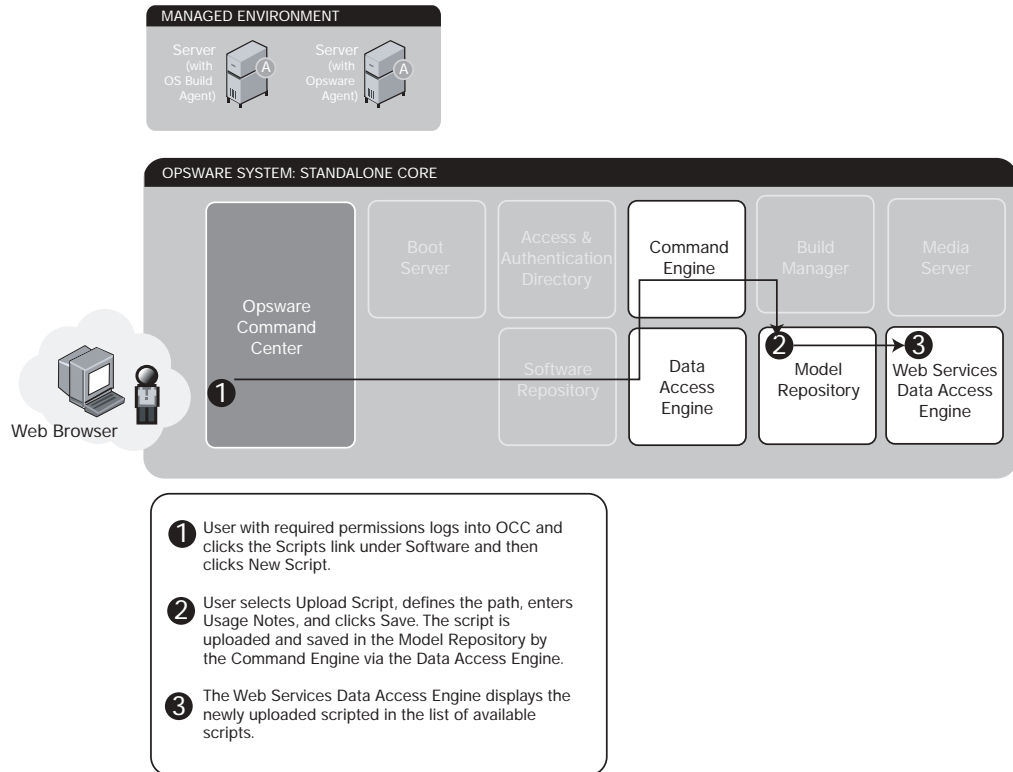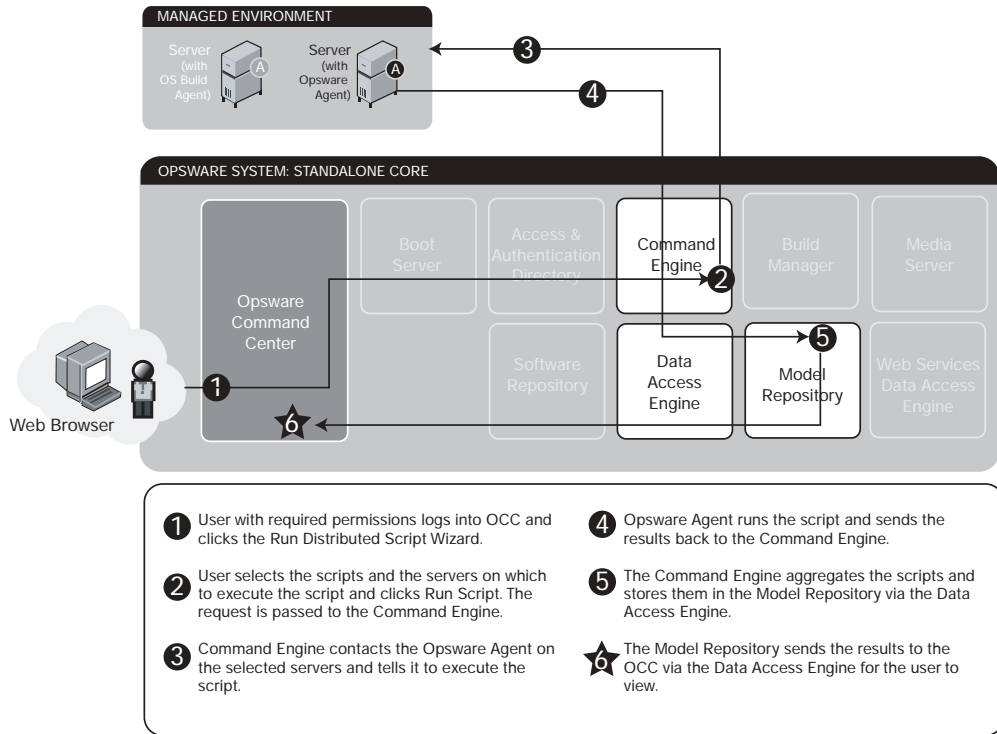*Figure 1-12: Scripting Subsystem: Upload Script*

❶ User with required permissions logs into OCC and clicks the Scripts link under Software and then clicks New Script.

❷ User selects Upload Script, defines the path, enters Usage Notes, and clicks Save. The script is uploaded and saved in the Model Repository by the Command Engine via the Data Access Engine.

❸ The Web Services Data Access Engine displays the newly uploaded scripted in the list of available scripts.

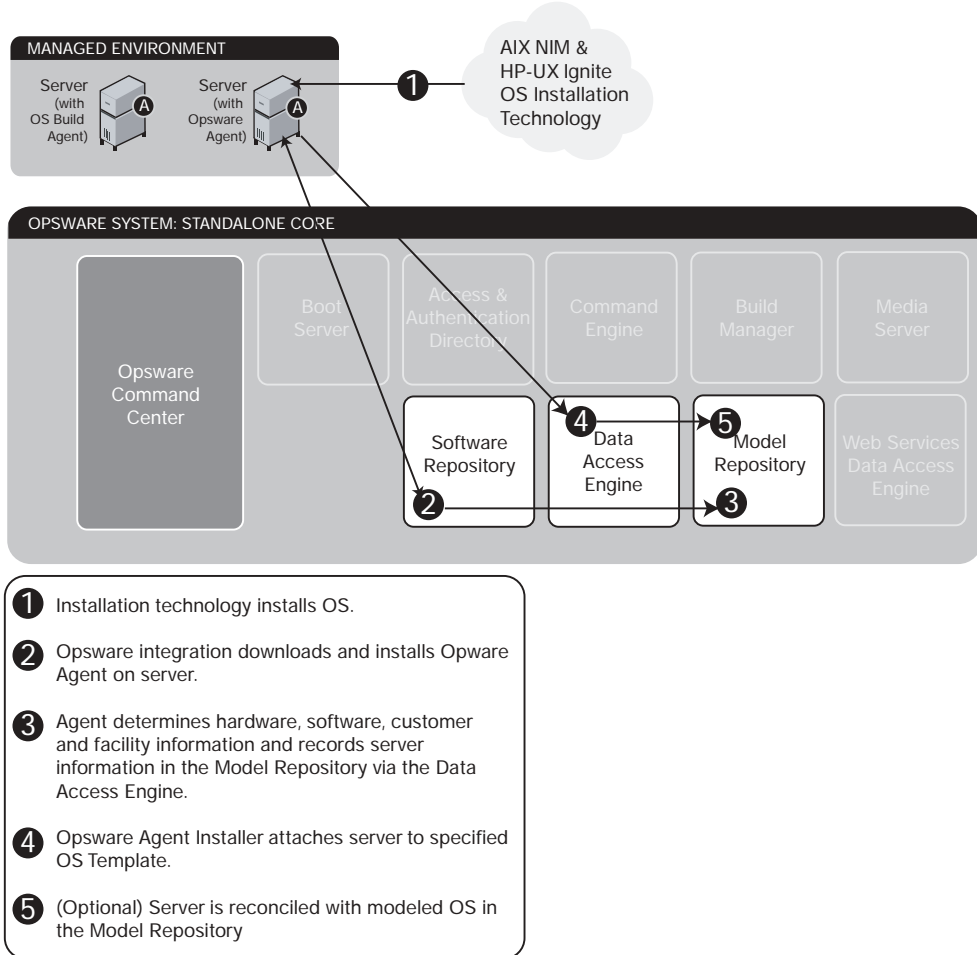*Figure 1-13: Scripting Subsystem: Execute Script*

Integration with AIX and HP-UX Installation Technology

Integrating the Opsware System with an OS installation technology enables installing an OS by using vendor utilities and automatically installing the Opsware Agent, which registers servers' initial configurations with the Model Repository.

See the *Opsware System 4.6 User's Guide* for information about how the Opsware System integrates with third-party OS installation technologies.

Figure 1-14 explains the interaction between the Opsware System components when the Opsware System is integrated with AIX and HP-UX OS installation technology. Opsware System installation integration with NIM and HP-UX occurs with the integration of the Opsware Installer. See the *Opsware System 4.6 User's Guide* for more information.

Figure 1-14: Opsware Integration with AIX and HP-UX OS Installation Technology

① Installation technology installs OS.

② Opsware integration downloads and installs Opware Agent on server.

③ Agent determines hardware, software, customer and facility information and records server information in the Model Repository via the Data Access Engine.

④ Opsware Agent Installer attaches server to specified OS Template.

⑤ (Optional) Server is reconciled with modeled OS in the Model Repository

Component Interaction in Multiple Facilities

Figure 1-15 and Figure 1-16 show how the Opsware System components interact when the Opsware System is running in multiple facilities. See "Opsware System Technology in Multiple Facilities" on page 69 in Chapter 2 for information about how to administer this Opsware System configuration.

Figure 1-15: Interaction Between Components in Multiple Facilities
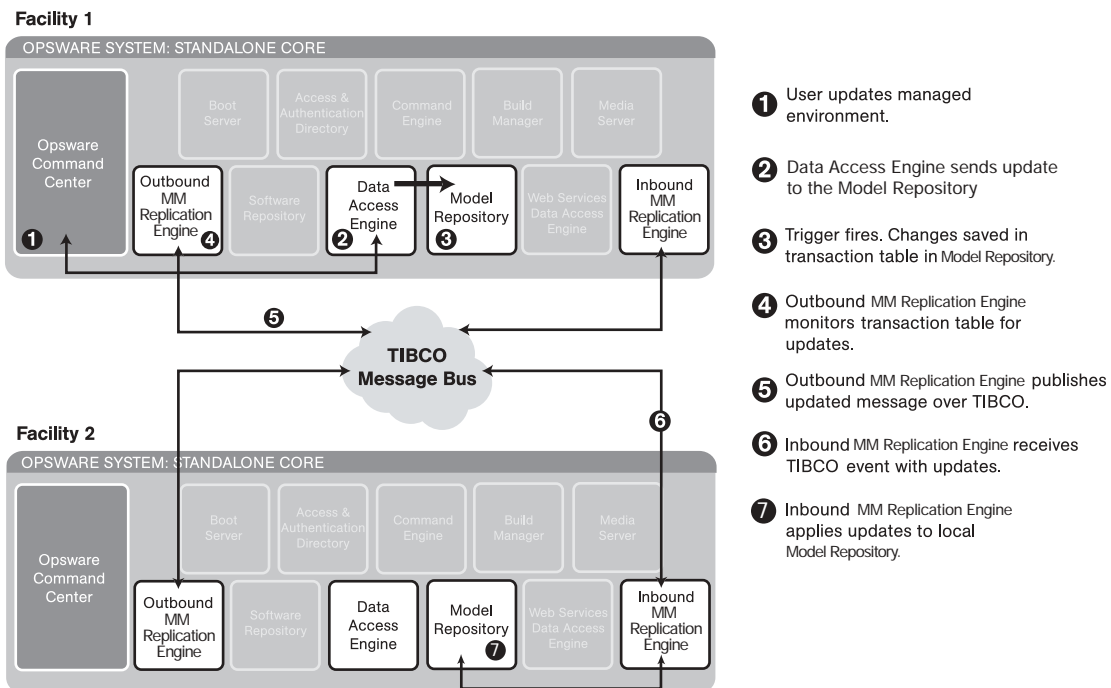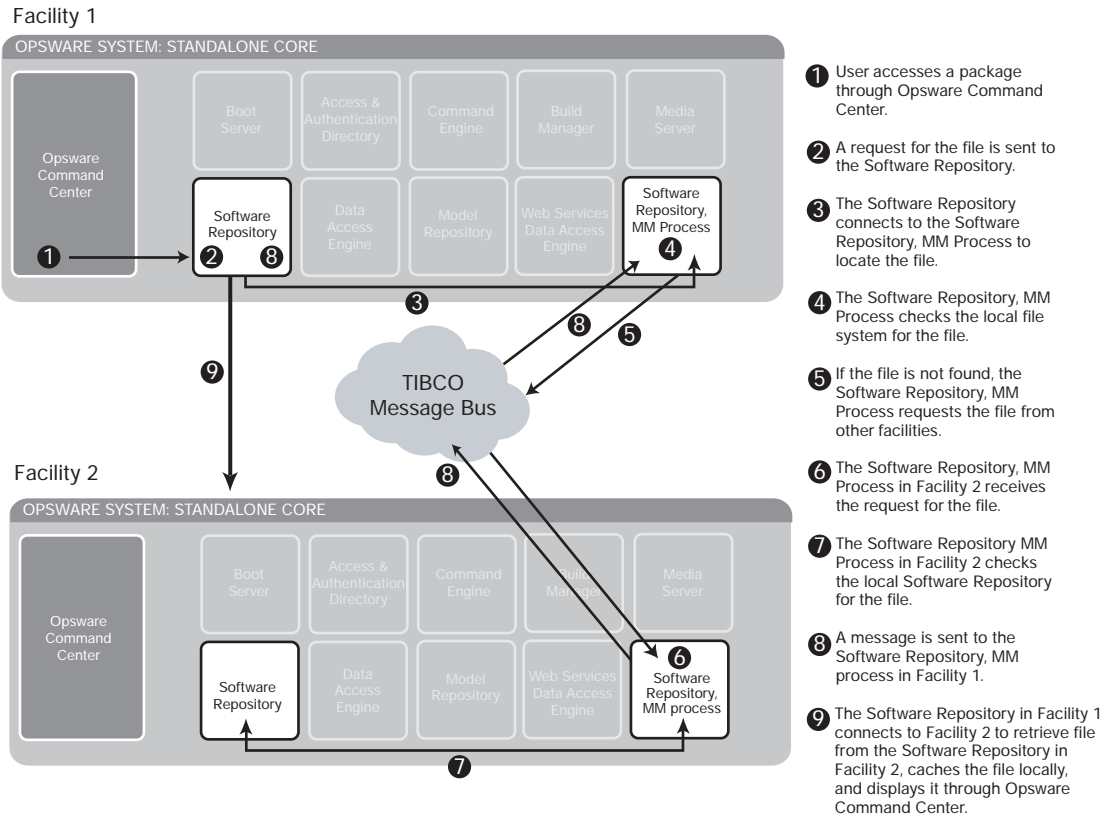
Data from Facility 1 to Facility 2

❶ User updates managed environment.

