Opsware System 4.7 User's Guide

Corporate Headquarters

599 North Mathilda Avenue Sunnyvale, California 94085 U.S.A.
T + 1 408.744.7300 F +1 408.744.7383 www.opsware.com

Table of Contents

Chapter 6: Application Provisioning Setup 307

Chapter 10: Script Execution Subsystem 455

Preface

Welcome to Opsware System 4.7 – 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 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 use Opsware System, starting with an introduction to the system and how to navigate the user interface. It provides information about managing servers, operating system provisioning, managing software packages, provisioning applications, managing patches, reconciling servers, automated script execution, automated configuration tracking, deploying and rolling back code, using the command line interface, and integrating with third-party OS installation technologies.

This guide is intended for system administrators who are responsible for all aspects of managing and provisioning the servers in an operational environment.

Contents of This Guide

This guide contains the following chapters and appendices:

Chapter 1: Opsware System Overview – provides a high-level overview of the entire Opsware system, including the automated subsystems, Web Service APIs, and multimaster. It includes information about supported operating systems and browsers, navigation of the user interface, and an explanation of each of the features found on the Opsware Command Center Home page.

Chapter 2: Server Management – provides information about all aspects of server management including server management in multiple facilities, server asset tracking, server histories and reports, and server properties. It includes information about the Opsware Agent on managed servers, server identification, IP address usage, server

network configurations, server life cycles, and server management tasks. It also discusses job details, time outs, server groups, custom attributes for servers, service levels, My Servers, server assimilation, and IP range groups and IP ranges.

Chapter 3: OS Provisioning Setup – provides a description of all tasks necessary to prepare for operating system provisioning including media management, operating system specific tasks, operating system definitions, build customization scripts, OS build process default definitions, operating system definitions in templates, and details of hardware support.

Chapter 4: Operating System Provisioning – provides information about supported environments for OS provisioning and an overview of the permissions and server life cycles associated with OS provisioning. It also describes the process for provisioning, an overview of the hardware preparation, information about booting new servers, and using the Opsware Command Center to install operating systems.

Chapter 5: Package Management – provides information about container and installable packages, and details about package types, metadata, and any prerequisites or scripts for each supported operating system. It also discusses how to view packages assigned to nodes, and how to upload, overwrite, edit, delete, deprecate, and download packages.

Chapter 6: Application Provisioning Setup – provides information about how to prepare for provisioning servers with software including a description of the Software Tree, managing nodes on the Software Tree, an overview of modeling software attached to nodes, configuration settings, viewing software attached to nodes, adding and removing software packages from nodes, and changing the installation order of software. It also discusses the concepts of inheritance and dependencies. Managing custom attributes is described, as well as an in-depth discussion of all aspects of creating templates for operating systems, patches, applications, and service levels.

Chapter 7: Application Provisioning – provides information about installing and uninstalling software, and using templates for installation.

Chapter 8: Patch Management Subsystem – provides information about managing patches including testing and installation standardization, operating systems and patch types, the roles of the patch administrator and system administrator in applying patches, and the permissions required for performing patch management. It also describes how to set up the patch management system, upload patches, the administration of patches using the Opsware Command Center, installing and uninstalling patches, and an overview of the Microsoft Patch Update Wizard.

Chapter 9: Reconcile – provides an overview of the server reconciliation process and how it works, including an in-depth discussion of how to perform a reconcile, reconcile and package metadata, installation and uninstallation order, “adopted” software, patches and reconcile, and reconcile preview. Types of reconcile are also discussed, along with reconcile on supported operating systems, reconcile and scripts, reconcile output, assigning servers to and removing servers from nodes, and an overview of the Reconcile Software Wizard.

Chapter 10: Script Execution Subsystem – provides a description of script types, permissions required for scripting tasks, creating, editing, and deleting scripts and information about script version history. It also discusses script execution for all script types, script execution results, and scripting error resolution.

Chapter 11: Automated Configuration Tracking – provides information about configuration tracking policies, the supported types of configuration files and databases, how changes are detected, node-based tracking policies, reconciling tracking policies, customizing tracking policies, and viewing the tracking policies for a specific server. It also discusses reconciling customized tracking policies, performing manual backups, viewing backup history, restoring backed up files, and enabling and disabling configuration tracking.

Chapter 12: Code Deployment & Rollback – provides information about uploading code and content to staging, and setting up and performing services, synchronizations, and sequences to deploy code and content to managed servers.

Appendix A: Opsware Command Line Interface – provides information about installing the Opsware CLI, file transfer commands, command syntax, command options, and supported operating systems and package types.

Appendix B: Agent Upgrade Tool – provides procedures for upgrading all the Opsware Agents in a facility to the latest version.

Appendix C: OS Installation Integration – provides a description of how the Opsware System handles OS installation integration with third-party installation technologies, how to perform integration including Opsware Agent Installer commands and options, and how to integrate with specific operating system installation technologies.

Appendix D: Communications Test Troubleshooting – provides troubleshooting information to diagnose Opsware Agent unreachability problems.

Appendix E: Glossary – defines terms that are unique to the Opsware System.

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 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 iconographic conventions.

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 4.7 User's Guide* is intended to be read by the system administrator who is responsible for performing the day-to-day functions of managing servers, provisioning operating systems, uploading packages, setting up the Software Tree 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 4.7 Administration Guide* is intended to be read by Opsware administrators who will be responsible for setting up accounts for users, creating user groups and additional Opsware administrators, assigning permissions for different levels of operation and access, adding customers and facilities, and monitoring and diagnosing the health of the Opsware System components.

The *Opsware System 4.7 Installation Guide* is intended to be used by system administrators who are responsible for the installation of Opsware System in a facility. It documents how to run the Opsware Installer and how to configure each of the components.

Types of Opsware Users

Depending on their access level, Opsware users are able to do tasks such as provision and configure servers.

The following table identifies the types of Opsware users and their responsibilities.

OPSWARE USER	RESPONSIBILITIES
Data Center and Operations Personnel	After manually racking and stacking servers, manage customer facilities and boot bare-metal servers over the network or from an Opsware boot floppy.
System Administrators	Install operating systems and applications (for example, Solaris 5.7 or WebLogic 6.0 Web Server), and upgrade servers; create operating system definitions and set up application provisioning.

OPSWARE USER	RESPONSIBILITIES
Site Engineers and Customer Project Managers	Deploy custom code on servers.

In addition to the Opsware users listed in table 1-3, this guide describes the following three types of users:

- **End Users** are the system administrators who are responsible for performing the day-to-day functions of managing servers, provisioning operating systems, and installing applications on servers.
- **Policy Setters** are the power users who are responsible for architecting what the Opsware System will do in the managed environment; for example, they determine which operating systems can be installed on your managed servers and how those operating systems will be configured during installation. Policy setters, for example, prepare specific subsystems in the Opsware System by defining the Software Tree, preparing Operating System Definitions, and acting as Patch Administrators to approve patches for installation in the operational environment.
- **Opsware administrators** are the users, with special training and information, who support the Opsware System by setting up accounts for users, creating user groups and additional Opsware administrators, assigning permissions for different levels of operation and access, adding customers and facilities, and monitoring and diagnosing the health of the Opsware System components. Opsware administrators need to understand how Opsware System features operate to support users and the Opsware System.

How to Read the User's Guide

This guide documents how to use the Opsware System to manage servers in the operational environment. This guide is intended for Opsware end users and policy setters.

This guide does *not* document how to install an Opsware System in a facility or how to set up the Opsware System so that a specific user can start managing servers in the Opsware Command Center. See the *Opsware System 4.7 Administration Guide* for information about how to create user accounts for the people in your organization.

This guide provides information for Opsware policy setters and Opsware end users.

Opsware policy setters, who are responsible for preparing specific subsystems in the Opsware System, should read the following chapters in this guide:

- To set up OS provisioning so that you can use the Opsware System to install operating systems on bare-metal servers, see Chapter 3, "OS Provisioning Setup" on page 175 of this guide.

  A policy setter must perform set up tasks so that the OS Provisioning Subsystem will install specific operating systems on managed servers.

- To set up automated patch management so that you can use the Opsware System to patch Opsware-managed servers, see Chapter 8, "Patch Management Subsystem" on page 403 of this guide.

  A policy setter must perform set up tasks so that the Patch Management Subsystem will patch specific operating systems and applications.

- To set up application provisioning with the Opsware System by defining the Software Tree and uploading packages, see Chapter 5, "Package Management" and Chapter 6, "Application Provisioning Setup" on page 307 of this guide.

  A policy setter must perform set up tasks so that the Application Provisioning Subsystem will install specific applications on Opsware-managed servers.

- To install the Opsware Command Line Interface to upload packages into the Software Repository, see Appendix A, "Opsware Command Line Interface."

- To set up the Code Deployment & Rollback Subsystem before using it to push code and content to managed servers, see Chapter 12, "Code Deployment & Rollback" on page 539 of this guide.

Opsware end users, system administrators who are responsible for managing servers running in the operational environment, should read the following chapters in this User's Guide:

- To understand how the Opsware System manages servers, see Chapter 2, "Server Management" on page 29 of this guide.

- To install Opsware Agents on the existing servers in your operational environment, Chapter 2, "Server Management" on page 29 of this guide.

- To install and uninstall applications on managed servers, see Chapter 7, "Application Provisioning" on page 385 of this guide.

- To patch managed servers, see Chapter 8, "Patch Management Subsystem" on page 403 of this guide.
- To run distributed scripts on managed servers simultaneously, see CChapter 10, "Script Execution Subsystem" on page 455 of this guide.
- To push code and content to managed servers, see Chapter 12, "Code Deployment & Rollback" on page 539 of this guide.
- To track configuration changes for managed servers, see Chapter 11, "Automated Configuration Tracking" on page 493 of this guide.

Chapter 1: Opsware System

IN THIS CHAPTER

This chapter provides an overview of the Opsware System and its core features. It also provides an introduction to the Opsware Command Center, the user interface to the Opsware System.

The topics covered in this chapter include:

- Opsware System Overview
- Opsware System's Model-Based Approach
- Supported Operating Systems
- Opsware Automation Subsystems
- Web Service APIs
- Multimaster Support
- Getting Started with the Opsware System

Opsware System Overview

Opsware System provides a core set of features that automate critical areas of server and application operations – including the provisioning, deployment, patching, and change management of servers – across major operating systems and a wide range of software infrastructure and application products.

Opsware System does not just automate your operations, it also allows you to make changes more safely and consistently because you can model and validate changes before you actually commit the changes to a server. Opsware System automation helps ensure that modifications to your servers work the first time that you attempt them, thereby reducing the risk of downtime.

Using Opsware System, you can coordinate many operations tasks, across many IT groups with everyone working with the same understanding of the state of servers, applications, and configurations. This coordination ensures that all IT administrators have full knowledge of the current state of the environment before further changes are made.

Opsware System allows you to incorporate and maintain operational knowledge that has often been gained through long hours of trial-and-error processes. After an administrator has found and tested a procedure or configuration, that knowledge can be translated into a model that is stored in a central repository. The ability to store this knowledge in a central repository allows you to continue to benefit from the operational knowledge gained by your system administrators even if they are no longer working in your organization.

Figure 1-1 provides an overview of how Opsware System automates server and application operations across all major platforms and a wide range of applications. Each feature that is shown in the diagram is discussed in the following sections.

Figure 1-1: Overview of Opsware System Features

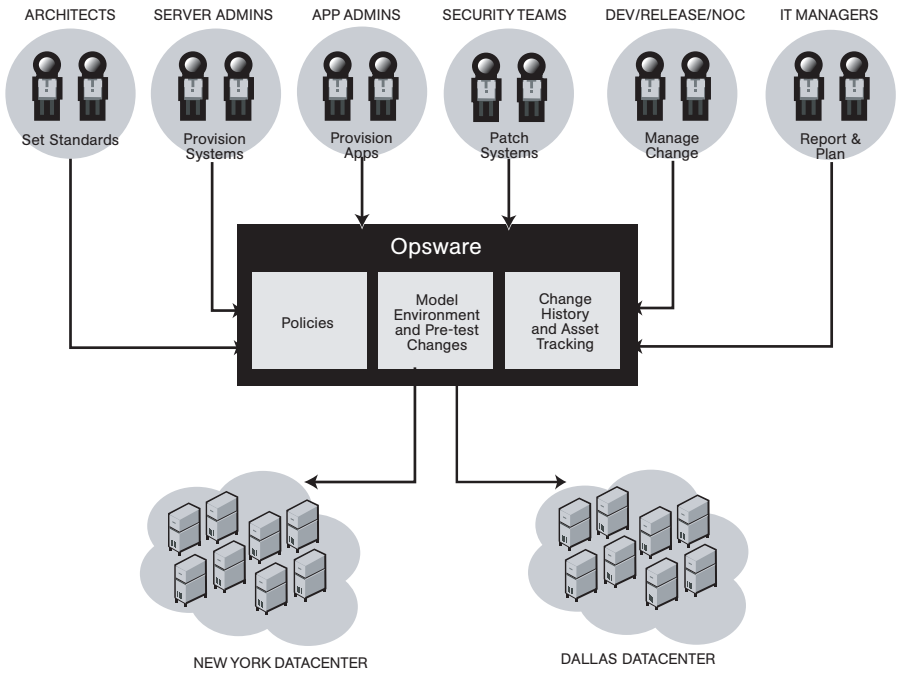

Opsware System's Model-Based Approach

Opsware System employs a methodical, model-based approach, using a centralized information system as the starting point for any changes made to the physical environment. This information system approach to initiating change into the environment means that when team members need to make a change, they model the change first to verify and validate the change. Next, they let the Opsware System propagate that change out to the environment in a secure, consistent, quick and reliable manner.

Figure 1-2 gives a high-level view of the steps involved when users interact with the model.

Figure 1-2: High-Level View of User-Model Interaction

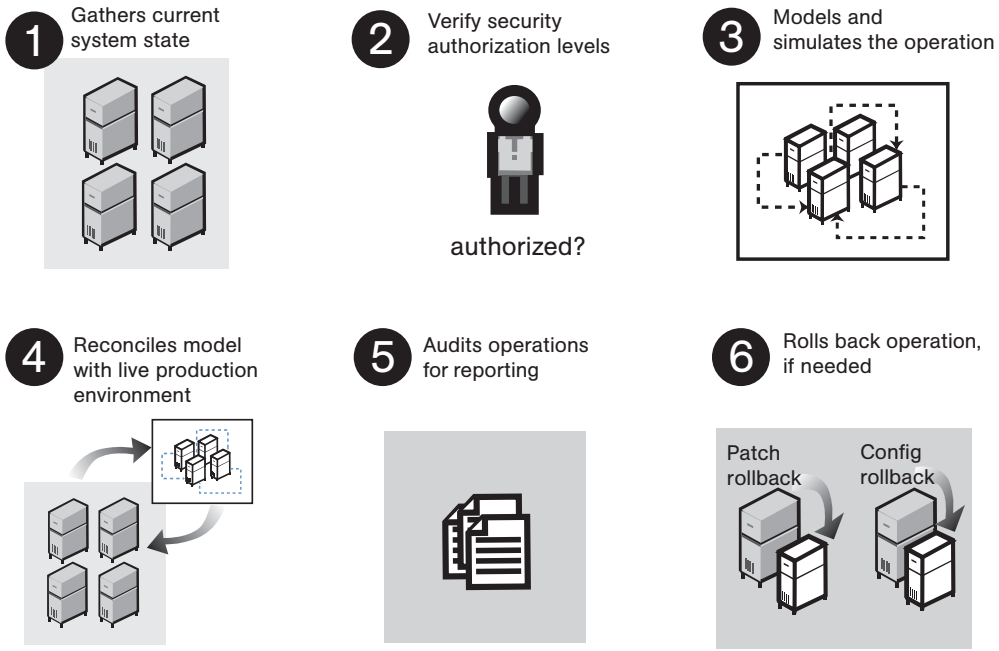

The Opsware Automation Platform includes three additional components to automate core functions. The role of these components is to provide policies, environment state information, and a comprehensive audit log for activity and asset reporting.

- **Software Tree**

  The Software Tree records a variety of information for software applications and operating systems, including, for example, data about how changes to a given software application might impact other existing applications.

- **Environment Tree**

  The Environment Tree manages characteristics about a customer's unique data center environment, including hardware, location of servers, network infrastructure, application names, business units and service levels assigned to servers and applications. The information contained in the Environment Tree, combined with the data contained in the Software Tree, is utilized by the Opsware Automation Platform to model and simulate operational changes before they are executed in the production environment.

- **Modeling and Change Simulation Engine**

  Opsware System enables users to first model and simulate proposed operational changes to their environment before propagating these changes to production servers and applications. Utilizing the information contained in the Software and Environment Trees, the Modeling and Change Simulation Engine maintains a model of the various hardware and software configurations and other customer characteristics associated with each of the production servers under Opsware System's control.

  Before committing any proposed changes to the production servers, this engine first conducts an impact analysis of the requested operation, enabling customers to test and validate changes prior to executing them in the production environment. This feature of Opsware System is designed to increase the success rates associated with initial deployments, improve the accuracy and security of these changes, and reduce application downtime.

Supported Operating Systems

Table 1-1 shows the supported operating systems and their version numbers.

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

OPERATING SYSTEM	VERSIONS
AIX	AIX 4.3 AIX 5.1 AIX 5.2 AIX 5.3

OPERATING SYSTEM	VERSIONS
HP-UX	HP-UX 10.20 HP-UX 11.00 HP-UX 11.11/11i
Sun Solaris	SunOS 5.6 SunOS 5.7 SunOS 5.8 SunOS 5.9
Fujitsu Solaris	Fujitsu 2.8 Fujitsu 2.9
Windows	Windows NT 4.0 Windows 2000 Server Family Windows Server 2003
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
SUSE Linux	SUSE Linux Enterprise Server 8.0 SUSE Linux Standard Server 8.0 SUSE Linux Enterprise Server 9.0

Opsware Automation Subsystems

Opsware System is made up of a set of Opsware Automation Subsystems. Opsware Automation Subsystems are the components that automate particular IT processes.

Automation Subsystems are designed to replace ad hoc, error-prone, manual processes. For example, by using the Opsware OS Provisioning Subsystem, users can set standards for different types of servers and automatically provision the servers, saving time and

ensuring that operating system builds are consistent. By using the Patch Management Subsystem, users can establish polices about how patches are installed. Opsware System uniformly enforces those polices.

The following Opsware Automation Subsystems are currently available as part of the Opsware System:

- Software Provisioning
- Operating System Provisioning
- Patch Management Automation
- Code Deployment and Rollback
- Configuration Tracking
- Script Execution
- Data Center Intelligence Reporting

All Opsware Automation Subsystems support cross-platform environments and are designed to automate both new and existing data center environments. See Figure 1-3.

Figure 1-3: Opsware System Automation Subsystems

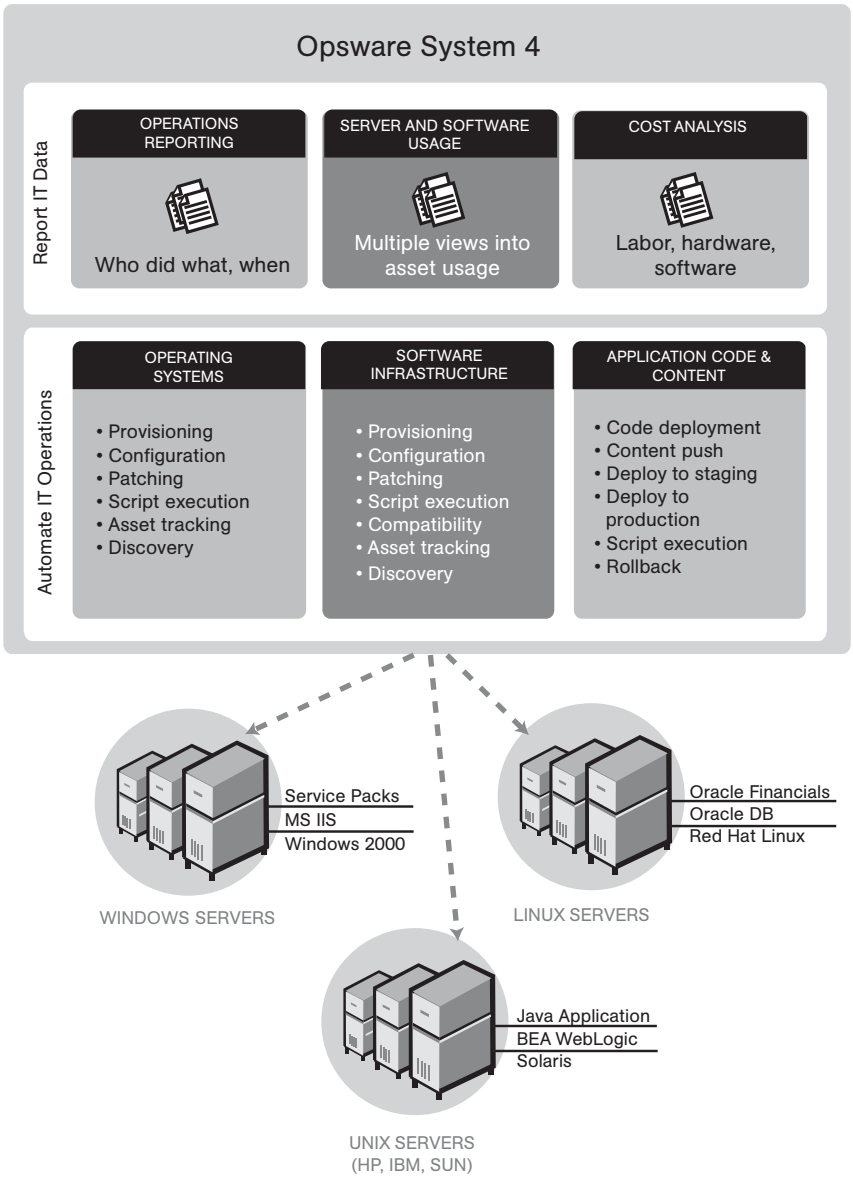

Software Provisioning

The Software Provisioning Subsystem gives system administrators a systematic way to install, configure, and remove packaged software across Windows, Unix, and Linux servers distributed across many different data centers. Opsware System's unique model-based approach enables many different teams, such as the system administration team, the database team, and the application development team, to manage the same set of servers. Each of these teams has a common view of the environment.

The Software Provisioning Subsystem leverages Opsware System's model-based approach, which provides the following unique capabilities and benefits:

- Detailed information about the latest system state and configurations

  The Software Provisioning Subsystem automatically creates and updates two lists: the list of software that users indicate should be installed on a server and the list of software that is actually installed on a server. By maintaining this detailed model of the server's current state, Opsware System helps keep different IT groups managing the same server in sync and ensures that all groups making server changes are working with the same knowledge of the current state of the environment.

  Using this model, Opsware System enables multiple groups to manage the same server without stepping over each other's changes. An accurate model of the software installed on a server, granular role-based access control, a unified audit trial, and the ability to roll back changes all contribute to Opsware System's ability to coordinate the activities of many different administrators managing the same server.

- Integration with other automation functions

  The Software Provisioning Subsystem is fully integrated with other Opsware Automation subsystems, enabling software provisioning to be performed automatically with other tasks, such as operating system provisioning. Because software provisioning shares the same environment model as the other functions, the state of the environment is always known. This means that different groups, such as OS administrators, application administrators, security administrators, and others, can work together and communicate more effectively.

- Simulation of software installation and removal

  Opsware System's provisioning engine simulates installation and uninstallation actions before it applies changes to production servers. Users can view the list of software packages to be added or removed before they authorize Opsware System to execute

the change. This ensures that all changes are pre-tested and validated before propagating changes to the production environment.

- An up-to-date model of the actual server environment

  Opsware System regularly refreshes its view of what is installed on a server, including both hardware and software. This real-time understanding of server state and configurations ensures that administrators provision the right software to the right servers at the right time. It also ensures that dependencies and pre-requisites are checked and installed as needed.

- Sophisticated role-based access control

  Opsware System enforces a security policy that allows only authorized users to install or remove particular types of software on a particular server. For example, companies can define an access control rule that permits only DBAs to add or remove database software from a server.

- A unified audit trail

  Opsware System maintains a comprehensive audit trail of the software that Opsware users install, configure, and remove from a server. When combined with the additional events that Opsware System tracks – including configuration updates, business application pushes and rollbacks, hardware upgrades, and executed scripts – organizations gain a complete view of server activity over time.

- The ability to roll back to a last known good state

  The Software Provisioning Subsystem allows users to back out of software provisioning operations. In the event an upgrade or installation goes awry, administrators can back out the change to return to the last known good state.

- Ability to store powerful name-value pairs

  Opsware System helps organizations increase software package re-use by enabling administrators to install the same software package on different servers. Server-specific configuration values are fetched from Opsware System (or calculated based on those values).

Operating System Provisioning

The OS Provisioning Subsystem gives administrators the ability to provision operating system baselines onto bare metal servers quickly, consistently, and with minimal manual intervention. Bare metal OS provisioning is a key part of the overall process of getting a server into production.

Benefits of the OS Provisioning Subsystem include the following items:

- Integration with the other subsystems of Opsware System

  Because the OS Provisioning Subsystem is integrated with the suite of Opsware System automation capabilities, including patch management, software provisioning, and distributed script execution, handoffs between IT groups are seamless. The Opsware System ensures that all IT groups are working with a shared understanding of the current state of the environment, which is an essential element of delivering high-quality operations and reliable change management.

- The ability to easily update server baselines without re-imaging servers

  Unlike many other OS provisioning solutions, systems provisioned with the Opsware System can be easily changed after provisioning to adapt to new requirements. Key to this benefit is Opsware System's use of templates and its installation-based approach to provisioning.

- Flexible architecture designed to work in many environment

  Opsware engineers carefully designed the OS Provisioning Subsystem to handle many different types of servers, networks, security architectures, and operational processes. Opsware System works well in floppy or CD- or network-boot environments, with scheduled or on-demand workflows, and across a large variety of hardware models. This flexibility ensures that you can provision operating systems to suit your organization's needs.

Opsware System automates the entire process of provisioning a comprehensive server baseline, which typically consists of the following tasks:

- Preparing the hardware for OS installation
- Installing a base operating system and default OS configuration
- Applying the latest set of OS patches, the exact list of which depends on what applications are going to run on the server

- Installing system agents and utilities such as SSH, PC Anywhere, backup agents, monitoring agents, or anti-virus software
- Installing widely-shared system software such as Java Virtual Machines
- Executing pre- or post-installation scripts that configure that system with values such as a root password

Patch Management Automation

The Patch Management Automation Subsystem provides two features critical to patch management: the ability to react quickly to newly-discovered threats and the degree of control required to ensure that a new patch has been properly tested and installed in a uniform way.

Opsware System has a deep understanding of native patch formats and structure. System administrators upload patches directly into Opsware System, which understands and respects the structure of those patches in their native forms. It treats Solaris patch clusters, for example, differently from Windows Hotfixes or AIX APARs. Native patch support greatly increases both the flexibility and reliability of patch installation.

The Patch Management Subsystem provides the following features:

- Scalable, cross platform patch deployment
- Reduced risk throughout automated patch rollback
- A central, shared patch repository to improve access
- Secure access control
- The ability to install patches on one server, or simultaneously on many servers
- The ability to schedule automated future installation (for example, to take advantage of maintenance windows)
- The inclusion of patches in the template for an operating system, so all newly provisioned servers receive the most up-to-date set of recommended patches

Code Deployment and Rollback

Opsware System automates code and content deployment to reduce the risk and time requirements associated with pushing new code to production. The Code Deployment & Rollback Subsystem (CDR) provides an automated system for deploying code (such as, ASP, JSP, JAR, Java, C++, and Perl files) and content (such as, HTML, JPEG, GIF, and PDF files). Specifically, CDR includes the following capabilities:

- Push code from staging or development environments to production environments
- Synchronize code and content across multiple servers and locations
- Automatically roll back to the previous version of code or content
- Sequence multiple, complex deployment steps into repeatable workflows
- Manage changes across heterogeneous operating systems

Configuration Tracking

The Configuration Tracking Subsystem tracks, backs up, and recovers critical software and system configuration information across Unix and Windows servers.

System administrators set up policies that describe which configuration files and databases to track, and what actions to take when a change in configuration is detected. Policies can be assigned to software, individual servers, groups of servers, and customers, and applied either locally or globally across data centers.

When the Opsware System notices a server configuration change, it can log the change, notify administrators about the change with email, or back up the configuration, depending on the policy set by the administrator.

When a bad configuration change forces administrators to roll back to a previous version, they can use Opsware System to restore the configuration file to the saved version of the configuration. By notifying users about configuration changes – and maintaining a version history of those changes – organizations can quickly diagnose problems related to configuration errors and roll back to a known good state. In addition, this capability helps teams plug security holes inadvertently created by bad server configurations.

Typically, system administrators define configuration-tracking policies on a per-application basis. So for example, a policy for BEA WebLogic might specify, "monitor the `weblogic.conf` file, notify app-server-admins@company.com of any changes, and maintain a version history of any changes that occur for 30 days." After a policy is defined in this fashion, administrators can apply the policy to all the WebLogic servers running in their environment, or to specific servers.

Script Execution

The Script Execution Subsystem enables system administrators to share and run ad-hoc or custom scripts across an entire farm of Opsware-managed servers. By executing scripts with Opsware System instead of manually, administrators benefit by using the following features:

- Parallel script execution across many Unix and/or Windows servers, saving time and ensuring consistency
- Role-based access control, ensuring only authorized administrators can execute scripts on hosts to which they have access
- The ability to control access to scripts by storing them in private or in public libraries
- The ability to see and download script output one server at a time or in a consolidated report, which captures output from all servers in a single place
- The ability for scripts to be mass-customized by accessing the information in Opsware System about the environment and state of servers which is critical to ensuring that the right scripts are executed on the right servers
- A comprehensive audit trail that reports who, what, when, and where a particular script was executed
- The ability to rollback changes (when used in conjunction with the Configuration Tracking Subsystem)
- Automatic backup of all private and shared scripts to all other Opsware-managed data centers (when used in conjunction with Opsware's Multimaster Replication Engine)

Because the Script Execution Subsystem is an integrated part of Opsware System, administrators enjoy unique benefits when compared to standalone script execution tools:

- Using known system state and configuration information to customize script execution, users can tailor each script by referencing and accessing the rich store of information in the Opsware System, such as the customer or business that owns the server, whether the server is a staging or production server, which data center the server is located in, and custom name-value pairs.
- Sharing scripts without compromised security. Users can share scripts with each other without compromising security because the Opsware System maintains strict controls on who can execute scripts on which servers and generates a comprehensive audit trail of script execution.

Data Center Intelligence Reporting

Every change made to your managed servers is recorded in Opsware System's Model Repository. The Model Repository maintains precise information about the state and configuration of every server under your management.

You can now take advantage of this information though Opsware System's Data Center Intelligence Reporting (DCI) component. The DCI provides accurate, detailed and up-to-date information about your operational environment. The DCI provides a new level of visibility into your operational environment that can help organizations make better decisions.

DCI reporting provides the following features and benefits:

- Exact information about the latest system state and configurations

  DCI reports display the most accurate and up-to-date information available, even during periods of frequent and substantial change. This level of accuracy reduces your risk of making the wrong decisions because of old data.

- Visibility across the data center environment

  Opsware System provides a comprehensive view across all operating systems and locations, allowing IT managers to generate on-demand snapshots driven from a single, high-quality data source. The ad-hoc capability allows you to view a variety of report types, filter by specific criteria, and display summary graphics or list views. In addition, a set of Quick Reports are pre-designed for one-click access to real-time information from the Reports Home page.

- Accurate and detailed change history information

  When a server's performance suddenly degrades, the best way to diagnose the cause is to learn what changes have been made to the server and who made the changes. Often, talking with the people who made the changes can help you understand what's causing the performance degradation.

  In most data centers, however, it's often difficult, if not impossible, to find out a server's exact change history, since records are not accurately kept. But Opsware System maintains a detailed record of each change: who made the change, what was the nature of the change, and when did the change occur. This record is presented in a comprehensive series of reports; these reports can significantly reduce the time and effort in debugging server and software problems.

- A comprehensive set of patch reports

  One of the most time-consuming aspects of patching servers is identifying the vulnerable servers. Data collection for this task typically involves manually logging onto each server to see if it contains a particular version of software, what patches are already installed on the server, and what patches are *not* installed on the server.

Opsware System helps administrators avoid this up-front effort by offering a comprehensive set of patch management reports.

- The ability to extend the DCI reports

  You can also create new reports or modify the reports that ship with Opsware System. The Opsware System provides the information about the database that you need to create your own reports.

  The Reports Home page checks for any new custom reports that you create, and presents them on the Reports Home page for easy access to all users. These reports are created by using the readily available Crystal Reports Designer 9.

  New reports can be extended to integrate with your own data sources (databases, spreadsheets, XML, and so forth), creating a powerful tool for more advanced data intelligence.

  See the *Opsware System Data Center Intelligence Administrator's Guide* for information about how to set up the DCI Reporting component.

  See the online Data Center Intelligence help and tutorial documentation for information about how to use and run the reports.

The Opsware Data Center Intelligence Reporting component is an optional component. By default, it is not installed with the Opsware System. If this reporting component is not available for your organization, contact your Opsware Support Representative for information about how to obtain it so that you can generate reports. The DCI component must be installed and running in order to access the online documentation.

Web Service APIs

The Opsware System exposes a new Web services interface to facilitate the integration of operations and business support systems with the Opsware System. The new Opsware Web Services APIs allow other IT systems, such as customers' existing monitoring, trouble ticketing, billing, and virtualization technology, to quickly exchange information with Opsware System. The Opsware Web Services APIs thereby extend the value of Opsware's DCI software across the IT organization.

Developers can now use any Web Services-enabled development environment, from Microsoft Visual Studio.NET to BEA WebLogic Workshop to simple Perl scripts, to develop monitoring and reporting applications that invoke procedures through this interface. Opsware System support for leading Web Services standards, such as XML, SOAP, and WSDL, ensures that enterprises are not locked into proprietary protocols that make integration more complex and costly.

Multimaster Support

With the Opsware Multimaster Replication Engine, customers can store and maintain a blueprint of software and environment characteristics of each data center in multiple locations so the infrastructure can be easily rebuilt in the event of a disaster. The Multimaster Replication Engine not only provides the ability to replicate an environment in case of a disaster, but can also assist in data center migration activities as well as knowledge sharing across the enterprise.

Through the Multimaster Replication Engine, the Opsware System provides the ability to easily rebuild server and application environments, provision additional capacity, distribute updates, and share software builds, templates and dependencies – across multiple data centers and from one user interface.

Getting Started with the Opsware System

The following section discusses getting started with the Opsware System and contains the following topics:

- Getting Access to Opsware Features
- Using the Opsware Command Center User Interface
- My Profile
- Search
- My Servers
- Mouseover Icon Tooltips
- Supported Browsers
- Configuring Your Browser

Getting Access to Opsware Features

To log in to the Opsware Command Center, your Opsware administrator must have created a login ID and password for you. The Opsware administrator also assigns permissions that control which features you can access and what actions you can take when you use them. Figure 1-4 shows these features.

Figure 1-4: Opsware System Navigation Panel, All Permissions View

Using the Opsware Command Center User Interface

Figure 1-5, which shows the Home page, appears when you log in or when you click the Home hyperlink in the navigation panel.

Figure 1-5: Opsware System Home Page

The time zone that appears in the upper right corner of the Home page is taken from the time zone preference that was defined for you when your profile was created. Consequently, the date and time information that displays throughout the Opsware Command Center is for that time zone. The occasional exceptions however are always labeled GMT.

Tasks

The Tasks area of the Home page displays links to the wizards that you have permissions for, a link to the Data Center Intelligence reports if you have that module installed, and a link to the Code Deployment page if you have that permission. If you do not have permissions to a task in this area, the task name still displays, but it is italicized and it is not an active link. Figure 1-6 shows the Tasks area with all permissions enabled.

Figure 1-6: Tasks Area of the Home Page

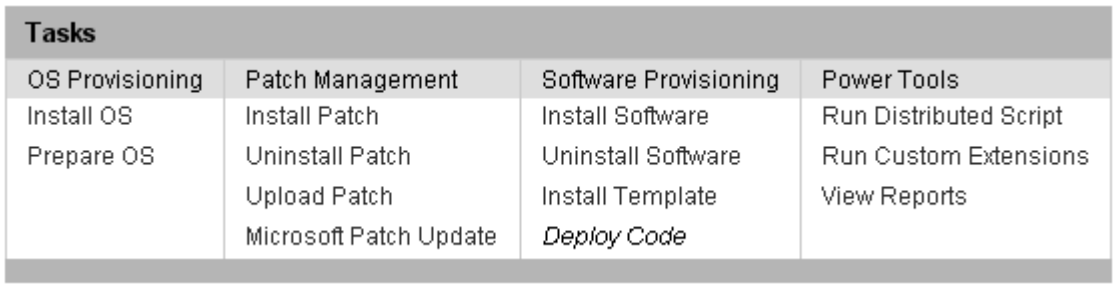

My Jobs

The My Jobs area of the Home page is populated with details of the jobs that you have run, jobs that are currently in progress, or jobs that you have scheduled to run, including the name of the job, the start time, the number of servers affected by the job, and the status of the job. If there are more than six jobs, you can see the rest of them by clicking the See All link, which also shows the total number of jobs in parentheses, as Figure 1-7 shows.

Figure 1-7: My Jobs Area of the Home Page

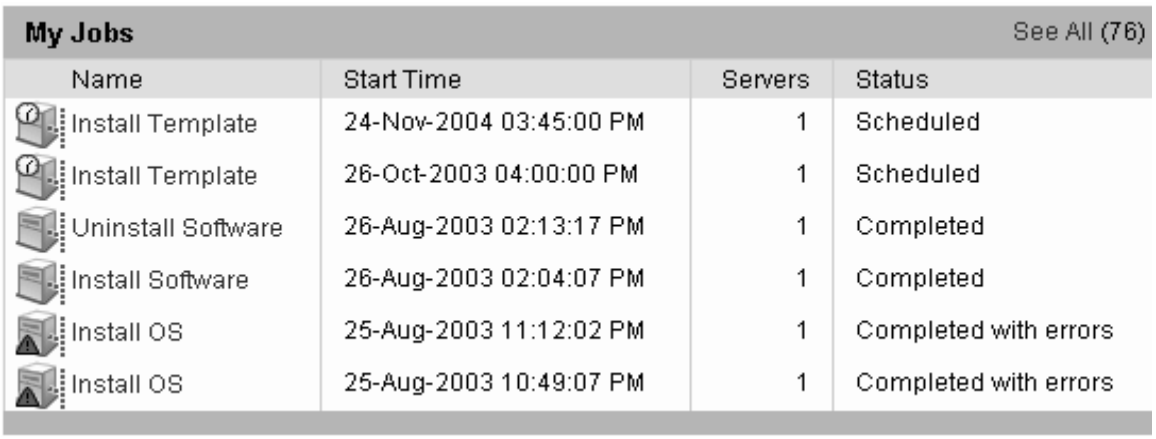

My Customers

The My Customers area of the Home page is populated with customer information, including unreachable servers associated with a customer and the total number of servers associated with that customer. To select the customers to display in the My Customers area of the Home page, click the Edit button and click the check box next to the customer name. See Figure 1-8.