❷ Data Access Engine sends update to the Model Repository

❸ Trigger fires. Changes saved in transaction table in Model Repository.

❹ Outbound MM Replication Engine monitors transaction table for updates.

❺ Outbound MM Replication Engine publishes updated message over TIBCO.

❻ Inbound MM Replication Engine receives TIBCO event with updates.

❼ Inbound MM Replication Engine applies updates to local Model Repository.

Figure 1-16: Interaction of Software Repository in Multiple Facilities

Data from Facility 1 to Facility 2

❶ User accesses a package through Opsware Command Center.

❷ A request for the file is sent to the Software Repository.

❸ The Software Repository connects to the Software Repository, MM Process to locate the file.

❹ The Software Repository, MM Process checks the local file system for the file.

❺ If the file is not found, the Software Repository, MM Process requests the file from other facilities.

❻ The Software Repository, MM Process in Facility 2 receives the request for the file.

❼ The Software Repository MM Process in Facility 2 checks the local Software Repository for the file.

❽ A message is sent to the Software Repository, MM process in Facility 1.

❾ The Software Repository in Facility 1 connects to Facility 2 to retrieve file from the Software Repository in Facility 2, caches the file locally, and displays it through Opsware Command Center.

Chapter 2: Opsware System Administration

To perform the tasks in this chapter, you must be an Opsware administrator and log in to the Opsware Command Center by using your Opsware administrator username and password.

How Users Access the Opsware Command Center

The Opsware Command Center comes installed with an initial Opsware administrator user account called *admin* for your operational environment. The Opsware administrator then adds users and assigns them to user groups that include specific Opsware System permissions. For example, an Opsware administrator could create and add users to a user group so that they can update server configurations in the Opsware System environment or deploy code and content for the operational environment.

The Opsware administrator creates groups as described later in this chapter, and after the groups are created, the Opsware administrator assigns users to the groups. Each user can belong to one or more user groups that control the operations a user can perform and the data that the user can view. In addition to the groups that the Opsware administrator created, the system also comes with two sets of predefined groups: the Basic, Intermediate, and Advanced user groups, and the Code Deployment groups. In both cases, the Opsware administrator assigns users to the groups that the Opsware administrator created, and assigns users to the pre-defined groups found on the Code Deployment tab of the Users and Groups page. These Code Deployment user groups grant permissions that allow members to set up or perform specific code deployment operations such as service startup and shutdown, and code and content file synchronization performed on staging or production servers. Also, the Opsware administrator cannot customize existing Code Deployment groups and cannot create new ones.

See *Creating a Customer* on page 58 in this chapter for more information. This topic contains information about how to create and update *Customers*, which lets you associate configuration and administration responsibility for different business units and operational environments within your organization. By creating additional customers, you can also assign specific permissions for these Customers to individual user groups.

Opsware System users are authenticated with a login name and password stored in the Opsware Access & Authentication Directory.

The Access & Authentication Directory is the central repository for information about people and user groups. It is an LDAP directory that contains information about Opsware System users, the user groups they belong to, and the authority they possess.

For example, when a user accesses the Opsware Command Center, a login screen prompts the person for a username and password. When the user clicks the Sign in button, the Opsware Command Center connects to the Access & Authentication Directory and presents the username and password. If the username exists in the Access & Authentication Directory and the password is correct, the user is granted access to the Opsware Command Center. Additionally, the Opsware Command Center retrieves user profile information and provides access to the Opsware Command Center features as appropriate based on the individual user's permissions.

User Account Administration

The Opsware Command Center creates an initial Opsware administrator account (username *admin*), along with a temporary password that is defined during Opsware System installation, for the designated Opsware administrator of your organization. The Opsware administrator can then use the administration features of the Opsware Command Center to perform the following tasks:

• Creating a User

• Updating User Information

• Deleting a User

• Creating Opsware Administrator Accounts

Access permissions, based on user groups to which users are assigned membership, determine the Opsware System resources and the Opsware Command Center features that individual users can access.

Creating a User

An Opsware administrator can create a user and grant permissions by adding the user to the Opsware System or Code Deployment user groups.

Perform the following steps to create a user:

1 From the navigation panel, click Administration ➤ Users & Groups. The Users & Groups: View Users page appears, as Figure 2-1 shows.

By default, the Users tab on the View Users page appears, which shows a list of all current users.

Figure 2-1: List of Current Users Displayed on View Users Page

2 Click the New User button to display the New User page.

3 In the Identification section, enter the new user's first and last name, as Figure 2-2 shows. Required fields are in bold text.

Figure 2-2: Identification Section of New User Page

4 In the Contact Information section, enter the user's address, telephone number, and email address, as Figure 2-3 shows. The email address is a required field.

Figure 2-3: Contact Information Section of New User Page

5 Enter the new user's username and password in the Login information section, and confirm the password, as Figure 2-4 shows. These fields are all required.

Figure 2-4: Login Information Section of New User Page

6 In the User Preferences section, select a time zone for the user, and choose a long and short date format for the user, as Figure 2-5 shows.

This date information is used throughout the Opsware System wherever a date stamp is used. The time zones that are shown on the drop-down list might appear in long form or as an acronym.

Figure 2-5: User Preferences Section of New User Page

Individual users can later update their personal information and change passwords by clicking the My Profile link that appears at the top of every Opsware Command Center page.

Daylight savings time does not appear as an option in the time zones list. For example, Pacific Standard Time (PST) appears on the list, but Pacific Daylight Time (PDT) does not. Select the standard time for your location, and the system automatically displays Daylight time, if Daylight Savings Time is in effect when you make the selection.

7 In the User's Privileges section, select the Group Membership that you want to grant to the user, as Figure 2-6 shows. Group membership defines access permissions available to specific Opsware System resources and operations in the Opsware Command Center.

Figure 2-6: Group Membership Section of the New User Page

You do *not* assign access privileges to individual users. You assign privileges to a user group and then assign users to that group. If a user is assigned to more than one user group, that person can access each of the permissions specified in each user group (as well as the permissions provided by the union of the multiple user groups). See "User Groups and Opsware System Permissions" on page 44 in this chapter for more information.

CDR groups for the New User and Edit User screen only apply to the default customer, which was specified during installation.

8 Click the Save button.

A confirmation page informs you that the user was successfully added.

9 Click the Continue button.

The Opsware Command Center displays the updated View Users page.

Updating User Information

As an Opsware administrator, you can update information for user accounts. Users can also update their own information, excluding group memberships, in their profiles when they log in with their usernames and passwords. See the *Opsware System 4.6 User's Guide* for more information.

Perform the following steps to update user information:

1 From the navigation panel, click Administration ➤ Users & Groups. The View Users page appears, which shows a list of your existing user accounts.

2 Click the hyperlinked ID of the user that you want to update.

The Users & Groups: Edit User page appears, which shows the user's information. The User ID field cannot be changed, the Group Memberships cannot be changed, and the list of Permissions cannot be changed.

3 Modify the information for the user as necessary.

You cannot change the Group Membership and Permissions assigned to the user by editing a user's account information. You must add or remove the user from the User Groups. See "Adding Users to a User Group" on page 48 in this chapter for more information.

4 Click the Save button. A confirmation page informs you that the user was successfully updated.

Figure 2-7: Users & Groups: Confirmation Page

5 Click the Continue button to return to the Users & Groups: View Users page, as Figure 2-1 shows.
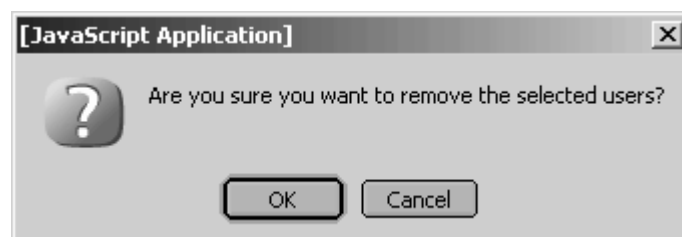
Deleting a User

Perform the following steps to delete a user:

1 From the navigation panel, click Administration ➤ Users & Groups. The Manage Users: View Users page appears.

By default, the Users tab displays. The Users & Groups: View Users page shows the users who can access the Opsware Command Center.

2 Select the check box next to each user who you want to delete.

3 Click the Delete button. A confirmation dialog box appears, as Figure 2-8 shows.

Figure 2-8: Delete User Confirmation Dialog Box

4 Click OK.

A confirmation page informs you that the user was successfully deleted, as Figure 2-9 shows.

Figure 2-9: Delete User Confirmation Page

Users & Groups: Confirmation

Confirmation
A user has been removed. User ID(s): avgjoe
Continue

5 Click the Continue button.

The Opsware Command Center displays the updated View Users page.

Creating Opsware Administrator Accounts

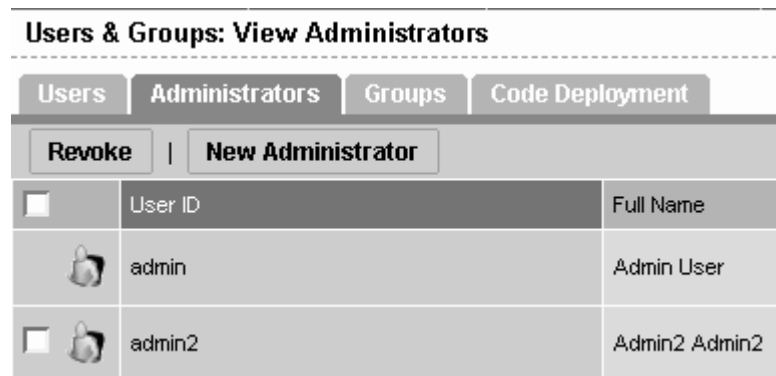After you create users, an Opsware administrator can optionally promote users to be Opsware administrators.

Perform the following steps to create Opsware Administrator accounts:

1 From the navigation panel, click Administration ➤ Users & Groups. The Manage Users: View Users page appears.

By default, the Users tab page displays.

2 Click the Administrators tab. The View Administrators page shows current Opsware administrators, as Figure 2-10 shows.

Figure 2-10: List of Current Administrators Displayed on View Administrators Page

Users & Groups: View Administrators

| Users | Administrators | Groups | Code Deployment |

| Revoke | New Administrator |

☐	User ID	Full Name
	admin	Admin User
☐	admin2	Admin2 Admin2

3 Click the New Administrator button. The Users & Groups: Add Administrators page appears, as Figure 2-11 shows.

Figure 2-11: Users & Groups: Add Administrators Page
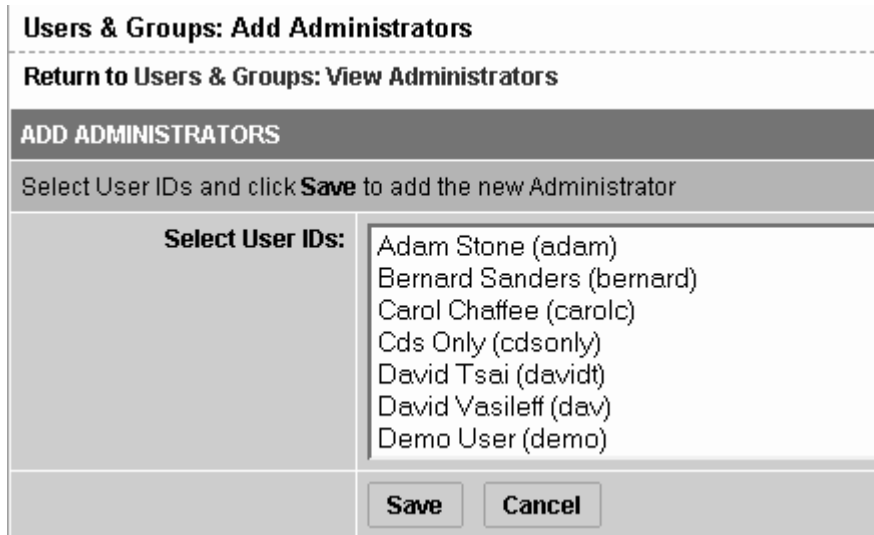
4 Select a user from the list.

5 Click the Save button.

The Opsware Command Center displays a confirmation message, as Figure 2-12 shows.

Figure 2-12: Add Administrator Confirmation Page

6 Click the Continue button.

The Opsware Command Center adds the user to the current list of Opsware administrators and displays an updated Users & Groups: View Administrators page.

User Groups and Opsware System Permissions

This section provides information on user groups and Opsware System permissions and contains the following topics:

- User Groups and Opsware System Permissions Overview

- Description of Permissions

- Creating User Groups

- Adding Users to a User Group

- Assigning Opsware System Permissions to a User Group

User Groups and Opsware System Permissions Overview

Opsware System permissions to access specific Opsware System resources and operations in the Opsware Command Center are assigned to user groups, not individual users. When you create a user group, you first specify the Opsware System permissions that the user group provides to members. Then, when you create a user, you select the user groups where that user should be a member.

When a user is assigned to more than one user group, that person can access each area of functionality specified in each user group (as well as the permissions provided by the union of the multiple user groups).

In the Opsware Command Center, you can define access control for the following functional areas:

- Features Permissions – provides access to specific Opsware Command Center features (available as menu selections in the navigation panel or as links to the functions in the Tasks section).

- Opsware System permissions – determines what the user's group can do after they have access to a feature, including running scripts as root, read/write access to customers and facilities, and the ability to lock specified functions.

- Code Deployment & Rollback (CDR) permissions – provides access to Code Deployment operations based on a user's group in deploying code and content changes. Permissions can be defined separately for deploying code and content changes to staging and production servers.

Description of Permissions

Permissions are defined in two areas of the system. The first area is Features Permission, which determines what areas of the Opsware Command Center are available to the user, and the second area is Opsware System Permissions, which determines what the user can do when accessing those features. The Opsware administrator cannot create new permissions.

Feature permissions that can be assigned include access to individual Opsware System Wizards, as well as configuration tracking, DNS, IP addresses, servers, applications, hardware, the Opsware System, service levels, operating systems, packages, patches, scripts, and templates.

The following list describes the different types of Opsware System permissions:

- General permissions – users in this group can deactivate servers. Users can deactivate servers in the Opsware Command Center; deactivating a server removes all its nodes and marks its status as Deactivated Server. Additionally, users in this group can edit shared scripts and run scripts as root.

- Customer permissions – allows you to specify which customers among those defined for your Opsware System installation are available to users in the associated user group. (See "Creating a Customer" on page 58 in this chapter for more information. This topic discusses creating customers for your operational environment.) In addition, these permissions specify whether users have read-only or read/write access to configuration settings and data for the customer.

- Facility permissions – lets you specify the facility or data center locations among those defined for your Opsware System installation that users in the associated user group have access to. In addition, these permissions specify whether users have read-only or read/write access to configuration settings and data for equipment installed at these locations.

  For example, if a user is in a user group that has None access selected for Customer A, Read access to Customer B, and Read & Write access to Customer C, then the user is unable to see any servers for Customer A, is able to see but not change servers for Customer B, and is able to see and change servers for Customer C. In addition, these access settings control visibility of associated objects in the Opsware Command Center as well as the ability to create, update, or delete related data.

- Locking permissions – provides the ability to lock and unlock the nodes associated with Application Servers, Database Servers, Service Levels, Web Servers, Other

Applications, OS Extras, and System Utilities. With locking rights, you can limit the modification rights for the specified permissions. Also, the lock icon is used to denote whether an object is locked as a result.

Creating User Groups

As an Opsware administrator, you can add users to any user group to make all functionality in the Opsware Command Center available to users, or you can define user groups that restrict access to specific areas of functionality for members of that group.

Perform the following steps to create a user group:

1 From the navigation panel, click Administration ➤ Users & Groups. The Users & Groups: View Users page appears.

By default, the Users tab displays. The View Users page shows the users who can access the Opsware Command Center.

2 Click the Groups tab. The Users & Groups: View Groups page appears that shows a list of all groups, as Figure 2-13 shows.

Figure 2-13: List of Current Groups Displayed on View Groups Page
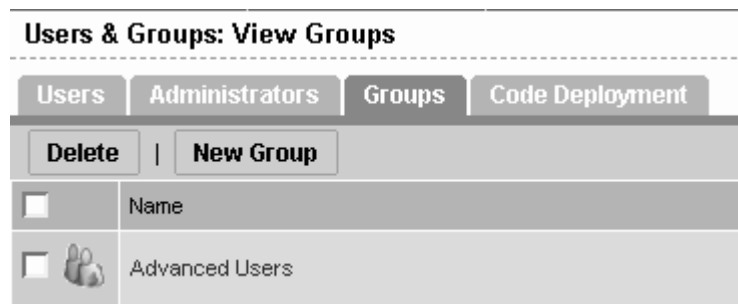
3 Click the New Group button. The Users & Groups: New Group page appears, which shows a form where you can define a user group, as Figure 2-14 shows.

Figure 2-14: Group Name Section of New Group Page

NEW GROUP

Enter the group's name and description, and select the appropriate channel permissions for group. Then click **Save** to create the new group

Group name:	
Group description:	

4 Enter the name of the user group. This field is required.

5 Optionally, enter a description for the user group.

6 Select the Feature permissions for the user group.

To edit the Feature permissions for the group after you create the group, click the hyperlinked group name to display the Users & Groups: Edit Group - [*group name*] page, and click the Features tab. The same list of features appears, which allows you to select or deselect features.

7 You assign General, Customer, Facility, and Locking permissions later, after the group is created.

8 Click the Save button. A confirmation page appears that indicates that the User Group was added, as Figure 2-15 shows.

Figure 2-15: Add Group Confirmation Page

Users & Groups: Confirmation

Confirmation

New user group is added.
Added user group Slightly Larger

Continue

9 Click the Continue button. The Users & Groups: View Groups page appears, which shows a list of all the groups, including the newly-added group.

To delete a user group, in the Users & Groups: View Groups page, select the check box next to the user group that you want to delete and click the Delete button. To edit a user group, click the hyperlinked name that corresponds to the user group that you want to edit. On the resulting page, Users & Groups: Edit Group - [*group name*], add or delete members of the user group.

Adding Users to a User Group

Permissions to access or perform specific Opsware System operations are granted by user groups that contain users, not by granting permissions to individual users. If a user is assigned to more than one user group, that person can access each of the resources specified in each user group (the union of the multiple user groups).

1 From the navigation panel, click Administration ➤ Users & Groups. The Users & Groups: View Users page appears.

2 Click the Groups tab. The Users & Groups: View Groups page appears, which shows a list of all groups.

3 Click the hyperlinked name that corresponds to the user group that you want to edit.

The Users & Groups: Edit Group - [*group name*] page appears with the User tab selected, which shows a list of users that you can choose from, as Figure 2-16 shows.

Figure 2-16: Edit Group User Assignment Page

4 In the Unassigned Members box, highlight the names of members whom you want to add to the group, and click the left-pointing arrow to move the names into the Assigned Members box.

5 When you finish selecting members, click the Save button. A confirmation page appears, as Figure 2-17 shows.

Figure 2-17: Edit Group Confirmation Page

6 Click the Continue button to return to the Assign Users page.

To remove an existing member from a user group, select the user's name in the Assigned Users list and click the right arrow.

Assigning Opsware System Permissions to a User Group

After you define a user group, you assign Opsware System permissions that determine the features and functionality that user group members have access to in the Opsware Command Center when they log in with their usernames and passwords. When users log in, only those features or areas of functionality that they have access to through their membership in user groups appears in the Opsware Command Center navigation panel.

Perform the following steps to assign Opsware System Permissions to a User Group:

1 From the navigation panel, click Administration ➤ Users & Groups. The Users & Groups: View Users page appears.

2 Click the Groups tab. The Users & Groups: View Groups page appears, which shows a list of all groups.

3 Click the hyperlinked name that corresponds to the user group that you want to edit.

4 Click the Other tab.

5 Select the General permissions items for the group by clicking the check box next to the permissions that you want to grant, as Figure 2-18 shows.

Figure 2-18: General Permissions Section of the Opsware Permissions Page

6 Select the Locking permissions for the group by clicking the check box next to the permissions that you want to grant, as Figure 2-19 shows.

Figure 2-19: Locking Permissions Section of Opsware Permissions Page

Locking Permissions

☑ Application Servers

☑ Other Applications

☑ Database Servers

☑ OS Extras

☑ Service Levels

☑ System Utilities

☑ Web Servers

7 Click the Save button. A confirmation message appears that confirms your permission change. Click Continue to return to the Permissions page.

8 Click the Customer tab.

9 Select the Customer permissions for the group by clicking the check box next to the permissions that you want to grant and select Read or Read & Write access for that customer, as Figure 2-20 shows.

Figure 2-20: Customer Permissions Section of Opsware Permissions Page

10 Click the Save button. A confirmation message appears that confirms your permission change. Click Continue to return the Permissions page.

11 Click the Facilities tab.

12 Select the Facility permissions for the group by clicking the check box next to the permissions that you want to grant and select Read or Read & Write access for that facility, as Figure 2-21 shows.

Figure 2-21: Facility Permissions Section of Opsware Permissions Page

13 Click the Save button. A confirmation message appears that confirms your permission change. Click Continue to return the Permissions page.

Code Deployment Permissions

Because Code Deployment & Rollback (CDR) is provided as an Opsware Command Center feature, it uses similar authentication to control users' access and ability to perform service operations, synchronizations, and sequences. Specific permissions to perform CDR operations are based on a user's membership in user groups predefined specifically for CDR.

The Opsware Command Center includes predefined user groups that have specific permissions to perform CDR operations. Opsware administrators create and add users to these user groups to grant them permissions to perform specific CDR operations, based on their role in an organization. When logged into the Opsware Command Center, a user sees only the services, synchronizations, and sequences that they have authorization to perform because of their user group membership. Users are assigned to these groups as part of the Create User process.

See the *Opsware System 4.6 User's Guide* for information about the process to deploy code and content to servers in the managed environment.

Some places in the interface might still show the legacy term CDS. However, all documentation references use the Opsware System 4.6 Code Deployment & Rollback term CDR.

Table 2-1 through Table 2-4 provide descriptions of the predefined Code Deployment user groups.

Table 2-1: Special Code Deployment User Groups

CODE DEPLOYMENT USER GROUP	DESCRIPTION
Super User	A user in this user group can define, request, or perform any code deployment operation on hosts designated for either staging or production.

Table 2-1: Special Code Deployment User Groups

CODE DEPLOYMENT USER GROUP	DESCRIPTION
History Viewer	A user in this user group can view a log of operations (service operations, synchronizations and sequences) that have been previously executed from the Code Deployment feature. Viewing this information can help you determine the status of particular deployment operations, and whether they completed successfully.

Table 2-2: Service User Groups

CODE DEPLOYMENT USER GROUP	DESCRIPTION
Service Editor	A user in this user group can define a service, and modify or delete service definitions.
Production Service Performer	A user in this user group can directly perform or request performance of service operations on hosts designated for use in production.
Staging Service Performer	A user in this user group can directly perform or request performance of service operations on hosts designated for use in staging.
Production Service Requester	A user in this user group can request performance of service operations on hosts designated for use in production.
Staging Service Requester	A user in this user group can request performance of service operations on hosts designated for use in staging.

Table 2-3: Synchronization User Groups

CODE DEPLOYMENT USER GROUP	DESCRIPTION
Synchronization Editor	A user in this user group can define a synchronization, and modify or delete the synchronization definition.

Table 2-3: Synchronization User Groups

CODE DEPLOYMENT USER GROUP	DESCRIPTION
Synchronization Performer	A user in this user group can directly perform or request performance of a synchronization action.
Synchronization Requester	A user in this user group can request performance of a synchronization action.

Table 2-4: Sequence User Groups

CODE DEPLOYMENT USER GROUP	DESCRIPTION
Sequence Editor	A user in this user group can define a sequence, and modify or delete the sequence definition.
Production Sequence Performer	A user in this user group can directly perform or request performance of a sequence of actions on hosts designated for use in production.
Staging Sequence Performer	A user in this user group can directly perform or request performance of a sequence of actions on hosts designated for use in staging.
Production Sequence Requester	A user in this user group can request performance of a sequence of actions on hosts designated for use in production.
Staging Sequence Requester	A user in this user group can request performance of a sequence of actions on hosts designated for use in staging.

When a user requests that a service operation, synchronization, or sequence be performed on their behalf, an email notification is sent to the individuals assigned to actually perform the requested service operation or synchronization.

Adding Members to a Code Deployment User Group

Permissions to perform specific Code Deployment operations are granted based on a user's membership in specific Code Deployment user groups. If a user is assigned to more than one user group, that person can access each of the resources specified in each user group (the union of the multiple User Groups).

1 From the navigation panel, click Administration ➤ Users & Groups. The Manage Users: View Users page appears.

2 Select the Code Deployment tab.

3 Select the code deployment user group that you want to modify by clicking the hyperlinked user group name.

The Users and Groups: Edit Code Deployment Group - [group name] page appears.

4 From the drop-down list, specify the customer whose group membership you want to modify.

5 Highlight the name in the left box of the user that you want to add to the group, then click the right-pointing arrow to put the name in the box on the right.

6 Click the Save button when you finish moving the user names to the box on the right.

A confirmation page appears.

7 Click the Continue button.

The Users & Groups: View Code Deployment Group page appears. You can continue modifying Code Deployment Groups, or you can select another function.

Customer Account Administration

This section provides information about customer account administration within the Opsware System and contains the following topics:

- Customer Account Administration Overview

- Creating a Customer

- Updating Customer Information and Settings

- Setting Custom Attributes

- Restrictions for Deleting Customers

- Deleting a Customer

Customer Account Administration Overview

As an Opsware administrator, you can add customers or update information and configuration settings for an existing customer.

See "Assigning Opsware System Permissions to a User Group" on page 50 in this chapter for more information. This topic discusses assigning specific administrative responsibilities to user groups for individual customers.

Creating a Customer

As an Opsware administrator, you can add customers to your Opsware System installation to create designations by business unit to provide management of Opsware System operations and configuration.

Perform the following steps to create a customer:

1 From the navigation panel, click Environment ➤ Customers.

The Customers page displays a list of all current customers, as Figure 2-22 shows.

Figure 2-22: List of Current Customers on Customers Page

2 Click the New Customer button.

The Customer: New Customer page appears where you can define the settings for a customer account, as Figure 2-23 shows.

Figure 2-23: Information Section of New Customer Page

3 In the Information section, enter a Name for the customer and a Short Name, or a nickname by which the customer name will display.

The short name is limited to 25 characters and must consist of uppercase letters, numbers, hyphens (-), and underscores (_).

4 If you have more than one facility, specify the associated facilities for the customer. To add a facility to the customer, select a facility from those displayed in the Available Facility list and click the left arrow button.

You only see the fields to add a facility for a customer when the Opsware System is running in multiple facilities. See "Opsware System Technology in Multiple Facilities" on page 69 in this chapter for more information.

5 When you finish defining the customer, click the Save button. A confirmation message appears that says that the customer was successfully created.

By default, no access permission of this customer is granted to any user groups. In order to permit users to access this customer, you must grant access permissions to the appropriate user groups.

6 Click the Continue button. The Manage Customer page appears.

You must log out and log in again to see the updated Customer list.

After you create a customer, you can define custom attributes for the customer. These custom attributes include setting miscellaneous parameters and named data values. Users can write scripts that use these parameters and data values when they perform a variety of functions, including network and server configuration, notifications, and CRON script configuration.

See "Setting Custom Attributes" on page 62 in this chapter for more information.

Updating Customer Information and Settings

As an Opsware administrator, you can update information or change configuration settings for existing customer accounts defined for your Opsware System installation.

Perform the following steps to update an existing customer:

1 From the navigation panel, click Environment ➤ Customers. The Customers page appears, which shows a list of your existing customer accounts.

2 Click the hyperlinked name of the customer whom you want to update.

The Customers: Edit Properties page appears, as Figure 2-24 shows.The list box on the left side includes facilities assigned to the customer. The list box on the right side includes all other Opsware System-managed facilities that are available to be added for a new or existing customer.

To change the name for the customer who appears in the Opsware Command Center, edit the name that appears in the Name field.

Figure 2-24: Assign Facility Section of the Customers: Edit Properties Page

3 To add a facility to the customer, select a facility from those displayed in the Available Facility list and click the left arrow button.