Figure 1-8: My Customers Area of the Home Page

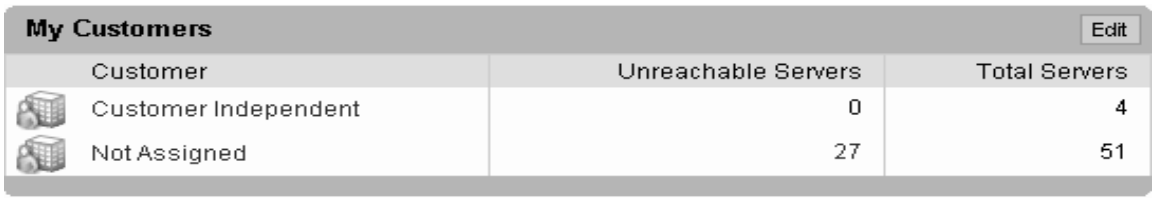

Navigation

Top-level navigation from the Home page is simple. To access any of the features in the navigation panel, click the feature name. To access any of the wizards or other features in the Tasks area of the Home page, click the name of the task.

After you select a task or a feature, other pages appear, which might have one or more of the following means of navigation:

- Clicking a hyperlinked name to display a page

  For example, if you select Software ➤ Applications (assuming that you have the correct permissions), the page that appears shows all of the applications that have been uploaded so far. Each application name is a hyperlink that takes you to another page with information about that application already displayed.

- Clicking a tab to display a page

  For example, when you select Applications and click one of the hyperlinked application names, the resulting page shows a row of tabs, like Figure 1-9 shows.

Figure 1-9: Example of Tabs

  Each tab displays a page, each with its own buttons and functionality.

- Clicking a button to display a page

  For example, when you click the Custom Attributes tab, a page appears with several buttons: Add, Delete, and Copy, which are common to each page called by these tabs,

and Add Custom Attribute, which is unique to this particular tab. You will find similar functionality on all tabbed pages in the Opsware System.

Navigation Panel

The items that appear in the navigation panel depend on the permissions the user has. Clicking an item on the navigation panel displays that feature in the main part of the Home page. For a user with all permissions, the following links appear:

Home – Displays the top level of the Opsware Command Center. The Home page is described in this section of the guide. Wizards are documented in their respective functional areas of the system.

My Jobs – Displays the My Jobs page, showing jobs completed during the previous 30 days, jobs currently in progress, and currently scheduled jobs. This page is an expanded view of the contents of the My Jobs area of the Home page and has the same effect as clicking the Show All button in the My Jobs area of the Home page.

Servers expands to display these selections:

- Server Pool – Use filters to display a list of servers, install operating systems on the servers, and delete the servers. See Chapter 4, “Operating System Provisioning” on page 237 of this guide.
- Managed Servers – Use filters to display a list of servers and perform operations on them such as edit server values, assign, reconcile, run scripts, and add servers to My Servers. See Chapter 4, “Operating System Provisioning” on page 237 of this guide.
- Server Groups – Define server groups and server types, assign servers to groups. See Chapter 2, “Server Management” on page 29 of this guide.
- Server Search – Search for specific servers using default criteria or user-defined criteria. See Chapter 2, “Server Management” on page 29 of this guide.

Software expands to display these selections:

- Operating Systems – Prepare operating systems for installation and delete existing operating systems. See Chapter 4, “Operating System Provisioning” on page 237 of this guide.
- Patches – Prepare patches for installation and define patch preferences. See Chapter 8, “Patch Management Subsystem” on page 403 of this guide.

- Applications – Define nodes in the Software Tree, identify packages to attach to nodes, assign servers to nodes. See Chapter 6, "Application Provisioning Setup" on page 307 of this guide.
- Templates – Define templates and associate operating systems, patches, applications and service levels. See Chapter 6, "Application Provisioning Setup" on page 307 of this guide.
- Packages – Upload packages, browse uploaded packages, and search for packages. See Chapter 5, "Package Management" on page 267 of this guide.
- Scripts – Run scripts, upload scripts, create new scripts. See Chapter 10, "Script Execution Subsystem" on page 455 of this guide.

Environment expands to display these selections:

- Customers – Create new customers and edit or delete existing customers. See the *Opsware System 4.7 Administration Guide*.
- Facilities – Create new facilities, edit facility properties, and assign and edit custom attributes for facilities. See the *Opsware System 4.7 Administration Guide*.
- Hardware – A read-only view of servers categorized by hardware manufacturer and model, and their related information. See Chapter 2, "Server Management" on page 29 of this guide.
- Service Levels – Define service levels and custom attributes. See Chapter 2, "Server Management" on page 29 of this guide.
- IP Ranges – Identify and create IP ranges and IP range types. See Chapter 2, "Server Management" on page 29 of this guide.
- IP Range Groups – Create IP Range Groups. See Chapter 2, "Server Management" on page 29 of this guide.

Code Deployment expands to display these selections:

- CDR Home – The exact CDR links that you see in the Code Deployment area are based on the permissions that you have for the customer you want to work with. See Chapter 12, "Code Deployment & Rollback" on page 539 of this guide.
- Service Management – Create, modify, and delete service definitions. Services define the location and commands to manipulate applications on hosts. See Chapter 12, "Code Deployment & Rollback" on page 539 of this guide.

- Run Service – Perform service operations on one or more hosts, or request that a service operation be performed on your behalf. Service operations include starting or stopping applications, cutting over or rolling back code, and backing up or restoring code. See Chapter 12, “Code Deployment & Rollback” on page 539 of this guide.
- Sync Management – Create, modify, and delete synchronization definitions. Synchronizations define the path for pushing code from a source host to one or more destination hosts. See Chapter 12, “Code Deployment & Rollback” on page 539 of this guide.
- Synchronize – Perform a synchronization to one or more hosts, or request that a synchronization be performed on your behalf. See Chapter 12, “Code Deployment & Rollback” on page 539 of this guide.
- Sequence Management – Create, modify, and delete sequence definitions. Sequences allow the grouping of service operations and synchronization operations to define higher level code deployment operations. Chapter 12, “Code Deployment & Rollback” on page 539 of this guide.
- Run Sequence – Perform a pre-defined sequence of service operations and synchronizations on one or more hosts, or request that a sequence be performed on your behalf. See Chapter 12, “Code Deployment & Rollback” on page 539 of this guide.
- View History – Get information about previously run Code Deployment operations. See Chapter 12, “Code Deployment & Rollback” on page 539 of this guide.

Reports – Displays the Data Center Intelligence top-level page (if this module is installed). See the *Administrator’s Guide to the Opsware Data Center Intelligence Reporting Server* document.

Administration expands to display the following features. For more information about these features, see the *Opsware System 4.7 Administration Guide*.

- Users & Groups – Administrators use this feature to create user groups, define permissions for those groups, create new administrators, and add users to groups.
- Sessions – Displays information about each user currently online, including the name of the user, how long they have been online, how long ago they last accessed the Opsware Command Center, and what features they accessed.
- Server Attributes – Define and edit server use attributes, enable them for code deployment, and define deployment stages. Also define and edit server deployment stage attributes.

- System Configuration – Contains the configuration parameters that define how the Opsware System works in your environment. This selection is only used at the direction of Opsware Inc.
- System Diagnosis – Runs a series of tests on Opsware components to make sure that they are functioning correctly.
- Opsware Software – Provides a view of the properties, custom attributes, installation order, and history of the software attached to the Opsware nodes in the system.

The Administration set of features is only available if you are logged in as an Opsware administrator.

My Profile

You can update your own personal user information without the assistance of the Opsware administrator with the My Profile link. You can change your first and last name, your contact information, your password, and your time zone and date display preferences.

Search

Click the Search link at the top of the Home page to open the dialog box that Figure 1-10 shows.

Figure 1-10: Opsware Command Center Search Function

You can search for servers, applications (nodes), packages, and jobs by making the selection from the drop-down list, and then entering an identifying string in the text box.

My Servers

The My Servers feature provides a convenient place to store a set of servers that have been selected and stored using the Add to My Servers function in Managed Servers. You might use it as a shortcut to the servers you work with most often, or as a way to gather a group of servers when you want to apply the same changes to all of them. All functions that are available in the Managed Servers page are also available from within My Servers.

Mouseover Icon Tooltips

When an icon appears on a page in the Opsware Command Center, a tooltip displays information about the icon when your mouse pointer hovers over it. For example, server icons display messages such as *Available* or *Build Failed* to describe the state of the server. Packages and patches display messages such as *Available*, *Managed*, *Unmanaged*, and so forth.

Supported Browsers

Table 1-2 shows the supported browsers.

Table 1-2: Supported Browsers

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

Configuring Your Browser

Make sure that your browser is configured as follows:

- You must configure your browser to accept cookies and use Java to function correctly with the Opsware System.
- Using a Web browser that provides 128-bit encryption is recommended.
- Using a pop-up blocker might prevent some functions in the Opsware System from working correctly. Either disable the pop-up blocker completely or use the supported browser's native pop-up blocking function instead of a third-party product.

Chapter 2: Server Management

IN THIS CHAPTER

This chapter provides the following information about server management in the Opsware System:

- Server Management Overview
- Server Asset Tracking Overview
- Server Identification
- Server Histories and Reports
- Server Life Cycle
- Agent Reachability Communication Test
- Server Locking
- Scheduling Server Management Jobs
- Communication Between Managed Servers and the Opsware System
- Viewing Hardware Information for Managed Servers
- IP Range Groups and IP Ranges
- Network Configuration
- Opsware Agent on Managed Servers
- Server Assimilation
- Custom Attributes for Servers
- Server Groups
- Service Levels

This chapter does *not* document how to install operating systems, patches, or applications on servers. However, it does discuss how those tasks fit into the overall server life cycle, and it does provide cross-references to the appropriate topics in other

chapters.

Server Management Overview

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

- Server Management Functions
- Agent-Server Architecture of Opsware Technology
- Required Permissions for Server Management
- How the Opsware Model Affects Server Management
- Distinquishing Among Packages, Nodes, and Templates
- Software Tree Nodes and Server Management
- Packages and Server Management
- Templates and Server Management
- Server Management in Multiple Facilities

Server Management Functions

The Opsware System manages servers in an operational environment in the following ways:

- Assimilating servers that are already functioning in the operational environment, which allows users to deploy and manage new applications installed on those servers
- Provisioning servers with Microsoft Windows, Red Hat Linux, and Sun Solaris operating systems by using vendor-provided operating system bootstrapping technologies. Additionally, the Opsware System integrates with AIX NIM and HP-UX Ignite installation technologies to provide a uniform method for OS provisioning across a heterogeneous environment.

  See "OS Provisioning Overview" on page 239 in Chapter 4 for information about how the Opsware System provisions Microsoft Windows, Red Hat Linux, and Sun Solaris operating systems on managed servers.

  See Appendix B, "Agent Upgrade Tool" for information about how the Opsware System integrates with AIX NIM and HP-UX Ignite installation technologies.

- Automating patch management on Opsware-managed servers by providing the ability to react quickly to newly-discovered security threats and the degree of control required to ensure that new patches are installed in a uniform way

  See “Opsware Patch Management” on page 403 in Chapter 8 for more information.

- Managing application provisioning, which enables users to deploy applications and databases across many servers simultaneously, and track what is deployed on each server

  See “Application Provisioning” on page 385 in Chapter 7 for more information.

- Providing configuration tracking, which allows users to monitor selected configuration files and databases and to take certain actions when change is detected

  See “Automated Configuration Tracking” on page 493 in Chapter 11 for more information.

Agent-Server Architecture of Opsware Technology

Server management with the Opsware System is possible because Opsware technology utilizes an agent-server architecture. The *server* portion of the platform consists of multiple, integrated systems, each with a unique purpose.

See the *Opsware System 4.7 Administration Guide* for information about a detailed description of the components that make up the Opsware System server portion of the architecture.

Each server that Opsware technology manages has an intelligent agent (the Opsware 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 Agents. Depending on the request, the Opsware Agent on a server might use global Opsware services in order to fulfill the request. For example, the Opsware Agent might often make requests to the Model Repository, the central database for all Opsware System components, and the Software Repository, the central repository for all software that the Opsware System manages.

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 Data Access Engine and registers itself. The Data Access Engine sends this data to the Model Repository, which allows the Model Repository to keep track of server status, and know when particular servers are disconnected from or reconnected to the network.

After you install an Opsware Agent on a server, users can manage the server by installing or upgrading software, patching the OS software, removing software, changing server properties, or decommissioning it.

See "Server Assimilation" on page 139 in Chapter 2 for information about how to install an Opsware Agent on a server so that the Opsware System can manage it.

Required Permissions for Server Management

Users who are authorized to administer servers and applications by using the Opsware System do not also require system privileges (such as root on Unix or administrator access on Windows) on the Opsware-managed servers. Using the Opsware System to manage servers reduces the number of people who need system privileges for servers.

Typical system administrators only need Opsware user accounts to perform their regular server management duties. Depending on the types of server management tasks you need to perform, you will need different Opsware permissions:

- Servers are associated with customers in the Opsware System. Therefore, to manage specific servers, you must have permissions to manage the resources of your customers' servers. For example, you want to manage the server with hostname `m0123.core3.xyzcustomer.com`, which is associated with customer XYZ in the Opsware Command Center. Therefore, you must have permission to manage customer XYZ's resources.

  Contact your Opsware administrator to obtain the required customer permission or see the *Opsware System 4.7 Administration Guide* for information about how to obtain permission to manage the servers associated with a specific customer.

- You must have permission to access the facility in which the servers you want to manage are located. Given the same example, the server with hostname `m0123.core3.xyzcustomer.com` is located in Facility 1. Therefore to manage this server, you must have permission to manage the resources in Facility 1.

Contact your Opsware administrator to obtain the required facility permissions, or see the *Opsware System 4.7 Administration Guide* for information about how to obtain permission to manage the servers in a specific facility.

The Opsware System provides two types of permissions for customers and facilities, read and read/write permissions. If you have read permission for a customer or facility, you can only view information about the server associated with the customer or located in the facility. If you have read/write permission, you can perform actions for the server, such as installing an application.

- Additionally, the different features and subsystems in the Opsware System are accessible only if you have the required permissions for that subsystem or feature. For example, if you want to use the OS Provisioning Subsystem to install operating systems on bare-metal servers, you must have the requisite OS provisioning permissions. Each chapter that documents how to use a particular Opsware feature includes a topic for the required permissions to access that feature.
- If you want to deactivate servers as part of your server management duties, you must have the Deactivate permission in the Opsware System.

How the Opsware Model Affects Server Management

The Opsware System employs a model-based approach to managing the operational environment. Users interact with the Opsware Command Center, the front-end application, to accomplish OS provisioning and application provisioning, patch management, and server asset tracking for the operational environment. When users work in the Opsware System, they work in the model and the Opsware System pushes the changes to the managed environment by reconciling the managed environment with the model. See Figure 2-1.

See "Opsware System's Model-Based Approach" on page 3 in Chapter 1 for information about how the Opsware System is model-based.

Figure 2-1: Opsware Model-Based Approach to Server Management

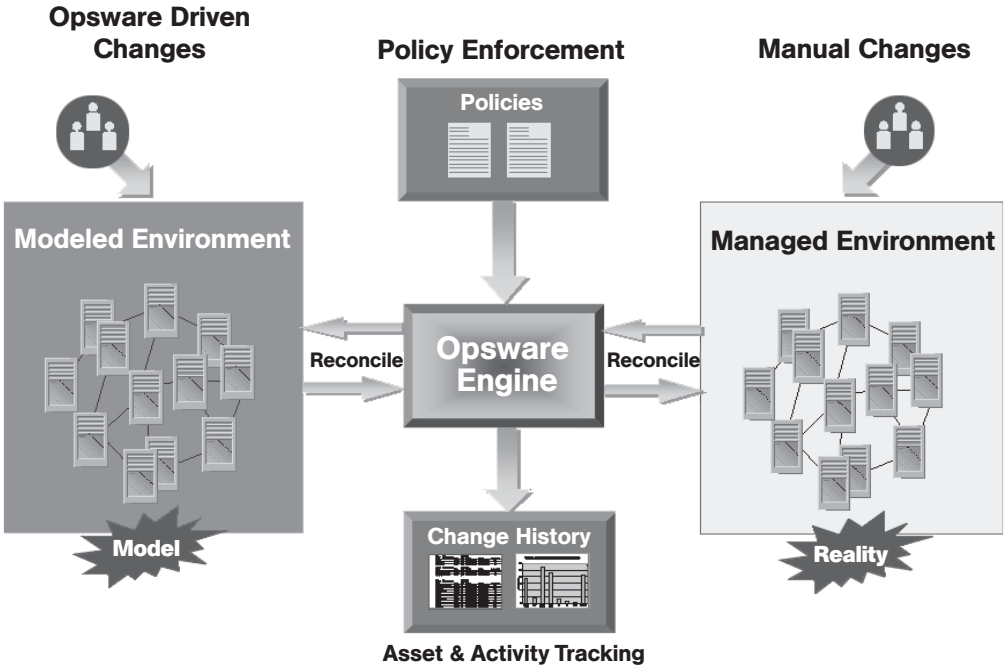

The Opsware Command Center visualizes the model as a tree called the Software Tree. By using the Software Tree, the Opsware System enables users to embed technical best practices by specifying software installation order and by creating templates.

Before system administrators can use the Opsware System to manage servers in the operational environment, the operational environment must be modeled in the Opsware Command Center.

A policy setter for your organization performs this goal by completing these setup tasks:

- Defining the Software Tree, which defines what applications are available to install on the servers running in your environment

  See "Application Provisioning Setup" on page 307 in Chapter 6 for more information.

- Uploading the packages that you will use in the operational environment

  See "Package Management" on page 267 in Chapter 5 for more information.

- Creating templates so that users can quickly install a complete software baseline on servers in the environment

See “Application Provisioning Setup” on page 307 in Chapter 6 for more information.

- Setting up the OS Provisioning Subsystem by preparing OS definitions for the operating systems that you need to install on your servers

  See “OS Provisioning Setup” on page 175 in Chapter 3 for more information.

- Setting up the Patch Management Subsystem by uploading the required patches for your environment

  See “Server Management” on page 29 in this chapter for more information.

After a policy setter completes these setup tasks, end-users are ready to manage the servers running in the operational environment. These setup tasks do *not* need to be repeated by end users for them to manage servers.

Distinquishing Among Packages, Nodes, and Templates

To understand the way that the Opsware System manages servers through the Opsware model, you need to understand the distinctions between certain elements in the model.

The *Software Tree* organizes all the technology building blocks of your organization’s environment, including the software and hardware. The Software Tree is made up of nodes and subnodes, modeling the interrelationships and dependencies among the software and customer accounts in the operational environment. Users navigate the Software Tree to perform specific operations.

The Software Tree provides the model for the operational environment, as Figure 2-2 shows.

Figure 2-2: Illustration of the Software Tree

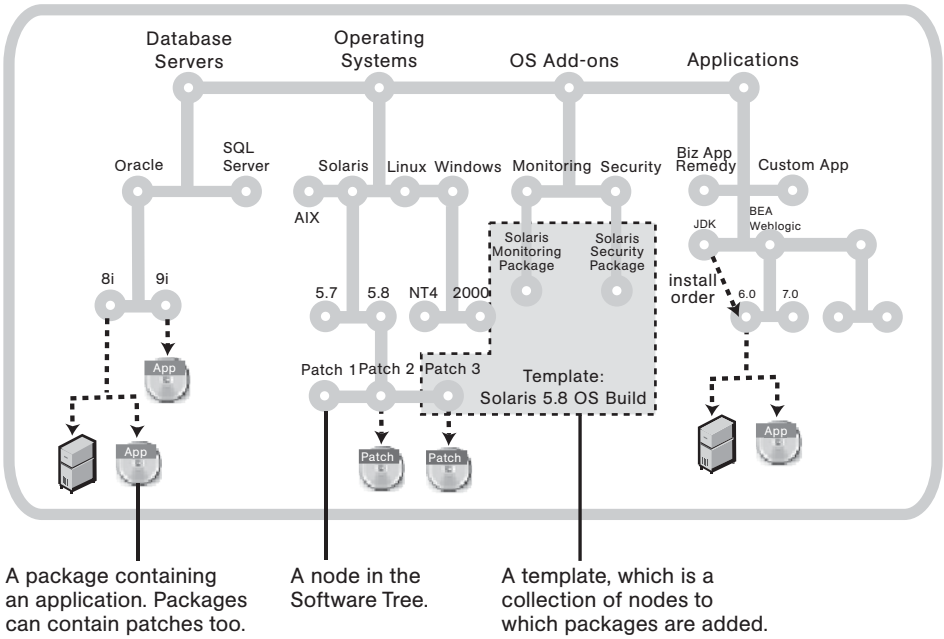

Software Tree Nodes and Server Management

Using nodes simplifies server management within the Opsware System and the management of the software applications and configurations associated with those servers. *Nodes* are a hierarchical set of categories or types that classify software, configuration, and other components running in the operational environment.

In the Opsware Command Center, you create nodes in the Software Tree. From the navigation panel, click Software ➤ Applications. The Applications page appears. As you navigate the Software Tree, you see the hierarchy of nodes in each category of applications, as Figure 2-3 shows.

Figure 2-3: Hierarchy of Nodes within an Application Category

The Software Tree has the following characteristics:

- Each point in the Software Tree is called a node.
- The node information at the top of the Software Tree is more general, and as you travel farther down the tree, each successive subnode contains more detailed and specific information that relates to the node above it.
- A node inherits properties or software from the nodes above it.
- Users can add nodes and subnodes within each category.
- Software or a server might be assigned to a node depending on its location in the Software Tree.

  See “Software Attached to Nodes” on page 334 in Chapter 6 for more information.

  This guide primarily documents how servers are automatically assigned to and removed from nodes when you use one of the Opsware wizards to install or uninstall software. Using an Opsware wizard is an efficient and easy method to assign or to remove servers from nodes.

  In certain situations, you might want to use the Software Tree and the reconcile operation to install or uninstall applications from managed servers.

  See “Assigning to and Removing Servers from Nodes” on page 451 in Chapter 9 for more information.

- Assigning a server to a node determines what software is installed on that server and how it is configured.

Users perform the following tasks for servers by using nodes:

- Locating servers and viewing the nodes to which they are assigned
- Copying the configuration (assigned to nodes) of a server to other servers
- Assigning servers to and removing servers from nodes (so that you can install or uninstall software by running the reconcile operation).

  See "Reconcile" on page 441 in Chapter 9 for more information.

  See "Application Provisioning Setup" on page 307 in Chapter 6 for information about how to create the Software Tree in the Opsware Command Center so that end users can install applications on Opsware-managed servers.

Packages and Server Management

In the Opsware System, *packages* contain software that is registered in the Software Repository. Packages contain software for operating systems, applications (for example BEA WebLogic, IBM WebSphere), databases, customer code, and software configuration information.

Packages are made available in the Opsware System by uploading the packages to the Software Repository with the Opsware Command Center or by using the Opsware Command Line Interface.

The term package describes the collection of executables, configuration, or script files that are associated with an Opsware-installable application or program.

Templates and Server Management

Opsware templates allow you to group related sets of software so that they can be installed in a single operation by using the Opsware wizards. The two basic types of Opsware templates are:

- Templates that include the installation of an operating system
- Templates that do *not* include an operating system

You can, for example, use Opsware templates to quickly bring new servers into production. A template might include an operating system for a new server, the latest security patches for the operating system, plus all the applications required to run a full-fledged Web service. Or, you can use templates to install a new service, made up of a set of applications and patches, on servers that are already in production.

Templates require little setup and you can create them and deploy them quickly. To create a template, select packages, patches, or operating systems that are already configured and tested for installation, and add them to the template. You can later edit the template. For example, if a new patch is released, you can edit the template and add the patch to the template.

Server Management in Multiple Facilities

The managed environment might span several facilities. A facility refers to the collection of servers that a single Opsware Model Repository manages, the database that stores information about the managed environment. For example, one facility might be dedicated to an organization's Intranet, while another facility might be dedicated to the Web services offered to the public.

Users can manage servers in any facility from an 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 Repositories located in all remote facilities.

When using Opsware technology in multiple facilities, users should follow these work process rules to reduce the chance of data conflicts between facilities:

- Users should not change data in one facility and then make the same change in another facility.
- More than one user should not change the same object in different facilities at the same time. For example, two users should not manage the same server from different facilities.

Server Asset Tracking

This section provides information on server asset tracking with the Opsware System and contains the following topics:

- Server Asset Tracking Overview
- Server Lists Overview
- About the Server Pool
- About the Managed Servers List
- Filtering Servers in the Server Lists
- My Servers Overview

- Adding Servers to My Servers
- Removing Servers from My Servers
- Searching for a Server By Using the Search Box
- Ways to Use Advanced Search
- Searching with Advanced Search
- Details About Advanced Searches
- Searching for Servers by IP Address
- Examples of Advanced Server Searches

Server Asset Tracking Overview

You can locate, list, and display servers in the Opsware Command Center in the following four ways:

- By searching when you know the name, hostname, or IP address of the server you want to provision or manage.
- By viewing the Managed Servers list and Server Pool list when you want to see a complete list of all your servers. You can refine the lists by using filters.
- By browsing nodes in the Software Tree when you want to determine on which servers specific software should be installed or find all servers that match a certain criteria; for example all servers with a specific application installed.
- By viewing servers sorted by hardware category (Click Environment ➤ Hardware in the navigation panel.) The Servers tab in the Hardware pages shows each manufacturer and model that you have running in the operational environment. See "Server Data That the Opsware Agent Tracks" on page 135 in this chapter for more information.

Every server that the Opsware System manages has the following properties:

- IP addresses, hostname, and the server ID
- Which nodes the server is assigned to in the Software Tree categories. Click a node link to display specific information about that node.
- What software *should be installed* on the server. Click the Install List tab from the Managed Servers: Server Properties page to display the list of software packages that should be installed on that server by virtue of that server's assigned nodes.

The Opsware System is able to determine what software *should be installed* on a server because of its model-based approach to server management. The software that should be installed is recorded in the Opsware Model Repository.

- All software that *is installed* on the server. Click the Installed Packages tab from the Managed Servers: Server Properties page to display the list of software that is reportedly installed on the server.

  The Opsware System is able to determine what software *is installed* on a server because the Opsware Agent communicates with the Opsware core and reports the installed hardware and software for the server.

  In some cases, Solaris packages might only be partially installed. In these cases, the partially installed Solaris package does not show up in the installed list.

See "Server Data That the Opsware Agent Tracks" on page 135 in this chapter for information about a complete list of all the hardware and software information that the Opsware System tracks for managed servers.

Server Lists Overview

The Opsware Command Center displays lists for two types of servers, as Figure 2-4 shows.

Figure 2-4: Servers Section in the Navigation Panel

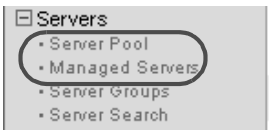

The **Server Pool** – Servers in the Server Pool have registered their presence with the Opsware System but do not have the target OS installed. An OS Build Agent is running on each server so that they can communicate with the Opsware System.

See "Operating System Provisioning" on page 237 in Chapter 4 for information about how to use the Server Pool when you install the target OS on a server.

Managed Servers – The Managed Servers list contains servers on which the Opsware System can perform management tasks because Opsware Agents are installed on them. However, the Opsware System might not have provisioned all the software running on the servers.

You begin the OS provisioning process by reviewing the servers in the Server Pool list. From the Server Pool, you can install a target OS by selecting a server and clicking the Install OS button. See Figure 2-5.

Figure 2-5: Server Pool List in the Opsware Command Center

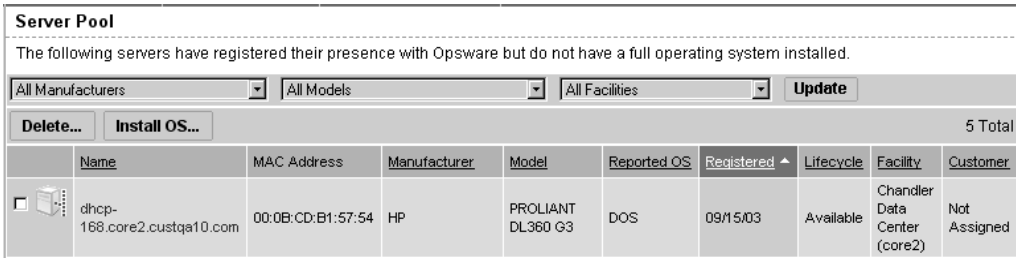

About the Server Pool

The Server Pool provides the following information about each server waiting to be provisioned with the target OS:

- The hostname set by booting the server for the first time over the network or by using an Opsware Boot Floppy
- The MAC address
- The manufacturer and model
- The OS that the OS Build Agent is running – DOS (Windows operating systems), Linux, or Solaris

  You use this information to select the target OS for servers. If the server is in the process of installing an OS, this value might change.
- The last date and time that the Opsware Agent on the server communicated with the Opsware System (by submitting the server's hardware and software information)

  If the server is in an unreachable state (that is, if the server icon has a red "x" on it), you can run a Communication Test to help you troubleshoot why that server is unreachable. See "Agent Reachability Communication Test" on page 80 in this chapter for more information.
- The life cycle value, such as whether the server is available to have a target OS installed on it
- The facility in which the server is located
- The customer association

- Additional hardware information (Click the server name to open a window that displays specific hardware information.)

About the Managed Servers List

The Managed Servers list contains servers on which the Opsware System can perform management tasks because Opsware Agents are installed on them. When an existing operational server is assimilated successfully into the Opsware System, it appears in the Managed Servers list and the server icon indicates that it is fully manageable, as Figure 2-6 shows.

Figure 2-6: The Managed Servers List in the Opsware Command Center

By default, servers in the Managed Servers list are sorted by the Name column. However, you can re-sort the list based on any of the column headings. For example, you can click the Hostname / IP Address column heading to re-sort the list by hostname or IP address.

If you have many servers that the Opsware System manages, the list of servers is grouped by pages. Click the page number links or the left arrow button at the bottom of the list.

The Managed Servers list provides the following information about each server:

- The name of the server

  By default, the server's hostname appears in this field. However, you can edit it so that it is more descriptive or useful.

- The hostname of the server determined by the Opsware Agent
- The IP address configured for the server, which users can edit by using the network configuration feature in the Opsware Command Center
- The reported OS, which is obtained by the Opsware Agent that is running on the server
- The stage of the server, which specifies the stages of deployment for servers

- The server's use
- The facility in which the server is located
- The customer association
- Additional hardware information (Click the server name to open a window that displays specific hardware information.)

Filtering Servers in the Server Lists

The Managed Servers list displays five filters that you can specify to qualify the servers that the Opsware Command Center displays, as Figure 2-7 shows.

Figure 2-7: Filters in the Managed Servers List

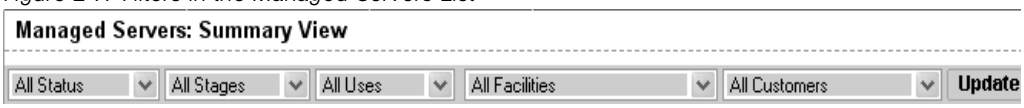

- **Status** – specifies the ability of the Opsware System to manage servers. The Opsware System automatically detects the status of servers; a server's status is OK or Not Reachable.
- **Stages** – specifies the stages of deployment for servers; for example, a server is live or offline. Users set this property for servers. The values in this list are customizable by the Opsware administrator.
- **Uses** – specifies how an organization is utilizing servers; for example, a server is a staging server. Users set this property for servers. The Opsware administrator can customize the values in this list.
- **Facilities** – specifies the location of servers. From an Opsware Command Center, users can manage servers located in any facility. For example, a user could log in to the Opsware Command Center running in facility A and manage the server located in facility B.
- **Customers** – specifies the customer associated with each server. Your Opsware administrator defines the options for customer selections by using the Administration pages.

Figure 2-8 shows the filters that are in the Server Pool list.

Figure 2-8: Filters in the Server Pool List

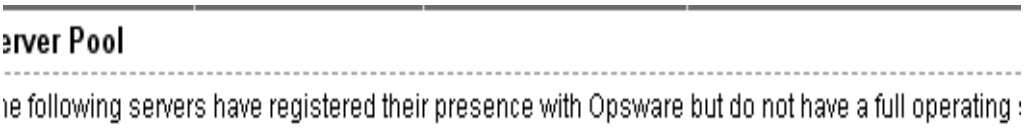

- **Manufacturers** – specifies the manufacturer for the server as reported by the OS Build Agent running on the server.
- **Models** – specifies the model of the server as reported by the OS Build Agent running on the server.
- **Facilities** – specifies the location of the server. Users can manage servers in any facility from an Opsware Command Center in any facility.

My Servers Overview

The My Servers feature provides an efficient way to manage servers when your operational environment contains hundreds or thousands of servers.

When you search for servers or browse the server lists, you can add servers to My Servers (similar to a shopping cart on an e-commerce site). Using My Servers allows you to view and perform actions on selected servers.

When you use the same browser and log in to the Opsware Command Center running in the same facility, servers stay in My Servers for one year or until you explicitly remove them. Each time that you log in to the Opsware Command Center, you see the servers that were in My Servers the last time that you logged in.

The My Servers feature is available only on a per-user basis. You cannot log in as an Opsware administrator to see the servers in the My Servers area of other Opsware users.

Adding Servers to My Servers

Perform the following steps to add a server to My Servers:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the servers that you want to add to the My Servers.

   Or

   Search for the servers that you want to add to My Servers.

See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

2. Select the servers that you want to add to My Servers.
3. Choose Servers ➤ Add to My Servers from the menu above the Managed Servers list. The Add To My Servers window appears, which indicates that you added the chosen number of servers to My Servers.
4. Click the Close button to close the window.
5. Next, select the My Servers link at the top of the page. You see the selected servers added to My Servers, as Figure 2-9 shows.

Figure 2-9: Servers in My Servers

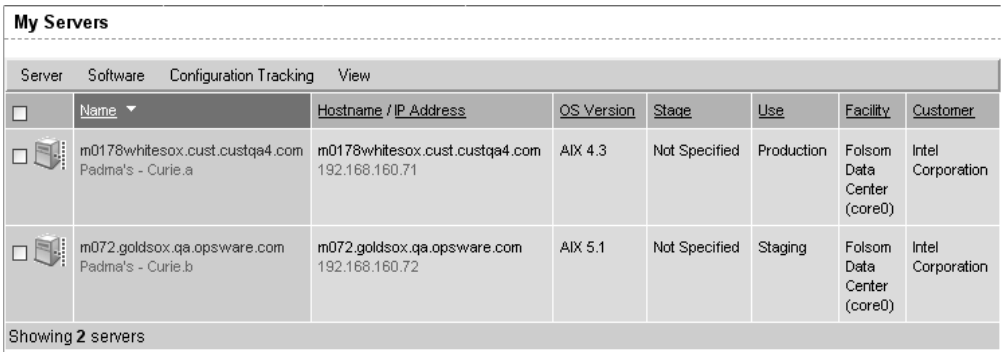

You can perform the same server management tasks on servers in My Servers as on the servers in the Managed Servers list.

Removing Servers from My Servers

Perform the following steps to remove servers from My Servers:

1. Click the My Servers link in the navigation panel of the Opsware Command Center. The My Servers page appears that shows the servers currently added to it.
2. Select the servers that you want to remove from My Servers and choose Server ➤ Remove from My Servers from the menu above the Server list.

   Or

   Choose Server ➤ Clear My Servers to remove all the servers from My Servers.

The My Servers page refreshes and displays the remaining servers in My Servers.

Searching for a Server By Using the Search Box

Perform the following steps to search for a server using the Search box:

1. On the Home page, click the down arrow in the navigation panel to open the Search box, as Figure 2-10 shows.

Figure 2-10: Search Text Box on the Opsware Command Center Home Page

2. Verify that the Servers option is selected in the list.
3. Type the server's IP address, hostname, or name in the Search box and click the Go button.

   The search text that you enter can include an asterisk (*) wildcard character. However, the search feature automatically pre-pends and appends an asterisk to the text. The search text is case sensitive.

   For example, you can type any of the following search queries:

   ```
   192.168.68.6
   host02.coredev-va1.sample.com
   192.168.*.192
   host1*.xyz.samplecompany.com
   ```

   The resulting page contains one or more servers, depending on the type of search query that you specified. If no servers are found, the Opsware Command Center displays a message that indicates that no servers were found that matched your query, as Figure 2-11 shows.

Figure 2-11: Search Results Page in the Opsware Command Center

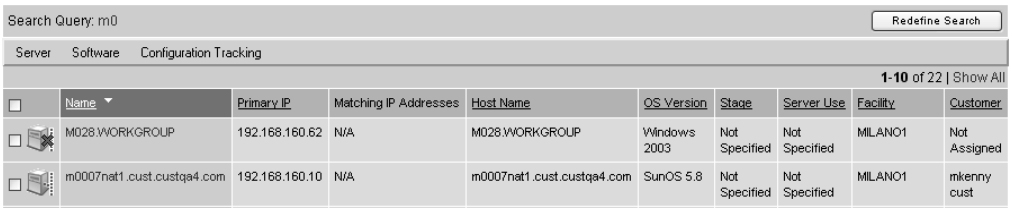

4. Refine your search by clicking the Redefine Search button.

See "Searching with Advanced Search" on page 48 in this chapter for information about how to formulate complex, multicriteria search queries.

Ways to Use Advanced Search

By using the Opsware Command Center, you can perform advanced searches in the following ways:

- From Opsware wizards

  While using the Opsware wizards from the Tasks panel, you are prompted to select (by browsing or searching) servers, operating systems, patches, applications, and templates at specific points in the process.

  What you can search for in the Opsware wizards is context sensitive to the type of operation that you are performing. For example, if you are using the Install Patch Wizard, you can click the Search tab to search for patches to install on servers.

- When adding operating systems, patches, or applications to templates

  Searching for an operating system, patch, or application to add to a template functions the same way as searching through the Opsware wizards.