4 Click the Save button.

A message appears that confirms that you want to save the changes you just made, as Figure 2-25 shows.

Figure 2-25: Confirm Save Message

5 Click the OK button to save the changes.

The Customers page appears, which allows you to continue making changes to customer properties.

Setting Custom Attributes

You can use the Custom Attribute function to apply special properties to customers.

Perform the following steps to apply special properties to customers:

1 From the navigation panel, click Environment ➤ Customers. The Customers page appears, which shows a list of your existing customer accounts.

2 Click the hyperlinked name of the customer who you want to update.

The Customers: Edit Properties page appears.

3 Click the Custom Attributes tab.

The Customers: Edit Custom Attributes page appears. If custom attributes have previously been applied to this customer, they appear on this page.

4 Click the New button.

The Opsware Command Center displays the Customers: New Custom Attributes page, as Figure 2-26 shows.

Figure 2-26: New Custom Attributes Page

Customers: New Custom Attribute

Return to Edit Custom Attributes

Edit a Name and Value suitable for the new custom attribute.

Name:	
Value:	

These named values are used to provide parameters to the Opsware System, for example, to customize displays or provide settings to use during installation or configuration of packaged software in the operational environment.

Do not use either an asterisk (*) or a question mark (?) in the name field.

Be careful when you update or remove existing attribute settings as it might affect or disrupt operation of the operational environment. Contact your Opsware System Support Representative to help you determine the appropriate changes to make when you update the information or settings for a specific customer.

5 When you finish entering names and values, click the Save button, or exit without entering any values by clicking the Return to Customers link.

Restrictions for Deleting Customers

You can only delete a customer account from the Opsware Command Center when the following conditions are true for the customer:

• No Nodes are attached to the customer account.

• The customer account does not own any software packages; the packages uploaded for the customer must be deleted or deprecated.

- All servers assigned to the customer are deactivated.

- No IP Range Groups are created for the customer.

- No IP Ranges are created for the customer.

- No server groups are created for the customer.

If any of these restrictions apply to this customer, a message appears and you cannot delete the customer.

If the restrictions do not apply, the user is prompted to move deactivated servers to the Not Assigned customer account.

Deleting a Customer

Deleting a customer removes the customer information from the Opsware Access & Authentication Directory, and moves deactivated servers assigned to the customer to the Not Assigned customer account or deletes the data about the servers from the Model Repository database.

See "Restrictions for Deleting Customers" on page 63 in this chapter for more information.

Perform the following steps to delete a customer:

1 From the navigation panel, click Environment ➤ Customers. The Customers page appears, which shows a list of your existing customer accounts.

2 Select the check box next to the customer whom you want to delete.

3 Click the Delete button.

A confirmation page appears, which verifies that the selected customer will be deleted, as Figure 2-27 shows.

Figure 2-27: Customers: Delete Confirmation Page

4 Click the Delete button.

A confirmation page appears, as Figure 2-28 shows.

Figure 2-28: Delete Customer: Confirmation Page

5 Click the Continue button.

The Opsware Command Center displays the updated Customers page.

Facilities Administration

This section provides information on facilities administration within the Opsware System and contains the following topics:

• Facilities Administration Overview

• Creating a Facility

• Updating Facility Information and Settings

• Deactivating a Facility in the Opsware Command Center

Facilities Administration Overview

In the Opsware Command Center, a facility refers to the collection of servers and server groups that are managed by a single Model Repository, the database that stores information about the managed environment. Users can manage servers in any facility from the Opsware Command Center in any facility. When a user updates data in a facility, the Model Repository for that facility is synchronized with the Model Repository databases located in all remote facilities.

In the Opsware Command Center, a facility is identified by a facility name and a facility ID. The facility name displays in the Opsware Command Center.

Creating a Facility

You can use the Opsware Command Center to create a facility by selecting Environment ➤ Facilities from the navigation panel. Creating a facility in the Opsware Command Center only *adds* the facility to the Opsware Command Center UI. Performing this procedure does *not* install the Opsware System in that facility.

Creating a facility in the Opsware Command Center is one part of a multistep process for installing and setting up a facility from which you can run the Opsware System.

This guide does *not* document how to install the Opsware System to run in a facility or how to create a facility. Consult your Opsware System Support Representative or see the *Opsware System 4.6 Installation Guide* for this information.

After you create a facility, an Opsware administrator must grant read or read/write permissions for all users who need to manage servers in the facility. Users (including the Opsware administrator who created the facility) do not see the facility in the Opsware Command Center until they have permissions for the facility. See "User Groups and Opsware System Permissions" on page 44 in this chapter for more information.

After you create a facility in the Opsware Command Center, you can define custom attributes for the facility. These custom attributes include setting miscellaneous parameters and named data values. Users can write scripts that use these parameters and data values when they perform a variety of functions, including network and server configuration, notifications, and CRON script configuration.

See "Updating Facility Information and Settings" on page 66 in this chapter for more information.

Updating Facility Information and Settings

Perform the following steps to update facility information and settings:

1 From the navigation panel, click Environment ➤ Facilities. The Facilities page appears and displays the names of the current facilities.

2 Click the hyperlinked name of the facility that you want to update. The Facilities: Edit Facility page appears with the Properties tab automatically selected, as Figure 2-29 shows.

Figure 2-29: Properties Tab of the Edit Facility Page

3 To change the name of the facility that appears in the Opsware Command Center, edit the Name field or click the Return to Facilities link to exit without making any changes.

Contact your Opsware System Support Representative if you need to make other changes to the facility properties.

4 Click the Save button.

The Opsware Command Center displays a message that confirms that the properties for the facility were updated.

5 Click the Custom Attributes tab.

The Custom Attributes page appears, which provides name-value pairs associated with this customer. These named values are used to provide parameters to the Opsware System, for example, to customize displays or provide settings to use during installation or configuration of packaged software in the operational environment.

6 Click the hyperlinked name of an attribute to display the Facilities: Edit Attribute for [*facility name*] and make changes to its associated value.

7 To add an attribute name and to specify a value to associate with the attribute, click the New button.

Be careful when you update or remove existing attribute settings as it might affect or disrupt operation of the operational environment. Contact your Opsware System Support Representative to help you determine the appropriate changes to make when you update the information or settings for a specific facility.

8 When you finish making updates to the facility properties or custom attributes, click the Return to Facilities link.

Deactivating a Facility in the Opsware Command Center

Performing this procedure does not shut down or uninstall the Opsware System in a facility. Performing this procedure *only* stops Model Repository updates from propagating to the deactivated facility and stops the Multimaster Tools from functioning for that facility. Deactivate facilities with care because this task cannot be undone.

When you deactivate a facility, the facility is still listed in the Opsware Command Center, however, it is marked with a D. After a short name is used, even if it is deactivated, that name cannot be reused.

Perform the following steps to deactivate a facility:

1 From the navigation panel, click Environment ➤ Facilities. The Facilities page appears, which shows a list of the existing facilities that are defined in the Opsware Command Center.

2 Select the check box next to the facility that you want to delete.

3 Click the Delete button. A confirmation dialog box appears.

4 Click OK.

The updated Facilities: Attributes for [*facility name*] page appears.

Opsware System Technology in Multiple Facilities

This guide does not document how to set up the Opsware System to run in multiple facilities (a multimaster configuration). Consult your Opsware System Support Representative or see the *Opsware System 4.6 Installation Guide* for information about how to set up the Opsware System to run in multiple facilities.

The managed environment might span several facilities when a distributed facility architecture is required. Distributed facility architecture can provide better service and fault tolerance, or a more efficient network design. For example, one facility might be dedicated to an organization's Internet, while another facility might be dedicated to the Web services offered to the public.

Running the Opsware System from multiple facilities has the following advantages:

• Redundancy - If a facility goes down, an Opsware Command Center is still usable from other facilities. Users in remote locations can have their own Opsware Command Center. Also, it provides the ability to move out of a facility and keep the Opsware System running in other facilities.

• Performance scalability – Write operations do not need to be proxied to a central location.

• Geographic scaling - International facilities can be independent and do not need to rely on a network connection across continents to a central facility.

Multimaster Mesh Administration

This section provides information on multimaster mesh administration within the Opsware System and contains the following topics:

• Multimaster Mesh Administration Overview

• Software Repository, Multimaster Component Conflicts

• Causes of Conflicts

• User Overlap

• User Duplication of Actions

• Connectivity Problems That Cause Out of Order Transactions

Multimaster Mesh Administration Overview

Customers who need to manage more than one facility should run the Opsware System in multimaster mode. Multimaster is a configuration for synchronizing copies of the Model Repository database located in different facilities. Any Model Repository database in any facility can be used as the updateable source at any time. In the multimaster architecture, there is no designated master for any individual data element.

Running the Opsware System in multimaster mode has the following characteristics:

• Model Repository databases are geographically dispersed.

• Each facility is independent of other facilities.

• Each Model Repository database is available for both read and write transactions.

• Data is updated locally and then propagated to every facility in the multimaster architecture.

• The multimaster mesh is virtually invisible to operations personnel.

Operating in a multimaster architecture involves the chance of conflicting updates being made to the same record in different Model Repository databases.

The multimaster Opsware System components detect conflicts and propagate alerts; however, the multimaster components do not resolve conflicts. Opsware administrators use the multimaster tools in the Opsware Command Center to resolve the conflicts at the target databases.

In a configuration with multiple facilities, one facility, known as multimaster central, is designated as the primary facility. It is responsible for generating the transaction table automatically, although the other facilities can also do so upon demand. Multimaster central was defined during installation. See the *Opsware System 4.6 Installation Guide* for more information.

Software Repository, Multimaster Component Conflicts

When an update from the source database arrives at the destination database, a conflict can be generated any time the data at the destination database is not what was expected—either the values are different or the row cannot be found.

The probability of multimaster conflicts occurring varies depending on the following factors:

• The number of servers under management

• The number of facilities

• The number of Opsware Command Centers

• The propensity for users to make changes in more than one facility by using different Opsware Command Centers

Data conflicts occur when the values of the objects in the local Model Repository do not match the values in a message from the Outbound Model Repository, Multimaster Component or a database constraint is violated. When a conflict is flagged, the Opsware System takes the following actions:

1 The transaction is canceled.

2 All rows affected by the transaction are locked, thereby preventing further changes to those rows.

3 The Outbound Model Repository, Multimaster Component propagates this change in a new transaction to all remote databases, thereby locking the rows in all facilities.

4 An alert message with the conflict information is emailed to the configured mailing list.

5 The Inbound Model Repository, Multimaster Component continues on to the next message.

If the Inbound or Outbound Model Repository, Multimaster Component encounters an exception that prevents it from going on to the next message, it sends an email and shuts itself down.

An Opsware administrator must manually resolve the problem by using the Opsware Command Center. Resolving the conflict unlocks the rows. See "Best Practices for Resolving Database Conflicts" on page 76 in this chapter for more information.

Causes of Conflicts

Conflicts can have the following causes:

- User overlap

- User duplication of actions

- Connectivity problems cause out of order transactions

User Overlap

Multiple users are working in the same area of data by using the Opsware Command Centers in different facilities. Conflicts occur when a user makes a change by using the Opsware Command Center in one facility and another user makes a change to the same object using the Opsware Command Center in another facility.

Partitioning the data space helps to reduce the number of conflicts that user overlap causes.

For example, this sequence of events occurs:

1 Alice removes Node A from a server in the Atlanta facility.

2 Bob removes Node A from the same server in the Boston facility.

3 The Opsware System propagates the change from the Atlanta facility to the Boston facility; however, the node has already been removed from the server in the Boston facility. The Opsware System generates a Model Repository, Multimaster Component conflict.

4 The Opsware System propagates the change from the Boston facility to the Atlanta facility; however, the node has already been removed from the server in the Atlanta facility. The Opsware System generates a second Model Repository, Multimaster Component conflict.

User Duplication of Actions

Conflicts occur when a user makes a change in one database, does not see the change reflected in another database, and makes the change again in the other database.

This situation involves a user bouncing back and forth between multiple Opsware Command Centers, or between an Opsware Command Center and some command line utilities in a facility.

For example, this sequence of events occurs:

1 From a server in the Seattle facility, Carol uses the Opsware Command Line Interface (OCLI) to upload the package `carol.conf`.

2 In the Phoenix facility, Carol logs into the Opsware Command Center to search for the package. She does not see the package because that data has not yet propagated from Seattle to Phoenix. Carol is unaware of the lag time for data propagation between facilities.

3 Carol uploads the package `carol.conf` by using the Opsware Command Center in Phoenix.

When the data arrives from Seattle, the Opsware System generates a conflict because the data already exists in Phoenix.

Connectivity Problems That Cause Out of Order Transactions

This situation causes conflicts when a user changes or inserts data at facility A (Model Repository database A). The transaction for that change propagates to facility B (Model Repository database B). The same data is modified again or somehow referenced at facility B (Model Repository database B). The transaction from facility B reaches facility C (Model Repository database C) before the transaction from facility A.

Transactions sent from a facility to another facility arrive in the order in which they were sent. However, the correct ordering is not guaranteed for transactions arriving from different facilities.

This type of conflict occurs only when the Opsware System is running from three or more facilities.

A common cause of this situation is a user uploading a package by using the OCLI, and then immediately adding the package to a node by using the Opsware Command Center in another facility. The delay in propagating data about the package to other facilities causes the data about the node attachments to arrive at other facilities out of order.

The occurrence of out of order transactions is aggravated by proximate updates in different facilities and unreliable inter-facility network connections.

For example, this sequence of events occurs:

1 From a server in the Denver facility, Henry uses the OCLI to upload the package `henry.conf`.

2 The Opsware System propagates data about the package to the Miami facility; however, it cannot propagate the data to the Paris facility because the network connection to the facility is down.

3 Henry updates the description of the package `henry.conf` by using the Opsware Command Center in Miami.

4 The Opsware System propagates data about the updated package description to the Denver facility; however, it cannot propagate the data to the Paris facility because the network connection to the facility is down.

5 Network connectivity to the Paris facility is restored and multimaster messages are propagated to the Paris facility.

6 The message about the updated package description arrives at the Paris facility before the message about the uploaded package. The Model Repository in the Paris facility does not contain data about the package, so a conflict is generated.

7 The message about the uploaded package arrives at the Paris facility and is processed without error. The package data exists in Paris but the package description differs from the other facilities.

Best Practices for Preventing Multimaster Conflicts

When you use the Opsware System in multiple facilities, try to keep the number of conflicts that can occur to a minimum. Educate users to consider the following factors when the Opsware System is running in multimaster mode:

• Users in multiple facilities are able to modify the same data at the same time.