- From the navigation panel (click Servers ➤ Server Search)

  The Server Search page allows you to search for managed servers that match specified criteria.

Searching with Advanced Search

In the Opsware wizards, you can browse for servers, operating systems, patches, applications, and templates, or use the Advanced Search feature to search for these items. Perform the following steps to search by using Advanced Search:

1 In an Opsware wizard, click the Search tab. An Advanced Search page appears. See Figure 2-12.

*Figure 2-12: Search Tab in the Select Servers Step of an Opsware Wizard**

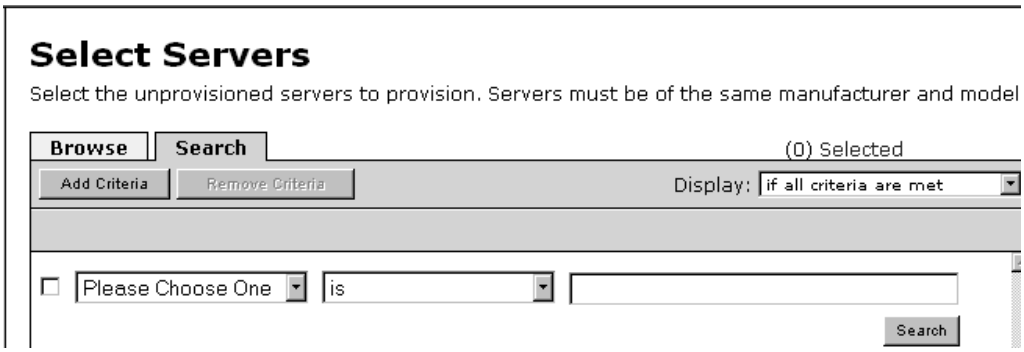

*(*Additionally, you can use the Search tab at other steps in the wizards to search for operating systems, patches, applications, and templates.)*

Or

From the navigation panel, click Servers ➤ Server Search. The Server Search page appears, as Figure 2-13 shows.

Figure 2-13: Server Search Page

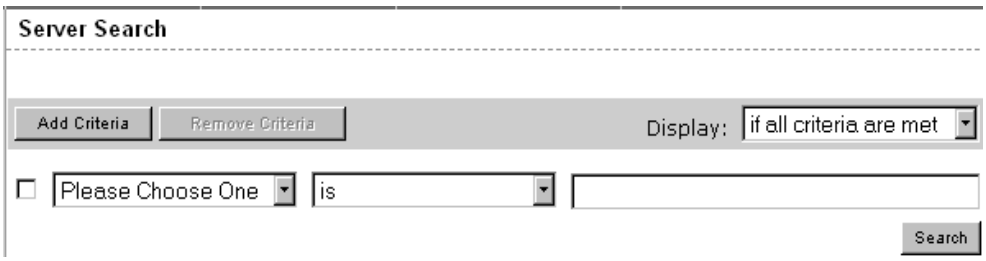

By default, one search criteria is added to the search.

2 Specify the attribute that you want to search for by selecting it from the first list, as Figure 2-14 shows.

Figure 2-14: Search Attribute List in the Search Page

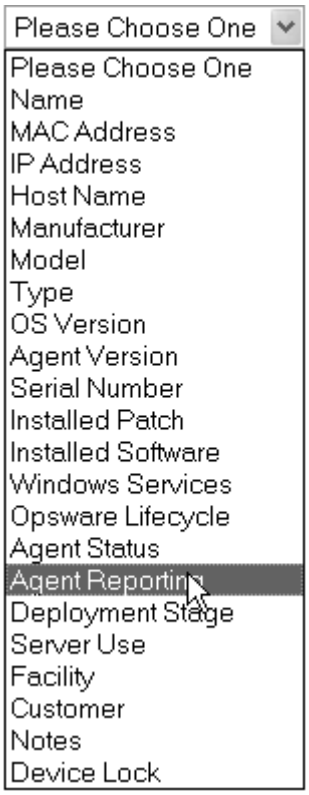

Depending on the attribute that you select, additional selection lists might appear in the page. For example, if you selected Manufacturer, an additional list of manufacturers appears in the page.

You cannot search in Notes that contain line breaks. See "Details About Advanced Searches" on page 53 in this chapter for more information and a workaround.

If you are searching while using an Opsware wizard, the first search attribute list might not have all the options that the previous example shows. The list only includes the options that are relevant for the Opsware wizard that is being used. For example, the Install OS Wizard does not include options to search for installed patches on the servers.

3 In the second list, specify how you want the Opsware System to search by selecting one of the following values. The operator selected defines how the search text is treated.

- Is – select when you want an exact match with the entered text. By default, all queries use the Is operator.
- Is not – select when you want to explicitly exclude specific results from the search query.
- Contains – select when you want a partial match. The search feature pre-pends and appends an asterisk (*) to the search query.
- Does not contain – select when you want to generally exclude certain text from the search results. The search feature pre-pends and appends an asterisk (*) to the search query.

Negative operators (Is not and Does not contain) might not be available in all cases.

4 Enter the text that you want to search for in the text box or choose from the list. The search text that you enter can include an asterisk (*) wildcard character. The search text is case sensitive. You can also use the SHIFT or CTRL key to select multiple criteria.

5 (Optional) Click the Add Criteria button as Figure 2-15 shows and repeat Steps 2 through 4.

Figure 2-15: Multiple Criteria in an Advanced Search

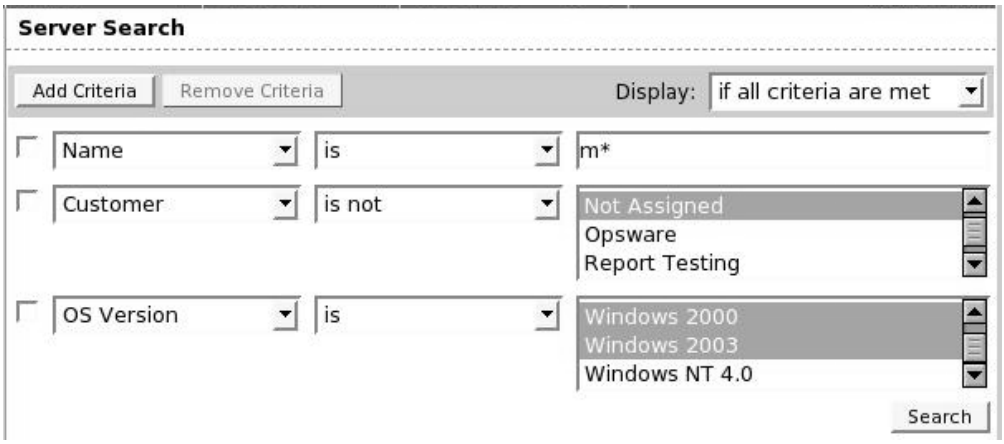

6 If you specified multiple criteria for the search, select whether you want search results only if all criteria are met or if any of the criteria are met, as Figure 2-16 shows.

Figure 2-16: Operator Controlling Advanced Server Search Results

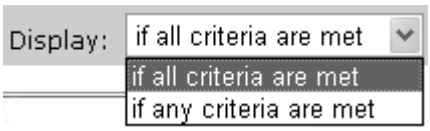

By default, search results appear for servers that match *all* the search criteria. If you are searching from an Opsware wizard, this field is *always* set to the value *if all criteria are met* and you cannot change it.

7 Click the Search button. The list of servers that match your search criteria appears in the page, as Figure 2-17 shows.

Figure 2-17: Displayed Search Results

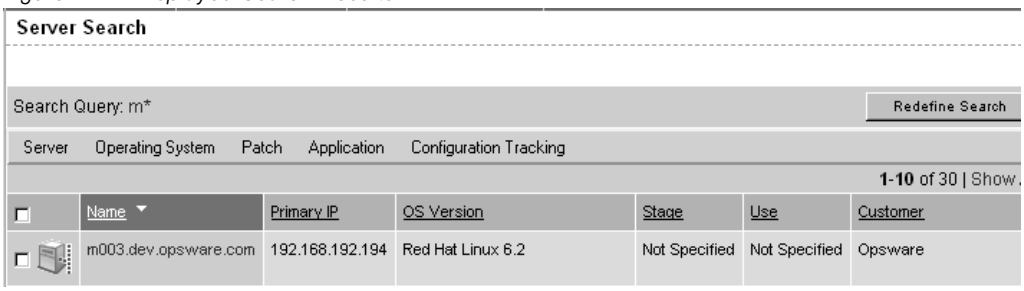

By default, the search results always include columns for Name, IP Address, OS Version, Stage, Use, and Customer. Depending on the search attribute that you select in Step 2, the search results might contain additional columns of data. For example, if you search for installed software, the search results contain an Installed Software column.

When you search for installed software or patches and include an asterisk in the search text, the Opsware System might take several minutes to display the search results.

If you searched for Installed Software or Installed Patches, move your cursor over a package name to display details about that installed software or patch, as Figure 2-18 shows.

Figure 2-18: Popup That Displays Details About an Installed Package

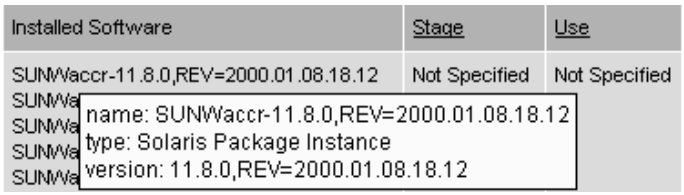

8 (Optional) Click the Redefine Search button to make changes to your search query and rerun it.

Details About Advanced Searches

You cannot search in notes that contain line breaks. For example, you cannot search for text in a note when it is this type of text:

```
line1 <line break>

line2
```

For example, the following query does not return any results:

```
Any/all notes contain line1 <line break> line2
```

However, the following query does return all servers that have notes:

```
Any/all notes is not line1 <line break> line2
```

To work around this limitation, include an asterisk (*) where the line break occurs, as Figure 2-19 shows. For example:

```
Any/all notes is line1*line2
```

Figure 2-19: Line Break Workaround in the Server Search Feature

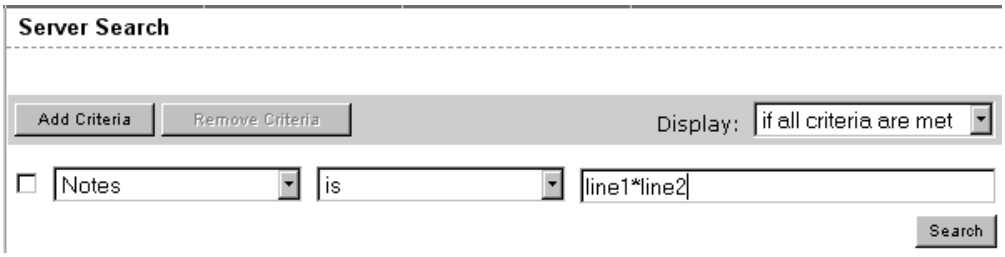

When you perform an advanced search with multiple criteria, the following conditions apply to the way that the Opsware System provides search results:

- When it evaluates criteria, the Search feature considers each criteria individually, finding servers (or operating systems, patches, and so forth) that match the individual criteria, and then the results of each criteria are combined.
- You must select at least one criteria and it must have a value. Default filter criteria count as criteria with a value. For example, in the MS Updates Wizard, when searching for a server, the Search feature automatically enters a default criteria to search only for servers running a Windows OS. Therefore, users are not required to specify any more criteria.
- Empty criteria are ignored in searches. Users do not have to manually remove them for the search to proceed.
- You cannot search for empty values. For example, you cannot search for all servers where the Notes field is empty.

At certain steps in the Opsware wizards, the Advanced Search feature provides default values based on previous selections in the wizard. When these default values are added automatically as search criteria, the search results are always relevant for all the criteria (the Advanced Search feature uses the AND operator to combine the criteria).

For example, in the Select Patches step of the Install Operating System Wizard, the search query automatically includes the values for customer and operating system that you specified in earlier steps. The search query finds only the patches that are associated with the specified customer and OS, as Figure 2-20 shows.

Figure 2-20: Default Values Entered in Advanced Searches

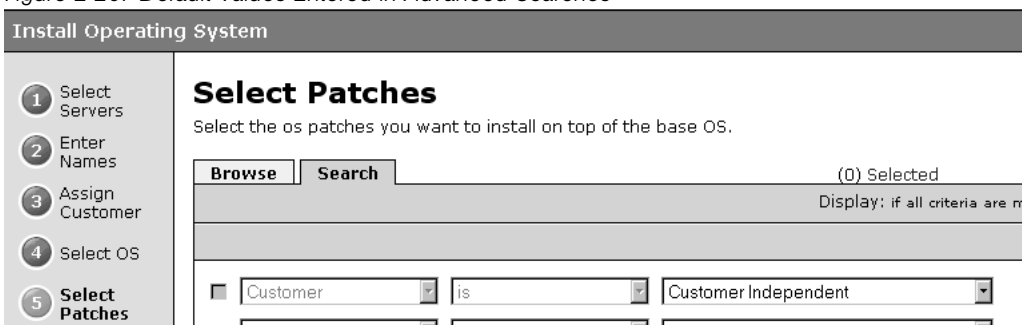

Searching for Servers by IP Address

Users can search for servers by entering a specific IP address in the Search Panel text box in the navigation panel or in the Advanced Search feature.

The Opsware System includes support for static Network Address Translation (NAT). This feature introduces the concept of a management IP, which might be different from any of the local IP addresses that the Opsware Agent reported for a server.

When searching for a server based on its IP address, the Advanced Search feature searches based on the server's primary IP address and based on the IP address for any interface that server has, including its management IP address. In the search results, an extra column is shown that lists all matching IP addresses for all interfaces. The management IP address is included if the server's networking is configured for static NAT.

See "Communication Between Managed Servers and the Opsware System" on page 111 in this chapter for information about how the Opsware System handles servers that are affected by static NAT.

Examples of Advanced Server Searches

Using the Advanced Search feature, a user creates a query with the following conditions:

- Installed Software contains *qa*
- Installed Software contains *man*
- If *all criteria are met* is selected

The results of this search will be all servers that have at least one installed package with *qa* somewhere in its name and at least one installed package (not necessarily the same one) with *man* in its name.

The search results are not limited to packages that contain both *qa* and *man* in the package name.

Find all servers that have some version of Apache or some version of Java installed:

- Installed Package contains *apache*
- Installed Package contains *java*
- If *any criteria are met* is selected

Server Identification

This section provides information on server identification within the Opsware System and contains the following topics:

- Server Identification Overview
- How Opsware System Identifies Servers
- Customer Accounts in the Opsware System
- Associated Servers with Customers

Server Identification Overview

The Opsware System uses the following IDs to track managed servers:

- **MID**: Machine ID. The unique identifier that the Opsware System uses to identify the server. The MID is usually equal to the server ID.

  The MID is stored in a file on a server's disk so that the MID can persist and be read by the Opsware Agent.

  The MID follows the hard disk, not the chassis, so system administrators can swap chassis for servers without affecting how the Opsware System tracks those servers.

  See "Example: How Opsware Handles Swapping a Server's Hard Disk" on page 57 in this chapter for information about swapping hard disks.

- **Server ID**: The primary key in the Opsware Model Repository (database) that represents a given server. The Server ID is used internally in the Opsware System.

Generally, users do not need this value for servers to manage them in the Opsware System.

- **MAC Address**: Media Access Control address, which is the network interface card's unique hardware number. The MAC is used as the server's physical address on the network.
- **Chassis ID**: A unique hardware-based identifier that the Opsware Agent discovers, typically derived from some property of the server's chassis. As a common source for this ID, the Opsware System uses an interface's MAC address or the hostid on Solaris servers, or the serial number for one of the interfaces.

How Opsware System Identifies Servers

Servers in the Managed Servers list are identified in the following ways when they register their hardware and software with the Opsware System:

- The Opsware System identifies each server by using the MID first.
- If the MID cannot be determined, the chassis ID is used to identify the server.
- If the server cannot be identified with the chassis ID, the MAC addresses are used to identify the server.

In the Server Pool, the MAC Address column displays values by which the Opsware System tracks the servers. The value used varies by platform:

- Intel x86 processor-based servers are identified by the MAC address of the server.
- Sun SPARC processor servers are identified by the host ID of the server.

  The host ID for Sun SPARC processor servers appears in the MAC Address column in the Server Pool list.

  To determine the value in the MAC Address column, the Opsware System uses the hardware address by which the server contacted the Opsware Build Manager (an OS Provisioning Subsystem component).

Example: How Opsware Handles Swapping a Server's Hard Disk

The following steps show how Opsware swaps a hard disk:

1. A system administrator swaps the hard disk of Server A (MID `1230001`, chassis ID `AB:08`) with the hard disk of Server B (MID `98730001`, chassis ID `XY:96`).
2. The Opsware Agent on Server A registers its hardware with the Opsware System. The MID for Server A equals `1230001` and the chassis ID equals `XY:96`.

3. The Opsware System locates Server A by using the MID.
4. The Opsware System updates the data it has for Server A in the Model Repository. It sets the chassis ID equal to `XY:96`.
5. The Opsware Agent on Server B registers its hardware with the Opsware System. The MID for Server B equals `98730001` and the chassis ID equals `AB:08`.
6. The Opsware System locates Server B by using the MID.
7. The Opsware System updates the data it has for Server B in the Model Repository. It sets the chassis ID equal to `AB:08`.

Customer Accounts in the Opsware System

Many enterprise customers have consolidated disparate IT operations into a single operation, yet they still need separate reporting, billing, and management for different business units or groups (for example, West Coast Office, East Coast Office, and London Office).

The Opsware System accommodates these requirements. Within the Opsware Command Center, users perform server provisioning and management by using customer accounts.

When your Opsware administrator creates a customer in the Opsware System, a value for that customer is automatically added to the customer filter in the Managed Servers list, as Figure 2-21 shows.

Figure 2-21: Customer Filter in the Managed Servers List

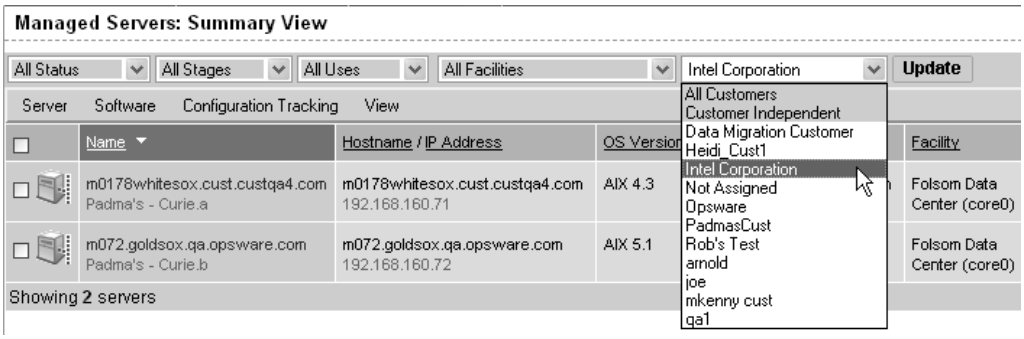

By using customer accounts in the Opsware Command Center, you can segregate servers that belong to different business units. By segregating servers, you can have separate accounting for each customer or different levels of security for different customers. You might want to segregate the servers based on the department or business unit to which they belong for many reasons.

By default, the Opsware System is shipped with the following two customers:

- Customer Independent – A global customer in the Opsware System. Resources (applications, patches, and templates) that are associated with "Customer Independent" can be installed on any managed server, no matter what customer it is associated with.
- Not assigned – The servers are not associated with a customer. You can install applications, patches, or templates that are *Customer Independent* on *Not Assigned* servers. However, you cannot install or use any resources associated with a customer on a server that is not assigned to a customer.
- When you assimilate a server into the Opsware System, the server is associated with the *Not Assigned* customer if IP ranges were not created to automatically associate assimilated servers with customers. See Figure 2-22.

Opsware Inc. recommends that you associate servers with customers, if necessary, by using the Server Properties pages. See "Editing the Properties of a Server" on page 75 in this chapter for more information.

Figure 2-22: Customers List Under Environment in the Opsware Command Center

Associated Servers with Customers

An Opsware user or an Opsware administrator can set up an IP range group so that servers are automatically associated with customers when users perform these server management tasks:

- Assimilate servers running in the operational environment by installing an Opsware Agent on the servers

See "Server Assimilation" on page 139 in this chapter for more information.

- Use the OS Provisioning Subsystem to install operating systems on bare-metal servers

  See "Operating System Provisioning" on page 237 in Chapter 4 for more information.

To set up this automatic customer association, you must create IP range groups for customers and specify the ranges of IP addresses that the groups contain.

In the Opsware Command Center, an IP range group is both a physical and logical list – an accounting way to group ranges of IP address and assign them to a particular customer. An IP range identifies a range of IP addresses within an IP range group.

When you set this up, IP addresses get their customer association through the IP range, which, in turn, gets its customer association from the IP range group.

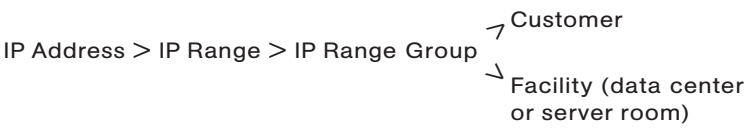

See "IP Range Groups and IP Ranges" on page 119 in this chapter for more information.

The loose relationship between server and IP address means that you can associate a server with a different customer from its IP address.

Even when IP range groups are set up for a customer, a server's IP address does not necessarily determine the customer to which the server is associated because a user can change the customer association in the Server Properties page.

See "Editing the Properties of a Server" on page 75 in this chapter for information about how to change the customer association for a server.

The customer association for a server is based on the management IP address of the server and not the primary IP address.

See "Communication Between Managed Servers and the Opsware System" on page 111 in this chapter for information about how the Opsware System uses management IP addresses for servers.

However, a server always belongs to the same facility (data center or server room) as its primary IP address. The Opsware System enforces the relationship between server and facility at hardware registration. See Figure 2-23.

Figure 2-23: Primary IP Addresses in the Opsware System

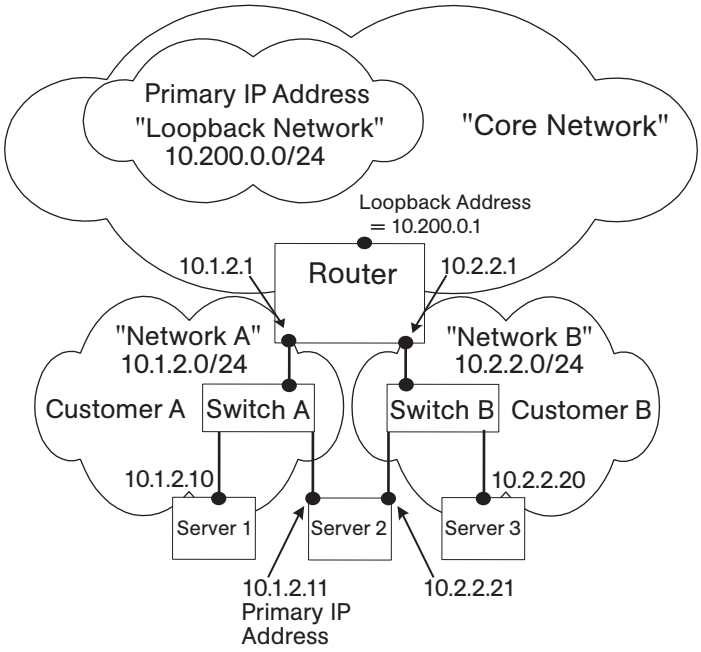

In this illustration, the following conditions apply:

- Server 1 belongs to Customer A.
- Server 2 belongs to Customer A but has IP addresses in Network A and Network B.
- Server 3 belongs to Customer B.
- The Router belongs to the Core Network but has IP addresses in Network A and Network B.

Server Histories and Reports

This section provides information on server histories and reporting with the Opsware System and contains the following topics:

- Server Histories and Reports Overview
- Viewing Server History

- Generating Server Reports

Server Histories and Reports Overview

By using the Opsware Command Center, you can view the history of changes made to a server. For example, you can see who has modified a server.

Entries are generated when actions are performed for managed servers in the Opsware Command Center. The History is read-only. See Figure 2-24.

Figure 2-24: Server History for a Server in the Managed Servers List

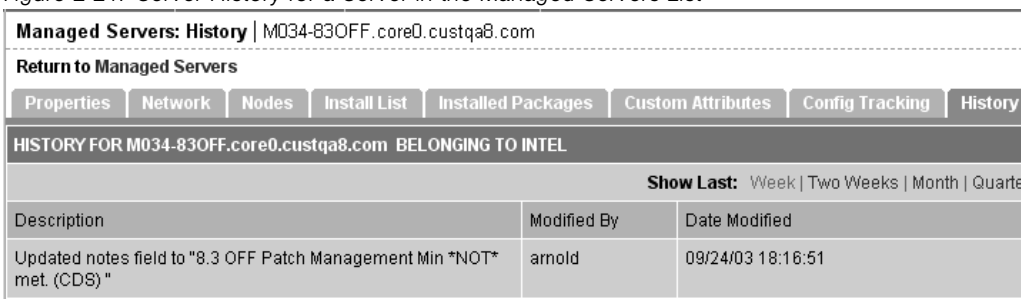

The Opsware System logs the following actions in the history for each managed server:

- Adds the server to a node
- Removes the server from a node
- Reassigns the server from one node to another

  This guide primarily documents how servers are automatically assigned to and removed from nodes when you use one of the Opsware wizards to install or uninstall software. Using an Opsware wizard is an efficient and easy method of assigning or removing servers from nodes.

  In certain situations, you might want to use the Software Tree and the reconcile operation to install or uninstall applications from managed servers.

  See "Assigning to and Removing Servers from Nodes" on page 451 in Chapter 9 for more information. See "Directly Reconciling Servers" on page 453 in Chapter 9 for information about how to install applications on managed servers by using the reconcile operation.

  See "Distinquishing Among Packages, Nodes, and Templates" on page 35 in this chapter for information about nodes.

- Adding a cloned server to a node (when cloning a server, node assignments occur in the Opsware System and are logged)
- Installing a template on the server
- When a preview reconcile fails or succeeds
- When a reconcile fails or succeeds

  See “Reconcile” on page 441 in Chapter 9 for information about how reconcile works to install software on or uninstall software from servers.

For each entry, the user who initiated the event is captured.

Data is maintained for servers in the Opsware System for the following periods of time:

- The Opsware Command Center maintains the history of changes for the last three-month interval.
- Command Engine session logs are retained for 30 days, except for the last reconcile session for a server, which is retained indefinitely.

  The Command Engine is the Opsware system component that enables distributed programs to run across many servers.
- Server node history is retained for 6 months.

If longer periods of time are required, Opsware Inc. recommends regular backups to enable offline storage of Opsware System data.

The Opsware System deletes old data from the Opsware Model Repository. It does not copy the data before it removes it. Customers can change the retention period for each type of data by using Oracle commands for manipulating scheduled jobs. Contact your Opsware support representative for assistance in changing these retention periods.

Each History entry contains three pieces of information, as Table 2-1 shows.

Table 2-1: Description of the Entries in the Server History Tab

HISTORY ENTRY	DESCRIPTION
Event Description	Description of the operation performed, for example: `Install Template (Job ID: 21870101L) completed successfully.` Locked with Reason: Changes to this server are not allowed
Modified By	The name of the Opsware user who made the change.
Date Modified	The date and time the change was made, for example: `Mon Aug 06 18:14:41 GMT+00:00 2001)`

Viewing Server History

Perform the following steps to view the server history:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server whose history you want to view.

   Or

   Search for the server whose history you want to view.

   See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

2. Click the name of the server whose information you want to see. The Managed Servers: Properties page appears for that server.

3. Click the History tab.

   By default, the view shows changes made within the past week.

Generating Server Reports

Every change made to managed servers is recorded in the Opsware Model Repository. The Model Repository maintains precise information about the state and configuration of every server under management.

You can take advantage of this information though the Opsware Data Center Intelligence (DCI) Reporting component. The DCI Reporting component provides dynamic and detailed information about the operational environment. The DCI Reporting component provides a new level of visibility into the operational environment that can help everyone, from system administrators to the CIO, to make better decisions.

DCI Reporting provides the following features and benefits:

- Exact information about the latest system state and configurations
- Visibility across the *entire* operational environment
- Accurate and detailed change history information
- A comprehensive set of patch reports
- The ability to extend the DCI reports

See the *Opsware System 4.7 DCI Administrator's Guide* for information about how to set up the DCI Reporting component.

See the online DCI reporting documentation for information about how to use and run the reports. The Data Center Intelligence component must be installed and running in the facility to view the online documentation.

The Opsware Data Center Intelligence Reporting component is an optional component. By default, it is not installed with the Opsware System. If this reporting component is not available for your organization, contact your Opsware Support Representative for information about how to obtain it so that you can generate reports for your managed servers.

Server Life Cycle

This section provides information on the server lifecycle within the Opsware System and contains the following topics:

- Server Lifecycle Overview
- Server Properties Overview
- Server Management Tasks Related to the Server Life Cycle
- Changing the Use and Stage Values for Servers

- Editing the Properties of a Server
- Deactivating a Server Overview
- Deactivating a Server
- Deleting a Server from the Opsware System
- Cloning a Server

Server Lifecycle Overview

The Opsware System is designed to enable multiple teams to work together to provision servers. The OS Provisioning Subsystem allows IT teams to separate the tasks of readying servers for provisioning (such as mounting servers in racks and connecting them to power and a network) from provisioning the servers with operating systems and applications.

For example, someone mounts a new server in a rack and connects it to the Opsware build network. Next they boot the server for the first time by using an Opsware Boot Floppy or by using the network.

At a later time, a different system administrator can select the available server from the Server Pool list and provision it with an OS. In the *available* state, servers do not have the target OS installed and might not have access to disk resources.

During OS provisioning, servers progress through the following Opsware life cycle state changes:

Unprovisioned (No OS Build Agent) ➤ Available ➤ Installing OS ➤ Managed

Table 2-2 describes the Opsware server life-cycle values.

Table 2-2: Opsware System Life-Cycle Values for Servers

OPSWARE LIFE CYCLE VALUE	DESCRIPTION
Server Pool Values	
Available	Indicates a server on which the OS Build Agent was installed and is running, but the target OS has not been installed on the server The OS Build Agent is a small agent that can run in the memory of the bare-metal server. See "Operating System Provisioning" on page 237 in this chapter for more information.
Installing OS	Indicates that a user is installing the target OS on the server The server stays in the Server Pool list until the installation process finishes successfully. Then the server moves to the Managed Servers list. See "Installing an OS by Using a Template" on page 255 in this chapter for more information. See "Installing an OS by Using a Custom Installation" on page 258 in Chapter 4 for more information
Build Failed	Indicates a server on which the OS Build Agent was installed and is running, but the installation of the target OS failed The server remains in the Server Pool list with this status for 7 days before the Opsware System deletes the entry. See "Recovering When an OS Build Agent Fails to Install" on page 252 in Chapter 4 for more information.

Table 2-2: Opsware System Life-Cycle Values for Servers

OPSWARE LIFE CYCLE VALUE	DESCRIPTION
Managed Server Values	
Managed	Indicates a server that the Opsware System is managing. The Opsware System performs periodic reachability checks on managed servers. After a server reaches this life cycle state, the entry for the server moves from the Server Pool list to the Managed Servers list. On managed servers, you can use the Opsware System to install applications and patches.
Deactivated	Indicates an Opsware-managed server that was removed from service. However, the server's history still exists in the Opsware System. Deactivated servers are not reachable.

Table 2-3 shows which server icons appear in the Opsware Command Center and explains what each icon indicates in regard to the server life cycle.

Table 2-3: Server Icons in the Opsware Command Center

SERVER ICON	DESCRIPTION
	Indicates a server that is *available* to have a target OS installed on it and on which an Opsware OS Build Agent is installed Appears in the Server Pool list
	Indicates a server on which the OS Provisioning Subsystem is in the process of installing the target OS Appears in the Server Pool list
	Indicates an *available* server on which an error occurred while the OS Provisioning Subsystem was installing a target OS Appears in the My Jobs panel in the home page, in the list in the My Jobs page, and in the Server Pool list

Table 2-3: Server Icons in the Opsware Command Center

SERVER ICON	DESCRIPTION
	Indicates a server that the Opsware Command Center is managing and that the Opsware System can communicate with. An Opsware Agent is running on the server Appears in the My Jobs panel in the home page, in the list in the My Jobs page, in the Managed Servers list, and in the server lists in the Opsware wizards
	Indicates a server that is scheduled for an operation (install software, uninstall software, and so forth) Appears in the My Jobs panel in the home page and in the list in the My Jobs page
	Indicates a server that the Opsware Command Center is managing but is currently locked Appears in the Managed Servers list and the Opsware wizards
	Indicates a managed server that the Opsware System cannot communicate with (it is Not Reachable) because the Opsware Agent on the server cannot connect to the Opsware System If you want to discover reasons why the managed server is unreachable, you can run a Communication Test. See “Agent Reachability Communication Test” on page 80 in this chapter for more information. Appears in the Managed Servers list
	Indicates a managed server that the Opsware System cannot communicate with (it is Not Reachable) and the server is currently locked Appears in the Managed Servers list and the Opsware wizards
	Indicates a managed server on which an error occurred while the Opsware System was installing or uninstalling software Appears in the My Jobs panel in the home page and in the list in the My Jobs page

Table 2-3: Server Icons in the Opsware Command Center

SERVER ICON	DESCRIPTION
	Indicates a managed server on which a warning occurred while the Opsware System was installing or uninstalling software Appears in the My Jobs panel in the home page and in the list in the My Jobs page
	Indicates a server that was deactivated in the Opsware System so that it is currently not managed and not reachable Appears in the Managed Servers list and in the server lists in the Opsware wizards (however, it is not selectable in the wizards)
	Indicates a group of servers. The same states that apply to single servers apply to groups. Appears in the Server Groups pages

Server Properties Overview

Figure 2-25 shows the Server Properties Columns. Table 2-4, Table 2-5, and Table 2-6 describe the Status, Stage, and Use properties for managed servers.

Figure 2-25: Server Properties Columns in the Managed Servers List

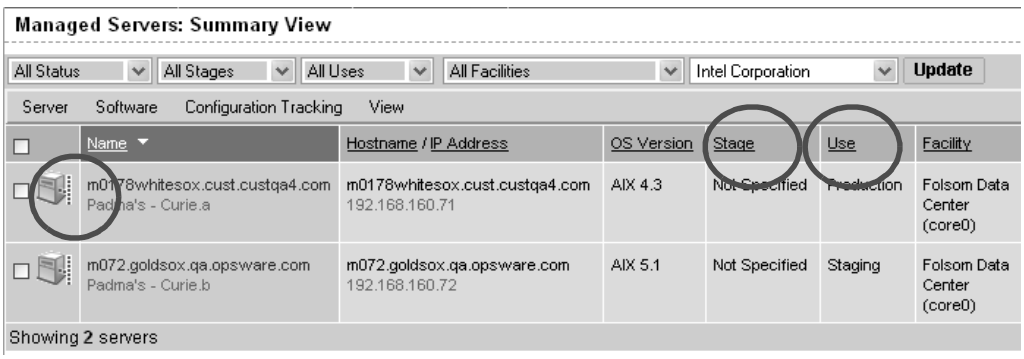

The Status property is represented by an icon in the first column in the Managed Servers list.

Status (short for Agent Status) is set automatically by the Opsware System.

The Opsware System toggles each server between OK and Not Reachable by reachability checks.

The Status value specifies the ability of the Opsware System to manage servers. The Opsware System automatically detects the status of servers. To verify the current status of a server, click the Update button in the Server Properties page for that server.

Table 2-4: Values for the Status Property for Managed Servers

STATUS VALUE	DESCRIPTION
OK	Server is manageable by the Opsware System. Represented as text (Ok) in the properties page for a server. Represented as an icon in the Managed Servers and Server Pool lists: 
Not Reachable	Server is unmanageable by the Opsware System due to error (for example, it cannot connect to the Opsware core); automatically set by the Opsware System. Represented as text (Not Reachable) in the properties page for a server. Represented as one of these icons in the Managed Servers list:  If you want to discover reasons why the managed server is unreachable, you can run a Communication Test. See “Agent Reachability Communication Test” on page 80 in this chapter for more information.

Stage (short for Deployment Stage) is set by a user.

The Stage value specifies the stages of deployment for servers; for example, a server is live or offline.

Your Opsware administrator can change the values for the Stage property. By default, the Opsware System is installed with the following Stage values.

Table 2-5: Values for the Stage Property for Managed Servers

STAGE VALUE	DESCRIPTION
In Deployment	Initial stage after being fully initialized.
Live	Your organization defines the meaning of this stage.
Not Specified	The default value for a server. Cannot be changed by the Opsware administrator.
Offline	Your organization defines the meaning of this stage.
Ops Ready	Your organization defines the meaning of this stage.

Use (short for Server Use) is set by a user.

The Use value specifies how an organization is utilizing servers. For example, a server is a staging server. Users set this property for servers.

By default, the Opsware System is installed with the following Use values. Except for the Staging, Production, and Not Specified values, an Opsware administrator can change the default values. The CDR Subsystem depends on the Staging and Production and Use values. Therefore, these default values cannot be changed or deleted.

Table 2-6: Values for the Use Property for Managed Servers

USE VALUE	DESCRIPTION
Development	A server that is not being used in production
Not Specified	The default value
Production	Fully live in-use servers (includes Opsware core servers)
Staging	A staging server for production

Server Management Tasks Related to the Server Life Cycle

Managing servers in the Opsware System involves the following standard tasks:

- Bringing a new server into the Opsware System so that it appears in the Server Pool

  See "Booting New Servers" on page 246 in Chapter 4 for more information.

- Installing an operating system on a server

  See “Ways to Install Operating Systems on Servers” on page 254 in Chapter 4 for more information.

- Installing a patch

  See “Overview of Installing and Uninstalling Patches” on page 430 in Chapter 8 for more information.

- Installing an application

  See “Installing and Uninstalling Software” on page 385 in Chapter 7 for more information.

- Reprovisioning a server with a new OS

  See “Reprovisioning a Solaris or Linux Server” on page 263 in Chapter 4 for more information.

  You can reprovision Solaris and Linux servers so that they are running another version of the same OS so long as the hardware supports that new version of the OS.

  You can reprovision servers built by the Opsware System and Opsware assimilated servers by using this feature.

You cannot reprovision a Linux server so that it runs a Windows OS.

- Deactivating a managed server so that the Opsware System no longer manages it.

  See “Deactivating a Server Overview” on page 77 in this chapter for more information.

You accomplish server management tasks by using the menus in the Managed Servers list, as Figure 2-26 shows.

Figure 2-26: Menus in the Managed Servers List

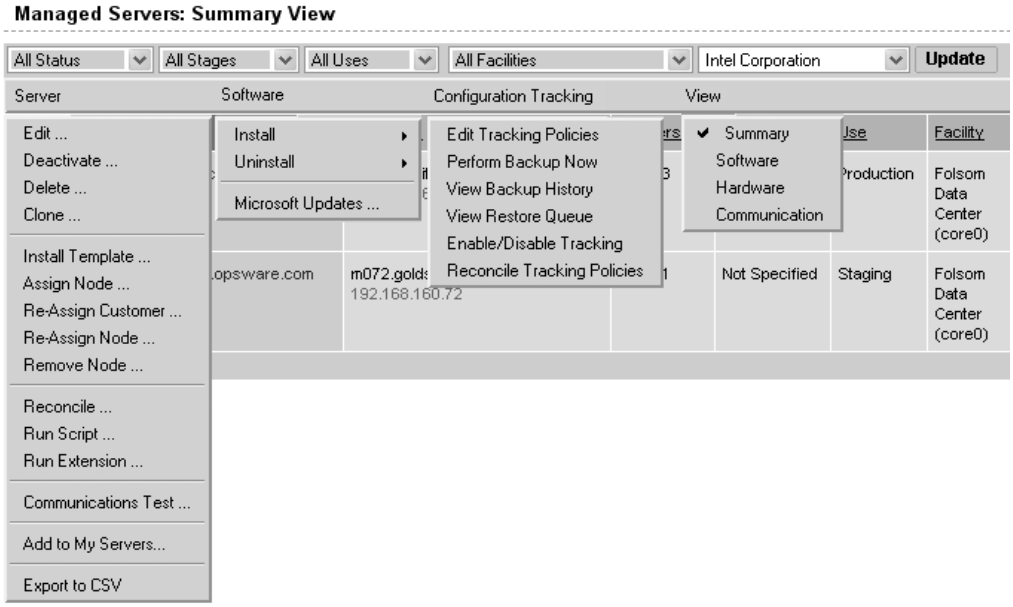

See "Application Provisioning" on page 385 in Chapter 7 for more information.

See "Reconcile" on page 441 in Chapter 9 for more information.

See "Script Execution Subsystem" on page 455 in Chapter 10 for more information about how to run a distributed script on a server.

See "About the Managed Servers List" on page 43 in this chapter for information about managed servers.

Changing the Use and Stage Values for Servers

Perform the following steps to change the Use and Stage values for multiple servers simultaneously:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server that you want to deactivate.

   Or

   Search for the server that you want to deactivate.

See “Searching with Advanced Search” on page 48 in this chapter for more information. See “Searching for Servers by IP Address” on page 55 in this chapter for more information.

2 Select the servers that you want different Use or Stage values for.

3 Choose Server ➤ Edit from the menu above the Managed Servers list. A window prompts you to select different values, as Figure 2-27 shows.

Figure 2-27: Edit Server Popup Window

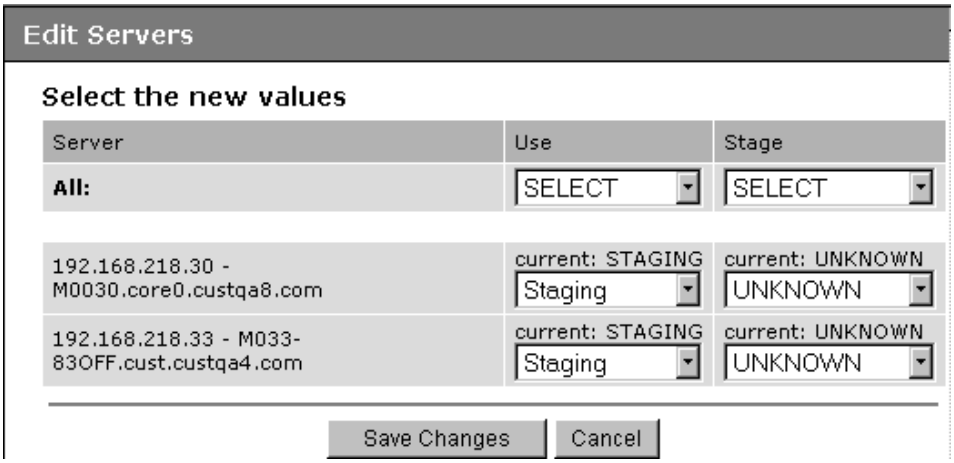

4 Select the Use and Stage values from the lists.

5 Click the Save Changes button. The window closes and the Managed Servers list refreshes with the updated values.

Editing the Properties of a Server

You can edit a server only if you have permission in the Opsware Command Center to access the customer to whom the server is associated.

When you edit the properties of a server, the server itself does not change; how the Opsware System views it changes. Perform the following steps to edit the properties of a server:

1 From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server whose properties you want to edit.

Or

Search for the server that you want to edit.

See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

2 Click the Server Display name. The Properties page for the server appears.

3 Change any of the following properties for the server:

- To change the name that appears in the Opsware Command Center, edit the text in the Name field.
- To change the description of the server, edit the Notes field.
- To change the customer associated with the server, select a different customer from the list. Your Opsware administrator defines the options for customer selections. Contact your Opsware administrator if the list does not contain the customer that you want to associate with this server.

You cannot change the customer associated with a server when the server is part of a CDR service, synchronization, or sequence. See "Defining CDR Services, Synchronizations, and Sequences" on page 559 in Chapter 12 for more information.

- To change the Use or Stage of the server, make your changes in either of those lists.

  See "Server Properties Overview" on page 70 in this chapter for more information.
- To change whether configuration tracking is enabled or disabled for the server, select a value from the list.

  See "Automated Configuration Tracking" on page 493 in Chapter 11 for more information.

4 To save your changes, click the Save button.

To change the custom attributes of the server, click the Custom Attributes tab. The Managed Servers: Custom Attributes page appears. See "Managing Custom Attributes" on page 157 in this chapter for more information.

Deactivating a Server Overview

You will want to deactivate a server when the Opsware System removes the server from management. For example, you are moving the server to a warehouse for storage. Additionally, you might choose to deactivate a server when you need to rebuild it from scratch, without using the OS Provisioning Subsystem.

When you deactivate a server, information about the server remains in the Opsware Model Repository for auditing purposes.

After you deactivate a server, you can reactivate it by re-installing an Opsware Agent with the Opsware Agent Installer and the `--clean` command line option.

See "Server Assimilation" on page 139 in this chapter for more information.

When you deactivate a server, you accomplish the following tasks:

- Reset the node assignments in the Software Tree to the defaults.
- Remove custom attributes from the server.
- Delete any configuration tracking policies from the server that are associated with backups.
- Set the server life cycle value to Deactivated.

You cannot deactivate a server when it is part of a CDR service, synchronization, or sequence. See "Defining CDR Services, Synchronizations, and Sequences" on page 559 in Chapter 12 for more information.

Deactivating a Server

Perform the following steps to deactivate a server:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server that you want to deactivate.

   Or

   Search for the server that you want to deactivate.

   See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

2 Select the servers that you want to deactivate.

3 Choose Server ➤ Deactivate from the menu above the Managed Servers list. A confirmation dialog box prompts you to confirm the deactivation.

4 Click OK. The Managed Servers list refreshes and the server appears with a deactivated icon.

Deleting a Server from the Opsware System

When you want to remove all record of a server from the Opsware System, you can delete it.

You must deactivate a server before you can delete it from the Opsware System.

When you delete a server from the Opsware System, it has these effects:

- Deletes all job information in the My Jobs feature
- Deletes the server from the Model Repository

Perform the following steps to delete a server:

1 From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server that you want to delete.

Or

Search for the server that you want to delete.

See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

2 Select the servers that you want to delete.

3 Choose Server ➤ Delete from the menu above the Managed Servers list. A confirmation dialog box prompts you to confirm the deletion.

4 Click OK. The Managed Servers list refreshes and the server disappears from the list.

Cloning a Server

The Opsware Command Center includes a feature that allows users to copy a server. Copying (referred to as *cloning* in the Opsware Command Center) is useful when you need to add more capacity to your operational environment.

A user selects a source server (the master server) and copies the configuration of that server to one or more other servers (the target servers). The other servers are assigned to every node to which the master server is assigned. The copied servers contain the same operating system, software applications, and configuration as the original server (though the customer association and facility location remain the same on the target servers). The copied servers also include any changes to the default server configuration made through the Opsware System.

The only restriction to cloning a server is that both servers need to be on the same platform. However, any existing user-configurable nodes are removed from the target server when you clone servers. The target server is made to look exactly like the master server in terms of node assignments.

Perform the following steps to clone a server:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the servers that you want to clone.
2. Select two or more servers. When copying servers, you must select at least two servers so that one server can be the master. The other servers are the ones that you want to copy the nodes to.

   Opsware Inc. recommends that you select servers with similar hardware architecture (based on the value in the Reported OS column in the Managed Servers list).
3. Choose Server ➤ Clone from the menu above the Managed Servers list.

A window prompts you to select the *master* server, as Figure 2-28 shows.

Figure 2-28: Clone Servers Popup Window in the Opsware Command Center

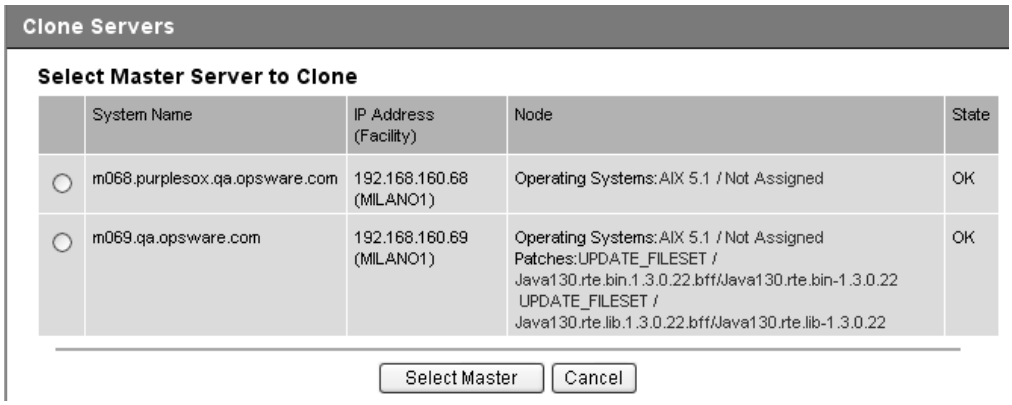

4. Select the option for the master server that you want to use and click the Select Master button.

   A confirmation page appears that shows how the nodes on the master server are copied to the server that you select.

5. Click the Commit Clone button.

   The Server List reappears and shows the updated target servers.

After you clone a master server to target servers, you must reconcile the target servers for the Opsware System to install the software on the target servers. See "Reconcile" on page 441 in Chapter 9 for information about reconciling servers.

Agent Reachability Communication Test

This section provides information on agent reachability Communication Tests within the Opsware System and contains the following topics:

- Communication Test Overview
- What Makes an Opsware Agent Unreachable?
- Communication Test Types
- Communication Test Errors

- Additional Information on a Communication Test
- Running a Communication Test on an Individual Server
- Running a Communication Test on Multiple Servers
- Viewing Servers by Communication Status
- Searching for Unreachable Servers
- Creating Communication Test DCI Reports
- Viewing My Jobs Communication Tests
- Exporting Unreachable Server Status List to CSV

Communication Test Overview

Sometimes an Opsware Agent can become unreachable, which means that the Opsware Command Center has difficulty communicating with the Opsware Agent. When an Opsware Agent is unreachable, the server it is installed on is considered unmanaged. This section explains how to use the Communication Test to find unreachable Opsware Agents and suggests ways that you can resolve these problems.

To help identify those managed servers that have unreachable agents, the Opsware Command Center runs periodic Communication Tests to verify that the Opsware Command Center can communicate with all servers under its management. You can always check the reachability of Opsware Agents by looking at the server's properties, or by viewing the current agent reachability status for all managed servers since the last Communication Test was run by choosing the Communication view from the managed servers list.

To determine the current reachability of a specific Opsware Agent, you can run a Communication Test to find those servers that have unreachable agents by using the Communication Test feature located in the Server menu of the managed server list. A Communication Test lists all servers with unreachable agents, returns specific errors associated with each unreachable Opsware Agent, and provides troubleshooting information to help you get the Opsware Agent back in working order.

You have the ability to check Opsware Agent reachability for individual servers, selected servers, or all servers under management of the Opsware Command Center. Each time that you run a Communication Test, this test is saved in the My Jobs panel, which allows

you to view a history of all the Communication Tests that you have run. You can even export the current reachability status of all managed servers to a Comma Separated Value (CSV) file.

What Makes an Opsware Agent Unreachable?

The Communication Test works by testing communication and data exchange between the specific components of the Opsware core and each managed server. The Opsware core is the entire collection of servers and services that provide Opsware services. In order to successfully manage servers, the Opsware core needs to be able to communicate with each Opsware Agent on all servers under Opsware management.

Communication Test Types

The Communication Test performs the following diagnostics to determine if an Opsware Agent is reachable:

- **Command Engine to Agent Communication (AGT)** – Determines if the Command Engine can communicate with the agent. The Command Engine is the Opsware system component that enables distributed programs to run across many servers. The Command Engine handles the entry of scripts into the Opsware Model Repository (the script storage location in the Opsware System) and the versioning of stored scripts.
- **Crypto Match (CRP)** – Checks that the SSL cryptographic files that the agent uses are valid.
- **Agent to Command Engine Communication (CE)** – Verifies that the agent can connect to the Command Engine and retrieve a command for execution.
- **Agent to Data Access Engine (DAE)** – Checks whether or not the agent can connect to the Data Access Engine and retrieve its device record. 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.
- **Agent to Software Repository Communication (SWR)** – Determines if the agent can establish an SSL connection to the Software Repository. The Software Repository is the central repository for all software that the Opsware system technology manages. It contains software packages for operating systems, applications, databases, customer code, and software configuration information.
- **Machine ID Mismatch (MID)** – Checks that the machine ID (MID) on the server matches the MID registered in the Model Repository for the agent.

When the test finishes, it returns results that show either success or failure for each test run on each server. For each failed test, the nature of the failure is indicated in the Test Summary column of the Communication Test results page. In some cases, the failure of one test might prevent other tests from being executed.

See “What an Opsware Agent Can Do on a Managed Server” on page 133 in this chapter for information about the Opsware Agent and its relationship to managed servers.

Communication Test Errors

After you run a Communication Test, three icons indicate the success or failure of agent reachability as Table 2-8 shows.

Table 2-7: Agent Unreachability Status Icons

STATUS ICON	DESCRIPTION
	Communication Test passed. Agent is reachable.
	Communication Test unable to be executed.
	Communication Test failed. Agent is unreachable.

Table 2-8 describes each type of Communication Test and all possible errors for each test.

Table 2-8: Communication Test Types with Possible Results

TEST	DESCRIPTION	RESULTS
Command Engine to Agent Communication (AGT)	Determines if the Command Engine can communicate with the agent	1 OK 2 Untested 3 Unexpected error 4 Connection refused 5 Connection timeout 6 Request timeout 7 Server never registered 8 Realm is unreachable 9 Tunnel setup error 10 Gateway denied access 11 Internal gateway error 12 Gateway could not connect to server 13 Gateway timeout
Crypto Match (CRP)	Checks that the agent's SSL cryptographic files are valid	1 OK 2 Untested 3 Unexpected error 4 Agent certificate mismatch 5 SSL negotiation failure

Table 2-8: Communication Test Types with Possible Results

TEST	DESCRIPTION	RESULTS
Agent to Command Engine Communication (CE)	Verifies that the agent can connect to the Command Engine and retrieve a command for execution	**1** OK **2** Untested **3** Unexpected error **4** Connection refused **5** Connection timeout **6** DNS does not resolve **7** Old agent version **8** Realm is unreachable **9** No gateway defined **10** Tunnel setup error **11** Gateway denied access **12** Gateway name resolution error **13** Internal gateway error **14** Gateway could not connect to server **15** Gateway timeout **16** No callback from agent

Table 2-8: Communication Test Types with Possible Results

TEST	DESCRIPTION	RESULTS
Agent to Data Access Engine (DAE)	Checks whether or not the agent can connect to the Data Access Engine and retrieve its device record	1 OK 2 Untested 3 Unexpected error 4 Connection refused 5 Connection timeout 6 DNS does not resolve 7 Old agent version 8 Realm is unreachable 9 No gateway defined 10 Tunnel setup error 11 Gateway denied access 12 Gateway name resolution error 13 Internal gateway error 14 Gateway could not connect to server 15 Gateway timeout

Table 2-8: Communication Test Types with Possible Results

TEST	DESCRIPTION	RESULTS
Agent to Software Repository Communication (SWR)	Determines if the agent can establish an SSL connection to the Software Repository	1 OK 2 Untested 3 Unexpected error 4 Connection refused 5 Connection timeout 6 DNS does not resolve 7 Old agent version 8 Server identification error 9 Realm is unreachable 10 No gateway defined 11 Tunnel setup error 12 Gateway denied access 13 Gateway name resolution error 14 Internal gateway error 15 Gateway could not connect to server 16 Gateway timeout
MID Match	Checks that the Machine ID (MID) on the server matches the MID registered in the Model Repository for the agent	1 OK 2 Untested 3 Unexpected error 4 MID Mismatch

See Appendix D for information about how to troubleshoot Communication Test Errors.

Additional Information on a Communication Test

In the case that the Communication Test cannot be performed on a server, you see an error named Unexpected Error with a small plus sign next to it. Click the plus sign to see the Additional Information window, which provides traceback information regarding the error. You can send this information to Opsware Customer Support to solve a problem of this nature. See Figure 2-29.

Figure 2-29: Additional Information on an Unexpected Error

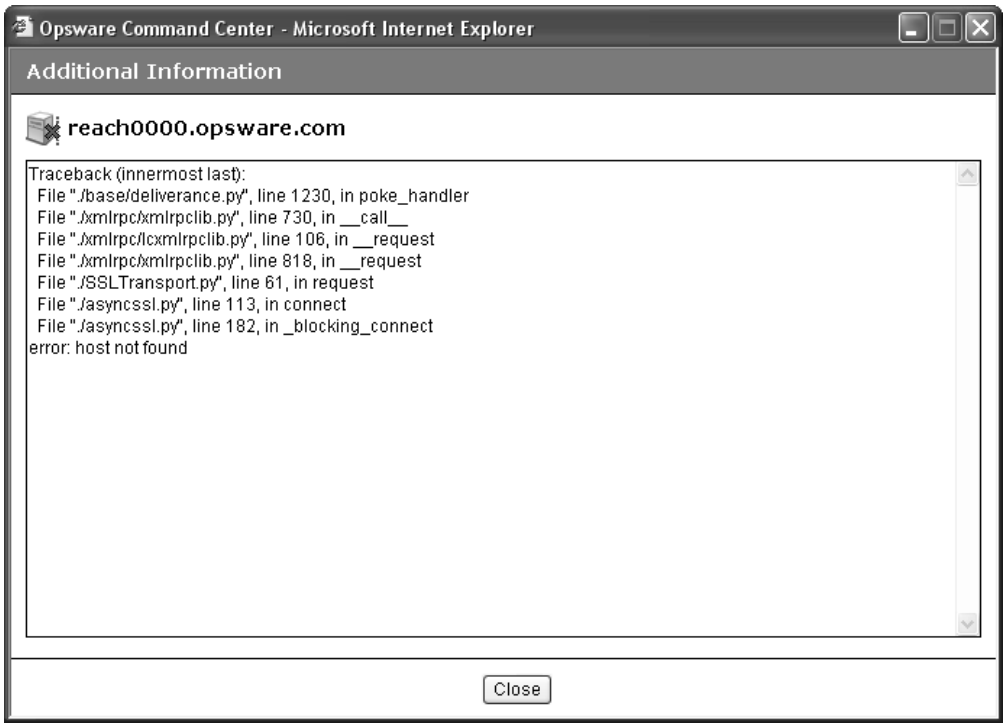

Running a Communication Test on an Individual Server

Perform the following steps to run a Communication Test on an individual server to find out if the Opsware Agent on that server is reachable:

1. From the navigation panel, click Servers ➤ Managed Servers.
2. From the Managed Servers list, click the display name of the server that you want to perform a Communication Test on.

   On the Server Property page, look in the Status field and notice that the server is either listed as Reachable or Not reachable. If listed as Not Reachable, a date indicates when the last regularly scheduled Communication Test was performed.

3 To see the results of the last Communication Test for this server, click the Details button. The Communication Test window for the server appears, as Figure 2-30 shows.

Figure 2-30: Communication Test Results on an Individual Server

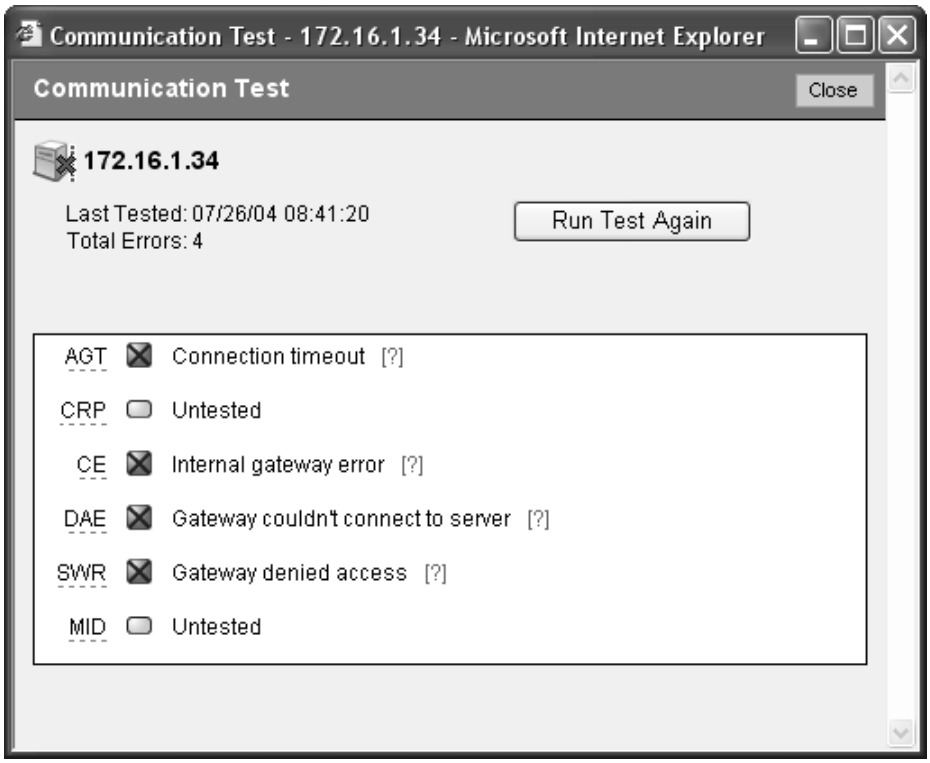

The results listed in this window show details from when the last regularly scheduled Communication Test was run.

4 To view troubleshooting information for any of the test errors, move your mouse over the error name (for example SWR). When your mouse cursor changes to a question mark, click the question mark to view the troubleshooting help.

5 To rerun the Communication Test, click the Run Test Again button. The new results display in the same window when the test finishes.

Running a Communication Test on Multiple Servers

Perform the following steps to run a Communication Test on multiple servers to find out which managed servers are not reachable:

1. From the navigation panel, click Servers ➤ Managed Servers.
2. In the Managed Servers Summary View page, select the servers for which you want a Communication Test.
3. Choose Server ➤ Communication Test. The Communication Test window opens and the test is initiated. The top of the test window shows the status report for the Communication Test, which indicates the time of the test, how many servers in the test were reachable, which servers were not reachable, and a progress bar. This summary information is shown in Figure 2-31.

Figure 2-31: Communication Test Summary

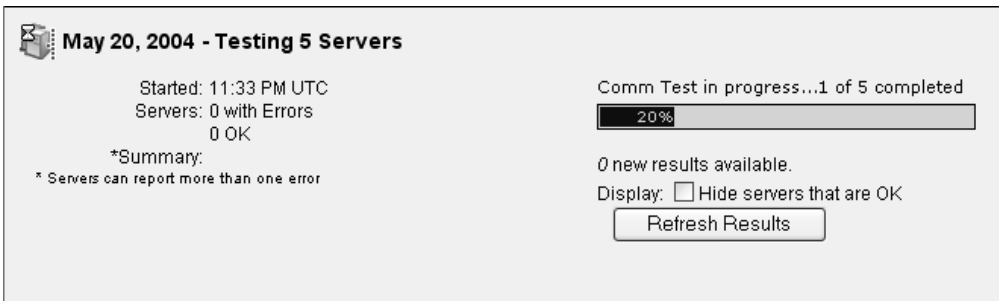

- **Date**: Provides the date of the test.
- **Statistics**: Shows start and finish time, total servers OK, total servers with unreachable agents, and a summary of errors.
- **Progress bar**: Provides live feedback of Communication Test progress. Progress data includes the number of servers completed, the total number of servers to be completed, and the list of servers completed so far.
- **Refresh Results button**: Refreshes the results screen with new results.

Below the summary section is a list of all Opsware Agents that were not reachable and their details, as Figure 2-32 shows.

Figure 2-32: Communication Test Results on an Unreachable Agent

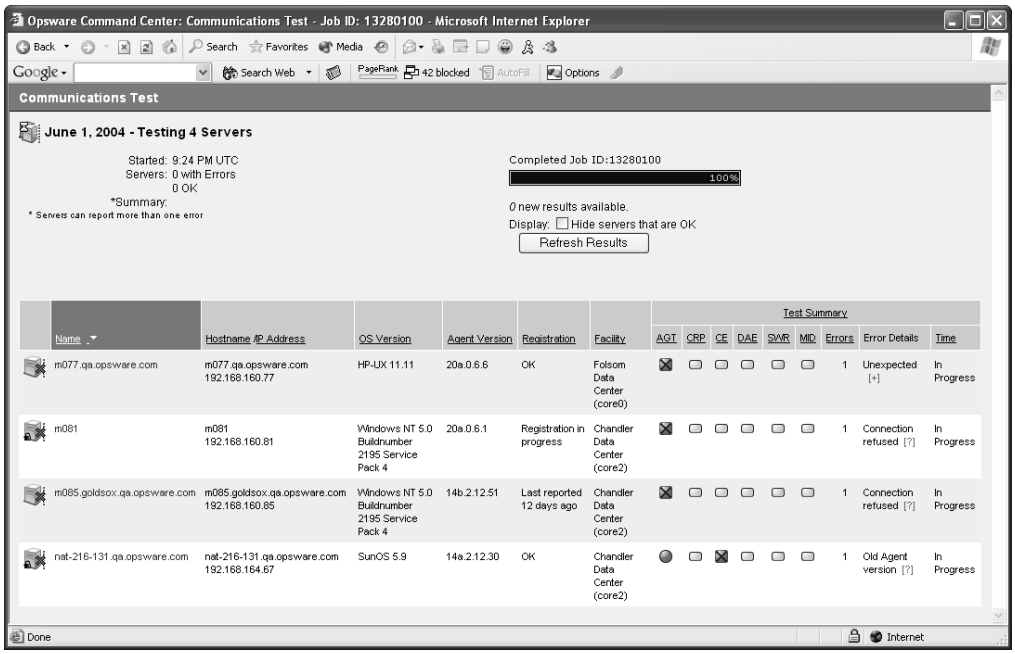

- The results section shows a table of all unreachable servers, detailing server name, hostname/IP address, OS, Opsware Agent version, registration (when the Opsware Agent last reported to the Opsware Command Center), and the time the test was completed.
- Click the title of each column to sort the test information by specific categories.
- The Test Summary section shows a list of all Communication Test types that were run, and which errors were returned. For information about each type of error and how to troubleshoot Opsware Agent reachability problems, click the link on each error to view online help.

Viewing Servers by Communication Status

Perform the following steps to view all managed servers with the most recent Communication Test results for each managed server:

1 From the navigation panel, click Servers ➤ Managed Servers.

2. From the View menu, choose Communication. A list of the most recent Communication Test that was run on all managed servers appears. All servers listed in this view are listed as unreachable or reachable since the last regularly scheduled Communication Test was run.
3. To sort the information in this view, click any of the column headings. For example, you might want to view the servers by number of errors (Click the Error column title) by name, by OS version, and so on).
4. To export this view to the CSV file format, choose Export to CSV from the Server menu.
5. To run a new Communication Test for some or all of the servers in this view, select the servers that you want to run a Communication Test on and then choose Communication Test from the Server menu.

Searching for Unreachable Servers

Another way you can discover unmanaged servers (those with unreachable agents) is to search for all servers that have a status of unreachable. See "Searching with Advanced Search" on page 48 in this chapter for more information.

Perform the following steps to search for unreachable agents:

1. From the navigation panel, click Servers ➤ Server Search. The Server Search page appears.

2 From the Server Search page, choose the Agent Reporting attribute from the first list, as Figure 2-33 shows.

Figure 2-33: Agent Reporting Server Search Attribute

More criteria displays for the search parameters.

3 From the far right list of search attributes, select Not Reporting. Your search parameters are *Agent Status is Not Reporting*.

4 Click the Search button. Wait a few moments (the speed of the search results depends on how many managed servers are being searched) for the results. If any of your managed servers are not reachable, these servers will be listed in the search results. All managed servers with unreachable agents are shown with an X next to the server icon, as this icon shows:

5 To run a new Communication Test on these servers, select the server (check box next to the server), and choose Communication Test from the Server menu.

Creating Communication Test DCI Reports

If you would like to view a Communication Test in a report format, you can use the Data Center Intelligence (DCI) Reporting tool to create a printable report of the Communication Test results. After you make a report of the test results, you can export the report to HTML, CSV, Microsoft Word, as well as other formats, so that you can exchange the report with others.

For information on how to create a DCI Communication Test report, view the DCI online help by clicking the Help button in the Opsware Command Center.

Viewing My Jobs Communication Tests

Each time that you run a Communication Test, the information is saved as a My Job. This feature automatically saves a history of all tests that you run. To view saved Communication Tests, perform the following steps:

1. From the navigation panel, click My Job.
2. From the My Job list, click the Communication Test job that you want to view.
3. In the Communication Test window, wait a few moments for the Communication Test information to load, then click the View Details button. You see the Communication Test window.

See the *Opsware System 4.7 Administration Guide* for information about My Jobs.

Exporting Unreachable Server Status List to CSV

Perform the following steps to export the list of all servers that have a status of unreachable to the CSV file format:

1. From the navigation panel, click Servers ➤ Managed Servers.
2. From the View menu, choose Communication. You see a list of all servers that are in an unreachable or reachable state.
3. To export a list of these servers to the CSV file format, select the check box next to each server that you want to include in the report, then from the Server menu choose Export to CSV.

Server Locking

This section provides information on server locking within the Opsware System and contains the following topics:

- Server Locking Overview
- Locking or Unlocking Multiple Servers
- Locking or Unlocking a Server
- Effects of Server Locking on the Opsware System

Server Locking Overview

The Opsware System provides a server locking feature that allows users with read/write access to lock a managed server, which prevents any server-modifying operations from being performed on that server until it is unlocked. Only users with read/write access to that specific server can unlock the server.

If a locked server is among those selected for a server-modifying operation, the Opsware Command Center prevents the operation from being performed, and displays an error message that says that the operation was prevented because the server was locked. The message includes the ID of the user who locked the server, the date that the server was locked, and any comments that the user entered.

Also, it is possible to flag scripts that have the potential to modify servers to indicate whether or not running the script will modify the server. The server-locking feature prevents server-modifying scripts from being run on locked servers if the script has been flagged.

You can view the current lock status of servers on the Properties tab of the Server Details page, and you can use the Server Search function to search for servers by lock status.

Locking or Unlocking Multiple Servers

Perform the following steps to lock or unlock multiple servers:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the servers that you want to lock or unlock.

   Or

   Search for the server that you want to lock or unlock.

See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

2. Click the check boxes next to each server you want to lock or unlock.
3. Select Edit from the Server menu. The selected servers appear in the Edit Server page.
4. To change the Locked value for all the servers on the list, use the Locking Status list at the top of the page or change the values for each server in the list.
5. Optionally, enter text in the Reason text box.

   After you save your changes, the text that you enter appears below the Locking Status field for each server as part of a sentence that reads something like "Last Locked on 16-Dec-2003 11:32:16 AM by [name of administrator] because [reason text you entered]."
6. Click the Save Changes button.

Locking or Unlocking a Server

Perform the following steps to lock or unlock a server:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server you want to lock or unlock.

   Or

   Search for the server that you want to lock or unlock.

   See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.
2. Click the name of the server that you want to lock or unlock. The Managed Servers Properties page appears.
3. In the Lock State field, click the radio button next to Locked or Unlocked.
4. Optionally, enter text in the Reason text box.

   After you save your changes, the text that you enter appears below the Locking Status field for each server as part of a sentence that reads something like "Last Locked on 16-Dec-2003 11:32:16 AM by [name of administrator] because [reason text you entered]."

5 Click the Save button at the bottom of the page.

Effects of Server Locking on the Opsware System

Locked servers prevent certain functions from being performed. The following tables represent areas of the system that are affected by server locking and show the impact of server locking on each function in that area. The affected areas are:

- Server Lists
- Tasks Panel of the Home Page
- Server Properties Page
- Distributed Script Execution
- Code Deployment – Run Service
- Code Deployment - Synchronize
- Code Deployment - Run Sequence

Server Lists

Table 2-9 describes which actions in the Server menu are allowable for locked servers.

Table 2-9: Menus in the Managed Servers List

MENU ITEM	BEHAVIOR IF LOCKED
Server – Edit	Allowed
Server – Deactivate	Not allowed
Server – Delete	Not a valid state. Only deactivated servers can be deleted. Deactivated servers cannot be locked.
Server – Clone	Allowed
Server – Install Template	Not allowed
Server – Assign Node	Allowed
Server – Reassign Customer	Allowed Regardless of whether the server is locked or unlocked, you *cannot* re-assign a server to another customer when the server is part of a CDR service, synchronization, or sequence.

Table 2-9: Menus in the Managed Servers List

MENU ITEM	BEHAVIOR IF LOCKED
Server – Reassign Node	Allowed
Server – Remove Node	Allowed
Server – Reconcile	Not allowed
Server – Run Script	If locked servers are selected to run a script, Ad-Hoc scripts are not available, and scripts that make changes to the server are not allowed.
Server – Run Extension	Using the Opsware Custom Extensions feature, you can run canned scripts that contain multiple commands. If one of the commands contains "changes server," and a locked server has been selected to run that script, an error message is generated when the system attempts to run that command.
Communication Test	Allowed
Server – Add to My Servers	Allowed
Server – Export to CSV (Available in server search only)	Allowed
Software - Install - Application	Not allowed
Software - Install - Patch	Not allowed
Software - Install - Operating System	Not allowed
Software - Uninstall - Application	Not allowed
Software - Uninstall - patch	Not allowed
Software - Microsoft Updates	Not allowed
Configuration Tracking – Edit Tracking Policies	Allowed
Configuration Tracking – Perform Backup Now	Allowed

Table 2-9: Menus in the Managed Servers List

MENU ITEM	BEHAVIOR IF LOCKED
Configuration Tracking – View Backup History	Allowed
Configuration Tracking – View Restore Queue	Allowed to view, but not allowed to do restores
Configuration Tracking – Enable/disable Tracking	Allowed
Configuration Tracking – Reconcile Tracking	Allowed

Tasks Panel of the Home Page

Table 2-10 describes which actions in the Tasks panel of the Home Page are allowable for locked servers.

Table 2-10: Links in the Tasks Panel

LINK	BEHAVIOR IF LOCKED
Install OS	Not a valid state, only server pool servers can be selected. Servers in a server pool cannot be locked.
Prepare OS	N/A – no servers listed
Install Patch	Locked servers not selectable
Uninstall Patch	Locked servers not selectable
Upload Patch	N/A – no servers listed
Microsoft Patch Update	Locked servers not selectable
Install Software	Locked servers not selectable
Uninstall Software	Locked servers not selectable
Install Template	Locked servers not selectable
Deploy Code	See Code Deployment tables
Run Distributed Script	When selecting a script, if you select Ad-Hoc, or your saved script shows a “Yes” in the “changes server,” column, then locked servers cannot be selected to run a script.

Table 2-10: Links in the Tasks Panel

LINK	BEHAVIOR IF LOCKED
Run Custom Extension	Using the new Opsware Custom Extensions feature, you can run canned scripts that contain multiple commands. If one of the commands contains "changes server," and a locked server is selected to run that script, an error message is generated when the system attempts to run that command.
View Reports	N/A

Server Properties Page

Table 2-11 describes which actions in the Server Properties page are allowable for locked servers.

Table 2-11: Fields in the Server Properties Page

FIELD OR SECTION	BEHAVIOR IF LOCKED
Properties Tab – Name	Change allowed
Properties Tab – Notes	Change allowed
Properties Tab – Customer	Change allowed
Properties Tab – Server Use	Change allowed
Properties Tab – Deployment Stage	Change allowed
Properties Tab – Agent Status	Change allowed
Properties Tab – Config Tracking	Change allowed
Network Tab – All	No changes allowed
Nodes Tab – Deactivate Button	Not allowed
Install List Tab	N/A
Installed Packages Tab	N/A
Custom Attributes Tab	Change allowed
Config Tracking Tab	Viewing allowed, backups allowed, restores not allowed
Services Tab	N/A
History Tab	N/A

Distributed Script Execution

Table 2-12 describes which actions in the Distributed Script Execution pages are allowable for locked servers.

Table 2-12: Fields in the Distributed Script Execution Pages

FIELD OR SECTION	BEHAVIOR IF LOCKED
My Scripts/Shared Scripts tab – Run button	If the selected script is flagged "changes server," selecting locked servers is not allowed.
Edit Script Contents Tab	New field - "Changes Server" – has been added, but behavior of the tab has not changed.

Code Deployment – Run Service

Table 2-13 describes which Code Deployment - Run Service actions are allowable for locked servers.

Table 2-13: Run Service Actions

ACTION	BEHAVIOR IF LOCKED
Perform service operations by service name	Locked servers are not selectable, the check box selection is disabled, and the server is marked as locked
Perform service operations by hostname	Locked servers are unlinked and marked as locked

Code Deployment - Synchronize

Table 2-14 describes which Code Deployment - Synchronize actions are allowable for locked servers.