• A slight time delay occurs before changes that a user makes arrive in other Opsware System facilities. (The length of delay varies depending on a number of factors, including network connectivity and bandwidth.)

Implement these best practices to reduce the chance of data conflicts between facilities:

• Ensure reliable network connections and sufficient network bandwidth between facilities. The risk of conflicts increases with degraded network connectivity between facilities.

  See "Network Administration for Multimaster" on page 85 in this chapter for more information.

For additional assistance, consult your Opsware System Support Representative or see the *Opsware System 4.6 Installation Guide* for information about network connectivity when running the Opsware System with a multimaster mesh.

• Educate users not to change data in one facility and then make the same change in another facility.

• Partition the data space so that more than one user does *not* change the same object in different facilities at the same time.

Have a user or a small group of coordinated users manage a given set of servers. Partitioning the data space ensures accountability of server ownership and prevents users from changing each other's data.

The Opsware System includes a mechanism for distributed access to data. Specifically, the Opsware Command Center includes permissions by customer, facility, and User Group Types.

See "User Groups and Opsware System Permissions" on page 44 in this chapter for more information.

Examining the State of the Multimaster Mesh

You can examine the state of the multimaster mesh by clicking the Multimaster Tools option, which is visible in the Opsware Command Center at all multiple facility installations.

When you select the Multimaster Tools option, the Multimaster Tools: State View page appears. In addition to a color-coded legend that shows possible transaction states (including red for Conflict, orange for Not Sent, yellow for Not Received, Gray for Unable to Connect, and green for Good), this page also:

• Presents an overview of the health of the multimaster mesh by automatically checking all facilities

• Shows the state of the last five transactions — a unit of change to a database that consists of one or more updates to rows and has a globally unique transaction ID — from each facility to each other facility, plus all conflicting transactions and all unpublished transactions

• Shows the time that the Opsware Command Center generated and cached the data. Click the Refresh button to refresh that cached data.

• Links each conflicting transaction ID to the Transaction Differences page for that transaction

Opsware administrators can also use the System Diagnosis tools in the Opsware Command Center to view information about the health of the multimaster components.

See "Opsware System Diagnosis" on page 105 in this chapter for more information.

Best Practices for Resolving Database Conflicts

Maintaining data consistency is complex and conflicts can occur even when implementation and work processes minimize them. This section contains the following topics:

• Types of Conflicts

• Guidelines for Resolving Each Type of Conflict

Types of Conflicts

The following types of conflicts can occur:

• Identical data conflict – the Multimaster Tools show a conflicting transaction but the data is the same between facilities. The data is the same because users made the same change in different facilities.

• Simple transaction conflict – the row exists in all facilities, but some columns have different values or the row does *not* exist in some facilities (missing objects).

• Unique-key constraint conflict – the object does not exist in a facility and cannot be inserted there because inserting it would violate a unique-key constraint.

• Foreign-key constraint conflict – the row does not exist in some facilities and cannot be inserted because the data contains a foreign key to another object that also does not exist in that facility.

In rare cases, you might encounter an additional type of conflict-a linked object conflict. The Opsware System includes business logic that links specific related objects in the Opsware System, such as a custom attribute name and value, and a customer created in the Opsware Command Center UI (appears in lists) and the associated Node for the customer in the node hierarchy. The Opsware System ensures that links between related objects are maintained. Resolving a linked object conflict can be complex because you

must attempt to preserve the intent of the transaction that caused the conflict. Contact your Opsware System Support Representative to help you resolve linked object conflicts.

Guidelines for Resolving Each Type of Conflict

In general, when you resolve conflicts, apply updates so that the target always reflects the most current data based on the time stamp of the originating changes.

When you cannot follow one of the preceding guidelines, attempt to preserve the intent of the transaction. Contact the users who are generating the transactions and determine what types of changes in the managed environment each user was trying to make.

Identical Data Conflict

All objects in a transaction contain exactly the same data across all facilities. This type of conflict includes the case where the objects do not exist in all facilities.

To resolve an identical data conflict, simply mark the conflict resolved.

Identical Data Conflict (Locked)

All objects in a transaction contain exactly the same data across all facilities but the objects in the transaction are still locked (marked conflicting).

To resolve this type of conflict, pick an arbitrary facility and synchronize all objects from it. Performing this action unlocks the objects. After synchronizing the data, mark the conflict resolved.

Simple Transaction Conflict

The data is different between facilities or some objects are missing from some facilities. None of the objects depends on the actions of other conflicting transactions. The results of synchronizing the objects does not result in a database foreign-key or unique-key constraint violation.

To resolve a simple transaction conflict, choose the facility that contains the correct data and synchronize from it. How you determine which facility contains the correct data varies depending on the type of transaction:

• If the conflict is the result of two users overriding each other's work, talk to the users and determine which user's change should be correct.

• If the conflict is the result of automated processes overriding each other's data, the most recent change is usually correct.

• If the conflict is the result of out-of-order transactions, the most recent change is usually correct.

After synchronizing the data, mark the conflict resolved.

Unique-Key Constraint Conflict

Resolving these conflicts results in a unique-key constraint violation.

For example, this sequence of events occurs:

1 From the Opsware Command Center in the London facility, John creates Node A1 as a subordinate node of Node A.

2 From the Opsware Command Center in the San Francisco facility, Ann performs the same action. She creates Node A1 as a subordinate node of Node A.

3 Node names must be unique in each branch of the node hierarchy.

4 The Opsware System propagates the node changes from the London and San Francisco facilities to the other facilities. Inserting the rows into the Model Repository databases at other facilities causes a unique-key constraint violation and a conflict.

Resolving this conflict by inserting the updates from the London facility in all facilities would fail with the same unique-key constraint violation.

Perform the following steps to resolve a unique-key constraint conflict:

1 Locate all the involved transactions and synchronize one transaction from a facility where the object does not exist, thereby deleting it in all facilities.

2 Synchronize the other transaction from a facility where the object exists, thereby inserting the object in all facilities. One of the two uniquely conflicting objects will take the place of the other.

Foreign-Key Constraint Conflict

Resolving these conflicts results in a foreign-key constraint violation.

For example, this sequence of events occurs:

1 Jerry creates Node B in facility 1.

2 Before that transaction has time to propagate to other facilities, Jerry creates Node C as a subordinate node of Node B.

3 When the first transaction arrives at facility 2, it generates a conflict for unrelated reasons.

4 When the second transaction arrives at facility 2, inserting the row for Node C causes a foreign-key constraint conflict because the parent Node (Node B) does not exist.

Resolving the second conflict first by inserting the update for Node C into all facilities would fail with the same foreign-key constraint violation.

Perform the following steps to resolve a foreign-key constraint conflict:

1 Resolve the conflicting transaction for Node B (the parent Node) by synchronizing the first transaction from the facility where the object exists.

2 Synchronize the second transaction (the Node C update) from the facility where the object exists.

Generally, resolving conflicts in the order in which they were created avoids generating foreign-key constraint conflicts.

Resolving Model Repository, Multimaster Component Conflicts

This section provides information on resolving model repository, multimaster component conflicts and contains the following topics:

• Resolving Model Repository, Multimaster Component Conflicts Overview

• Resolving a Conflict by Object

• Resolving a Conflict by Transaction

Resolving Model Repository, Multimaster Component Conflicts Overview

Opsware administrators can view and resolve multimaster conflicts in any Opsware Command Center by using the Multimaster Tools. The Multimaster Tools are available in all Opsware Command Centers.

Before you resolve conflicts, notify the subscribers of the email alert alias. Notifying these users helps to prevent other Opsware administrators from undoing or affecting each other's conflict resolution efforts.

If you see a large volume of conflicts that you cannot resolve by using the Multimaster Tools, contact your Opsware System Support Representative for assistance synchronizing databases.

Resolving a Conflict by Object

When you resolve a conflicting transaction that contains multiple objects, you must resolve the conflicts for all objects in the transaction. Generally, you resolve the conflicting objects in the order that they appear in the transaction. To undo the effects of a transaction, resolve the objects in the reverse order that they appear in the transaction (last to first).

Perform the following steps to resolve conflicting transactions:

1 From the navigation panel, click Administration ➤ Multimaster Tools. The Multimaster Tools: State View page appears. By default the State View tab displays. Each conflicting transaction ID links to the Transaction Differences page for that transaction. See Figure 2-30.

Figure 2-30: Transaction Table That Shows Conflicts

Different types of conflicts are indicated by the color of the box next to the Transaction ID. A key that explains the significance of the colors, like the one that Figure 2-31 shows, is at the top of the page.

Figure 2-31: Conflict Color Key

Key				
Problem	Potential Problem			Good
■ Conflict	▣ Not Sent	☐ Not Received	☐ Unable To Connect	■ Received

2 From the State View page or the Conflict View page, click the hyperlinked transaction ID for the conflicting transaction that you want to resolve.

The Multimaster Tools: Transaction Differences page appears, as Figure 2-32 shows.

Figure 2-32: Transaction Differences Page That Shows Partial Details of Transaction Conflict

Multimaster Tools: Transaction Differences | 298600003

Return to State View

Synchronize all objects from [C01 ▾] **Update**

Account 420003

DB Field	C01	C06
ACCT_ID	420003	Not Found
ACCT_NAME	TESTER	
AUTH_DOMAIN	TESTER.COM	
BUSINESS_ACCT_ID	-1	
CONFLICTING	0	
DISPLAY_NAME	TEST	
INTERNAL_ACCT_FLG	0	
STATUS	ACTIVE	
TRAN_ID	298600003	
USE_ROLE_CLASSES_FLG	1	
	Synchronize From	**Synchronize From**

The page shows the time that the data was generated, displays the objects in the order that they were created, and shows the differences among all objects that are involved in a transaction across all the facilities.

The page might show a conflict where the data is the same in both facilities but a conflict exists because the same change was made in both facilities. Even though the data is correct, the conflict still exists and must be resolved. See "Best Practices for Resolving Database Conflicts" on page 76 in this chapter for more information.

3 Choose a facility as the correct source of data by clicking the Synchronize From button at the bottom of an object.

The Multimaster Tools insert or delete objects in the transaction where necessary, and then propagate the change to every facility in the multimaster mesh.

The Multimaster Tools: Object Synchronization Results page appears. The page displays the results of the transaction synchronization and contains a link back to the Transaction Differences page.

4 Click the Return to Transaction Differences link. The Multimaster Tools: Transaction Difference page appears.

The data in the object that you just synchronized might not appear resolved. The Opsware System must propagate the conflict resolution to other facilities.

5 Continue synchronizing the conflicts for objects in the transaction until all conflicts in the transaction are resolved. (Repeating Steps 3 and 4) When all conflicts in the transaction for all objects are resolved, a Mark Resolved button appears at the bottom of the page, as Figure 2-33 shows.

Figure 2-33: When All Conflicts Are Resolved, the Mark Resolved Button Appears

_RoleClassAccount 210002-3220002		
DB Field	C06	C01
ACCT_ID	Not Found	Not Found
CONFLICTING		
ROLE_CLASS_ID		
TRAN_ID		
Mark Resolved		

6 Click the Mark Resolved button. The Multimaster Tools: Mark Conflict Resolved page appears, as Figure 2-34 shows. The page displays the results of marking a transaction resolved.

Figure 2-34: Page That Shows All Conflicts Successfully Marked Resolved

Multimaster Tools: Mark Conflict Resolved | 310130002

Return to State

All conflicts successfully marked resolved.

Facility	Conflict ID	Status
C01	2440003	OK
C06	No conflicts found.	OK

After it is marked resolved, the transaction disappears from the State and Conflicts views after the Opsware System refreshes the data in the Multimaster Tools.

7 Click the link to return to the State view.

Resolving a Conflict by Transaction

Perform the following steps If you know that synchronizing all objects from one facility will resolve the conflict:

1 From the navigation panel, click Administration ➤ Multimaster Tools. The Multimaster Tools: State View page appears. By default, the State View tab displays. Each transaction ID links to the Transaction Differences page for that transaction.

2 From the State View page or the Conflict View page, click the hyperlinked ID for the transaction that has the conflict that you want to resolve.

The Multimaster Tools: Transaction Differences page appears. The page shows the time that the data was generated and the differences between all objects involved in a transaction across all the facilities.

The page might show a conflict where the data is the same in both facilities but a conflict exists because the same change was made in both facilities. Even though the data is correct, the conflict still exists and must be resolved. See "Best Practices for Resolving Database Conflicts" on page 76 in this chapter for more information.

3 From the drop-down list at the top of the page, select the facility to use as the correct source of data, as Figure 2-35 shows.

Figure 2-35: Resolving Conflicts by Transaction Uses the Facilities Drop-Down List and the Update Button

Multimaster Tools: Transaction Differences | 298600003

Return to State View

Synchronize all objects from C01 ▾ **Update**

See "Best Practices for Resolving Database Conflicts" on page 76 in this chapter for more information.

4 Click the Update button beside the drop-down list. The Multimaster Tools: Transaction Synchronization Results page appears. The page shows the results of the synchronization and prompts you to mark the conflicts for all objects resolved.

5 Click the Mark Resolved button. The Multimaster Tools: Mark Conflict Resolved page appears. The page displays the results of marking a transaction resolved.

6 Click the link to return to the State view.

After it is marked resolved, the transaction disappears form the State and Conflicts views after the Opsware System refreshes the data in the Multimaster Tools.

Network Administration for Multimaster

Opsware Inc. does *not* require that a multimaster configuration meet specific guidelines on network uptime. A multimaster configuration functions acceptably in a production environment that experiences *temporary* inter-facility network outages.

However, as the duration of a network outage increases, the probability of multimaster conflicts increases. Extended network outages between facilities can cause the following problems:

• Multimaster messages fail to propagate between facilities.

• The Multimaster Tools stop functioning.

• Opsware Command Centers cannot contact the multimaster central Data Access Engine.

Production experience for multimaster configurations supports the performance data that Table 2-5 shows.

Table 2-5: Performance Data for Multimaster Configurations

# FACILITIES	DURATION NETWORK OUTAGE	# MULTIMASTER CONFLICTS *
8 facilities (Opsware core installed in each facility)	12 hour outage (1 facility looses network connectivity to the other facilities)	12 to 24 conflicts (average number generated)
* The propensity of users to manage servers in the disconnected facility with Opsware Command Centers in other facilities increases the number of conflicts.		

Network connectivity issues include TIBCO or multicast routing problems.

Multimaster Alert Emails

When multimaster conflicts occur or multimaster components experience problems, the Opsware System sends an email to the configured multimaster email alias.