Table 2-14: Synchronization Action

ACTION	BEHAVIOR IF LOCKED
Perform/request synchronization	Locked servers can be selected as the source server for a synchronization, but not as the destination server. Locked servers are not selectable, the check box selection is disabled, and the server is marked as locked. Preview and List buttons do not appear.

Table 2-15: Synchronization Action

ACTION	BEHAVIOR IF LOCKED
Perform/request synchronization	Locked servers can be selected as the source server for a synchronization, but not as the destination server. Locked servers are not selectable, the check box selection is disabled, and the server is marked as locked. Preview and List buttons do not appear.

Code Deployment - Run Sequence

Table 2-16 describes which Code Deployment - Synchronize actions are allowable for locked servers.

Table 2-16: Run Sequence Action

ACTION	BEHAVIOR IF LOCKED
Perform/request sequence	If any of a sequence's defined servers are locked (unless that server is the source server for a synchronization), the sequence itself will not be selectable, and the check box for the sequence is disabled. The locked servers are marked as locked, and the check box next to each locked server is disabled.

Scheduling Server Management Jobs

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

- Scheduling Server Management Jobs Overview
- Viewing Job Details
- Scheduling and Notifying Server Management Tasks
- Time-Outs for Server Management Jobs

Scheduling Server Management Jobs Overview

The My Jobs feature provides information about the following Command Engine scripts:

- Reconcile (install software, uninstall software, install a template, patching servers, and reconcile)
- OS provisioning
- CDR requests
- Distributed script execution
- Custom Extensions

The My Jobs information is available only on a per-user basis. You cannot log in as an Opsware administrator to see the jobs that other Opsware users have run.

The My Jobs information appears in two places in the Opsware Command Center:

- A panel on the Opsware Command Center home page that lists your most recent six jobs
- A page (accessed by clicking My Jobs in the navigation panel) that lists all the jobs that you have run

  The Opsware System maintains information about the server operations that you have run for the last 30 days in the My Jobs list. By default, the jobs are deleted from the Opsware Model Repository after 30 days. (The bottom of the My Jobs page indicates how long this interval is set for the Opsware System installed at your organization.)

For each job, the My Jobs lists display the following information:

- The name of the job, which is a link to a page that displays more detailed information about the job
- The date and time the job started or is scheduled to start (using your preference for time display)
- The number of servers that the job affects
- The status of the job:
  - Scheduled
  - In Progress
  - Completed
  - Completed with errors
  - Completed with warnings
- You can search for an existing Job by the Job's ID. On the Home page, click the down arrow in the navigation panel to open the Search box. Select Job from the drop-down list and enter a Job ID and click Search.

Viewing Job Details

Perform the following steps to view the job details:

1. From the Opsware Command Center home page, click the link in the My Jobs panel for the job that you want to view, as Figure 2-34 shows.

Figure 2-34: My Jobs Panel in the Opsware Command Center Home Page

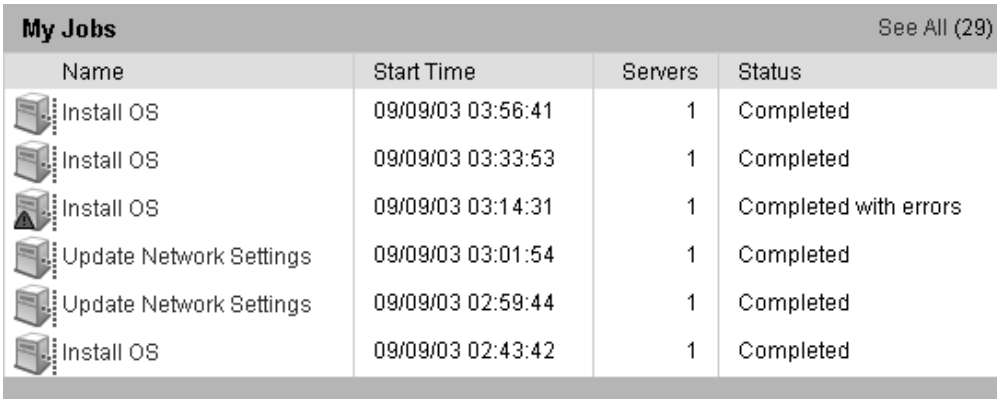

Or

From the navigation panel, click My Jobs and then click the link for the job to open a window that shows the details of the job, as Figure 2-35 shows.

Figure 2-35: My Jobs Page Accessed from the Navigation Panel

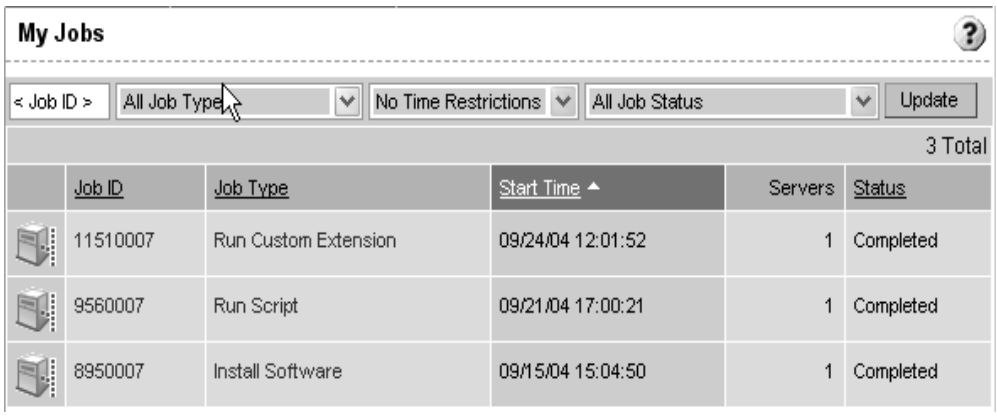

The My Jobs page displays the operations that you performed. See Figure 2-36.

Figure 2-36: Details Page for a Job in the My Jobs Panel

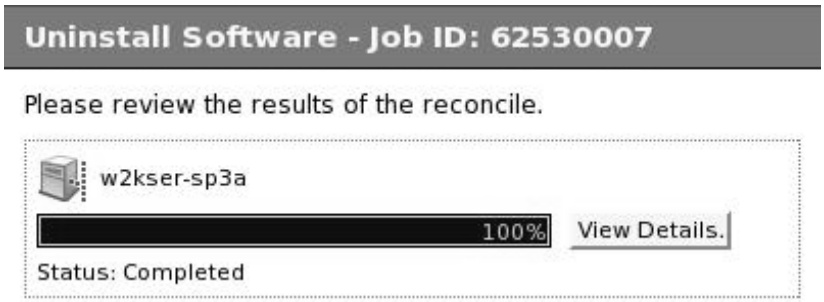

2 Click the View Details button to see detailed information about the job. The My Jobs information contains a build log for the job. This build log contains any error messages that the Opsware System generates. See Figure 2-37.

Figure 2-37: Reconcile Details Page in the My Job Window

See "Reconcile" on page 441 in Chapter 9 for information about how to interpret the reconcile operation output.

Scheduling and Notifying Server Management Tasks

Scheduling a Job

Perform the following steps to schedule server management tasks:

1. In the Confirmation step of an Opsware wizard, click the Schedule Job link at the top of the wizard window. (If you are installing or uninstalling software, you are prompted to schedule the job in the Preview step of the Opsware wizard.) See Figure 2-38.

Figure 2-38: Scheduling an Operation in an Opsware Wizard

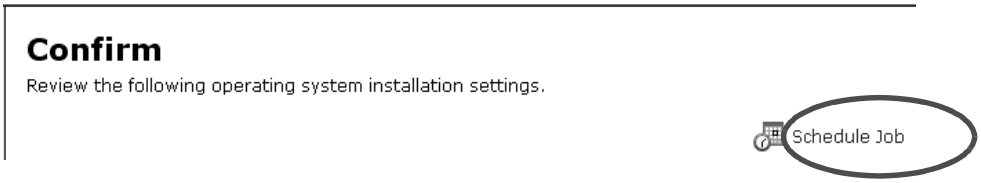

The page refreshes and controls appear that allow you to specify the date and time that you want the operation to run. See Figure 2-39.

Figure 2-39: Scheduling Options in an Opsware Wizard

2. Specify the date and time that you want the operation to run and click the Schedule button at the bottom of the wizard window. The Opsware wizard displays a message that the operation was successfully scheduled.
3. Click the Close button to end the Opsware wizard.
4. Additionally, you can view a scheduled job in the My Jobs page and change the time for the job to run, or cancel the job entirely. (Click the name of a scheduled job to open a window to change the date or time that the job will run or cancel it.)

Email Notification

The email notification feature provides you with an option to receive an email summarizing the job details when a job is over, and to notify others at the address they have registered with the Opsware System.

You have the option of sending the notification

- on the job success only
- on the job failure only
- on any result

You can also send notification to various people based on the condition of the job. For example, you can select to send the notification to your manager only on the job success, select to send the notification to yourself on any result of the job, and select to send the notification to support only on the job failure.

To send an email about the job details, choose the Condition option on the Schedule and Notify page and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field. See Figure 2-41.

Figure 2-40: Notifying about a Job in the Opsware Command Center

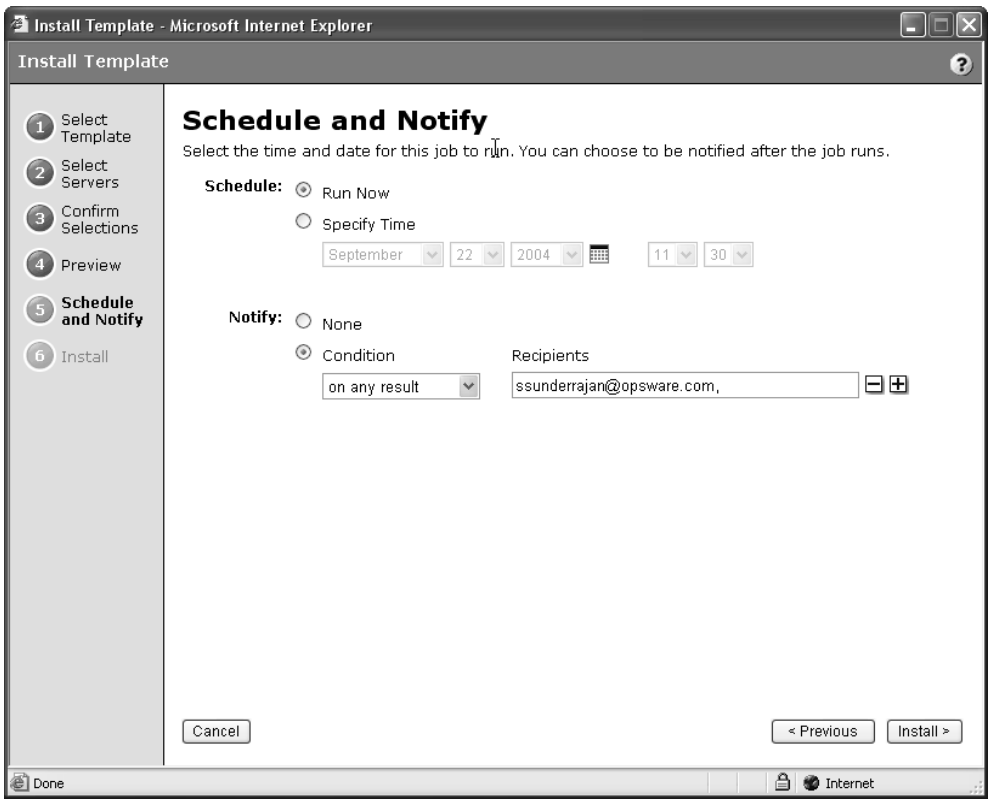

Time-Outs for Server Management Jobs

Table 2-17 describes the time-out values that apply to the server management operations in the Opsware System.

Table 2-17: Time-Outs in the Opsware System

TIME OUT (MINUTES)	OPSWARE SYSTEM OPERATION
420 (7 hours)	Reconciling software (installation and uninstallation)
2	Starting Command Engine sessions in response to a command. If the Opsware Agent does not start executing the command within this time, the command will time out and the Command Engine script will continue.
30	Responding to a command (for example, after a reboot, the maximum time to wait until a server responds) or sending a message to the Command Engine from the Opsware Agent. If the Opsware Agent does not respond to the Command Engine at least once during this interval, the command will time out and the Command Engine script will continue. The Opsware System polls the Opsware Agent every 15 minutes and if the Opsware Agent fails to respond two consecutive times, the command will time out.

Customizing Monitor Time-Out Duration

If you would like to set a different monitor time-out duration, you can create a custom attribute named OPSW_reconcile_monitor_timeout and change the number of minutes before a time out occurs. For each type of hardware running in your operational environment, you can set a custom attribute with the time-out duration that you want. To set a time-out duration for a type of hardware, click Environment ➤ Hardware in the navigation panel. Then navigate to the type of hardware that you want to add a custom attribute for.

During reconcile, a periodic heartbeat occurs between the Opsware Agent and the Command Engine to ensure that the agent is still responsive. This setting controls the maximum amount of time that can pass between these heartbeat messages. Typically,

you only need to increase this setting if you install software that reboots the server, and the time that it takes for the server to reboot and for the agent to restart exceeds the default value.

See "Custom Attributes Set for the Environment Overview" on page 350 in Chapter 6 for more information.

Communication Between Managed Servers and the Opsware System

This section provides information about communication between managed servers and the Opsware system and contains the following topics:

- Managed Server and Opsware Communication Overview
- Viewing the Management IP Address for a Server
- Code Deployment and Static NAT
- Setting the Primary IP Address of a Server
- How Changing NAT Tables Affects Managed Servers

Managed Server and Opsware Communication Overview

To manage a server, the Opsware System requires the server to have a unique IP address that is routable from the Opsware System. However, in a large operational environment, all servers might not have unique IP addresses. In this case, the Opsware System supports static network address translation (NAT) for managed servers.

Static NAT maps public IP addresses to hosts inside the internal network, which allows the Opsware System to manage all servers in the environment.

Unlike dynamic NAT, the mapping between the Opsware System and the servers under management is set ahead of time, not dynamically at runtime.

To understand how the Opsware System communicates with managed servers, you must understand these three terms:

- **Management IP** – The IP address that the Opsware System uses to communicate with the Opsware Agent on the server.

  During hardware registration for a server, the Opsware Agent opens a TCP/IP connection to the Opsware System. The connection contains the source IP (called peer IP) address of the server. By default, the Opsware System uses this peer IP address as the management IP for the server.

- **Management Interface** – When a server has more than one network interface, you can designate one of them as the *management interface*.

- **Primary IP** – The IP address of the management interface. When you change the management interface, the primary IP changes to the IP address of that interface. The primary IP address is a locally-configured IP address.

During synchronizations, the Code Deployment and Rollback Subsystem uses the primary IP address to communicate with the server.

See "Code Deployment and Static NAT" on page 114 in this chapter for more information.

The Opsware Agents on servers communicate with each other by using the primary IP addresses, even though the Opsware System uses management IP addresses to communicate with the servers.

The Opsware System does not support managed servers that have IPv6 addresses.

When static NAT is being used, the management IP address for a server will *not* be the same as the primary IP address. When static NAT is being used, the management IP is the NAT-translated IP address for the server. When static NAT is *not* being used, the management IP address is always the same as the primary IP address.

Viewing the Management IP Address for a Server

In the Opsware Command Center, you can find the management IP address of a server and check whether it is using static NAT. You might need this information for troubleshooting any servers marked Not Reachable and to determine whether your NAT configuration is correct. The Opsware Command Center displays the management IP address of a managed server in the following two places:

- The Network tab of the Server Details page
- The Hardware view of the Managed Servers list

The Network tab shows (and allows the user to set) the management interface for the server by selecting it from a drop-down list, as Figure 2-41 shows.

Figure 2-41: Management IP Address Information in the Network Tab

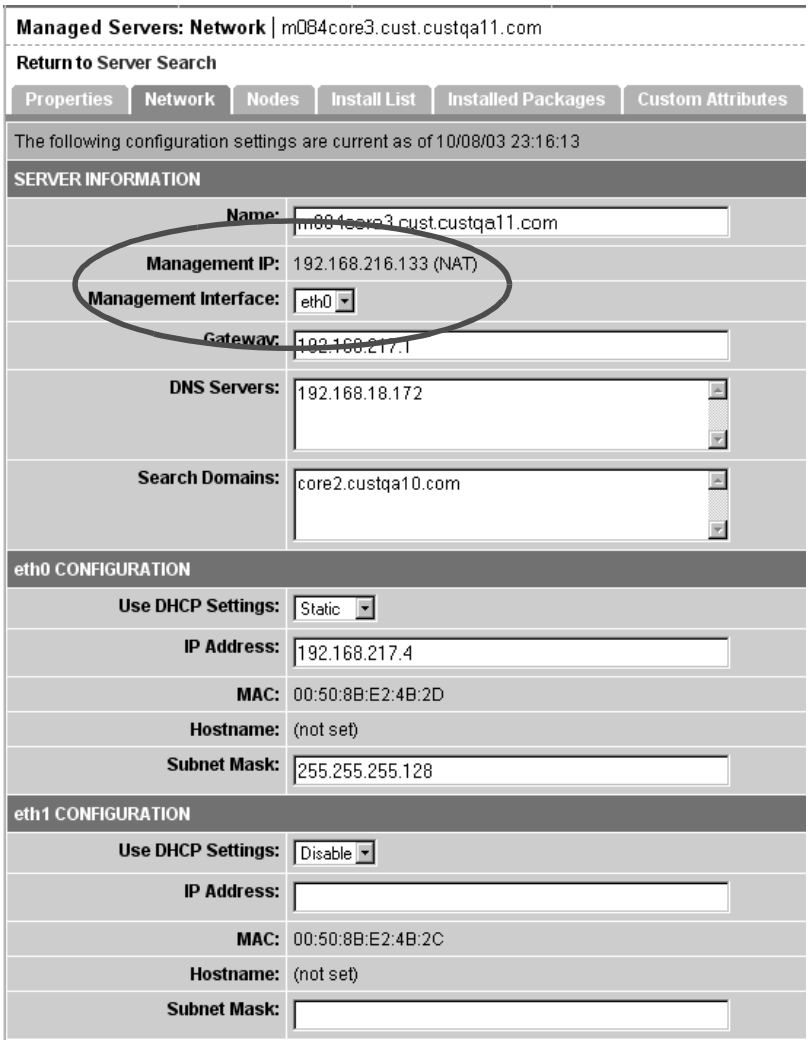

Figure 2-42 shows the Hardware view in the Managed Servers list, which displays the management interface for the server in the Network Info column. (To access the Hardware view, choose Hardware from the View menu.)

Figure 2-42: Hardware Tab in the Managed Servers List

The Network Info column shows the IP addresses and interfaces configured for each server in the list. If a server is using static NAT, the management IP is the first entry in this list and (NAT) appears after the IP address. If it is not using static NAT, the management IP is the same as the management interface, so it is already shown.

Code Deployment and Static NAT

Code Deployment and Rollback (CDR) synchronizations can only occur between Opsware Agents in the same NAT domain. Synchronizations cannot be performed between Opsware Agents in different NAT domains.

The Opsware System uses the primary IP address (instead of the management IP address) during synchronization because it is assumed that static NAT is not occurring between the servers in the synchronization. During CDR synchronizations, two Opsware Agents must communicate directly. Users can override the IP address that the Opsware System determined and designate a specific network interface as the management interface.

See "Code Deployment & Rollback" on page 539 in Chapter 12 for information about how to use CDR.

Setting the Primary IP Address of a Server

When a server has more than one network interface, users can specify one of them as the management interface and the IP address for this interface is designated the primary IP address. The primary IP address is used for Opsware Agent-to-Opsware Agent communication.

If static NAT is *not* being used, the management and primary IP addresses are the same. If static NAT is being used, the management IP is unaffected when a user changes the management interface.

Perform the following steps to set the primary IP address of a server:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server whose management IP address you want to view.

   Or

   Search for the server whose management IP address you want to view.

2. Click the server name. The Managed Servers: Server Properties page appears.

3. Click the Network tab. The network information for the server appears.

   The Network tab shows (and allows you to set) the server's management interface.

4. Set the management interface by selecting it from the Management Interface field. The IP address for this interface is designated the primary IP address.

5. Click the Update button.

See "Searching with Advanced Search" on page 48 in this chapter for more information.
See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

How Changing NAT Tables Affects Managed Servers

Static, one-to-one NAT tables map routable IP addresses between the Opsware System and managed servers. Network administrators configure and maintain these NAT tables. After the static NAT tables are configured, you do not have to perform any additional setup for the Opsware System.

The Opsware System does not control these NAT tables and errors can occur if they are modified after Opsware-managed servers register their hardware information. The following errors can occur if the IP address mapping of a server changes:

- If the IP address on the Opsware side of the NAT mapping is modified, the server becomes unmanageable and might be marked Not Reachable on the Managed Servers: Status page. It stays in this state until the Opsware Agent requests another hardware registration and the server's management IP is updated.

- If an IP address mapped to a particular server is mapped to a different server, both servers become unmanageable and might be marked Not Reachable on the Managed Servers: Status page. This problem is resolved when one of the two servers reports its IP address during hardware registration. The other server remains unmanageable until the server registers with the Opsware Agent. Both servers eventually become manageable.

Viewing Hardware Information for Managed Servers

The hardware link of the Opsware Command Center provides a read-only view of all servers in your managed environment categorized by hardware manufacturer and model. The hardware link provides hardware related information for each server, such as:

- Manufacturer
- Model number
- MAC ID
- Serial number
- CPUs used on the server
- Memory
- Storage capacity

For more information, see See "Server Asset Tracking" on page 39 in this chapter for more information.

Viewing Managed Server Hardware Information

Viewing hardware information allows you to see all the servers in your managed environment by hardware vendor, and view such infromation as MAC ID, CPUs, memory, and so on.

To view hardware of managed servers:

From the navigation panel, click Environment ➤ Hardware. You see the top level of the Hardware node in the managed environement, as shown in Figure 2-43.

Figure 2-43: Top level hardware node

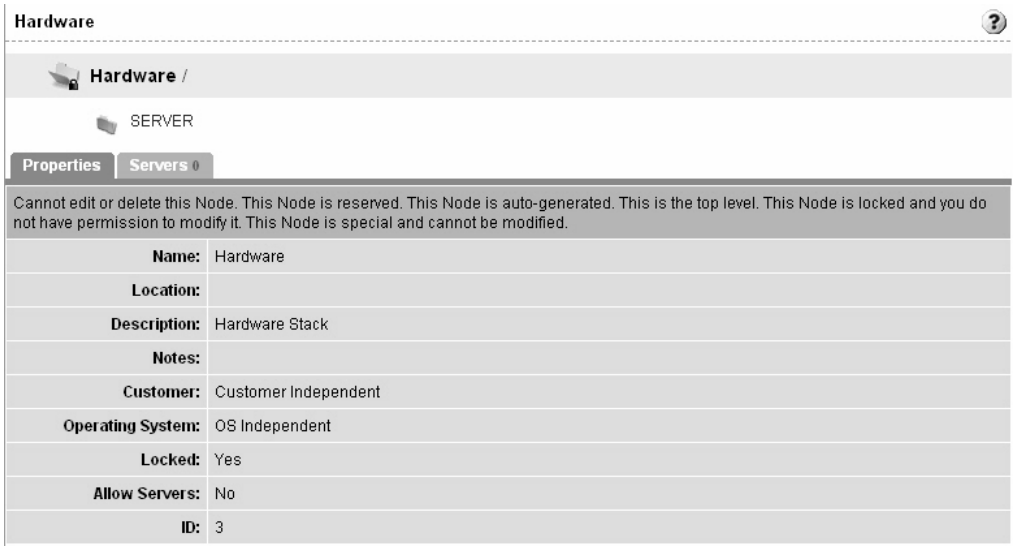

6 To view the servers in your managed environment, click the Servers link.

7 Drill down to the type of server you want to look at. For example, you may want to look at all Dell POWEREDGE 650s, as shown in Figure 2-44.

Figure 2-44: Hardware home page for Dell POWEREDGE 650s

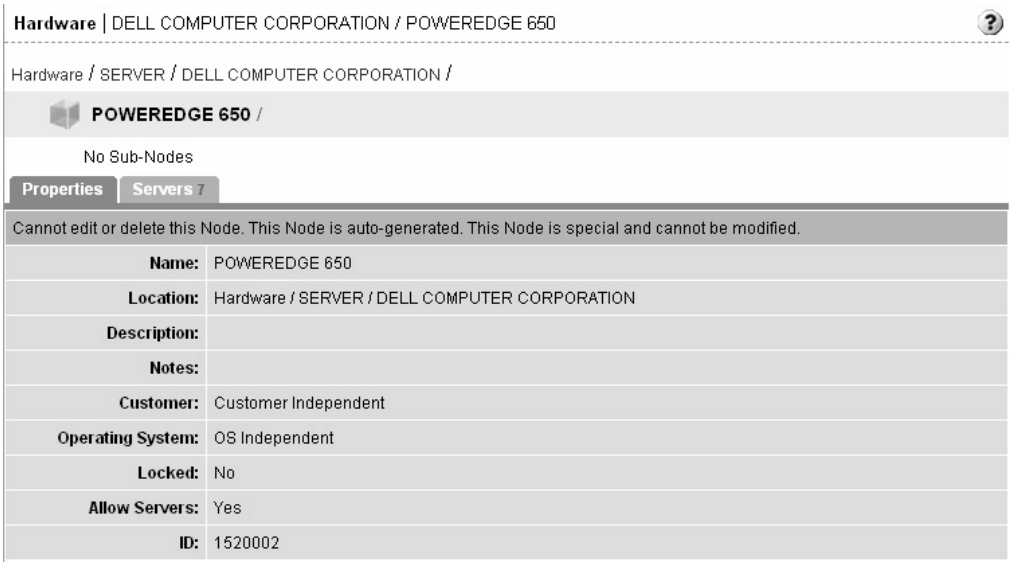

8 Next, click the Servers tab. You will see a list of all the Dell POWEREDGE 650s in your managed envrionment.

9 To view specific hardware information, from the View menu, choose Hardware. You now see more detailed information about all the Dell POWEREDGE 650s computers being managed by the Opsware System, as shown in Figure 2-45.

Figure 2-45: Detailed hardware information for Dell POWEREDGE 650s

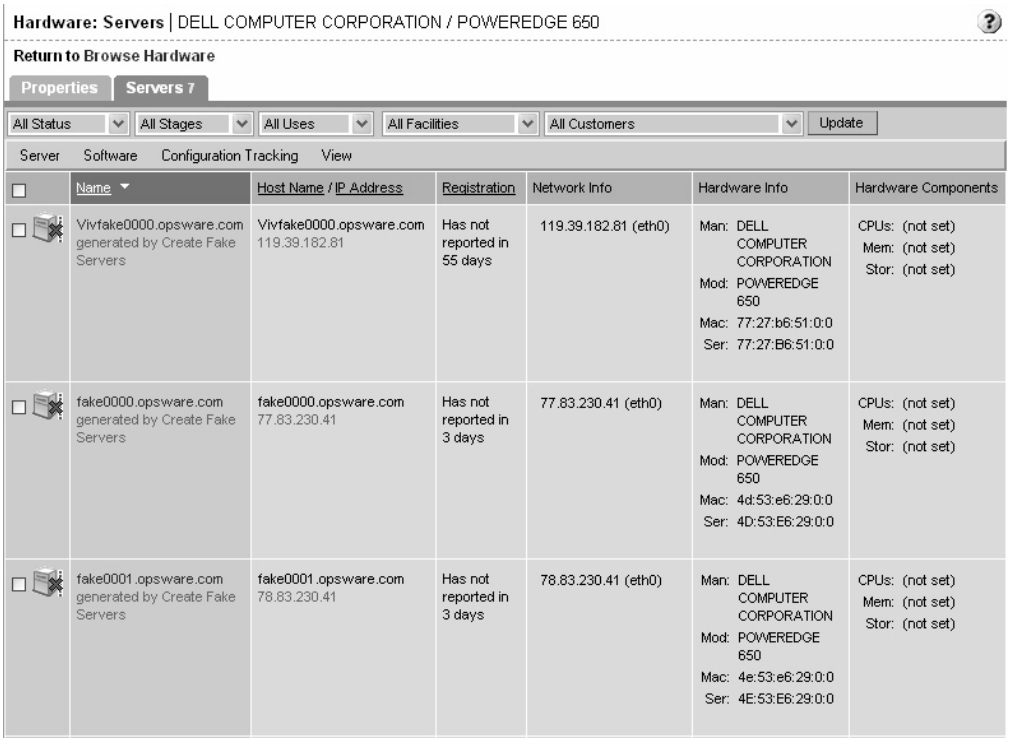

IP Range Groups and IP Ranges

This section provides information on IP range groups and IP ranges within the Opsware System and contains the following topics:

- IP Range Groups and IP Ranges Overview
- Creating an IP Range Group
- Creating an IP Range
- Changing Address Ranges on IP Ranges
- Increasing and Decreasing the Prefix Length

- Changing the Status of an IP Address in an IP Range

IP Range Groups and IP Ranges Overview

An Opsware user or an Opsware administrator can set up IP range groups and IP ranges so that servers are automatically associated with customers when users perform the following server management tasks:

- Assimilate servers running the operational environment by installing an Opsware Agent on the servers

  See "Server Assimilation" on page 139 in this chapter for more information.

- Use the OS Provisioning Subsystem to install operating systems on bare-metal servers

If you do not assign an IP range group to a customer, by default, a server is not assigned to a customer (Not Assigned appears in the Customer column of the server list) when you assimilate it.

An IP Range Group is a group of IP ranges that belong to a customer. It is both a physical and logical list – an accounting way to group IP ranges and assign them to a specific customer.

In the Opsware Command Center, an IP range identifies a range of IP addresses (in the OSI model – layer 3 IP address ranges). Each IP range can contain many IP addresses. The range of IP addresses is dependent on the subnet specified.

There is no direct association of an IP range with a specific customer; an IP range inherits its association to a customer from the IP Range Group it is created in.

See "Associated Servers with Customers" on page 59 in this chapter for more information.

Several types of IP Ranges are available in the Opsware Command Center, as Figure 2-46 shows.

Figure 2-46: Types of IP Ranges in the Opsware Command Center

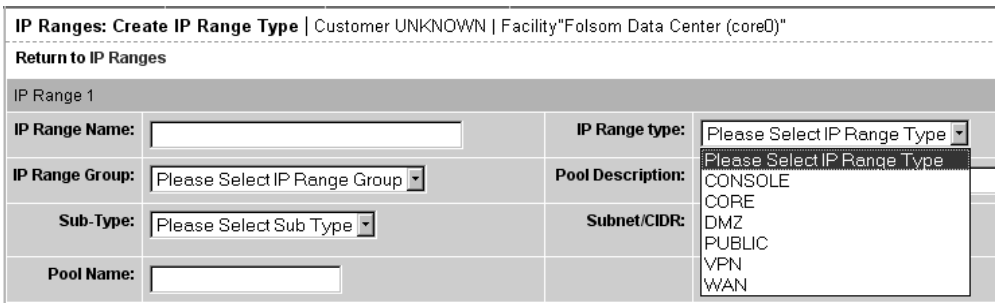

Creating an IP Range Group

You perform this task to create a group of IP ranges for a specific customer. After you create the group, you can designate the IP ranges that you want in that group.

Perform the following steps to create an IP Range Group:

1. From the navigation panel, click Environment ➤ IP Range Groups. The IP Range Groups page appears, as Figure 2-47 shows.

Figure 2-47: IP Range Groups Page in the Opsware Command Center

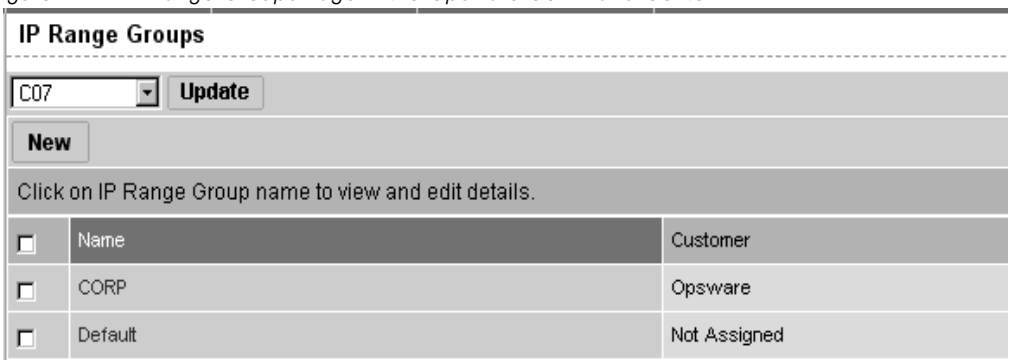

2. From the list, select the facility in which you want to create the IP range group and click the Update button. The list of IP range groups for that facility appears.
3. Click the New button at the top of the page. The IP Range Groups: Create IP Range Group page appears.
4. Enter a name for the new IP range group.
5. Select the customer from the drop-down list.
6. Click the Save button.

Creating an IP Range

Perform the following steps to create an IP Range:

1 From the navigation panel, click Environment ➤ IP Ranges. The IP Ranges: View IP Ranges page appears, as Figure 2-48 shows.

Figure 2-48: IP Ranges for the Default Customer ("Not Assigned")

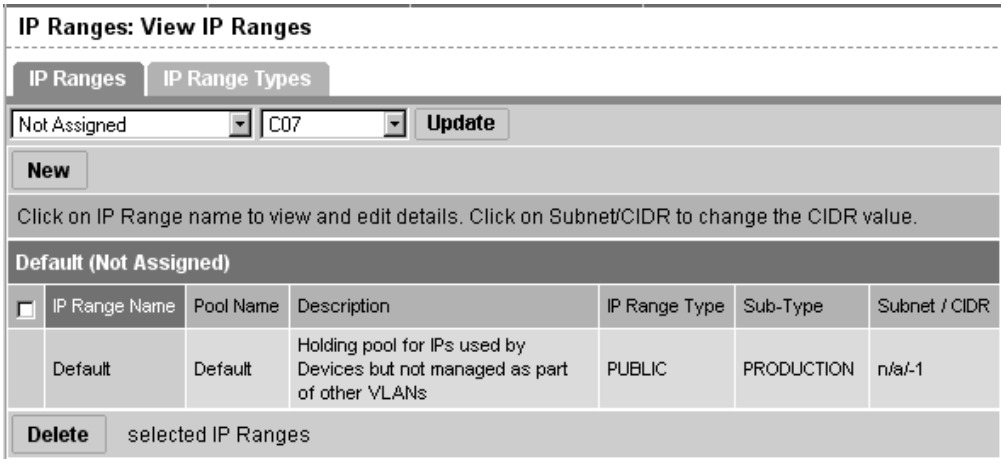

2 From the list, select the customer and facility in which you want to create the IP range and click the Update button. The list of IP ranges for that customer and facility appears.

3 Click the New button at the top of the page. The IP Ranges: Create IP Range Type page appears, as Figure 2-49 shows. You can add up to five new IP ranges at a time.

Figure 2-49: Creating an IP Range

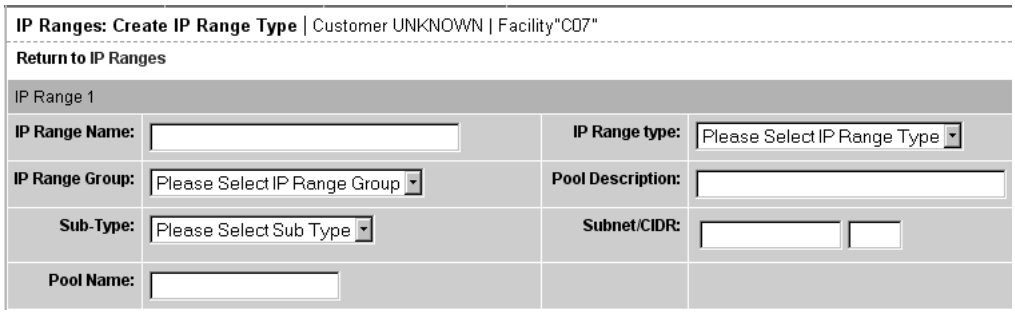

4 Define the following properties for each IP range:

- IP Range Name (for example, VLAN999 or SERVER100)
- IP Range Group – a customer might have several IP range groups and you must select one for the IP range that you are creating
- Sub-Type (for example, Development, Production, Staging, and so forth)

- Pool Name (for example, SAMPLE CUSTOMER SERVER pool)
- IP Range Type (for example, SERVER, PUBLIC, CONSOLE, TRANSIT, CORE, and so forth)
- Pool Description – provides detailed information about the IP range
- Subnet (for example, 10.2.0.0)
- Mask or Prefix Length – enter the prefix length or netmask (for example, 24, for /24, a netmask 255.255.255.0) in the CIDR field

You must complete all fields for each new IP range, which assumes that you have specific knowledge of your network's configuration and know the correct entries to include.

5 After you complete all entries, click the Save button at the bottom of the page.

Changing Address Ranges on IP Ranges

Classless Inter-Domain Routing (CIDR) in the Opsware System provides a way of specifying a range of IP addresses to include in an IP range.

The Opsware System might take several minutes to display an IP range with many IPs. For example, an IP Range with CIDR 19 (which has 8,192 IP addresses) might take 5 minutes to display in the Opsware Command Center.

Perform the following steps to change address ranges on IP ranges:

1 From the navigation panel, click Environment ➤ IP Ranges. The IP Ranges: View IP Ranges page appears.

2 From the list, select the customer and facility whose IP range you want to update and click the Update button. The list of IP ranges for that customer and facility appears.

3 Click the SUBNET/CIDR link at the end of the row for the IP range that you want to change, as Figure 2-50 shows.

Figure 2-50: IP Ranges in the Opsware Command Center

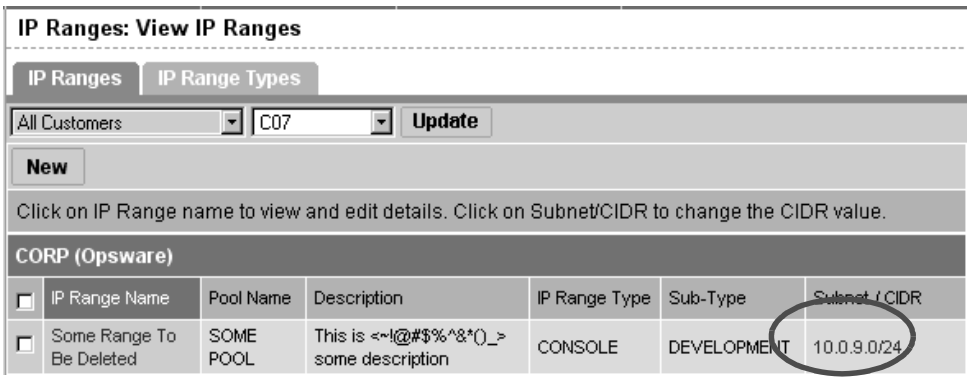

The last two digits in the Subnet/CIDR column make up the current prefix length for a particular IP range. The IP Range: Change CIDR page appears.

4 To change the current CIDR setting, select a new value from the list in the New CIDR column.

5 Click the Change button.

Changing the prefix length in the Opsware Command Center does not automatically change the net masks of servers for the servers themselves.

Increasing and Decreasing the Prefix Length

You can use the Opsware Command Center to increase or decrease the length of an IP range.

Increasing the Prefix Length

Increasing the prefix length reduces the IP range size. For example, if you have an IP range with prefix length 24 and it has 256 IP addresses in it, changing the prefix length to 25 results in the creation of two IP ranges with prefix length 25, each containing 128 IP addresses.

Example:

```
Network A – 10.1.0.0/24
```

Becomes:

```
Network A – 10.1.0.0/25
```



```
Network B: 10.1.0.128/25 (new network)
```

Decreasing the Prefix Length

Decreasing the prefix length expands the IP range size. Take the two CIDR 25 IP ranges from above. On the first IP range, changing the prefix length to 24 results in one IP range that contains twice as many IP addresses as before. The two original CIDR 25 IP ranges are combined to make one larger CIDR 24 IP range.

Example:

```
Network A – 10.1.0.0/24
```

```
Network B – 10.1.1.0/24
```

Becomes:

```
Network AB: 10.1.0.0/23 (one network)
```

This change only works if the two CIDR 25 IP ranges occupy contiguous blocks in the same IP range group. If they do not occupy contiguous blocks, you get an error message.

Changing the Status of an IP Address in an IP Range

You can use the IP Range feature to change the status of that IP address in the Opsware Command Center. For example, you might want to reserve an available IP address because you will assign it to a specific server in the next few days.

The status of an IP address automatically changes from available to assigned when a server with that IP address registers its hardware with the Opsware System.

Perform the following steps to change the status of an IP address in an IP range:

1. From the navigation panel, click Environment ➤ IP Ranges. The IP Ranges: View IP Ranges page appears.
2. From the list, select the customer and facility for which you want to assign IP addresses and click the Update button. The list of IP ranges for that customer and facility appears.
3. Click the name for the IP range in which you want to assign IP addresses. The IP Range: View IP Range page appears. By default, the View tab displays.

   The bottom of the page contains the IP addresses within that range. For each assigned or reserved IP address in the range, you can see its status.

4 Click an individual IP address link.

From the page that appears, you can change the status of the IP address (to ASSIGNED, AVAILABLE, RESERVED, and so forth). See Figure 2-51.

Figure 2-51: Editing the Properties of an IP Address in an IP Range

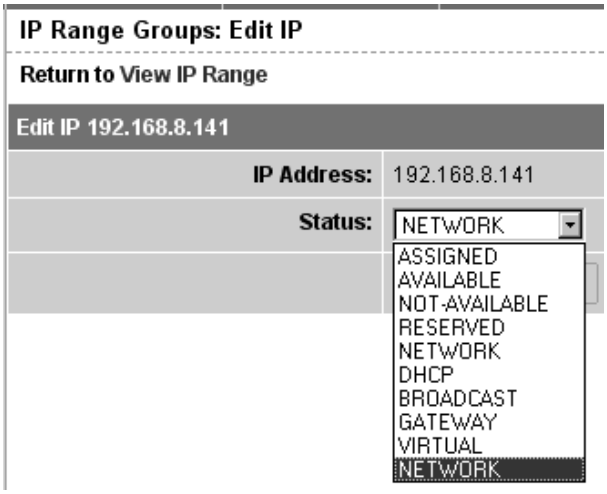

5 Select the status from the list. IP addresses can have one of the following statuses:

- ASSIGNED: A server is registered with this IP address.
- AVAILABLE: Available IP address.
- NOT AVAILABLE: Used to *reserve* an IP address for future use. For example, you might want to build a new server but need an IP address for the server prior to it being plugged into the network. Setting the status of an IP address to NOT AVAILABLE reserves it, so that another user does not take that IP address before the server is racked, stacked, and plugged into the network.
- RESERVED: The first couple of IP addresses after the first IP address is reserved
- NETWORK: Always assigned to the first IP address in a subnet
- DHCP: IP addresses reserved for use by a DHCP server
- BROADCAST: A special IP address reserved for sending a message to all stations
- GATEWAY: An IP address that acts as an entrance to another network
- VIRTUAL: Indicates a virtual IP address, such as

  www.samplecustomer.com

which is an IP address associated with a load balancer. The IP address does not correspond to any server, but yet the active load balancer responds to this request and forwards it to the appropriate Web server.

6 Click the Save button.

Network Configuration

This section provides information about network configuration within the Opsware System and contains the following topics:

- Network Configuration Overview
- Configuring Networking for an Opsware Managed Server

Network Configuration Overview

You can use the Opsware System to automatically configure network settings for a server after you install the OS.

The OS Provisioning Subsystem provisions servers with an OS by using DHCP addresses. Because DHCP servers often assign temporary IP addresses to servers that boot over a network, system administrators typically need to assign static IP addresses (and other network properties) before the servers can be put into service. The Opsware System enables system administrators to do this through the Opsware Command Center rather than logging onto the server manually after OS provisioning is complete.

The Opsware System does not support managed servers that have IPv6 addresses.

The Server Network Configuration functionality allows you to configure the following settings on a server that are related to its network configuration:

- Host name
- Domain Name System (DNS) servers
- Management interface (the interface that the Opsware System should use when managing the server)

See "Viewing the Management IP Address for a Server" on page 112 in Chapter 2 for more information.

- Gateway
- DNS search domains
- WINS (Windows Internet Naming Service) Servers
- Configuration for each network interface, including whether the interface is configured statically or with a Dynamic Host Configuration Protocol (DHCP) IP address, hostname, and subnet mask

You can alter any of these options and then apply the settings to the managed server. The Opsware System updates the server and reboots it to cause the new settings to take effect.

Configuring Networking for an Opsware Managed Server

You can only use the Network Configuration feature for servers running Sun Solaris, Red Hat Linux, and Microsoft Windows operating systems.

Perform the following steps to configure networking for an Opsware Managed Server:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server that you want to configure networking on.

   Or

   Search for the server that you want to configure networking on.

   See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

2. Click the name of the server that you want to configure networking for. The Managed Servers: Properties page appears for that server.
3. Click the Network tab. The network information for the server displays.
4. Modify any of the following settings to configure the server networking.

   For all fields, the default value is the one currently configured on the server.

- **Hostname**: The hostname configured on the managed server. This field only sets the name by which the server knows itself and does not update DNS records for the server.
- **Management Interface**: Instructs the Opsware System to use a particular network interface when contacting the server. This is useful, for example, when a server has multiple network interfaces, but not all of them are reachable by the Opsware core. Designating a particular interface as the management interface allows the Opsware System to know which interface to use for managing the server.
- **Gateway**: The IP address of the default router
- **DNS Servers**: A list of DNS nameserver IP addresses
- **Search Domains**: A list of DNS domains to search when attempting to resolve hostnames
- **WINS Servers**: Set for Windows only; a list of WINS server IP addresses
- **Interface Configuration** (for each network interface in the system):
  - **DHCP**: If DHCP is enabled for an interface, the system uses DHCP to configure this network interface. In this case, static configuration settings (IP address, hostname, and subnet mask) are not relevant for this interface, and the Opsware Command Center makes those fields not editable. If DHCP is not enabled, then static settings are required.
  - **IP Address**: The IP address for this interface (unless DHCP is enabled).
  - **Hostname**: The local hostname for the server. This item is only required for servers running Solaris. Like the Computer Name field, this setting only affects the name by which the managed server knows itself, and does not update DNS records.
  - **Subnet Mask**: The IP network mask to use for this interface

  In addition, the Opsware System displays the management IP address and MAC address for the server; however, you cannot change these values reported by the Opsware Agent.

5 Click the Update Server button at the bottom of the page.

(If you click the Revert button, it causes any changes that you made to the fields to be discarded.)

A confirmation dialog box appears that shows the changes that will be made to the server. The confirmation dialog box includes a check box that allows you to indicate that the server should revert to its old network configuration when it cannot contact the Opsware core after you save the new network configuration. By default, the Revert check box is selected.

The Opsware Command Center does not validate the network configuration changes that you make in the Network tab. Therefore, it is possible to provide a malformed IP address in the IP address field for an interface.

6 To have the server revert to its previous network configuration if an error occurs, ensure that the check box is selected in the confirmation dialog box.

7 Click OK to proceed with the configuration changes.

A progress dialog box appears that shows the progress of the operation. The process of setting a new network configuration involves rebooting the managed server. The operation might take several minutes.

You can wait for the operation to complete or close the progress dialog box and perform other work in the Opsware Command Center. The status of the task is available in the My Jobs user interface if you want to check the status of the network configuration update.

Details About Changing the Domain for Windows Servers

You cannot use the DNS Domain field to change the domain name for a Windows server.

The Opsware System does not change the domain name of a Windows server because changing the domain name of a server requires password authentication. Changing the domain name of a Windows server is a manual operation. See Figure 2-52.

Figure 2-52: DNS Domain Field Displays in the Network Tab for a Server

Opsware Agent on Managed Servers

This section provides information about the Opsware Agent on managed servers and contains the following topics:

- Opsware Agent on Managed Servers Overview
- Security for Opsware Agents Running on Managed Servers
- What an Opsware Agent Can Do on a Managed Server
- Server Data That the Opsware Agent Tracks

Opsware Agent on Managed Servers Overview

The Opsware Agent regularly performs the following management tasks on each server autonomously:

- On a regular interval, the Opsware Agent gathers a hardware and software inventory of each managed server. It opens a secure communication channel to the Opsware core, presenting its IP address and public-key certificate for authentication purposes. If properly authenticated, the Opsware Agent is permitted to write its updates about the server to the Opsware Model Repository.

Every 12 hours, the Opsware Agent submits hardware information for the managed server on which it is running. Hardware registration also occurs during Opsware Agent installation (server assimilation) or software installation.

Every 24 hours, the Opsware Agent submits software information for the managed server to the Opsware core.

The Reporting field indicates the status of the Opsware Agent's reporting capability and tells you whether or not the Opsware Agent is reporting regularly and successfully. The four possible reporting states for the Opsware Agent are as follows:

- **OK**: Opsware Agent is reporting properly.
- **Registration in progress**: Opsware Agent is currently registering server hardware information.
- **Reporting error**: Opsware Agent encountered error while trying to report hardware information.
- **Last reported days ago:** Indicates when the Opsware Agent last reported.

You can access Opsware Agent reporting information in Server Properties by using advanced search, and by viewing managed servers by Communication status. When viewing by Communication status, the Opsware Command Center user interface displays this information in the registration column.

If the Opsware Agent experiences an error in reporting, or has not reported within 24 hours, you can run a Communication Test to troubleshoot the problem. See "Agent Reachability Communication Test" on page 80 in this chapter for more information.

If you modify the server hardware, it could take up to 12 hours for the change to appear in the Opsware Command Center user interface, depending on the time that the Opsware Agent for that server contacted the Opsware core.

If you install or uninstall software on a managed server outside of the Opsware System, it could take up to 24 hours for the change to appear in the Opsware Command Center user interface. For example, if you update the Microsoft Patch database, it could take up to 24 hours for all managed servers to display whether they need new patches based on the updated Microsoft Patch database.

In some cases, not all of a server's hardware information is reported. For example, if the Opsware Agent was installed with its default settings, not all hardware information is reported to the Opsware Command Center until an hour after the agent is installed. Or there might be a problem retrieving certain hardware information, such as a disk failure,

that could prevent some hardware information from being reported. In these cases, the server's property page lists unreported information as *not set*.

- If configuration tracking is enabled for a server, the Opsware Agent sweeps through the managed server on a regular interval to see if any of the configurations being tracked are changed. If a tracked configuration is changed, the Opsware Agent performs the action specified by the tracking policy; namely, writing the information to a log file, generating a backup, or sending an email message through SMTP to the email address specified in the tracking policy.

Security for Opsware Agents Running on Managed Servers

The Opsware Agents act as both clients and servers when they communicate with Opsware core components. All communication is encrypted, integrity-checked, and authenticated using X.509v3 client certificates using SSL/TLS.

A small number of core components can issue commands to the Opsware Agent over a well-defined TCP/IP port. The Opsware Agent can also call back to Opsware core components, each with its own well-defined port.

The Code Deployment and Rollback Subsystem uses agent-to-agent communication for performance reasons. In particular, CDR synchronizations (the process of copying changed files and directories from one server to another) happen when an Opsware Agent connects to another Opsware Agent and sends across the network files that have changed since the last time the two Opsware Agents connected.

To further safeguard the SSL/TLS-based communication channel, the two Opsware Agents participating in the code deployment also need to have a common shared secret provided by the Command Engine. Before one Opsware Agent can begin a file transfer to another Opsware Agent, the two agents must verify that they have a common shared secret provided on a per-session basis by the Command Engine. This safeguard prevents unauthorized users from copying files from one managed server to another.

What an Opsware Agent Can Do on a Managed Server

The Opsware Agent is designed such that:

- It can only discover information about its own managed server (and no others).
- It cannot make changes on a server unless explicitly instructed to do so by an Opsware core component.

The Opsware System runs with administrator privileges (root on Unix servers and Local System on Windows servers) because it performs tasks that require administrator privileges, such as installing patches and rebooting servers.

The Opsware core performs client authentication and, additionally, checks to see if the presenting certificate belongs to that particular server. The Opsware System does this by comparing the certificate to the server's IP address that the Opsware System generates when the Opsware Agent is initially described. If the certificate is not valid or the originating IP address does not match the IP address stored in the Opsware Model Repository, authentication fails and the Opsware Agent cannot continue communication with the Opsware System.

If an unauthorized user were able to log on to a managed server with administrator privileges and compromise a server's security, the user would have only limited access to the following information in the Opsware Model Repository:

- The server's own hardware inventory (already available to someone logged on with administrator privileges)
- The server's own software inventory (already available to someone logged on with administrator privileges)
- The set of assignments from itself to the nodes in the Software Tree
- The custom attributes contained in those nodes

Server Data That the Opsware Agent Tracks

For each managed server, the Opsware Agent reports server and hardware information, as Figure 2-53 and Figure 2-54 show.

Figure 2-53: Manager Servers: Properties Page-Locking Status and Server Information

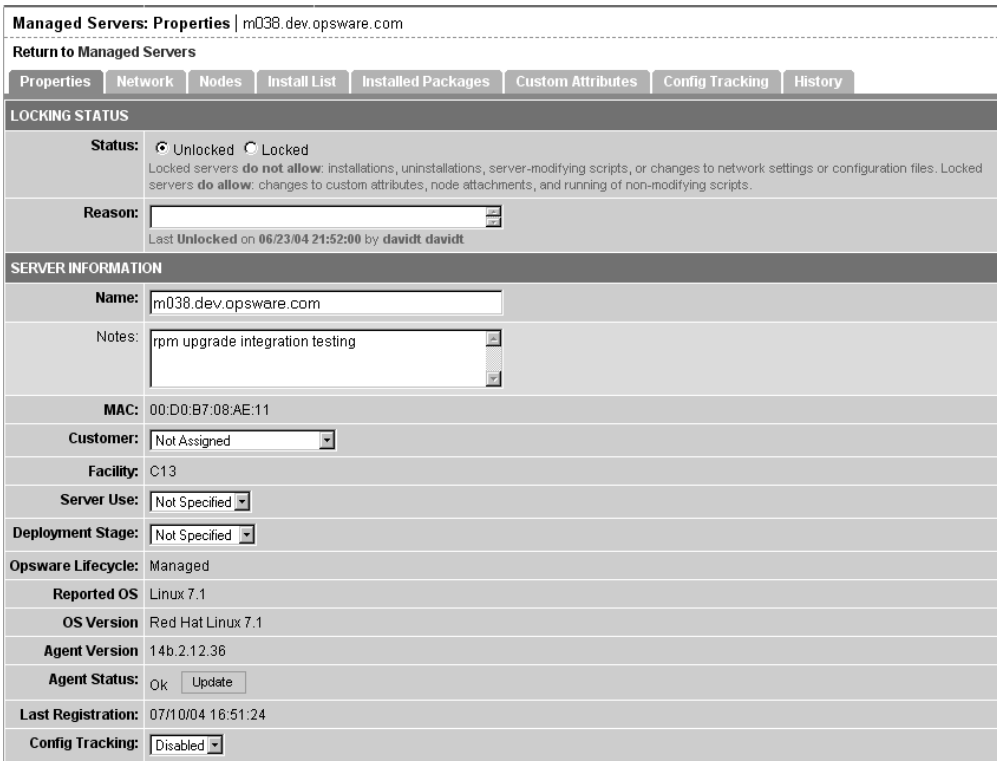

Figure 2-54: Server Properties-Hardware and Additional Information

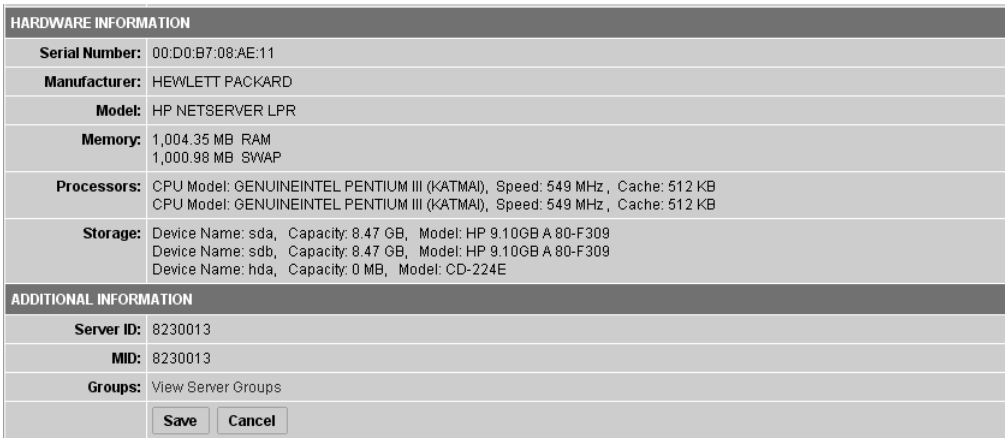

Additionally, the Opsware Agent reports networking information. See "Network Configuration" on page 127 in this chapter for information about descriptions of the networking information reported and how you can modify it by using the Network tab in the Opsware Command Center.

The Opsware Agent reports software information for managed servers. Click the Install List or Installed Packages tabs to view what software *should be installed* on the server or all software that *is installed* on the server.

- Click the Install List tab from the Managed Servers: Server Properties page to display the list of what software packages should be installed on that server by virtue of that server's assigned nodes.

  The Opsware System is able to determine what software *should be installed* on a server because of its model-based approach to server management. The software that should be installed is recorded in the Opsware Model Repository.

- Click the Installed Packages tab from the Managed Servers: Server Properties page to display the list of software that is reportedly installed on the server.

  The Opsware System is able to determine what software *is installed* on a server because the Opsware Agent communicates with the Opsware core and reports the installed hardware and software for the server.

  Partially-installed Solaris packages do not show up in the Installed Packages list, even though the package was partially installed.

See "Application Provisioning" on page 385 in Chapter 7 for information about how to set up the Software Tree. Table 2-18 shows how the Opsware Agent obtains the server and hardware information about each managed server. i

Table 2-18: Hardware Information That the Opsware Agent Reports for Servers

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
Name	The user-configurable name for the server. By default, the Opsware System uses the configured hostname of the server until a user edits it.	Windows – Uses the fully qualified DNS name of the server. Linux, Solaris, AIX, HP-UX – Uses the current hostname of the server that the `hostname` command returns.

Table 2-18: Hardware Information That the Opsware Agent Reports for Servers

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
Reported OS	The version number of the server's operating system	Windows – Uses the Windows version number as reported by the operating system. This information includes the major version number, the minor version number, the Windows build number, and the Service Pack level. Linux, Solaris, AIX, HP-UX – Uses the operating system version that the `uname` command returns.
OS Version	The OS version specified for the OS definition	Specified by the user who prepared the OS with the Prepare Operating System Wizard See "Defining an Operating System" on page 212 in Chapter 3 for more information.
Serial Number	The serial number of the system. The Opsware System attempts to report a chassis ID if possible.	Windows, Linux – Obtained from the system BIOS. Solaris, AIX, HP-UX – Obtained from the system ROM.
Manufacturer	The manufacturer of the server if available	Windows, Linux – Obtained from the system BIOS. Solaris, AIX, HP – Obtained from the system ROM.
Model	The model of the server if available	Windows, Linux – Obtained from the system BIOS. Solaris, AIX – Obtained from the system ROM. HP-UX – Output of model command (which is read from the system ROM).

Table 2-18: Hardware Information That the Opsware Agent Reports for Servers

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
Memory	The amount of physical RAM and the total amount of virtual memory paging space configured	Windows – Uses the Windows 2000 API `GlobalMemoryStatus()`. Linux – Obtained from information in the file `/proc/meminfo`. Solaris – Obtained from the `sysconf` and `swapctl` APIs. AIX – Uses the `lsattr` command for memory information and the `lsps` command for paging space. HP-UX – Uses the `pstat` system call.
Processors	Information about each of the processors in the system	Windows – If WMI is available, iterates over all instances of `Win32_Processor`. If WMI is not available, parses the registry key `HARDWARE\DESCRIPTION\System\CentralProcessor`. There is one sub-key for each processor. Linux – Obtained from information in the file `/proc/meminfo`. Solaris, HP-UX – Uses system APIs to enumerate the processors in the system. AIX – Uses the `lscfg` command.
Storage	Information about each installed disk drive or RAID array	All Platforms – Uses system APIs to discover and probe ed disk drives and RAID arrays.
Server ID	The internal ID the Opsware System uses to identify the sever	In most cases, the server ID is the same as the MID.

Table 2-18: Hardware Information That the Opsware Agent Reports for Servers

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
MID	The MID (Machine ID) is a unique number that the Opsware System assigns when the server first registers. The server stores the MID and reports it each time the server registers.	Windows – The MID is stored in the file `%ProgramFiles%\Common Files\Loudcloud\cogbot\mid` if present. Linux, Solaris, AIX, HP-UX – The MID is stored in the file `/var/lc/cogbot/mid`.

Server Assimilation

This section provides information on server assimilation within the Opsware System and contains the following topics:

- Server Assimilation Overview
- Preparation for Server Assimilation
- Preassimilation Checklist
- Installing an Opsware Agent on a Server
- Opsware Agent Installer Options
- Examples of Opsware Agent Installer Command and Options
- Starting an Opsware Agent on a Server
- Verifying Opsware Agent Functionality
- Augmenting the Information for an Assimilated Server
- Uninstalling an Opsware Agent (Unix and Windows)
- Uninstalling Earlier Versions of Opsware Agents on Unix
- Uninstalling Earlier Versions of Opsware Agents on Windows

Server Assimilation Overview

When you assimilate existing operational servers into the Opsware System, you should synchronize the local time on the servers with an external time-server that uses a network time protocol (NTP).

Assimilating servers makes existing operational servers known to the Opsware System so that they can be managed. Assimilating servers into the Opsware System is appropriate when many servers are already functioning in the operational environment and need to be managed (for example, when Opsware technology is initially deployed in a facility).

Assimilating a server with a pre-built OS into the Opsware System enables:

- Baseline discovery of the operating system on the server
- Managing the baseline operating system, including patch management, when the operating system is defined in the Opsware System with the Prepare Operating System Wizard
- Full provisioning and management capabilities for any new applications deployed on the server

The Opsware Agent assimilates a server by registering it with the Opsware Model Repository. The Opsware System assigns the server to a generic operating system that corresponds to the operating system that the Opsware Agent discovered during

assimilation. The server is assigned to a placeholder OS node. For each operating system, the Opsware Command Center contains a node <*operating_system_version*>/ Not Assigned, as Figure 2-55 shows.

Figure 2-55: Nodes Tab for an Assimilated Server

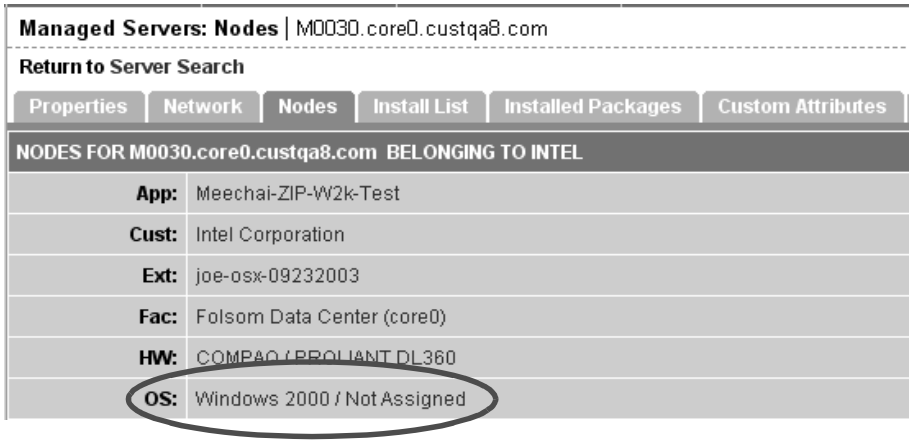

The Opsware Agent Installer can install Opsware Agents when the Opsware System core is not available to a server. If a newly-installed Opsware Agent cannot contact an Opsware System core, the Opsware Agent runs in a dormant mode. While dormant, it periodically attempts to contact the Opsware System core.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

The server is tracked in the Opsware Command Center. However, the server operating system *cannot* be managed while the server is assigned to the generic operating system node. You must reassign the server to the operating system that was defined with the OS Provisioning Subsystem. (From the Managed Servers list, choose Servers ➤ Re-Assign Node.)

The server is associated with the default facility for the local instance of the Opsware System.

If the assimilated server's IP address does not fall within a specified IP range, the server is associated with the default IP range group (Default). The default group is associated with the customer Not Assigned.

See "Overview of Server Association with Customer Accounts" and "How Servers Are Associated with Customers" for more information.

Users assimilate servers by installing an Opsware Agent on each server. Running an Opsware Agent on a server allows the Opsware System to manage the server. To install an Opsware Agent, run the Opsware Agent Installer.

The Opsware Agent Installer is an application that has the following features:

- Invokable from the command line or within a script
- Installs an Opsware Agent
- Logs its decisions and actions
- Can be operated unattended because user interaction is not required

The Opsware Agent Installer installs the Opsware Agent, retrieves cryptographic material, retrieves configuration information, and writes a configuration file.

Preparation for Server Assimilation

Opsware Inc. recommends that you set up a Windows file share to make the Opsware Agent Installer for various operating systems available from one place.

Setting up a file share allows you to install Opsware Agents on servers quickly and easily. If this is not possible, the Opsware Agent Installer needs to be moved by using an alternate file transfer mechanism, such as SFTP.

At the completion of the Opsware Agent installation process, a managed server is assimilated and the hardware and software data that the Opsware Agent discovered is stored in the Model Repository.

To use the Patch Management features on Windows NT 4.0 and Windows 2000 servers, you must install Internet Explorer (IE) 6.0 or later on the server first because a patch utility depends on it. If you do not install IE 6.0 on the server first, the Opsware Agent Installer warns you that the Patch Management feature does not work as expected because it checks for IE 6.0 on all Window servers. This prerequisite is not required for Windows 2003 because IE 6.0 is pre-installed for this operating system.

Preassimilation Checklist

Prior to installation, you should perform the following tasks on the server where the Opsware Agent is to be installed. Performing these tasks is vital to installing the Opsware Agent quickly within maintenance windows.

1. If DNS is not being utilized for servers or name resolution is not occurring through the use of the `/etc/hosts` file, you must add entries to the hosts file on the server so that it can resolve the Opsware System core service names. Perform this task for each server that you want to assimilate into the Opsware System.

   Set up entries in the hosts file to allow the server to resolve Opsware System core IP addresses:

   - `spin`
   - `way`
   - `theword`

2. Verify that the following ports are accessible from the server to the Opsware System core by entering the following commands from a terminal window:

   - `telnet spin 1004`
   - `telnet way 1018`
   - `telnet theword 1003`

3. For the Code Deployment and Rollback Subsystem, verify that the following port is accessible between the server you will push code from and the server where you will push code to:

   - `telnet <staging_server> 1002`
   - `telnet <production_server> 1002`

4. Because the Opsware Agent runs on port 1002, verify that no other applications are using this port.

   - On a Unix server, enter this command from a terminal window:

     ```
     netstat -an | grep 1002 | grep LISTEN
     ```

   - On a Windows server, enter this command from a terminal window:

     ```
     netstat -an | find "1002" | find "LISTEN"
     ```

5. Check for sufficient disk space for Opsware Agent installation on the server.

The Opsware Installer checks for the following amounts of free disk space in these directories:

- 30MB in `/opt/OPSW/installdir` (Unix)
- 100MB in `/var/lc/vardir` (Unix)
- 30MB in `%SystemDrive%\\Program Files\\Loudcloud\\installdir` (Windows)
- 100MB in `%SystemDrive%\\Program Files\\Common Files\\Loudcloud\\vardir` (Windows)

(These default directories can be overridden with parameters at installation time.)

These space requirements might not be enough. The `vardir` directory is used for dynamic content like logs and downloaded packages. If there is not enough disk space for the packages during a reconcile, reconcile will fail.

6. On the Solaris operating system, check for legacy sun4m architecture. Currently, the Opsware Agent works only for sun4u architecture.

7. For Windows, check the following items:

   - At a minimum, NT 4.0 Service Pack 6a must be installed on the server.
   - Verify that the Windows Registry has the correct settings:

   1. Start regedit and locate the following registry key:

      `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem`

   2. Select the `NtfsDisable8dot3NameCreation` entry.
   3. On the Edit menu, click DWORD and verify that the value is set to 0. The value *must* be set to 0. If necessary (because the value is set to 1), change the value by following your organization's IT policies and reboot the server.

8. To install an Opsware Agent on a server running Solaris, you must also install the following Solaris packages:

   For Python:

   ```
   SUNWtoo
   SUNWtoox
   ```

   For showrev:

   ```
   SUNCadm
   SUNWlibC
   ```



```
SUNWlibCx
SUNWadmfw
```

9 Before you install an agent on a server, the server must meet certain patch requirements that vary by operating system, as Table 2-19 shows.

Table 2-19: Required Patches for Opsware Agent Installation

SERVER OPERATING SYSTEM	REQUIRED PATCHES
AIX 4.3 AIX 5.1	APAR IY39444 APAR IY39429 NOTE: If AIX 4.3.3.388, 4.4.4.89, or 5.1.0.3 is installed, the Opsware Agent Installer displays an error message that indicates the correct APAR to install on the server.
HP-UX (10.20, 11.00, 11.11/11i)	For HP-UX 10.20, PHCO_21018 Additionally, SW-DIST should be upgraded to the HP recommended patch level. You should continue to upgrade this package when HP recommends new versions.
Linux AS 3.0	openssl096b NOTE: This package must be included in the `%packages` section of the Kickstart configuration.
Solaris 9, 8, 7, and 2.6	SUNWADMC SUNWcsl SUNWcsu SUNWesu SUNWlibms SUNWswmt
Windows 2000, NT 4.0	Service Pack 6a Internet Explorer 6.0 or later

Installing an Opsware Agent on a Server

The Opsware Agent needs administrator-level privileges (root on Unix servers and Local System on Windows servers) to manage a server. Therefore, Opsware Agent installation needs to be performed as root on Unix operating systems and as administrator on Windows operating systems.

You can install an Opsware Agent on any server that is running an Opsware-supported operating system.

See "Supported Operating Systems" on page 4 in Chapter 1 for information about the complete list of supported operating systems for managed servers.

Perform the following steps to install an Opsware Agent on a server:

1. Log on to the server that you want to assimilate by using a remote shell.
2. For Unix operating systems, change the user login to root (`su - root`) and for Windows operating systems, log in as administrator.
3. From the Opsware Command Center, download the package that contains the Opsware Agent Installer to a directory on the server you want to assimilate:
   1. Search for the package `opsware-agent`. From the navigation panel, enter `opsware-agent` in the Search box, select the Packages option, and click the Go button. The Manage Packages: Search Packages page appears.

      Each operating system and operating system version has different packages for the Opsware Agent Installer.

      Unix:

      ```
      opsware-agent-<version>-<system_name>-<system_version>
      ```

      Windows:

      ```
      opsware-agent-<version>-<system_name>-<system_
      version>.exe
      ```
   2. Click the package name for the Opsware Agent Installer that you want to download. The Packages: Edit Properties page appears.
   3. Click the Download button to save the package locally.
4. From the directory where the Opsware Agent Installer was copied, run the Installer by entering the correct executable and options for the installation environment.

See “Examples of Opsware Agent Installer Command and Options” on page 151 in this chapter for more information.

Opsware Agent Installer Options

When you use the Opsware Agent Installer CLI, you can include the options that Table 2-20 shows to control the way that the Opsware Agent is installed on a server.

Table 2-20: Opsware Agent Installer Command Line Options

OPTION	DESCRIPTION
`--clean` `(-c)`	Removes any machine-specific identifying material from the server. Specifically, removes the machine ID file (MID), and all machine-specific cryptographic material. Use this option when a server is deactivated and deleted from the Opsware Command Center and needs to be returned to service at a later time.
`-f`	Forces Opsware Agent installation and removes the target installation directory if it exists. REQUIREMENT: When you use the `-f` option, you must run the Opsware Agent Installer as root on Unix and as administrator on Windows.
`-h`	Displays help for the Opsware Agent Installer options.
`--installdir <path>`	Unix only. Agent installation directory. Default is: `/opt/OPSW`
`--logfile`	Specifies the path to the Opsware Agent Installer log file. By default, the current directory is set as the path. By default, the log file has the following filename: `opsware-agent-installer-<date>.log`

Table 2-20: Opsware Agent Installer Command Line Options

OPTION	DESCRIPTION
`--loglevel <level>`	Sets the log level for log messages. With this option, specify one of the following levels: `error`, `warn`, `info`, `trace`, or `none`. The level `error` logs the least amount of detail. The level `trace` logs all messages. By default, the log level is set to the log level `info`.
`-o`	Logs all output to `stdout` instead of a log file. This option is invoked automatically if the default log file or the log file passed with the `--logfile` option cannot be created, for example, when running the Opsware Agent Installer from non-writeable media, such as a CD-ROM.
`--reconcile <type>`	Reconciles the server against any nodes assigned to the server. The <*type*> can be `full` or `addonly`. `full` – All nodes in a category are selected and reconcile removes software that the Opsware System did not install. `addonly` – Software installed outside of the Opsware System is not removed. WARNING: When assimilating a server that is already functioning in the operational environment, use caution when you specify the option `--reconcile`. If you specify this option, you might inadvertently uninstall software from the server.
`--rpmbin <path>`	Specifies the path to the RPM binary to use for RPM operations. Use this option when RPM is already installed on the server to point the Opsware Agent at the RPM binary. Use the `--withrpm` option to install RPM if a usable instance of RPM is *not* already installed. **NOTE** It is unnecessary to use this option with the `--withrpm` option.

Table 2-20: Opsware Agent Installer Command Line Options

OPTION	DESCRIPTION
`-s`	Starts the Opsware Agent after installing it. By default, the Opsware Agent Installer does not start the Opsware Agent. NOTE; If you do *not* include the `-s` option on the command line when you install an Opsware Agent, you need to start the Opsware Agent on the server manually. See "Starting an Opsware Agent on a Server" on page 151 in this chapter for more information.
`--template <ID>`	Assigns the nodes contained in the template to the server. *<ID>* can be an ID or a full name of a template. If you specify this option with the `--reconcile` option, the Opsware System assigns the nodes in the template to the server before reconciling the server. WARNING: When assimilating a server that is already functioning in the operational environment, use caution when you specify the option `--template`. If you specify this option, you might inadvertently uninstall software from the server.
`--withmsi`	Installs MSI 2.0 along with the Opsware Agent. If MSI 2.0 is already installed, this option has no effect. Works with Windows NT 4.0 Service Pack 6a, Windows 2000, and Windows 2003.

Table 2-20: Opsware Agent Installer Command Line Options

OPTION	DESCRIPTION
`--withrpm`	Installs the RPM handler with the Opsware Agent. By default, an Opsware Agent is not installed with this option. Opsware Inc. recommends that you always include the `--withrpm` option when you install Opsware Agents on Solaris servers. NOTE: Use the `--withrpm` option only with the Opsware Agent Installers for these operating systems: Solaris 5.6, 5.7, 5.8, and 5.9, and AIX 4.3 and AIX 5.1. On Solaris, RPM 3.0.6 is installed in the directory `/opt/OPSWrpm` and the RPM database is installed in the directory `/var/opt/OPSWrpm/lib/rpm`. On AIX, RPM 3.0.5 is installed in the directory `/opt/freeware` and the RPM database is installed in the directory `/var/opt/freeware/lib/rpm`.
`--workdir <path>`	Specifies the path to the Opsware Agent Installer temporary working directory. Use this option if the default working directory causes problems with installation.
`--opsw_gw_addr_list`	If an agent is being installed on a server in a satellite, the Opsware Gateway `host:port` settings need to be specified during agent installation.
`--del_opsw_gw_addr_list`	If an agent is being upgraded on a server, and it has had Opsware Gateway `host:port` settings configured previously that are no longer needed, this option clears the settings. An example could be if a server is being moved from a satellite to the main data center.
`--force_full_hw_reg`	By default, the agent installer reports minimal server information during installation. This option tells the installer to report the full hardware information.
`--force_sw_reg`	By default, the installer does not report software information during install. This option tells the installer to report installed software information.

Table 2-20: Opsware Agent Installer Command Line Options

OPTION	DESCRIPTION
`--no_check_ reachability`	By default, the installer triggers the core to check whether the server is reachable. This option disables this check during install. This option is useful because checking reachability can be a resource intensive operation, and installing or upgrading many agents at the same time impacts the performance of the core.

Examples of Opsware Agent Installer Command and Options

Enter the following command and options to install the Opsware Agent for Solaris 5.7 in the default directories and log the results of the installation in the log file:

```
% opsware-agent-14.2.12.5-solaris-5.7 --logfile opsware-agent-installer.log --loglevel info
```

Enter the following command and options to install the Opsware Agent for Windows NT 4.0 in the default directories and log the results of the installation in the log file:

```
% opsware-agent-14.2.12.5-win32-4.0.exe --logfile opsware-agent-installer.log --loglevel info
```

Starting an Opsware Agent on a Server

If you do *not* include the `-s` option on the command line when you install an Opsware Agent, you have to start the Opsware Agent on the server manually.