This email address is configured when the Opsware System is installed in a facility. For assistance changing this email address, contact your Opsware System Support Representative or see the *Opsware System 4.6 Installation Guide*.

The subject line of the alert email specifies:

• What type of error occurred when a transaction was being applied to a Model Repository database

• What type of Opsware System problem occurred that caused problems with the multimaster operation

Contact your Opsware System Support Representative for assistance troubleshooting and resolving Opsware System problems that affect the multimaster operation.

See Table 2-6 for error messages.

Table 2-6: Multimaster Error Messages

SUBJECT LINE	TYPE OF ERROR	DETAILS
vault.ApplyTransactionError	Multimaster Transaction Conflict	The local database was not successfully updated with the changes from the other database. Each update must affect only one row and not result in any database errors.
vault.configValueMissing	Opsware System Problem	No value was specified for a given configuration parameter. Log in to the Opsware Command Center and provide the value for this configuration parameter. Contact your Opsware System Support Representative for assistance setting Opsware System configuration values.
vault.DatabaseError	Multimaster Transaction Conflict	An error occurred while querying the database for updates to send to other databases or while applying updates from other databases. Restart the Model Repository, Multimaster Component.

Table 2-6: Multimaster Error Messages

SUBJECT LINE	TYPE OF ERROR	DETAILS
vault.InitializationError	Opsware System Problem	An error occurred when the Model Repository, Multimaster Component process started. The application returned the message specified. The thread that encountered the error stopped running. This error occurs when running the Opsware System in multimaster mode. Resolve the error condition. Restart the Model Repository, Multimaster Component.
vault.ParserError	Multimaster Transaction Conflict	An error occurred when parsing the XML representation of the transaction. The application returned the message specified. This error occurs when running the Opsware System in multimaster mode. Run the Opsware Admin Multimaster Tools and verify that the transaction data does not contain special characters that the XML parser might be unable to interpret.

Table 2-6: Multimaster Error Messages

SUBJECT LINE	TYPE OF ERROR	DETAILS
vault.SOAPError	Multimaster Transaction Conflict	An error occurred while using SOAP libraries to marshal or un-marshal transactions into XML. The application returned the message specified. This error occurs when running the Opsware System in multimaster mode. Run the Opsware Admin Multimaster Tools and verify that the transaction data does not contain special characters that SOAP might be unable to interpret.
vault.TibcoError	Opsware System Problem	The TIBCO transport raised an error. The application returned the message specified. The thread that encountered the error stopped running. This error occurs when running the Opsware System in multimaster mode. Resolve the TIBCO transport error. See the TIBCO User's Guide for information. Restart the Model Repository, Multimaster Component.
vault.UnknownError	Opsware System Problem	The Model Repository, Multimaster Component process encountered an unknown error. Contact technical support and provide the database name and Opsware System component's log file.

The Opsware System User Sessions

The Sessions selection from the Administration menu displays a list of users who are currently online. This function is used primarily as a monitoring tool to view, for example, whether anyone is logged into a server that you need to shut down. See Figure 2-36.

It is possible to use the End function to end a user's session, but it is best only to use it for that purpose in extreme situations, because abruptly ending a user's session terminates whatever the user was doing and could conceivably create error conditions.

Figure 2-36: User Session Information

In addition to the number of sessions currently running, the Sessions page shows the user's name, statistics about their usage, such as how long they have been online, the time they logged in, and how long ago they last accessed the Opsware Command Center. Also displayed is a list of the Opsware System services invoked by their activities during the session.

Ending a User Session

Use this function with extreme caution.

Perform the following steps to end a user session:

1 From the navigation panel, click Administration ➤ Sessions. The Sessions page appears, which shows all users currently on the system.

2 Click the check box next to the session that you want to end. To select all sessions, click the check box at the top of the column.

3 Click the End button. A confirmation message appears.

4 Click OK to end the selected sessions. The Sessions page refreshes and displays any remaining sessions.

Server Attributes

This section provides information about server attributes within the Opsware System and contains the following topics:

• Server Attributes Overview

• Creating Server Use Values

• Editing Server Use Values

• Deleting Server Use Categories

• Creating Deployment Stage Values

• Editing Deployment Stage Values

• Deleting Deployment Stage Values

Server Attributes Overview

The Server Attribute function is used to identify broad categories of servers and what they are used for, as well as to describe their various stages of life cycle deployment.

The Opsware System comes with three use categories already defined (and which cannot be changed or deleted): Not Specified, Production, and Staging. It also has two deployment stages pre-defined - Not Specified and Decommissioned - that also cannot be changed or deleted.

The attributes defined here, along with the default attributes, populate two lists in the Server Management function: Server Use and Deployment Stage. See the *Opsware System 4.6 User's Guide* for more information.

91

Creating Server Use Values

Perform the following steps to create server use values:

1 From the navigation panel, select Administration ➤ Server Attributes. The Server Attributes page appears with the Server Use tab displayed, which shows all previously defined server use categories, as Figure 2-37 shows.

Figure 2-37: Server Use Tab of the Server Attributes Function

Server Attributes

Name ▼	Code Deployment	Description
DevLab	Enabled	Development Lab Servers
DevLab2	Enabled	Development Lab, South Building
Development	Disabled	Site development

Server Use | Deployment Stage

Delete New Value

2 Click the New Value button. The Create Server Use Value page appears, as Figure 2-38 shows.

Figure 2-38: Create Server Use Value Page

Create Server Use Value

Return to Server Use

Name:	
Description:	
Code Deployment:	☐
	Save Cancel

3 Enter the name of the Server Use category that you want to create. This name appears in the Server Use list in the Managed Servers area of the system. This field is required.

4 Enter a description of the server use value that you are defining.

5 Click the Code Deployment check box if you want this server use category to also appear in the Code Deployment list.

6 Click the Save button.

Editing Server Use Values

Perform the following steps to edit server use values:

1 From the navigation panel, select Administration ➤ Server Attributes. The Server Attributes page appears with the Server Use tab displayed, which shows all previously-defined server use categories.

2 The names of the server use categories that you already defined, as well as the default categories names, are hyperlinks. Click the link to edit the values of the categories. The Edit Server Use Value page appears, as Figure 2-39 shows.

Figure 2-39: Edit Server Use Value Page

For the default categories of Not Specified, Production, and Staging, only the description can be modified.

For the server use categories that you have defined, all values can be modified.

3   Make any necessary changes to the Name, Description, or Code Deployment fields.

4   Click the Save button.

Deleting Server Use Categories

Perform the following steps to delete the server use categories:

1   From the navigation panel, select Administration ➤ Server Attributes. The Server Attributes page appears with the Server Use tab displayed, which shows all previously-defined server use categories.

2   Click the check box next to each of the server use categories that you want to delete. You cannot delete the ones with no check boxes - Not Specified, Production, and Staging.

3   Click the Delete button. A confirmation window appears. You can view or hide the details of the server uses that you are about to delete.

4   Click the Delete button. The server use values are deleted, and the Server Attributes page refreshes, which shows the remaining serve use values.

Creating Deployment Stage Values

Perform the following steps to create deployment stage values:

1   From the navigation panel, select Administration ➤ Server Attributes. The Server Attributes page appears with the Server Use tab displayed, which shows all previously-defined server use categories.

2   Click the Deployment Stage tab. The list of previously defined deployment stages appears, as Figure 2-40 shows.

Figure 2-40: Deployment Stage Tab of the Server Attributes Function

	Name ▼	Description
☐	Decommissioned	Site has been completely shut down and is no longer in use.
☐	In Deployment	Server is actively being built.
☐	Live	Code is successfully deployed and tested, site is ready for live traffic.
	Not Specified	The stage of the server is not yet known.
☐	Offline	Server may still be operating, but has been removed from active management.
☐	Ops Ready	Build tasks have been completed and server is ready for Ops Ready testing.

3 Click the New Value button. The Create Deployment Stage Value page appears, as Figure 2-41 shows.

Figure 2-41: Create Deployment Stage Value Page

4 Enter the Name of the deployment stage. This field is required.

5 Enter a description of the deployment stage.

6 Click the Save button.

Editing Deployment Stage Values

Perform the following steps to edit deployment stage values:

1 From the navigation panel, select Administration ➤ Server Attributes. The Server Attributes page appears with the Server Use tab displayed, which shows all previously defined server use categories.

2 Click the Deployment Stage tab. The list of previously defined deployment stages appears.

3 Each of the deployment stages on the list is a hyperlink. Click the name of the deployment stage whose values you want to edit. The Edit Deployment Stage Value page appears, as Figure 2-42 shows.

Figure 2-42: Edit Deployment Stage Value Page

4 Change the Name or Description as necessary.

5 Click the Save button.

Deleting Deployment Stage Values

Perform the following steps to delete deployment stage values:

1 From the navigation panel, select Administration ➤ Server Attributes. The Server Attributes page appears with the Server Use tab displayed.

2 Click the Deployment Stage tab. The list of previously defined deployment stages appears.

3 Click the check box next to the deployment stage that you want to delete. The one with no check box, Not Specified, cannot be deleted.

4 Click the Delete button. A confirmation page appears. You can view or hide the details of the deployment stages that you are about to delete.

5 Click Delete. The Server Attributes page with the Deployment Stage tab appears, which shows the remaining deployment stages.

System Configuration

During the installation of an Opsware System core the Opsware Installer sets specific system configuration parameters. In addition to the parameters that are set during installation, there are also many default values for the various system configuration parameters that should not be changed unless expressly directed to do so by Opsware Inc.

For information about how to use this function when you install an Opsware core, see the *Opsware System 4.6 Installation Guide*

The Opsware Agent reads the system configuration values at installation time only. If any of the configuration values change, the agent configuration must be updated manually. Contact Opsware Technical Support for help making these changes, or in making any other changes in the System Configuration area of the Opsware System.

Opsware Software

The Opsware Software function is populated during the Opsware System installation.

Each component of the Opsware System is shown by its internal name, and the display is graphically similar in appearance to the way that nodes appear in the Software Tree. You cannot add or delete components or nodes in this area of the Opsware System.

Table 2-7 shows the internal and external names of the Opsware System components.

Table 2-7: Opsware Internal and External Component Names

INTERNAL NAME	EXTERNAL NAME
Agent	Opsware Agent
buildmgr	OS Build Manager
cast	Access and Authentication Directory
occ	Opsware Command Center
spin	Data Access Engine
truth	Model Repository
twist	Web Services Data Access Engine
vault	Multimaster Replication Engine
way	Command Engine
word	Software Repository

Some of the functionality available in the Server Management area of the system is also available to be applied to the servers that appear on the server tab. Take care in applying changes to the core servers. In particular, do not assign or unassign servers to these nodes or install or uninstall software or change networking unless directed to do so during the installation process by the *Opsware System 4.6 Installation Guide*.

To view the servers on which each component is installed, click the component's hyperlinked name, then click the Servers tab. The number of servers associated with that component appears on the tab itself, and detailed information about those servers shows when you click the tab.

The *Opsware System 4.6 Installation Guide* has the following two topics that deal with the Opsware Software function:

• Chapter 4, "Reassigning the Data Access Engine to a Secondary Role"

• Chapter 6, "Designating the Multimaster Central Data Access Engine"

Overview of Mass Deletion of Backup File

The Opsware System includes a script that you can run as a cron job for performing mass deletions of backup files. Backup files are created by configuration tracking. They can accumulate quickly and take up disk space. Consequently, performance when viewing backup history in the Opsware Command Center can be sluggish, and the information that displays might be cluttered with out-of-date configuration tracking data.

When the backup deletion script is run, it deletes all backed up files with the exception that it always keeps one copy of the latest version of every file ever backed up. If you want to delete those files, use the process for deleting backups individually or a few at a time that is covered in the *Opsware System 4.6 User's Guide.*

The script is called `backup_delete.pyc.` It is located on the server where the Data Access Engine resides, in the following directory:

`/cust/usr/blackshadow/spin/util`

The script is run using a configuration file that contains the script arguments such as host name, port number, whether you want full or incremental backups, the backup retention period, the name of the log file to use, email addresses for notifications, and the email server to use. See Table 2-8 on page 100, Configuration File Options, for the arguments, their values, and their descriptions.

Syntax for the Command

`backup_delete.pyc` [*options*]

Usage: `backup_delete.py` [`-c` *<conf_filename>*]

Using the Mass Deletion Script to Delete Backup Files

Perform the following steps to use the mass deletion script to delete backup files:

1 Log in as root to the server where the Data Access Engine is installed.

2 Make sure that `/lc/blackshadow` is in your PYTHONPATH environment variable.

3 Create a file that contains the arguments and values that you want the Opsware System to use with the mass deletion script. See Table 2-8 on page 100, Configuration File Options, for the available arguments.

For example, the following file specifies that a host called spin.yourcore.example.com, on port 1004 will have incremental backups that are three months old deleted. In addition, a log file called run.log, located in /tmp will be used to capture events, and email will be sent to user@example.com from user1@example.com reporting that the mass deletion was performed successfully.

```
host: spin.yourcore.example.com

port: 1004

inc: 1

time: 3m

logfile: /tmp/run.log

emailto: user@example.com

emailserver: smtp.example.com

emailfrom: user1@example.com

emailsuccess: 1
```

Table 2-8: Configuration File Options

ARGUMENTS	VALUES	DESCRIPTION
`host`	host: hostname, for example `host: spin.yourcore.example.com`	Hostname of the Data Access Engine
`port`	port: port number, for example `port: 1004`	Port of the Data Access Engine (defaults to 1004)
`full`	Set value to 1 to enable, for example `full:1`	Delete full backups. You must specify Either `full` or `inc`.
`inc`	Set value to 1 to enable, for example `inc:1`	Delete incremental backups. You must specify either `full` or `inc`.
`time`	time: [digits][dmy] For example, 6d equals six days. 3m equals three months. 1y equals one year.	Retention period beyond which backups should be deleted.

ARGUMENTS	VALUES	DESCRIPTION
`hostsfile`	hostsfile: [filename] The hostsfile should contain the name of each host on a line by itself, for example <hostname> <hostname>	The script deletes backups on every managed server in your system, unless you provide a hostsfile that contains a specific list of servers on which to perform the mass backup deletion.
`logfile`	logfile: filename, for example `logfile: /tmp/run.log`	File to use for log events.

`emailto`	emailto: email address, for example
	`emailto:`
	`user@example.com`
		Optional email notification recipient.

`emailserver`	emailserver: server name, for example
	`emailserver:`
	`smtp.example.com`
		SMTP Server to send email through. Optional if `emailto` not specified, otherwise required.