For Solaris, enter the command:

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

For Linux, enter the command:

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

For AIX, enter the command:

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

For HP-UX, enter the command:

```
/sbin/init.d/cogbot start
```

For Windows, enter the command:

```
net start shadowbot
```

Verifying Opsware Agent Functionality

Perform the following steps to verify Opsware Agent functionality:

1. From the navigation panel in the Opsware Command Center, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server whose Opsware Agent installation you want to verify. If necessary, select the correct customer and facility for the server and click the Update button.

   Or

   Search for the server whose Opsware Agent installation you want to verify.

   Or

   If you want to discover reasons why a server is unreachable, you can run a Communication Test. See "Agent Reachability Communication Test" on page 80 in this chapter for more information.

2. Verify that the server appears in the Managed Servers list and has the correct properties.

   See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

Augmenting the Information for an Assimilated Server

Use caution when you augment the discovery process for a server that is assimilated and functioning in the operational environment. You might inadvertently

software from the server. During the test reconcile, verify what software will be uninstalled from the server before you perform the actual reconcile.

Perform the following steps to augment the information for an assimilated server:

1. Model the OS and other applications running on the server in the Opsware System by defining the OS with the Prepare Operating System Wizard and by creating nodes and templates for applications running on the assimilated server.

See “Operating System Definitions” on page 191 in Chapter 3 for more information.

See “Application Provisioning Setup” on page 307 in Chapter 6 for information about modeling an operating system and applications in the Opsware Command Center.

2. Move the server to the appropriate nodes for the OS and installed applications.

   The server is tracked in the Opsware Command Center; however, the server operating system *cannot* be managed while the server is assigned to the generic operating system node. You must reassign the server to the operating system that was defined with the OS Provisioning Subsystem. (From the Managed Servers list, choose Servers ➤ Re-Assign Node.)

3. Reconcile the server.

   See “Directly Reconciling Servers” on page 453 in Chapter 9 for more information.

4. If an IP range group was set up, servers are automatically associated with customers when users install an Opsware Agent on the servers. Otherwise, the servers are associated with the *Not Assigned* customer. To change the customer associated with a server, See “Editing the Properties of a Server” on page 75 in this chapter for more information.

5. To specify the server’s use, stage, and state, edit the server’s properties. See “Editing the Properties of a Server” on page 75 in this chapter for more information.

Discovery is complete. The Opsware System assumes that the server should always be running the specific OS build it has been associated with. Any changes to the OS outside of the Opsware System are not captured in the model.

Users can deploy and manage new applications on the server, just as if the Opsware System initially provisioned the server. Users can also deploy OS level patches on the server, or rebuild the OS by using the OS build with which the server was associated.

Uninstalling an Opsware Agent (Unix and Windows)

Perform the following steps to uninstall an Opsware Agent on Unix or Windows:

1. Log in to Unix as root user. Log in to Windows as Administrator.
2. Change directories to any directory other than the agent's installation directory.
3. On Unix, enter the following command:

   ```
   <installation_directory>/bin/agent_uninstall.sh
   ```

By default, for Solaris and AIX, the Opsware Agent uninstaller will not remove the Opsware RPM package. For command line options for the agent uninstaller, including how to activate removal of the Opsware RPM package, See "Opsware Agent Uninstaller Options" on page 154 in this chapter for more information.

4 On Windows, enter the following command:

```
msiexec /x <installation_directory>\bin\agent_uninstall.msi
```

5 As the uninstall proceeds, the Unix platform `stdout` shows the uninstallation progress. The Windows uninstall does not show uninstallation progress.

Opsware Agent Uninstaller Options

When you use the Opsware Agent Uninstaller, you can include the options that Table 2-21 and Table 2-22 show.

Table 2-21: Agent Uninstallation Unix Options

OPTION	DESCRIPTION
`--uninstallerVersion`	Show uninstaller version.
`--help`	Show this help.
`--no_deactivate`	Do not deactivate server; default is deactivated.
`--force`	Do not prompt for confirmation before deactivating.
`--delete_opsw_rpm`	Remove the OPSW RPM package (AIX, SunOS only). Use the following commands to remove the RPM package: SunOS: `pkgrm -n OPSWrpm` AIX: `installp -u rpm.rte`

Table 2-22: Agent Uninstallation Windows Options

OPTION	DESCRIPTION
`NO_DEACTIVATE="1"`	Do not deactivate server; default is deactivated.
`FORCE="1"`	Do not prompt for confirmation before deactivating.

Uninstalling Earlier Versions of Opsware Agents on Unix

Perform the following steps to uninstall Opsware Agents versions 4.7

and earlier:

1. Stop the Opsware Agent on the server by running the following command as root:

   (Linux) `% /etc/rc.d/init.d/cogbot stop`

   (Solaris) `% /etc/init.d/cogbot stop`

   (HP-UX) `% /sbin/init.d/cogbot stop`

   (AIX) `% /etc/rc.d/init.d/cogbot stop`

2. Deactivate or delete the server by using the Opsware Command Center Server menu.

3. For Linux servers only, run `chkconfig` to de-register the Opsware Agent initialization script:

   `% /sbin/chkconfig -del cogbot`

4. As root, delete the following files and directories to remove the Opsware Agent files from the server:

   (Linux) `/etc/rc.d/init.d/cogbot`

   (Solaris) `/etc/init.d/cogbot`

   (Solaris) `/etc/rc2.d/S79cogbot`

   (Solaris) `/etc/rc0.d/K44cogbot`

   (HP-UX) `/sbin/init.d/cogbot`

   (HP-UX) `/sbin/rc2.d/cogbot`

   (AIX) `/etc/rc2.d/init.d/cogbot`

   (AIX) `/etc/rc.d/S79cogbot`

   (All Unix) `/opt/OPSW`

   (All Unix) `/var/lc`

Uninstalling Earlier Versions of Opsware Agents on Windows

Perform the following steps to uninstall earlier versions of Opsware Agents on Windows:

1. Stop the Opsware Agent by running the following command as administrator:

   `C:\> net stop shadowbot`

2. Deactivate or delete the server by using the Opsware Command Center Server menu.

3. Deregister the Opsware Agent service by running the following command as administrator:

```
C:\> "%SystemDrive%\Program
Files\Loudcloud\blackshadow\watchdog\watchdog.exe" -x
```

4. As administrator, delete the following directories to remove the Opsware Agent:

```
"%SystemDrive%\Program Files\Loudcloud"
```

```
"%SystemDrive%\Program Files\Common Files\Loudcloud"
```

Custom Attributes for Servers

This section provides information on custom attributes for servers within the Opsware System and contains the following topics:

- Custom Attributes for Servers Overview
- Managing Custom Attributes
- Adding Server Custom Attributes
- Editing Server Custom Attributes
- Deleting Server Custom Attributes

Custom Attributes for Servers Overview

Users often need to store specific miscellaneous information in the Opsware Model Repository to facilitate server or application installation and configuration, scripting, or other purposes.

The Opsware Command Center provides a data management function by allowing users to set custom attributes for servers. These custom attributes include setting miscellaneous parameters and named data values. Users can write scripts that use these parameters and values when performing a variety of functions, including network and server configuration, notifications, and CRON script configuration.

Custom attributes can be accessed by software packages at installation time to configure settings that might be unique to the installation.

For information about how to set custom attributes required by the software running on a specific server, contact the group responsible for packaging your applications, as Figure 2-56 shows.

Figure 2-56: Custom Attributes Set for a Server in the Opsware Command Center

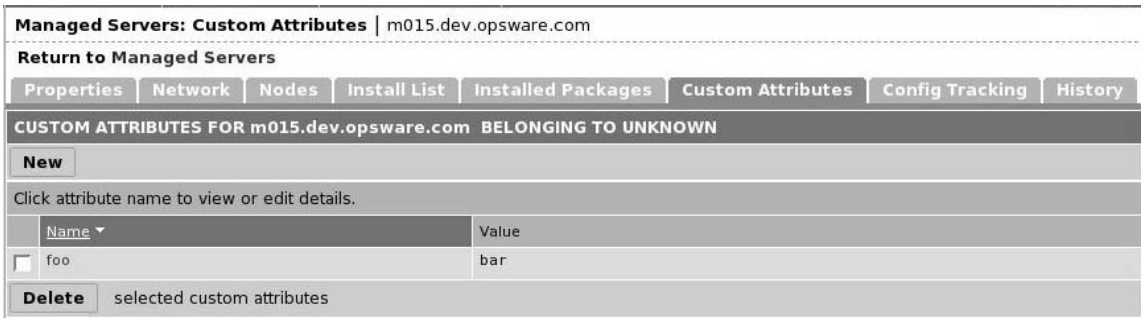

Managing Custom Attributes

Do not edit or remove custom attributes without verifying that the change you are making does not impact other users or critical Opsware operations.

To set custom attributes that affect a specific server, use the Managed Servers list. After locating the server and displaying the server properties, click the Custom Attributes tab. The Opsware Command Center displays the currently defined custom attributes for the selected server.

See “Adding Server Custom Attributes” on page 158 in this chapter for more information.

To set custom attributes for all the servers assigned to a node in the Software Tree, navigate to the node where you want to set attributes, click the Custom Attributes tab, and add attributes for all the servers assigned to the node. Inheritance applies when you set custom attributes for nodes in the Software Tree.

See “Custom Attributes Set for the Environment” on page 349 in Chapter 6 for more information.

Additionally, you can set custom attributes that affect every server associated with a specific customer or for every server in a facility. Navigate to the customer or facility where you want to set attributes (select Environment ➤ Customers or Facilities in the navigation panel, and click the correct name in the list), click the Custom Attributes tab, and add

attributes for all the servers associated with the customer or located in the facility. When you use this option, you define custom attributes at a customer- or facility-specific level. The procedure to add custom attributes to a customer or facility is the same as adding custom attributes to individual servers.

Additionally, you can add custom attributes for a server group by viewing a server group, and then clicking the Custom Attributes tab for that group. The procedure to add custom attributes to a server group is the same as adding custom attributes to individual servers.

See "Modifying a Server Group" on page 164 in this chapter for more information.

Adding Server Custom Attributes

For information about how to set custom attributes required by the software running on a specific server, contact the group responsible for packaging your applications.

Perform the following steps to add a custom attribute for a server:

1. From the navigation panel, click Servers ➤ Managed Servers or Server Pool. The Managed Servers page appears. Browse the list to find the server for which you want to add custom attributes.

   Or

   Search for the server for which you want to add custom attributes.
2. Click the display name of the server. The Managed Servers: Server Properties page appears.
3. Click the Custom Attributes tab. The Managed Servers: Custom Attributes page appears.
4. Click the New button.
5. Enter the name and value for the custom attribute that you want to add.
6. Click the Save button.

See "Searching with Advanced Search" on page 48 in this chapter for more information.
See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

Editing Server Custom Attributes

If you want to change the name of a custom attribute entry, you need to create a new custom attribute and delete the old custom attribute.

Perform the following steps to edit custom attributes for a server:

1. From the navigation panel, click Servers ➤ Managed Servers or Server Pool. The Managed Servers page appears. Browse the list to find the server for which you want to edit custom attributes.

   Or

   Search for the server for which you want to edit custom attributes.

2. Click the display name of the server. The Managed Servers**:** Server Properties page appears.
3. Click the Custom Attributes tab to change the custom attributes of the server. The Managed Servers: Custom Attributes page appears.
4. Click the attribute name link for the custom attribute that you want to change.
5. Update the value of the custom attribute.
6. Click the Save button to save your changes. The Managed Servers: Custom Attributes page reappears with the updated value.

See “Searching with Advanced Search” on page 48 in this chapter for more information.
See “Searching for Servers by IP Address” on page 55 in this chapter for more information.

Deleting Server Custom Attributes

Perform the following steps to delete custom attributes for a server:

1. From the navigation panel, click Servers ➤ Managed Servers or Server Pool. The Managed Servers page appears. Browse the list to find the server from which you want to remove custom attributes.

   Or

   Search for the server from which you want to remove custom attributes.

See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

2. Click the display name of the server. The Managed Servers: Server Properties page appears.
3. Click the Custom Attributes tab. The Managed Servers: Custom Attributes page appears.
4. Select the check box for the custom attribute that you want to delete.
5. Click the Delete button. The Opsware Command Center displays a confirmation page.
6. Click OK to delete the custom attribute.

See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

Server Groups

This section provides information on server groups within the Opsware System and contains the following topics:

- Server Groups Overview
- Creating a Server Group Type
- Creating a Server Group
- Viewing the Servers in a Server Group
- Modifying a Server Group
- Deleting a Server Group

Server Groups Overview

You can use the Server Group feature to categorize clusters of servers. Using the Server Groups feature is particularly useful when you have many servers deployed in your operational environment.

Not everyone thinks of their operational environment in terms of Web, application, and database tiers. Therefore, you can designate your own groups of servers in the Opsware Command Center. You can use Server Groups to associate servers for a specific purpose or role.

You can create a server group on which you set specific custom attributes. You can write scripts that use the group custom attributes when performing a variety of functions, including network and server configuration, notifications, and CRON script configuration.

You create your own groups. The Opsware Command Center does not create groups for you.

The Server Group feature provides the following two ways to display and define information:

- **Groups** – Groups contain servers and groups assigned to groups.

  When you create a group, you assign servers to the group or other existing groups to the new group.

  For example, you might assign existing groups to a new group when you have a customer for whom you are hosting different applications and each application's servers are defined in a separate group.

- **Servers** – The Servers tab in the Server Groups pages displays a list of the managed servers, which you can filter by customer. For each server, you can view the groups to which it belongs.

Creating a Server Group Type

You can create Server Group Types as a way to identify the different types of server groups in the managed environment.

Perform the following steps to create a server group type:

1. From the navigation panel, click Servers ➤ Server Groups. The Server Groups: View Server Groups page appears. By default, the Server Groups tab is selected.

2 Click the Server Group Types tab. The Server Groups: View Server Group Types for <customer> page appears. See Figure 2-57.

Figure 2-57: Server Group Types Tab

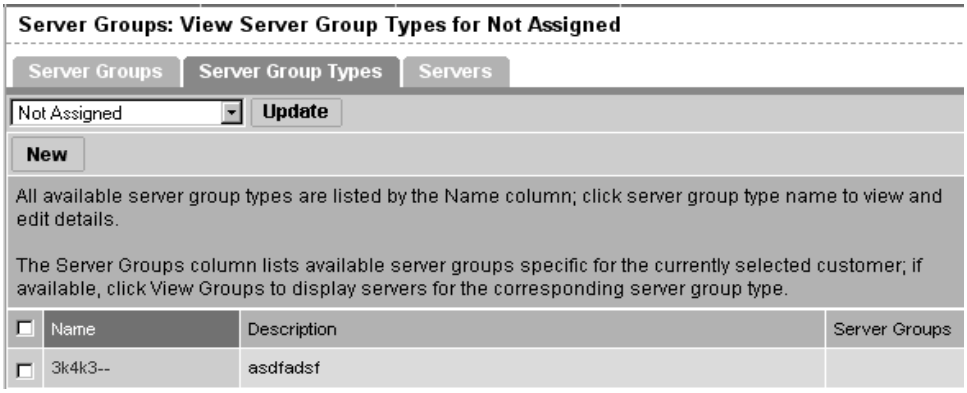

3 If necessary, select the customer for whom you want to create the server group type. Click the Update button. The server group types for that customer appear in the list.

4 Click the New button.

5 Enter a name and description for the new group type.

6 Click the Save button.

Creating a Server Group

Perform the following steps to create a server group:

1 From the navigation panel, click Servers ➤ Server Groups. The Server Groups: View Server Groups for <customer> page appears.

2 Select the customer for whom you want to create the server group and click the Update button. The server groups for that customer appear in the list.

3 Click the New button. The Server Groups: New Server Group page appears, as Figure 2-58 shows.

Figure 2-58: New Server Group Page in the Opsware Command Center

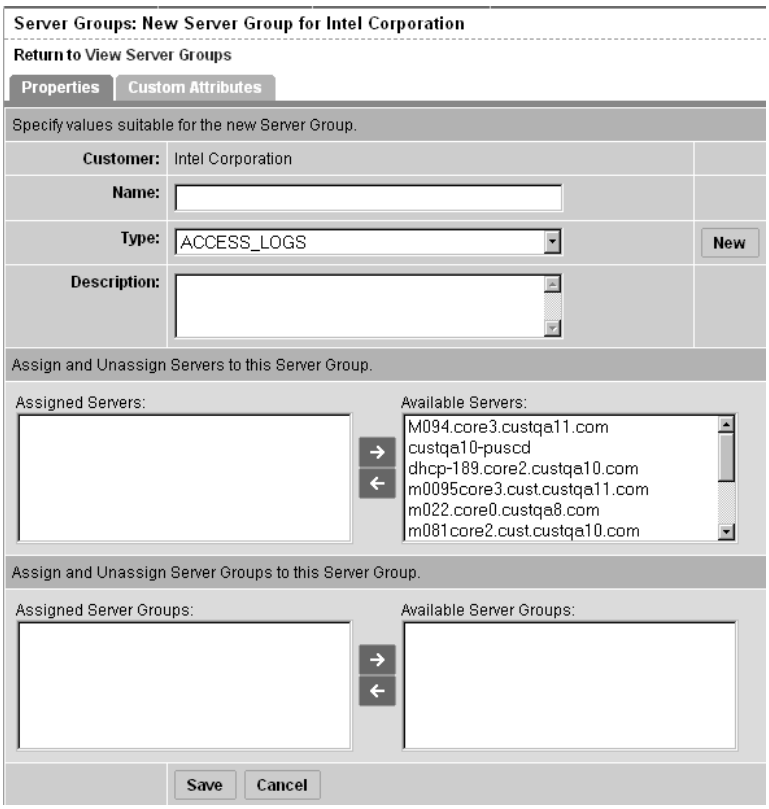

4 Enter a name and description for the new group and select a group type from the list.

5 Add servers or existing server groups to the new group by selecting them in the available list and clicking the left arrow. The servers or server groups move to the assigned list.

6 Click the Save button.

Viewing the Servers in a Server Group

Perform the following steps to view the servers in a server group:

1 From the navigation panel, click Servers ➤ Server Groups. The Server Groups: View Server Groups for <customer> page appears.

2 Select the customer for whom you want to see the servers in a group and click the Update button.

3 From the Server Groups tab, click the View Servers link in the Assigned Servers column for the group that you want to view. The servers in that group appear as a list in the column, as Figure 2-59 shows.

Figure 2-59: Servers Displayed for a Server Group

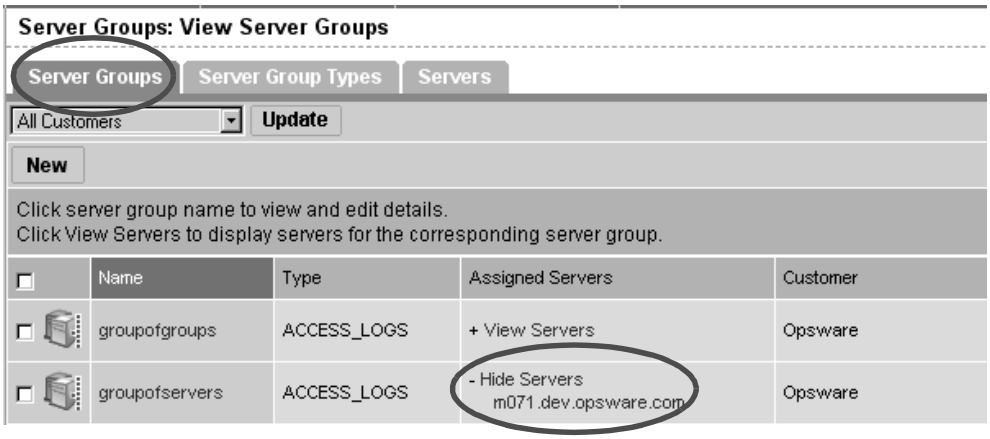

Or

Click the Servers tab to see all the managed servers in the operational environment. The Assigned Groups column displays any groups in which each server is included, as Figure 2-60 shows.

Figure 2-60: Servers Tab Showing the Groups for the "Not Assigned" Customer

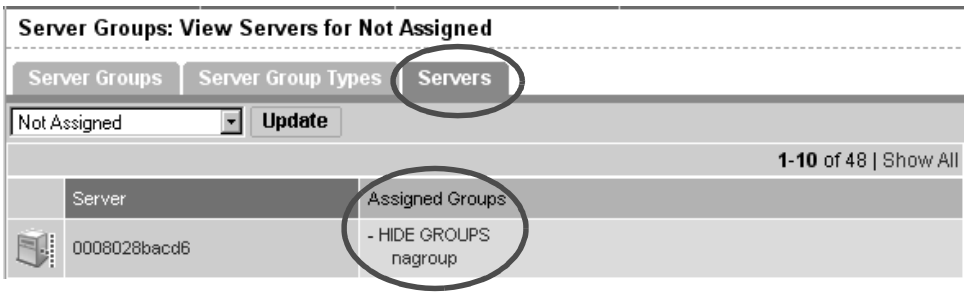

Modifying a Server Group

Perform the following steps to modify a server group:

1 From the navigation panel, click Servers ➤ Server Groups. The Server Groups: View Server Groups for <customer> page appears.

2. Select the customer for whom you want to modify the server group and click the Update button.
3. Click the name of the group that you want to modify. The Server Groups: Edit <group name> page appears.
4. Modify the group name, description, type, or add or remove servers from the group.
5. To modify the custom attributes for the server group, click the Custom Attributes tab.
6. Click the Save button.

See "Editing Server Custom Attributes" on page 159 in this chapter for more information.

Deleting a Server Group

Deleting a server group just deletes the group and does not affect the servers in that group.

Perform the following steps to delete a server group:

1. From the navigation panel, click Servers ➤ Server Groups. The Server Groups: View Server Groups for <customer> page appears.
2. Select the customer for whom you want to delete the server group and click the Update button.
3. Select the server group that you want to delete.
4. Click the Delete button. A popup dialog box prompts you to confirm that you want to delete the group.

Service Levels

This section provides information about service levels within the Opsware System and contains the following topics:

- Service Levels Overview
- Adding a Service Level to the Opsware Command Center
- Adding a Hierarchy of Service Levels
- Editing a Service Level
- Ways to View the Service Level for Servers
- Assigning a Server to a Service Level

- Removing a Server from a Service Level

Service Levels Overview

Service levels are user-defined categories that you can use as an organizational tool. Using service levels allows you to group servers in an arbitrary way. Initially, this category will be fairly empty.

Be default, when an Opsware Agent is installed on a server in the operational environment, the server is added to the UNKNOWN Service Level.

Users often like to organize their servers by functionality (finance, engineering, and so forth) or tier (Web, application, and database) or by ontogeny (development, staging, and production).

Using the Stage and Use properties, you can categorize servers in those ways in the Opsware Command Center. Using the Service Level feature offers users more flexibility for designing their own organizational schemes.

You can create service levels to indicate the Service Level Agreement (SLA) for the servers that your IT organization manages. For example, you might create service levels to denote Silver, Gold, and Platinum services.

Assigning servers to service levels does not cause the Opsware System to operate any differently with respect to those servers.

Adding a Service Level to the Opsware Command Center

Perform the following steps to add a service level to the Opsware Command Center:

1. From the navigation panel, click Environment ➤ Service Levels. The Service Levels page appears.

2 Navigate the hierarchy of service levels until you reach the point in the hierarchy where you want to add a new service level, as Figure 2-61 shows.

Figure 2-61: Service Level Hierarchy

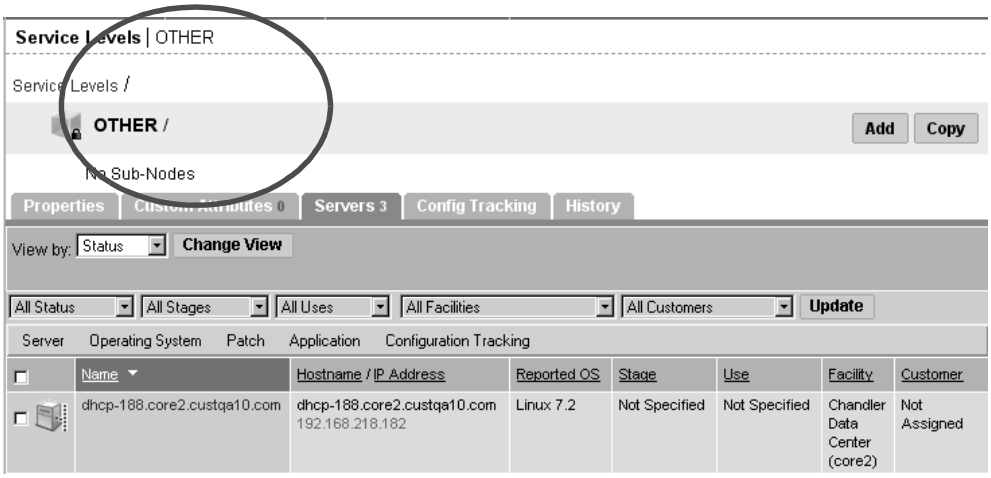

3 Click the Add button. The Service Levels page refreshes and the ADD SUB-NODE TO Service Levels form appears in the page.

4 Enter a name for the service level (required), and (optionally) enter notes and a description for the service level.

5 Click the Save button. The service level is added to the hierarchy of service levels. The Edit Service Level page appears, where you can change the properties of the service level, such as the customer association.

Adding a Hierarchy of Service Levels

Perform the following steps to add a hierarchy of service levels:

1 From the navigation panel, click Environment ➤ Service Levels. The Service Levels page appears, as Figure 2-62 shows. You are at the top level of the hierarchy of service levels. You can only add a hierarchy of service levels at this point.

Figure 2-62: Service Levels Page in the Opsware Command Center

2 Click the Add Many button. The Service Levels page refreshes and the CREATE MODEL in Service Levels form appears in the page, as Figure 2-63 shows.

Figure 2-63: Create Model Page to Add Hierarchy of Service Levels

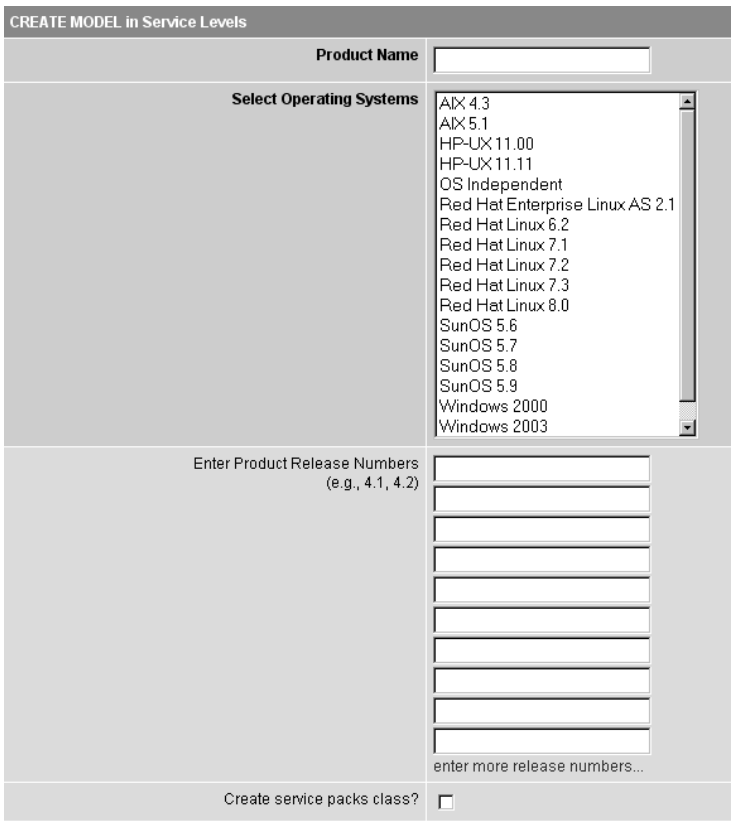

3 Complete the following entries to define the hierarchy of service levels:

- Product Name (Required)

  Product names appear at the top of the hierarchy for the node navigation path. The product name is limited to 25 characters; product names for nodes under the same category must be unique.

- Operating Systems (Required)

  Creates branches within the hierarchy for the operating systems selected. Also, specifies the operating system that each node, by default, is associated with. You can later edit this field for individual nodes in the Opsware Command Center.

- Product Release Numbers

Creates nodes within the hierarchy under each operating system for each product version specified. You can later edit the nodes for versions in the Opsware Command Center.

4 Click the Show Model button. The hierarchy of service levels is added to the Opsware Command Center.

Editing a Service Level

Perform the following steps to edit a service level:

1 From the navigation panel, click Environment ➤ Service Levels. The Service Levels page appears.

2 Navigate the hierarchy of service levels until you reach the point in the hierarchy where you want to edit an existing service level.

3 Click the Edit button in the Properties tab. The page refreshes and an editable form appears for the service level properties.

4 Make changes to the service level name, description, notes, whether servers are allowed to be assigned to the service level, the associated customers and operating systems.

Unlike nodes in the Software Tree, service levels can have multiple customers and operating systems associated with them.

5 Click the Save button.

Ways to View the Service Level for Servers

Find the server whose service levels you want to view by searching or browsing the Managed Servers list.

- If you are browsing the Managed Server list, you can find the service level for a server by locating the value in the Environment column, as Figure 2-64 shows.

Figure 2-64: Service Level Node Appearing in the Software Tab of the Server List

- If you searched for the server, click the server name and then click the Nodes tab. You can find the service levels, as Figure 2-65 shows.

Figure 2-65: Nodes Tab That Shows the Service Level to Which a Server is Assigned

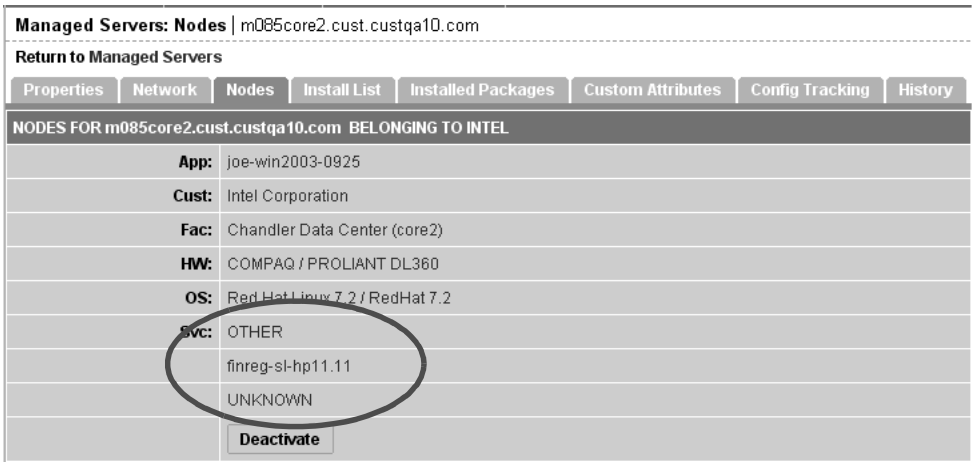

- To view all the servers assigned to a particular service level, click Environment ➤ Service Levels in the navigation panel. Navigate the hierarchy of service levels until you

reach the one for which you want to see which servers are assigned. Click the Servers tab, as Figure 2-66 shows.

Figure 2-66: Managed Server Assigned to a Service Level

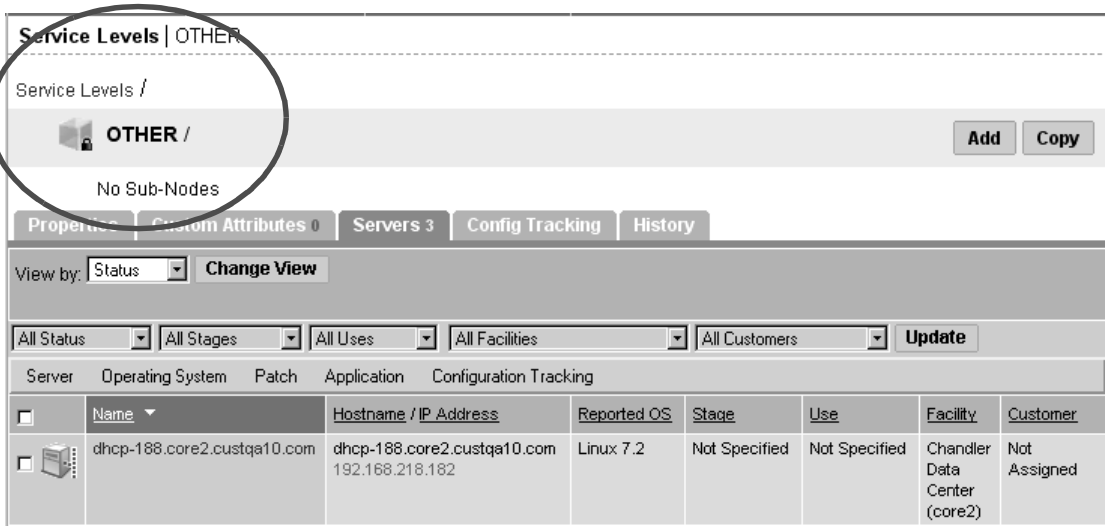

Assigning a Server to a Service Level

Perform the following steps to assign a server to a service level:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server you want to assign to a service level.

   Or

   Search for the server that you want to assign to a service level.

2. Select the servers that you want to assign to a service level.

3 Choose Server ➤ Assign Node from the menu above the Managed Servers list. A window displays the categories of nodes, as Figure 2-67 shows.

Figure 2-67: Assign Nodes Popup Window

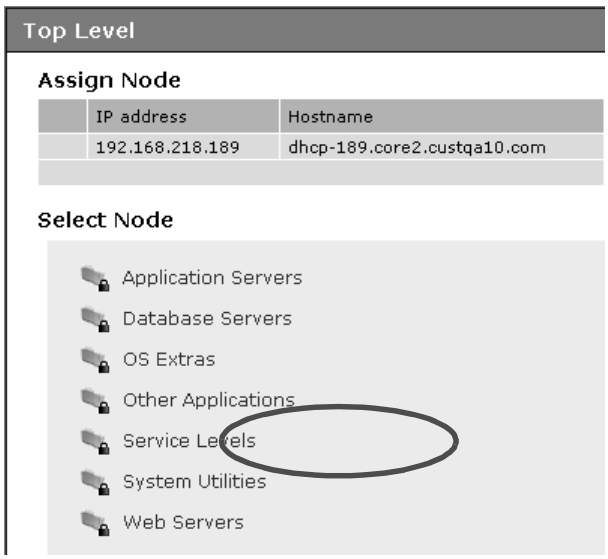

4 Click the Service Levels link. The window refreshes to show the service levels created for your operational environment.

5 Navigate to the service level to which you want to assign the server.

6 Click the Assign button. The window closes and you are returned to the Managed Servers list.

See “Searching with Advanced Search” on page 48 in this chapter for more information.
See “Searching for Servers by IP Address” on page 55 in this chapter for more information.

Removing a Server from a Service Level

Perform the following steps to remove a server from a service level:

1 From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears. Browse the list to find the server that you want to remove from a service level.

Or

Search for the server that you want to remove from a service level.

See "Searching with Advanced Search" on page 48 in this chapter for more information. See "Searching for Servers by IP Address" on page 55 in this chapter for more information.

2. Select the server that you want to remove from a service level.
3. Choose Server ➤ Remove Node from the menu above the Managed Servers list. A window displays the nodes to which the server is assigned, as Figure 2-68 shows.

Figure 2-68: Remove Nodes Popup Window

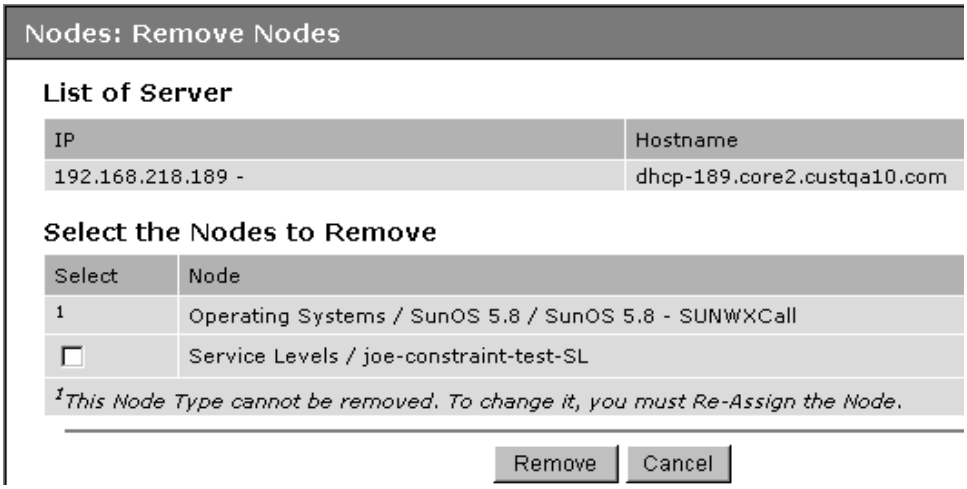

4. Select the service level node from which you want to remove the server and click the Remove button. You are prompted to confirm that you want to remove the server from the service level.
5. Click the Confirm Remove button. The window closes and you are returned to the Managed Servers list.

Chapter 3: OS Provisioning Setup

IN THIS CHAPTER

This chapter provides the following information about how to set up operating system (OS) provisioning in the OS Provisioning Subsystem:

- OS Provisioning Setup
- OS Media Management
- Additional Windows NT Media Setup Tasks
- Operating System Definitions
- Build Customization Scripts
- Working with OS Definitions
- Default Values for the OS Build Process
- Including OS Definitions in Templates Overview
- Hardware Support in OS Provisioning

Before you set up the OS Provisioning Subsystem, the OS provisioning components must have been installed in the local facility with the Opsware Installer and configured correctly. Contact your Opsware administrator for information about the installation and configuration of the Opsware OS provisioning components.

The OS Provisioning Subsystem does not provision HP-UX or AIX operating systems out of the box; however, the Opsware System can be integrated with Network Installation Management (NIM) to provision AIX and Ignite-UX to provision HP-UX. See the *Opsware System 4.7 Installation Guide* for information about how OS provisioning is configured during Opsware System installation.