`emailfrom`	emailfrom: email address, for example
	`emailfrom:`
	`user1@example.com`
		Email address to appear in the From: line. Optional if `emailto` not specified, otherwise required.

`emailsuccess`	Set value to 1 to enable, for example `emailsuccess: 1`	Send email even if no errors occurred deleting backups and more than one backup was deleted.

4 Optionally, if you want to run the script as a cron job, create a crontab entry.

For example, to run the job at 3:00 AM daily, create the following entry:

```
0 3 * * * env PYTHONPATH=/lc/blackshadow /lc/bin/python/
cust/usr/blackshadow/spin/util/backup_delete.pyc -c <path>/
<your_backup_filename.conf>
```

The crontab entry *must* be all on one line.

5 If you do not plan to run the script as a cron job, enter the following command at the prompt:

```
# python /cust/usr/blackshadow/spin/util/backup_
delete.pyc\-c /[conf_filename]
```

Chapter 3: Opsware System Health

Possible Opsware System Problems

This section provides information about possible Opsware System problems and contains the following topics:

- Possible Opsware System Problems Overview

- Troubleshooting Opsware Components

- Contacting Opsware Inc. Support

Possible Opsware System Problems Overview

While maintaining the Opsware System, you might encounter the following types of problems:

- Operational problems: processes failing or becoming unresponsive (Data Access Engine, Command Engine, Software Repository)

- Failure of an Opsware component, which causes other components to fail

The following examples describe the effects of some component failures:

- If the Data Access Engine fails, the Opsware Command Center, the Command Engine, the Access & Authentication Directory, and the Software Repository components will fail.

- If the Software Repository fails to contact the Data Access Engine, downloads from the Software Repository are impossible.

- If the Model Repository fails, the Data Access Engine fails.

- The Software Repository fails to contact the Data Access Engine without either a functioning DNS, or a properly-configured `/etc/hosts` file.

- Unreachable servers existing in the managed environment.

Many problems with the Code Deployment & Rollback (CDR) subsystem are caused by errors with the CDR configuration and setup. See the *Opsware System 4.6 User's Guide* for information about CDR configuration.

Troubleshooting Opsware Components

The following mechanisms for troubleshooting the Opsware System are available:

- Running the Opsware System Diagnosis tool (a tool for debugging common problems with Opsware components). See "Opsware System Diagnosis" on page 105 in this chapter for more information.

- Reviewing error logs for Opsware components. See "Logs for Opsware Components" on page 112 in this chapter for more information.

- Contacting Opsware Inc. Support

Contacting Opsware Inc. Support

When you contact Opsware Inc. Support, have the following information available to help you with your support call:

- Be at your computer and have network access to the servers running the Opsware core.

- Have your Opsware guides available.

- Write down the steps followed prior to the problem occurring.

• Write down the exact text of the error that appears on your screen or print out the page on which the error appears.

• Be able to describe the problem in detail.

Contact Opsware Technical Support:

Phone: +1 877 677-9273 (1-877-Opsware), in the United States

International Phone: 1 408-212-5300
Email: support@opsware.com

Opsware System Diagnosis

This section provides information about how to diagnose Opsware System problems and contains the following topics:

• Opsware System Diagnosis Overview

• System Diagnosis Testing Process

• What the System Diagnosis Tests Check

• Data Access Engine Tests

• Software Repository Tests

• Web Services Data Access Tests

• Command Engine Tests

• Model Repository, Multimaster Component Tests

Opsware System Diagnosis Overview

By using the System Diagnosis tool, you can check the functionality of the Opsware components and the ability of servers running in the managed environment to interact with the Opsware core.

You can troubleshoot most of the errors that occur within the Opsware core by running the Opsware System Diagnosis tool.

System Diagnosis Testing Process

The System Diagnosis tool tests the Opsware components first, and then, optionally, tests the servers that you specify, which are running in the managed environment.

The System Diagnosis tool performs intensive tests of the Opsware components, which check the functionality of the Opsware components:

• Stand-Alone Tests – The first suite of tests, which tests as much of the functionality of that component as possible without the use of other Opsware components. The Stand-Alone Tests are run to verify a base level of functionality and the component's ability to respond to an XML-RPC call.

• Comprehensive Tests – The second suite of tests, which tests the full functionality of each component.

On completion of the Comprehensive Tests, the System Diagnosis tool displays the success of each test, the results, and error information for the tests that failed.

The components are not tested in a specific order; however, the tests generally occur in this order:

• Opsware Agent Stand-Alone Tests

• Opsware Agent Comprehensive Tests

• Component Stand-Alone Tests

• Component Comprehensive Tests

What the System Diagnosis Tests Check

The tests for the components simulate all the functionality that each component represents. In addition to errors, the tests verify that each component is functioning within certain conditions (for example, whether database connections are near maximum on the Data Access Engine).

The System Diagnosis tool tests the following components:

• Data Access Engine

• Software Repository

• Web Services Data Access Engine

• Command Engine

• Opsware Agents on Opsware core servers

• Optionally, Opsware Agents on specified IP addresses or DNS hostnames

• Model Repository, Multimaster Component

When using the System Diagnosis function in an environment with multiple facilities, System Diagnosis can only be run on one facility at a time.

In addition to component functionality, the System Diagnosis tool can check the ability of servers running in the managed environment to interact with the Opsware core by specifying the hostnames or IP addresses of servers. Performing this test checks the functionality of the Opsware Agent on the servers and checks the servers for problems that relate to the Opsware components.

The System Diagnosis tool does not test the Build Manager.

Data Access Engine Tests

The following section describes two types of Data Access Engine diagnostic tests: Stand-Alone and comprehensive.

Stand-Alone Tests

- Check for the current Data Access Engine version

- Check for the current Model Repository database version

- Obtain a Device object

- Obtain a MegaDevice object

- Verifies advanced query functioning

- Verify a Device object

- Obtain the list of facilities

- Obtain the names of the Data Access Engine cronbot jobs

- Check whether the usage of database connections is below the acceptable level

- Check whether any database connection has been open more than 600 seconds

- Check whether the Data Access Engine and Model Repository are in the same facility

- Verify that all Model Repository garbage-collectors are running when the Model Repository is running in multimaster mode

- If the Data Access Engine is configured as the central multimaster Data Access Engine:

  - Check whether multimaster transactions are being published

  - Check whether multimaster transactions are showing up at remote facilities

  - Check for multimaster transaction conflicts

Comprehensive Tests

- Test connectivity to the Model Repository on the configured port

- Test connectivity to the Access & Authentication Directory on the configured port

- Test connectivity to the Command Engine on the configured port

- Test connectivity to the Software Repository on the configured port

Software Repository Tests

The following section describes two types of Software Repository diagnostic tests: stand alone and comprehensive.

Stand-Alone Tests

None

Comprehensive Tests

- Test whether a file that is not a package can be uploaded to the Software Repository process that serves encrypted files. This test verifies whether the file is present in the Software Repository file system and that the file size matches the source.

- Verify that a file can be downloaded from the Software Repository

- Verify whether the Software Repository process that serves unencrypted files is running and serving files

- Try to download a file without encryption

- Verify that a package can be uploaded to the Software Repository and that the package is registered with the Model Repository

- Verify that a package can be deleted from the Software Repository and removed from the Model Repository

Web Services Data Access Tests

The following section describes two types of Web Services Data Access diagnostic tests: stand-alone and comprehensive.

Stand-Alone Tests

• Connect to the Web Services Data Access Engine and retrieves its version information

Comprehensive Tests

• Connect to the Web Services Data Access Engine

• Read a server record from the Model Repository and thereby checks connectivity to the Model Repository

Command Engine Tests

The following section describes two types of Command Engine diagnostic tests: stand alone and comprehensive.

Stand-Alone Tests

• Check the state machine

• Check session tables

• Check lock-down status

• Check for signature failures

• Check command and service tables

• Check the facility cache

Comprehensive Tests

• Check Data Access Engine connectivity

• Check security signatures

• Check lock operation

• Run an internal script

• Run an external script

Model Repository, Multimaster Component Tests

The following section describes two types of Model Repository, Multimaster Component diagnostic tests: stand alone and comprehensive.

Stand-Alone Tests

• Check the ledger state by examining the ledger file

• Report the total number of messages sent, number of messages still in the ledger file (for example, not confirmed by all listeners), and the sequence number of the last message confirmed by each listener

• Check the sender health by examining the state of the Outbound Model Repository, Multimaster Component

• Check the receiver health by examining the state of the Inbound Model Repository, Multimaster Component

Comprehensive Tests

None

Running a System Diagnosis of Opsware Components

To access the System Diagnosis tool, you must have Opsware administrator privileges. See "Opsware System Administration" on page 33 in Chapter 2 for information about how to assign user privileges. The Opsware Command Center has access to all the Opsware Agents running on the Opsware component servers.

Perform the following steps to run a system diagnosis of the Opsware Components:

1 From the navigation panel, click Administration ➤ System Diagnosis. The System Diagnosis: Begin Diagnosis page appears.

2 Select the components that you want to test. By default, all components are selected (the Data Access Engine, the Software Repository, Command Engine, and Web Services Data Access Engine; in multiple core environments, there is also a selection for the Multimaster Replication Engine component). See Figure 3-1.

Figure 3-1: System Diagnosis Page That Shows Opsware Components Selected for Testing on the Indicated Facility

3 Click the Run Diagnosis button.

The System Diagnosis: Performing Diagnosis window appears, which displays a progress bar while the tests are running, as Figure 3-2 shows.

Figure 3-2: System Diagnosis Progress Bar

When all the tests are complete, the window closes and the System Diagnosis: Failed Tests page appears in the main Opsware Command Center window. If all tests passed, the System Diagnosis: Successful Tests page appears.

4 To review the results of a test, click the linked test name in the Test column. The System Diagnosis: Test Information page appears. If the test contained an error, error information appears at the bottom of the page.

Logs for Opsware Components

This section provides information on Opsware component logs and contains the following topics:

• Opsware Component Logs Overview

• Access & Authentication Directory Logs

• Boot Server Logs

• Build Manager Logs

• Command Engine Logs

• Data Access Engine Logs

• Media Server Logs

• Model Repository Logs

• Model Repository, Multimaster Component Logs

• Opsware Agents Logs

• Opsware Command Center Logs

• Software Repository Logs

• Software Repository Replicator Logs

• Software Repository, Multimaster Component Logs

• Web Services Data Access Engine Logs

Opsware Component Logs Overview

Opsware components log events to a log-event handler and write human-readable summaries of log records, which are useful for troubleshooting.

By default, Opsware Inc. has configured the logging levels for the highest value (indicating higher priority) and the default log file size before rotation.

To increase or decrease the log levels for debugging or log file sizes in your organization's operational environment, contact your Opsware Support Representative for assistance.

You can access the logs for the Opsware components as files on each host running an Opsware component. The components use standard logging libraries for access and error logs, so the logs are output in a standard format. Use shell commands to view the error log files.

Refer to your own internal documentation where you have noted the Opsware component hostnames configured for your environment.

You can access the logs for the following components:

• Access & Authentication Directory

• Boot Server

• Build Manager

• Command Engine

• Data Access Engine

• Media Server

• Model Repository

• Model Repository, Multimaster Component

• Opsware Agents

• Opsware Command Center

• Software Repository

• Software Repository Replicator

• Software Repository, Multimaster Component

• Web Services Data Access Engine

Access & Authentication Directory Logs

Perform the following steps to access the Access & Authentication Directory Logs:

1 Access the host running the Access & Authentication Directory by opening a terminal session.

2 Locate the error log file in the directory:

```
/cust/usr/netscape/server4/slapd-cast/logs/errors
```

Boot Server Logs

The Boot Server does not create its own logs. Installing the boot server component enables TFTP with INETD and the NFS server components. Additionally, the ISC DHCPD is installed. All these components log with syslog. Consult your vendor documentation for more information on these log files.

Build Manager Logs

Perform the following steps to access the build manager logs:

1 Access the host running the Build Manager by opening a terminal session.

2 Locate the following files:

`/var/lc/buildmgr/buildmgr.log`

Command Engine Logs

Perform the following steps to access the command engine logs:

1 Access the host running the Command Engine by opening a terminal session.

2 Locate the file:

`/var/lc/waybot/waybot.err`

Data Access Engine Logs

Perform the following steps to access the data access engine logs:

1 Access the host running the Data Access Engine by opening a terminal session.

2 Locate the file:

`/var/lc/spin/spin.err`

Two Data Access Engine logs are in a multihost core.

Media Server Logs

Perform the following steps to access the media server logs:

1 Access the host running the Media Server by opening a terminal session.

2 Locate the following files relevant to Windows OS provisioning:

`/var/opt/OPSWsamba/log.smbd`

`/var/opt/OPSWsamba/log.nmbd`

Solaris and Linux OS provisioning makes use of vendor-provided services such as NFSD. These services typically log through syslog. Consult your vendor documentation for more information on these log files.

Model Repository Logs

The Model Repository is an Oracle database. The location of any logs that the database produces is specific to your installation.

Model Repository, Multimaster Component Logs

Perform the following steps to access model repository, multimaster component logs:

1 Access the host running the Model Repository, Multimaster Component by opening a terminal session.

2 Locate the files:

`/var/lc/vault/log.*`

*Where the files are `log.1`, `log.2`, `log.3`, and so forth. The size of the log file is configurable. Additional logs are created when the specified maximum file size is reached.

Opsware Agents Logs

Perform the following steps to access Opsware Agents Logs:

1 Use a terminal session to access the host from which you need the Opsware Agent error log.

2 Locate the following file:

`/var/lc/cogbot/cogbot.err` (UNIX-like systems)

`%ProgramFiles%Common Files\loudcloud\cogbot\cogbot.err` (Windows systems)

Opsware Command Center Logs

The Opsware Command Center does not generate specific error logs. You can access the WebLogic access log, Web server logs, and the standard out log by performing the following steps:

1 Access the host running Opsware Command Center by opening a terminal session.

2 Locate the following files:

/var/lc/owm/logs/owm_weblogic*.log (new file nightly)

/var/lc/owm/logs/owm_access.log (new file nightly)

/var/lc/owm/logs/stdout.log (one file)

The stdout.log file essentially contains debug output. The file does not conform to a specific format. The file can contain stack traces that were trapped on an Opsware Command Center page and sent to the file, or data such as init Spinclient.

Software Repository Logs

Perform the following steps to access the software repository logs:

1   Access the host running the Software Repository by opening a terminal session.

2   Locate the following files:

/var/lc/mm_wordbot/wordbot.err*

/var/lc/mm_wordbot-clear/wordbot-clear.err*

Software Repository Replicator Logs

Perform the following steps to access the software repository replicator logs:

1   Access the host running the Software Repository, Multimaster Component by opening a terminal session.

2   Locate the following files:

/var/lc/replicator/replicator.err

/var/lc/replicator/daemonbot.out

/var/lc/replicator/replicator.log.*

/var/lc/replicator/1

*Where the files are log.1, log.2, log.3, and so forth. The size of the log file is configurable. Additional logs are created when the specified maximum file size is reached.

Software Repository, Multimaster Component Logs

Perform the following steps to access the software repository, multimaster component logs:

1 Access the host running the Software Repository, Multimaster Component by opening a terminal session.

2 Locate the following files:

`/var/lc/mmword/log.*`

*Where the files are `log.1, log.2, log.3`, and so forth. The size of the log file is configurable. Additional logs are created when the specified maximum file size is reached.

Web Services Data Access Engine Logs

The Web Services Data Access Engine logs any errors or exceptions in its standard out log. Weblogic-specific messages are generated in the Weblogic server log.

1 Access the host running the Web Services Engine by opening a terminal session.

2 Locate the following files:

`/var/lc/twist/stdout.log`

`/var/lc/twist/twist.log`

`/var/lc/twist/access.log`

The `stdout.log` file contains debug output and logging of every exception that the server generates. The file does not conform to a specific format.

The `twist.log` file contains Weblogic-specific error or informational messages.

The `access.log` file contains access information in common log format.

Restarting Opsware Components

This section provides information on restarting Opsware components and contains the following topics:

• Restarting Opsware Components Overview

• Restarting the Access & Authentication Directory

• Restarting the Boot Server

• Restarting the Build Manager

• Restarting the Command Engine

- Restarting the Data Access Engine

- Restarting the Media Server

- Restarting the Model Repository

- Restarting the Model Repository, Multimaster Component

- Restarting an Opsware Agent

- Restarting the Opsware Command Center

- Restarting the Software Repository

- Restarting the Software Repository, Multimaster Component

- Restarting the Web Services Data Access Engine

Restarting Opsware Components Overview

While maintaining the Opsware System, you might encounter operational problems or an Opsware component might fail. A server might become unreachable because the Opsware Agent on the server has stopped responding.

The Opsware System components are installed on both Linux and Solaris servers. The processes run in the same directories on both platforms.

As part of maintaining the Opsware System, you might need to stop and restart specific components or the Opsware Agent on an Opsware-managed server.

Opsware components are configured to restart upon process failure, meaning that if a component's process terminates, nothing needs to be done to restart that process. However, there might be times when a manual restart is necessary. There are some restart ordering dependencies to take into consideration when you manually restart Opsware components.

You must restart the following components in the following order:

1 Make sure that the machine on which the Model Repository resides is running, and that Oracle is also running.

2 If you restart the Model Repository, you must also restart the Data Access Engine and, in a multiple facility environment, restart the Multimaster Replication Engine.

3 If you need to restart the Web Services Data Access Engine, the Data Access Engine, and the Model Repository must be running first.

4   If you need to restart Build Manager, these components must be started first: the Data Access Engine, the Command Engine, and the Web Services Data Access Engine.

5   In a multiple facility environment, if you need to stop and restart the TIBCO Rendezvous process, make sure to also stop and restart the Opsware components (such as Multimaster Replication Engine, Web Services Data Access Engine, and Software Repository, Multimaster Component) that are also running on the same host. If you leave the components running, they spawn a new TIBCO process, but it is the rvd process, not the rvrd process. To launch the right process, be sure to first stop the Opsware components, then stop TIBCO. Then restart TIBCO, followed by the Opsware components.

6   In a multiple facility environment, if you need to restart the Web Services Data Access Engine, in addition to the Data Access Engine, and the Model Repository, you must also be sure that the Multimaster Replication Engine is running first.

Restarting the Access & Authentication Directory

To stop the Access & Authentication Directory, enter the following command as root at the prompt on the server running the Access & Authentication Directory:

```
/etc/init.d/cast-updown stop
```

To start the Access & Authentication Directory, enter the following command as root at the prompt on the server running the Access & Authentication Directory:

```
/etc/init.d/cast-updown start
```

Restarting the Boot Server

You should not need to restart the Boot Server. However, you might need to restart the DHCP server that runs on the Boot Server.

To stop the DHCP server, enter the following command as root at the prompt on the server running the Boot Server:

```
/etc/init.d/dhcpd stop
```

To start the DHCP server, enter the following command as root at the prompt on the server running the Boot Server:

```
/etc/init.d/dhcpd start
```

Restarting the Build Manager

The Build Manager was installed when the Build Scripts ran during installation. To stop the Build Manager, enter the following command as root at the prompt on the server running the Build Manager:

```
/etc/init.d/buildmgr stop
```

To start the Build Manager, enter the following command as root at the prompt on the server running the Build Manager:

```
/etc/init.d/buildmgr start
```

Restarting the Command Engine

Use caution when stopping and restarting the Command Engine. If processes are active when the Command Engine stops, data might be lost.

Before you stop the Command Engine, check the Way Administration page:

```
https://[hostname:1018]/way/bidniss/activeScripts.py
```

If none of the processes has a status of Active, it is safe to stop the Command Engine. There are also links to Command Engine processes at the following URL:

```
https://[hostname:1018]/way/bidniss/
```

To access the Way Administration page, you must have an Opsware certificate of authority. Contact Opsware Support.

Click the Write Machine link. If the resulting table shows messages in the Write Queue column, there are active processes, and stopping the Command Engine results in a loss of data.

To stop the Command Engine, enter the following command as root at the prompt on the server running the Command Engine:

```
/etc/init.d/waybot stop
```

To start the Command Engine, enter the following command as root at the prompt on the server running the Command Engine:

```
/etc/init.d/waybot start
```

Restarting the Data Access Engine

To stop the Data Access Engine, enter the following command as root at the prompt on the server running the Data Access Engine:

```
/etc/init.d/spin stop
```

To start the Data Access Engine, enter the following command as root at the prompt on the server running the Data Access Engine:

```
/etc/init.d/spin start
```

Restarting the Media Server

For Windows, check to see if the Media Server process is running by entering the following command as root at the prompt on the server running the Windows Media Server:

```
ps -e|grep -w smbd
```

To stop the Windows Media Server, enter the following command as root at the prompt on the server running the Windows Media Server:

```
/etc/init.d/samba.server stop
```

To start the Windows Media Server, enter the following command as root at the prompt on the server running the Windows Media Server:

```
/etc.init.d/samba.server start
```

To restart the Windows Media Server, enter the following command as root at the prompt on the server running the Windows Media Server:

```
/etc/init.d/samba.server restart
```

Restarting the Model Repository

To make sure that the Oracle database is running on Solaris and Linux, enter the following command:

```
ps -fu oracle | grep pmon
```

To make sure that the Oracle listener is running on Solaris, enter the following command:

```
ps -fu oracle | grep -i listen
```

To make sure that the Oracle listener is running on Linux, enter the following command:

```
ps axwww |grep -i listen
```

In both cases, the correct output is one or more lines, and no output at all is an error.

To stop and restart the Model Repository, contact your company database administrator.

Restarting the Model Repository, Multimaster Component

The Model Repository, Multimaster Component is automatically installed on the same server as the Model Repository. To stop the Model Repository, Multimaster Component, enter the following command as root at the prompt on the server running the Model Repository:

```
/etc/init.d/vaultdaemon stop
```

To start the Model Repository, Multimaster Component, enter the following command as root at the prompt on the server running the Model Repository:

```
/etc/init.d/vaultdaemon start
```

Before you restart the Model Repository, Multimaster Component, perform these verification steps: Access the Model Repository, Multimaster Component logs to determine if the Model Repository, Multimaster Component shut down. When shutdown, the Model Repository, Multimaster Component logs contain the statement, "Vault has been shut down." Check that the server running the Model Repository, Multimaster Component has stopped the Model Repository, Multimaster Component process.

Restarting an Opsware Agent

To stop the Opsware Agent on the Solaris platform, execute the command:

```
/etc/init.d/cogbot stop
```

To restart the Opsware Agent on the Solaris platform, execute the command:

```
/etc/init.d/cogbot start
```

For Linux 6.2, to stop the Opsware Agent, execute the command:

```
/etc/rc.d/init.d/cogbot stop
```

For Linux 6.2, to restart the Opsware Agent, execute the command:

```
/etc/rc.d/init.d/cogbot start
```

For Linux 7.2 and higher, to stop the Opsware Agent, execute the command:

```
/etc/init.d/cogbot stop
```

For Linux 7.2 and higher, to restart the Opsware Agent, execute the command:

```
/etc/init.d/cogbot start
```

For AIX, to stop the Opsware Agent, execute the command:

```
/etc/rc.d/init.d/cogbot stop
```

For AIX, to restart the Opsware Agent, execute the command:

```
/etc/rc.d/init.d/cogbot start
```

For HP-UX, to stop the Opsware Agent, execute the command:

```
/sbin/init.d/cogbot stop
```

For HP-UX, to restart the Opsware Agent, execute the command:

```
/sbin/init.d/cogbot start
```

On Windows, enter the following commands:

```
net stop shadowbot

net start shadowbot
```

Restarting the Opsware Command Center

To stop the Opsware Command Center, enter the following command as root at the prompt on the server running the Opsware Command Center:

```
/etc/init.d/owm.server stop
```

To start the Opsware Command Center, enter the following command as root at the prompt on the server running the Opsware Command Center:

```
/etc/init.d/owm.server start
```

Restarting the Software Repository

To stop the Software Repository, enter the following commands as root at the prompt on the server running the Software Repository:

```
/etc/init.d/mm_wordbot stop
```

```
/etc/init.d/mm_wordbot-clear stop
```

To start the Software Repository, enter the following commands as root at the prompt on the server running the Software Repository:

```
/etc/init.d/mm_wordbot start
```

```
/etc/init.d/mm_wordbot-clear start
```

Restarting the Software Repository, Multimaster Component

To stop the Software Repository, Multimaster Component, enter the following commands as root at the prompt on the server running the Software Repository:

```
/etc/init.d/mmworddaemon stop
```

To start the Software Repository, Multimaster Component, enter the following commands as root at the prompt on the server running the Software Repository:

```
/etc/init.d/mmworddaemon start
```

Restarting the Web Services Data Access Engine

The Web Services Data Access Engine was installed automatically with the Opsware Command Center during installation, and it runs on the same server. To stop the Web Services Data Access Engine, enter the following command as root at the prompt on the server running the Web Services Data Access Engine:

```
/etc/init.d/twist stop
```

To start the Web Services Data Access Engine, enter the following command as root at the prompt on the server running the Web Services Data Access Engine:

```
/etc/init.d/twist start
```

The Web Services Data Access Engine start is asynchronous. It could take up to ten minutes to restart, so the Opsware Command Center is unavailable during that time.

Opsware Access and Authentication Directory Administration

This section provides information on access and authentication directory administration and contains the following topics:

- Access and Authentication Directory Administration Overview

- Backing up the Access & Authentication Directory

- Restoring the Access & Authentication Directory

- Checking the Status of the Access & Authentication Directory

- Re-initializing an Access & Authentication Directory Replica

Access and Authentication Directory Administration Overview

The Opsware Access & Authentication Directory is the repository of information about users, groups, Opsware resources, and the permissions that users have. Opsware resources, especially Opsware Command Center, consult the Opsware Access & Authentication Directory to determine to which resources each user has access.

The Opsware Access & Authentication Directory is an instance of the iPlanet Directory Server 4.x, an industry-standard Lightweight Directory Access Protocol (LDAP) server. The iPlanet Directory Server package includes the following iPlanet applications:

- Directory Server

- Administration Server

- Server Console

The Server Console, a Java application that administers the Directory Server, communicates with the Administration Server to perform its administration tasks.

See the *iPlanet Directory Server 4.1 Administrator's Guide* for information about how to administer and run these applications.

Backing up the Access & Authentication Directory

You can back up the Opsware Access & Authentication Directory while the server is running and serving requests.

Perform the following steps to back up the Access & authentication Directory:

1 Log in to the Opsware Access & Authentication Directory host.

2 Run the following commands:

```
cd /cust/usr/netscape/server4/slapd-cast
```

```
./db2bak
```

A backup of the database is stored in a subdirectory of the following directory:

```
/cust/usr/netscape/server4/slapd-cast/bak
```

The directory is named by the date and time the backup was taken.

See the *iPlanet Directory Server 4.1 Administrator's Guide* for information about how to back up an instance of the iPlanet Directory Server 4.x.

Restoring the Access & Authentication Directory

Perform the following steps to restore the Access & Authentication Directory:

1 Before you restore the data, shut down the iPlanet Directory Server with the command:

```
/etc/init.d/cast-updown stop
```

2 Log in to the Access & Authentication Directory host.

3 Run the following commands:

```
cd /cust/usr/netscape/server4/slapd-cast

./bak2db backup-directory
```

Where *backup-directory* is the full path name to the directory that contains the backup that you want to restore.

4 Restart the Access & Authentication Directory after the restoration is complete with this command:

```
/etc/init.d/cast-updown start
```

See the *iPlanet Directory Server 4.1 Administrator's Guide* for information about how to restore the data for an iPlanet Directory Server 4.x.

Checking the Status of the Access & Authentication Directory

This guide does not document how to set up the Opsware Access & Authentication Directory to run in multiple facilities (creating Access & Authentication Directory replicas in a multimaster architecture). Consult your Opsware Support Representative for information about how to set up replicas of the Opsware Access & Authentication Directory.

You can use the Netscape Console application to check on the status of replication to each Opsware Access & Authentication Directory replica.

See the *iPlanet Directory Server 4.1 Administrator's Guide* for information about how to monitor replication status for iPlanet Directory Servers 4.x.

Re-initializing an Access & Authentication Directory Replica

Opsware User Directories can get out of sync with the master Opsware Access & Authentication Directory when network connectivity is lost for long periods.

If a replica of an Opsware Access & Authentication Directory gets out of sync with the master Opsware Access & Authentication Directory, you can re-initialize its data from the master Opsware Access & Authentication Directory.

See the *iPlanet Directory Server 4.1 Administrator's Guide* for information about how to initialize consumers for iPlanet Directory Servers 4.x.

This guide does not document how to set up the Opsware Access & Authentication Directory to run in multiple facilities (creating Access & Authentication Directory replicas in a multimaster architecture). Consult your Opsware Support Representative for information about how to set up replicas of the Opsware Access & Authentication Directory.

Index

W