OS Provisioning Setup

This section provides information on OS Provisioning setup within the Opsware System and contains the following topics:

- OS Provisioning Setup Overview
- Permissions Required to Set Up OS Provisioning
- Process for Setting up OS Provisioning
- Setting Up for Sun Solaris OS Provisioning
- Setting Up for Linux OS Provisioning
- Setting Up for Microsoft Windows OS Provisioning

OS Provisioning Setup Overview

Setting up the OS Provisioning Subsystem is an ongoing process. Before you can provision servers with a new OS, you must set up the OS Provisioning Subsystem to install that OS on the servers in your environment.

Additionally, you should continue to update existing operating systems with the latest patches and security fixes by updating the templates used to install the operating systems.

See "Including OS Definitions in Templates Overview" on page 225 in this chapter for more information.

See "Patch Administration Using the Opsware Command Center" on page 424 in Chapter 8 for information about how to set up patch management in the Opsware System.

The OS Provisioning Subsystem supports installation-based provisioning using Red Hat Linux Kickstart, SUSE Linux YaST2, Sun Solaris JumpStart, and Microsoft Windows unattended installation. Image-based provisioning requires customization that Opsware Professional Services can perform for your environment.

Contact your Opsware Support Representative for information about using image-based provisioning with the Opsware System.

Because the OS Provisioning Subsystem supports installation-based provisioning, your organization can keep its OS installations very 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).

Permissions Required to Set Up OS Provisioning

In the Opsware Command Center, users access only the areas of functionality relevant to their responsibilities in the managed server environment. If access is allowed to a functional area in the Opsware Command Center, the link for that function displays in the navigation panel and on the home page.

To set up OS provisioning in the Opsware Command Center, you must have the following permissions to perform the tasks described in this chapter:

- Access to operating systems so that you can create, edit, and delete OS definitions
- Access to templates so that you can view and edit templates to include OS definitions
- Access to the Prepare Operating System Wizard so that you can use this Wizard to define new operating systems

To obtain the required permissions to perform OS provisioning set up, contact your Opsware administrator.

Process for Setting up OS Provisioning

An OS standards setter records in the OS Provisioning Subsystem the standard configuration of an OS and its required utilities, drivers, and agents. System administrators can then use the OS Provisioning Subsystem to install the OS, configure networking, and install other software required for smooth operation of the server.

Before you perform the tasks to set up OS provisioning, you must have a licensed copy of the OS installation media, which typically comes as a CD-ROM or DVD.

You must perform the following tasks to set up the OS Provisioning Subsystem to install an OS:

1. Make the media for that OS available on the Media Server by performing these tasks:
   1. Copy the OS media to the Media Server.
   2. Create a Media Resource Locator (MRL) for the OS media by using the Opsware Import Media tool.

      See "OS Media Management" on page 185 in this chapter for more information.
2. Create a configuration file with a text editor to specify how the OS will be installed.
3. Prepare a definition in the Opsware Command Center for the OS by performing these tasks:
   1. Indicate the location of the OS media by specifying the correct MRL.
   2. Upload the configuration file into the OS Provisioning Subsystem.

   See "Defining an Operating System" on page 212 in this chapter for more information.

Setting Up for Sun Solaris OS Provisioning

The OS Provisioning Subsystem includes a DHCP-based JumpStart configuration that hides the complexity of JumpStart from the end user. Unlike typical JumpStart systems, the OS Provisioning Subsystem does not require configuration updates to the JumpStart server for each installation that you provision.

Instead, you prepare an OS definition in the OS Provisioning Subsystem for each version of the Solaris OS that you want to install on servers in your environment.

The setup process for Solaris OS provisioning follows the general process for OS provisioning setup. However, you must perform certain setup tasks specifically for each Solaris OS. See the topics that Figure 3-1 lists.

Table 3-1: Setting Up Tasks for Solaris OS Provisioning

	TO PERFORM THIS TASK...	SEE THESE TOPICS...
1	Copying the Sun Solaris OS media to the Media Server by using the scripts included on the Sun Solaris installation CD-ROM or DVD	See "Prerequisites for Creating an MRL" on page 187 in this chapter for more information.
2	Creating an MRL for the Solaris media by using the Import Media tool	See "Creating an MRL with the Import Media Tool" on page 187 in this chapter for more information.
3	Creating a Solaris profile with a text editor	See "About Sun Solaris Profiles" on page 194 in this chapter for more information.
4	Preparing an OS definition for the Solaris OS in the Opsware Command Center by specifying the location of the Solaris OS media (with the MRL) and uploading the profile	See "Defining an Operating System" on page 212 in this chapter for more information.
5	Optionally, specifying a list of software packages or clusters to install after the base OS installation is complete by adding the packages directly to the OS definition	See "About Conditional Packages for Solaris" on page 211 in this chapter for more information. See "Overview of Installation Order for Solaris and Linux" on page 211 in this chapter for more information. See "Modifying Which Packages an OS Definition Installs" on page 219 in this chapter for more information.

Table 3-1: Setting Up Tasks for Solaris OS Provisioning

	TO PERFORM THIS TASK...	SEE THESE TOPICS...
6	Customizing the default build process that the OS Provisioning Subsystem uses to install the version of Solaris on servers	See "About the Solaris Build Customization Script" on page 203 in this chapter for more information. See "Requirements for Solaris Build Customization Scripts" on page 204 in this chapter for more information.
7	Editing the OS definition for the version of Solaris after you have created it so that it passes specific information to the Solaris build script to configure aspects of the installation process	See "Default Values for the OS Build Process" on page 221 in this chapter for more information. See "Custom Attributes for Sun Solaris" on page 222 in this chapter for more information.

Setting Up for Linux OS Provisioning

The OS Provisioning Subsystem includes a Kickstart and YaST2 system that hides the complexity of Kickstart and YaST2 from the end user.

Unlike typical Kickstart or YaST2 systems, mapping a specific installation client to a particular configuration is a simple procedure in the OS Provisioning Subsystem. In the OS Provisioning Subsystem, each Linux OS (and template) have a single configuration associated with them.

The setup process for Linux OS provisioning follows the general process for OS provisioning setup. However, you must perform certain setup tasks specifically for the Linux OS. See the topics that Figure 3-2 lists.

Table 3-2: Setting Up Tasks for Linux OS Provisioning

	TO PERFORM THIS TASK...	SEE THESE TOPICS...
1	Copying thet Linux OS media to the Media Server	See "Prerequisites for Creating an MRL" on page 187 in this chapter for more information.
2	Creating an MRL for the Linux media by using the Import Media tool	See "Creating an MRL with the Import Media Tool" on page 187 in this chapter for more information.

Table 3-2: Setting Up Tasks for Linux OS Provisioning

	TO PERFORM THIS TASK...	SEE THESE TOPICS...
3	Creating a configuration file with a text editor	See “About Red Hat Linux Configuration Files” on page 195 in this chapter for more information. See “About SUSE Linux Configuration Files” on page 195 in this chapter for more information.
4	Preparing an OS definition for the Linux OS in the Opsware Command Center by specifying the location of the Linux OS media (with the MRL) and uploading the configuration file	See “Defining an Operating System” on page 212 in Chapter 3 for more information.
5	Optionally, specifying a list of software packages to install after the base OS installation is complete by adding the packages directly to the OS definition	See “Overview of Installation Order for Solaris and Linux” on page 211 in this chapter for more information. See “Modifying Which Packages an OS Definition Installs” on page 219 in this chapter for more information.
6	Customizing the default build process that the OS Provisioning Subsystem uses to install the version of Linux on servers	See “About Linux Build Customization Scripts” on page 207 in this chapter for more information. See “Requirements for Linux Build Customization Scripts” on page 208 in this chapter for more information.
7	Editing the OS definition for the version of Linux after you have created it so that it passes specific information to the Linux build script to configure aspects of the installation process	See “Default Values for the OS Build Process” on page 221 in this chapter for more information. See “Custom Attributes for Linux” on page 223 in this chapter for more information.

Table 3-2: Setting Up Tasks for Linux OS Provisioning

	TO PERFORM THIS TASK...	SEE THESE TOPICS...
8	Additionally, if necessary, adding new hardware support to a Linux build image The OS Provisioning Subsystem includes build images that install the target OS on servers for Linux.	See "Adding Hardware Support to a Linux Build Image" on page 233 in this chapter for more information.

Setting Up for Microsoft Windows OS Provisioning

To prepare a Windows OS definition, you must set up a Windows unattended installation. You need to provide the following items when you set up Windows provisioning in the OS Provisioning Subsystem:

- A licensed copy of the Windows OS installation media, which typically comes as a CD-ROM or DVD
- Mass storage drivers and network interface card (NIC) drivers. The latest drivers can usually be downloaded from the hardware vendor's website.
- A Windows setup response file

The setup process for Windows OS provisioning follows the general process for OS provisioning setup. However, you must perform certain setup tasks specifically for the Windows OS. See the topics that Figure 3-3 lists.

Table 3-3: Setting Up Tasks for Windows OS Provisioning

	TO PERFORM THIS TASK...	SEE THESE TOPICS...
1	Copying the Windows OS media to the Media Server	See "Prerequisites for Creating an MRL" on page 187 in this chapter for more information.
2	For the Windows NT media, modifying the media from the vendor to install Service Pack 6a and applying Microsoft patch Q143473 to the media	See "Additional Windows NT Media Setup Tasks" on page 190 in this chapter for more information.

Table 3-3: Setting Up Tasks for Windows OS Provisioning

	TO PERFORM THIS TASK...	SEE THESE TOPICS...
3	Creating an MRL for the Windows media by using the Import Media tool	See "Creating an MRL with the Import Media Tool" on page 187 in this chapter for more information.
4	Creating a Windows response file with a text editor	See "About Microsoft Windows Response Files" on page 195 in this chapter for more information.
5	Preparing an OS definition for the Windows OS in the Opsware Command Center by specifying the location of the Windows OS media (with the MRL) and uploading the response file	See "Defining an Operating System" on page 212 in this chapter for more information.
6	In the OS definition, uploading hardware-specific files for the hardware you expect to provision by mapping a signature for that hardware to the correct hardware-specific profile The OS Provisioning Subsystem will select the correct Hardware Signature file at build time based on the hardware signature of the server that is about to be provisioned.	See "About Hardware Signature Files for Windows" on page 212 in this chapter for more information.
7	Optionally, specifying a list of software packages to install after the base OS installation is complete by adding the packages directly to the OS definition	See "Modifying Which Packages an OS Definition Installs" on page 219 in this chapter for more information.

Table 3-3: Setting Up Tasks for Windows OS Provisioning

	TO PERFORM THIS TASK...	SEE THESE TOPICS...
8	Customizing the default build process that the OS Provisioning Subsystem uses to install the version of Windows on servers	See "About Windows Build Customization Scripts" on page 210 in this chapter for more information.
9	Editing the OS definition for the version of Windows after you have created it so that it passes specific information to the Windows build script to configure aspects of the installation process For a Windows OS definition, you can set a value for the timeout custom attribute. Setting this value controls the timeout value after an error.	See "Default Values for the OS Build Process" on page 221 in this chapter for more information. See "Custom Attribute for Microsoft Windows" on page 224 in this chapter for more information.
10	If you need to boot x86-process based servers from a floppy (perhaps you cannot boot servers over the network), creating a Windows boot floppy	See "Creating a Windows Boot Floppy" on page 231 in this chapter for more information.
11	Additionally, if necessary, adding new hardware support to the Windows boot images The default boot images for Windows include common NIC drivers for many hardware makes and models. The Opsware System uses these NIC drivers to boot new x86-processor-based servers for the first time.	See "Adding NIC Support to a Windows Floppy Image" on page 229 in this chapter for more information.

Table 3-3: Setting Up Tasks for Windows OS Provisioning

TO PERFORM THIS TASK...	SEE THESE TOPICS...
1 2 Updating the Windows PXE image after adding hardware support When the Opsware System was installed with the Opsware Installer, a image was added to the PXE system by default. You only need to update the PXE image when you have added support for additional NIC drivers to the image.	See "Updating the PXE Image for Windows" on page 233 in this chapter for more information.

OS Media Management

This section provides information on OS media management within the Opsware System and contains the following topics:

- OS Media Management Overview
- Prerequisites for Creating an MRL
- Creating an MRL with the Import Media Tool
- Editing an MRL
- Deleting an MRL

OS Media Management Overview

OS media consists of the installation software for an OS from the software vendor. Typically, OS media is distributed on CD-ROM, DVD, or by downloading the software distribution from the vendor's FTP site. The OS media can contain binaries for installing the OS, packages of different types, metadata about the packages, and other information.

So that the OS Provisioning Subsystem can access the media, you must copy it to the Opsware Media Server. The Media Server provides access to the OS media over the network by using NFS for Linux and Solaris OS provisioning, and by using SMB for

Windows OS provisioning. After copying the OS media to the Media Server, you must import it into the Opsware System by running the Opsware Import Media tool (a utility script included with the Opsware System).

Running the Import Media tool creates an Opsware-generated string called a Media Resource Locator (MRL) for each OS media that you want to provision.

An MRL is a network path (in URI format) to the installation media for an OS on the Opsware Media Server. When a server is being provisioned with an OS, the server mounts the network path for the OS media by using NFS (for Linux and Solaris), SMB (for Windows). The MRL is registered with the Opsware System. An MRL should resolve to the Media Server in the local facility where the Opsware System is installed.

To create an MRL, run the Media Import tool. Running the Import Media tool automatically performs the following functions:

- Mounts the media at the specified network path by using NFS or SMB
- Detects the OS (Solaris, Linux, or Windows) and version of the media
- Based on the server name and path that you specify, creates that MRL in the Opsware System so that you can use it in OS definitions
- Extracts vendor-provided metadata (such as the package list and dependencies between packages) from the OS software and stores this data in the Opsware System
- For Sun Solaris and Linux, uploads all packages to the Software Repository so that the OS Provisioning Subsystem can install them after initial OS provisioning

  For Solaris, an MRL represents or contains a path to the media for JumpStart purposes, a hierarchy of metaclusters, clusters, and packages, and information about package dependencies and installation order.

  For Linux, an MRL contains a path to the media for Kickstart or YaST2 and information about package dependencies and installation order.

  Re-running the Import Media tool with the same server and path as an existing MRL updates the MRL, but does *not* re-upload duplicate Linux or Solaris packages.
- For Linux and Microsoft Windows, modifies portions of the OS media to integrate the OS Provisioning Subsystem with the vendor provisioning boot process.

Prerequisites for Creating an MRL

Before you run the Import Media tool, the OS media that you want to import must be available through the network on the Media Server. If necessary, contact your Opsware administrator for the hostname of the Media Server.

Before you perform the tasks to set up OS provisioning, you must have a licensed copy of the OS installation media, which typically comes as a CD-ROM or DVD.

You must know what locations were specified for the OS media. When the Opsware System was installed, the Opsware Installer prompted for the pathnames of the root directories for the Windows, Solaris, and Linux OS media on the Opsware Media Server. If necessary, contact your Opsware administrator for this information.

Perform the following tasks to set up OS provisioning:

1. On the Media Server host, create the directory structure for the versions of the OS that you plan to use for server provisioning.

   Create the directory structure based on the root directories specified for the OS media during Opsware System installation. If necessary, contact your Opsware administrator for the locations of the OS media root directories.

2. The media for each OS that you want to provision needs to be available on the Media Server.

   - For Microsoft Windows, copy the OS media files to the correct location on the Media Server.
   - For Linux, copy the OS media files to the correct location specified on the Media Server. The OS media needs to be NFS exported read write.

     For SUSE Linux, see http://www.suse.de/~nashif/autoinstall/multiplesource.html for information on how to deal with multiple sources.
   - For Sun Solaris, use the Sun Solaris scripts included on the CD-ROM or DVD to copy the OS media files to the correct location on the Media Server.

Creating an MRL with the Import Media Tool

Perform the following steps to create an MRL with the Import Media Tool:

1. Log in to the Software Repository host as root.
2. For Sun Solaris and Linux media, NFS mount the OS media on the Media Server from the Software Repository host.

You must know the correct location for the OS media.

For example, enter the following command to NFS mount Solaris and Linux media:

```
theword# mount mediaserver:/usr/local/solaris/5.8 /mnt
```

3 On the Software Repository host, run the `import_media` script in the following directory:

```
/cust/usr/blackshadow/mm_wordbot/util/
```

To write-protect the Windows media share, a password was set for the root user (parameter: `media_server.windows_share_password`) when the Opsware System was installed. The Opsware Import Media Tool prompts for the password each time you run it. Contact your Opsware administrator for this password.

When running the `import_media` script, specify as an argument the directory where the OS media is mounted. For Windows, you must specify the directory of the Windows OS media by using UNC style with the following syntax:

```
//<server_name>/<sharename>/I386
```

The path must end at the `/I386` directory.

For example, enter the following Import Media tool command for Solaris and Linux:

```
theword# /cust/usr/blackshadow/mm_wordbot/util/import_media
/mnt
```

For example, enter the following Import Media tool command for Windows:

```
import_media //mediasrv.corp.lionscapital.com/PUB/WIN2000/
SERVER/I386
```

Running the Import Media tool writes progress to the log file `import_media.log`. The log file is located on the server where you are running the Import Media Tool script in the directory from which you invoke the script.

Editing an MRL

Perform the following steps to edit an MRL:

1 Log in to the Opsware Command Center. The Opsware Command Center home page appears.

2 From the navigation panel, click Software ➤ Operating Systems. The Operating Systems page appears.

3 Click the OS Media tab. A list of Media Resource Locators appears.

Each MRL represents media available for installation. See Figure 3-1.

Figure 3-1: OS Media Page in the Opsware Command Center

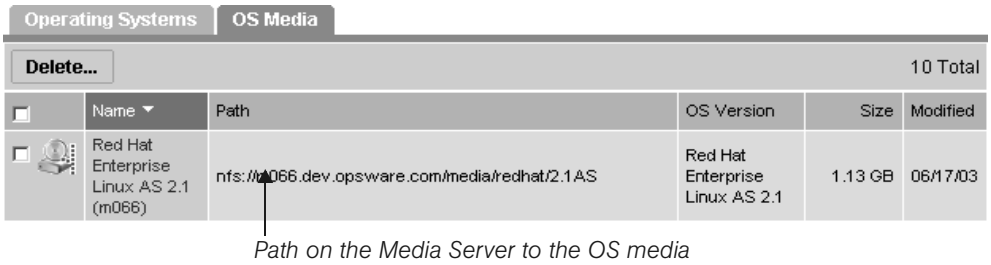

4 Click the display name for the MRL that you want to edit. The Edit OS Media page appears, as Figure 3-2 shows.

Figure 3-2: Edit OS Media Page in the Opsware Command Center

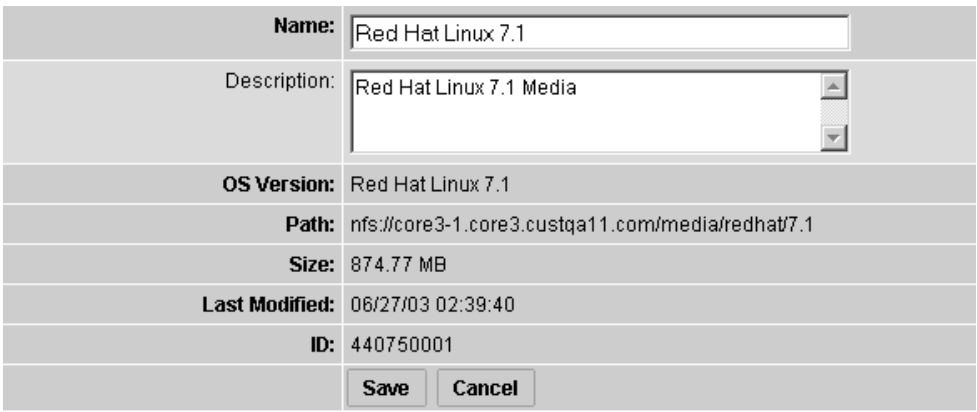

5 Modify the name or description of the MRL.

You cannot edit the OS media path with the Opsware Command Center. If the path for the OS media changes on the Media Server, create a new MRL with the Import Media tool. Then delete the out-of-date MRL by using the Opsware Command Center.

See “Creating an MRL with the Import Media Tool” on page 187 in this chapter for more information.

6 Click the Save button.

Deleting an MRL

You cannot delete an MRL with the Opsware Command Center when the MRL is specified in an OS definition. To delete an MRL specified in an OS definition, you must first delete the OS definition or specify another MRL in the OS definition.

See "Defining an Operating System" on page 212 in this chapter for more information.

Perform the following steps to delete an MRL:

1 Log in to the Opsware Command Center. The Opsware Command Center home page appears.

2 From the navigation panel, click Software ➤ Operating Systems. The Operating Systems page appears.

3 Click the OS Media tab. The list of media available for installation appears.

4 Select the OS Media that you want to delete.

5 Click the Delete button. (If the MRL is specified in an OS definition, a warning message appears.)

The list of Media Resource Locators re-appears.

Additional Windows NT Media Setup Tasks

You must modify the Windows NT media from the vendor before you can use it to provision servers. Perform the following tasks to setup OS provisioning for Windows NT:

- Setting Up Installation of Service Pack 6a
- Applying Microsoft Patch Q143473 to the Windows NT Media

Setting Up Installation of Service Pack 6a

Before you provision Windows NT servers, you must set up the OS Provisioning Subsystem to install Service Pack 6a with the OS. Use the `cmdlines.txt` feature of Windows setup to install Service Pack 6a along with the setup.

Perform the following steps to set up installation of Service Pack 6a:

1. Obtain the Service Pack 6a executable file `sp6i386.exe` from the Microsoft FTP site that contains product updates and copy the file `sp6i386.exe` into the Windows NT `I386\$OEM$` directory on the Media Server.
2. Create a file named `cmdlines.txt` in the `I386\$OEM$` directory that has the following contents:

```
[Commands]
"sp6i386.exe -u -o -z -q"
```

By performing these tasks, Service Pack 6a is silently installed on servers during Windows setup.

For more information about how to install Service Pack 6a with `cmdlines.txt`, see the Microsoft Knowledge Base Article 168814, Installation Option 3, on the Microsoft Web site.

Applying Microsoft Patch Q143473 to the Windows NT Media

Before you provision Windows NT servers, you must apply Microsoft Patch Q143473 to the Windows NT media that you copied to the Media Server.

Without the patch, Windows NT unattended setup stops and prompts you to press any key to shut down. The Windows NT media requires this patch for unattended builds to function properly.

Perform the following steps to apply Microsoft Patch Q143473:

1. Download the patch Q143473 from the Microsoft FTP site that contains patches.
2. Copy the file into the Windows NT `I386\$OEM$` directory on the Media Server.

For more information about applying patch Q143473 to the Windows NT media, see the Microsoft Knowledge Base Article Q143473 on the Microsoft Web site.

Operating System Definitions

This section provides information on operating system definitions within the Opsware System and contains the following topics:

- Operating System Definitions Overview
- About Specifying Software in OS Definitions
- About Configuration Files
- About Sun Solaris Profiles

- About Red Hat Linux Configuration Files
- About Microsoft Windows Response Files
- About Microsoft Windows Response Files
- Sample Response File for Windows 2000
- Sample Response File for Windows NT

Operating System Definitions Overview

To provision a server with an OS, the OS must first be defined in the OS Provisioning Subsystem.

See "Process for Setting up OS Provisioning" on page 177 in this chapter for information about the overall process of setting up OS provisioning.

OS definitions store all relevant information needed to provision an OS. You create OS definitions by using the Prepare Operating System Wizard in the Opsware Command Center.

Perform the following tasks to define an OS:

- Specify properties for the OS
- Specify the OS media from which to perform the installation by selecting an MRL

  See "OS Media Management" on page 185 in this chapter for information about MRLs.
- Upload the following installation resources used during unattended installation:
  - A standard configuration file for the OS

    See "About Configuration Files" on page 194 in this chapter for more information.
  - A build customization script, which can modify the installation process at certain points

    See "Build Customization Scripts" on page 199 in this chapter for more information.
  - For Microsoft Windows only, a Hardware Signature, which contains hardware specific information

    See "About Hardware Signature Files for Windows" on page 212 in this chapter for more information.

Table 3-4 compares the installation resources across operating systems.

Table 3-4: Installation Resources for OS Definitions

INSTALLATION RESOURCE	WINDOWS	SOLARIS	LINUX
Configuration File	Required File name: `unattend.txt`	Required `profile`	Required `profile`
Build Customization Script	Optional Executable file: `run.bat`	Optional Executable file: `run`	Optional Executable file: `run`
Hardware Signature File	Optional *`filename`*`.txt`	Not required	Not required

The configuration file that you upload for each OS can have any filename. When the file is uploaded, the OS Provisioning Subsystem renames the file so that it has the correct name for that OS.

You can edit an OS definition later to add support for new hardware or to change the way the OS is installed.

See “About Editing OS Definitions” on page 216 in this chapter for more information.

About Specifying Software in OS Definitions

Solaris and Linux are package-oriented operating systems. In other words, you can define a particular OS build as a set of Solaris or RPM packages.

An OS definition can contain a list of software packages or clusters to install after the base OS installation is complete. You specify the software packages to install during OS provisioning in the following ways:

- Uploading a configuration file that specifies to the vendor installation program what software packages to install

  For example, a JumpStart profile contains a list of clusters (and optionally packages) to be installed by JumpStart. A Kickstart configuration file specifies to Kickstart the RPMs

to be installed. When you upload a configuration file, the OS Provisioning Subsystem extracts the list of software packages that will be installed by the vendor's installer.

Extracting the packages allows the Opsware System to manage the software so that you can upgrade or remove software from an OS definition later.

- By adding packages directly to an OS definition

  You can select packages from the list of packages already uploaded to the Opsware System. The Opsware Agent installs the selected packages after the vendor installation program installs the initial OS and the packages specified in the configuration file.

About Configuration Files

A configuration file is required for each OS definition:

- For Solaris, you must create and upload a JumpStart profile.
- For Red Hat Linux, you must create and upload a Kickstart configuration file.
- For SUSE Linux, you must upload a YaST2 configuration file.
- For Windows, you must create and upload a response file.

The purpose of these configuration files is described in the following topics.

About Sun Solaris Profiles

When preparing a Solaris OS definition, the OS Provisioning Subsystem requires that you upload a JumpStart profile. The OS Provisioning Subsystem extracts the list of software to be installed from the uploaded profile by examining the cluster and package specifications. If the profile specifies an invalid cluster or package name, the OS Provisioning Subsystem generates an error. No other profile validation occurs when the profile is uploaded.

The Solaris profile must meet the following requirements:

- Be a valid profile that you would use with a JumpStart server
- Specify that the installation type is an initial installation and not an upgrade
- Specify a package-based installation by listing the clusters and packages to install
- Specify disk partitioning information

See "About Conditional Packages for Solaris" on page 211 in this chapter for information about how the OS Provisioning Subsystem handles Solaris conditional packages.

About Red Hat Linux Configuration Files

The Red Hat Linux configuration file instructs the Kickstart server on what packages to install, how to partition the drive, and how to configure the runtime network post-installation.

When preparing a Red Hat Linux OS definition, the Opsware System validates the Kickstart configuration file. When the configuration file is uploaded, the OS Provisioning Subsystem parses the file in order to extract the package list.

The Red Hat Linux configuration file must meet the following requirements:

- Be a valid configuration file that you would use with a Kickstart server
- Specify the RPM packages to install
- Must include the reboot option

In the Red Hat Linux configuration file, do not enable firewalls. The Opsware Agent must communicate with the Opsware System on port 1002.

About SUSE Linux Configuration Files

The SUSE Linux configuration file instructs YaST2 on what packages to install, how to partition the drive, and how to configure the resulting machine.

When preparing a SUSE Linux OS definition, the Opsware System validates the YaST2 configuration file. When the configuration file is uploaded, the OS Provisioning Subsystem parses the file in order to extract the package list.

The SUSE Linux configuration file must meet the following requirements:

- Be a valid YaST2 configuration file
- Under the general options, the reboot and confirm properties in the mode resource needs to be set to true and false respectively

For SUSE Linux, see http://www.suse.de/~nashif/autoinstall/8.0/html/index.html and http://www.suse.de/~nashif/autoinstall/sles8/html/index.html for information on Linux installations.

About Microsoft Windows Response Files

For a Windows OS definition, the configuration file must meet the following requirements.

- Be an unattended installation response file that contains the following settings:
  - Sets the `OemPreInstall` key to yes. If this key is not set, the OS Provisioning Subsystem will set it automatically.
  - Specifies a network configuration so that the OS boots for the first time with a valid IP address.
  - Suppresses any dialog boxes that might appear during the Text and GUI mode portions of Windows setup.

When uploading an `unattend.txt` file, the Opsware System validates the response file and rejects incomplete response files.

See "Sample Response File for Windows 2000" on page 196 in this chapter for information about examples of valid Windows response files. See "Sample Response File for Windows NT" on page 197 in this chapter for information about examples of valid Windows response files.

Sample Response File for Windows 2000

The following sample response file shows how to create a valid response file for a Windows 2000 installation. This sample response file contains the required settings for Windows 2000 provisioning with the OS Provisioning Subsystem.

```
; Minimal unattend.txt for installing Windows 2000 Professional,
; Server, and Advanced Server
;
; All parameters listed in this file are required for Windows
; 2000 setup and Opsware OS provisioning to be completely
; unattended.
;
; Values between <> are values that you must provide.
; For more information, see the unattend.doc file in the
; Support\Tools folder in the Deploy.cab file on the Windows
; 2000 CD-ROM.
;
[Unattended]
UnattendMode=FullUnattended
TargetPath=*
OemSkipEula=Yes
; The OemPreInstall key is automatically provided by Opsware
; OS provisioning.
OemPreinstall=Yes
```



```
[GuiUnattended]
AdminPassword=<*>
OEMSkipRegional=1
OEMSkipWelcome=1
TimeZone=<085>

[UserData]
; The ComputerName parameter is automatically provided by
; Opsware OS provisioning.
ComputerName=*
FullName=<Your User Name>
OrgName=<Your organization name>
ProductID=<License key provided by Microsoft>

; For server installs only
[LicenseFilePrintData]
AutoMode = <PerServer>
AutoUsers = <5>

; Installs TCP/IP on network interfaces. Interfaces are
; configured for DHCP.
[Networking]

[Identification]
JoinWorkgroup = <Workgroup>
```

Sample Response File for Windows NT

The following sample response file shows how to create a valid response file for a Windows NT installation. This sample response file contains the required settings for Windows NT provisioning with the OS Provisioning Subsystem.

```
; Minimal unattend.txt for installing Windows NT Workstation,
; Server, and Enterprise Server.
;
; All parameters listed in this file are required for Windows NT
; setup and Opsware OS provisioning to be completely unattended.
;
; Values between <> are values that you must provide.

[Unattended]
ConfirmHardware = no
TargetPath = *
NoWaitAfterTextMode = 1
NoWaitAfterGuiMode = 1
```

```
OEMSkipEula = yes

; The OemPreInstall key is automatically provided by Opsware
; OS provisioning.
OemPreinstall = yes

[UserData]
; The ComputerName parameter is automatically provided by
; Opsware OS provisioning.
ComputerName = *
FullName=<Your User Name>
OrgName=<Your organization name>
ProductID=<License key provided by Microsoft>

; For server installs only
[LicenseFilePrintData]
AutoMode = <PerServer>
AutoUsers = <5>

[GuiUnattended]
AdvServerType = <SERVERNT>
OEMSkipWelcome = 1
OEMBlankAdminPassword = 1
TimeZone = <"(GMT) Monrovia, Casablanca">

[Display]
ConfigureAtLogon = 0
BitsPerPel = 16
XResolution = 1024
YResolution = 768
VRefresh = 70
AutoConfirm = 1

; Installs TCP/IP on network interfaces. Interfaces are
; configured for DHCP.
[Network]
JoinWorkgroup = <Workgroup>
DetectAdapters = ""
InstallProtocols = ProtocolsSection

[ProtocolsSection]
TC = TCParameters

[TCParameters]
DHCP = Yes
```


Build Customization Scripts

This section provides information on build customization scripts within the Opsware System and contains the following topics:

- Build Customization Scripts Overview
- Sun Solaris Build Process
- About the Solaris Build Customization Script
- Requirements for Solaris Build Customization Scripts
- Sample Solaris Build Customization Script
- Linux Build Process
- About Linux Build Customization Scripts
- Requirements for Linux Build Customization Scripts
- Microsoft Windows Build Process
- About Windows Build Customization Scripts

Build Customization Scripts Overview

To control the way each OS is installed on servers, the OS Provisioning Subsystem utilizes OS-specific build scripts. Build scripts manage each OS installation from the point where bare-metal hardware is connected to the network and booted for the first time through the point where the OS and an Opsware Agent are installed on the server.

Customers need flexibility to customize how operating systems are installed in their environments. Therefore, the OS provisioning build scripts provide hooks into the build process to modify OS installations at specific points. These hooks call a single build customization script at the appropriate time in the OS installation process. Examples of what you can customize by using build customization scripts are given in the following topics.

Because each build script is specific to the OS it installs, how to create a build customization script varies by OS. Additionally, the way a build customization script can modify the OS installation process varies by OS.

To use a build customization script, follow this general process:

1. Upload the file that contains the build customization script (with the correct filename) in the Opsware Command Center by clicking Software ➤ Packages in the navigation panel. (The build script for an OS looks for a build customization script that has a specific name.) When you upload the file, specify "Installation Hooks" as the type of package.

   See "Uploading a Package" on page 293 in Chapter 5 for more information.

2. While preparing an OS definition with the wizard, select the build customization script during Step 2 (Define Installation). The uploaded build customization scripts appear in a list when you click the Select button.

   See "Defining an Operating System" on page 212 in this chapter for more information.

Sun Solaris Build Process

The following table describes in detail the exact steps that occur in the OS Provisioning Subsystem to provision an installation client with Solaris.

A user initiates the build process with Steps 1 and 5. The rest of the build process steps happen automatically in the OS Provisioning Subsystem.

It is important to understand the Solaris build process before you include a build customization script in a Solaris OS definition. See Table 3-5.

Table 3-5: Sun Solaris Build Process

PHASE	BUILD PROCESS STEPS
Pre-installation	1 A user boots the installation client over the network by entering the command in a console attached to the server: `boot net:dhcp - install` 2 The installation client boots from the network by using a provided Solaris 9 JumpStart miniroot (included as part of the Opsware OS Provisioning Subsystem), eventually running a JumpStart `begin` script. The `begin` script is used to start the Opsware OS Build Agent. 3 The OS Build Agent registers with the OS Build Manager. 4 The Solaris `build` script probes the hardware configuration of the installation client and registers it with the Opsware System. The installation client then appears in the Server Pool list in the Opsware Command Center.

Table 3-5: Sun Solaris Build Process

PHASE	BUILD PROCESS STEPS

Phase One	**5** In the Opsware Command Center, a user chooses to install an OS on an available installation client.
	6 The Solaris build script mounts the Solaris installation media indicated by the MRL in the OS definition that the user selected.
	7 The Solaris build script retrieves the profile associated with the selected OS definition and copies it to `$SI_PROFILE`, the standard JumpStart location for dynamic JumpStart profiles.
	8 The Solaris build script executes the build customization script:
	`/sbin/sh run Pre-JumpStart`
	9 The Solaris build script validates the profile by using the JumpStart installer (`pfinstall`) in ted mode.
	10 The Solaris build script causes the OS Build Agent to run in the background, allowing the JumpStart `begin` script to complete.
	11 The JumpStart installer `pfinstall` is invoked by the JumpStart installer script and Solaris is installed. Concurrently, the OS Build Agent monitors the installation process. Feedback is displayed in the Opsware Command Center.
	12 The JumpStart installer `pfinstall` completes and runs the JumpStart `finish` script, which indicates to the OS Provisioning Subsystem that the OS installation is complete.
	13 The build script executes the build customization script a second time:
	`/sbin/sh run Post-JumpStart`
	14 The installation client reboots.

Table 3-5: Sun Solaris Build Process

PHASE	BUILD PROCESS STEPS

Phase Two	**15** On entering multiuser mode, the OS Build Agent is invoked and it contacts the OS Build Manager.
	16 The Solaris build script executes the build customization script:
	`/sbin/sh run Pre-Agent`
	17 The Solaris build script installs the Opsware Agent.
	18 The Solaris build script executes the build customization script:
	`/sbin/sh run Post-Agent`
	19 The Solaris build script exits and Phase Two finishes.

The OS Provisioning Subsystem takes over, causing a reconcile of the selected software to be installed onto the installation client.

See "Overview of Reconcile" on page 442 in Chapter 9 for information about how reconcile works to install software on servers.

About the Solaris Build Customization Script

You can customize a Solaris installation at multiple points; therefore, the Solaris build customization script runs more than once:

- A pre-installation hook for the first stage (Pre-JumpStart)

  During phase one, the build customization script runs in the JumpStart environment. The script can use all the standard JumpStart environment variables, such as `SI_PROFILE`. All the environment variables associated with the standard JumpStart probe keywords and values are set (for example, `SI_DISKLIST`, `SI_HOSTADDRESS`, and `SI_MEMSIZE`).

  When the `run` script is invoked at the Pre-JumpStart point, it can perform any actions that a JumpStart `begin` script would perform. For example, the script could modify the downloaded profile before the OS installation begins. At this point, the Solaris profile is downloaded from the OS Provisioning Subsystem but the profile has not been passed to the JumpStart server.

  For the complete list of the environment variables, see the *Solaris 9 Installation Guide*.

- A post-installation hook for the first stage (Post-JumpStart)

When the `run` script is invoked at the Post-JumpStart point, it can perform any actions that a JumpStart `finish` script would perform. One example would be to set custom `eeprom` settings. The installation client's file systems are available for modification at this point and are mounted on the `/a` partition for the `finish` script environment.

- A pre-installation hook for the second stage (Pre-Agent)
- A post-installation hook for the second stage (Post-Agent)

  During Phase Two, the `run` script is executed after the installation client has rebooted, at a point after the system is up and running in multi-user mode with most services started.

The last 4K of output produced by the build customization script (`stdout` and `stderr`) appears in the Opsware Command Center output details for the OS.

Requirements for Solaris Build Customization Scripts

To use a build customization script for Solaris, you must meet the following requirements:

- Create the script as a Bourne shell script and name it `run`.
- Include the `run` script in an archive file in `tar.Z` format. Include the script at the top level of the archive. During OS provisioning, the `tar.Z` archive is unpacked on the installation client and the script is processed by `/sbin/sh`.
- The `run` script is unpacked in its own directory with the other files in the archive. This directory serves as the current working directory when the `run` script is invoked. Based on this fact, correctly refer to the other files in the archive. For example, unpacking and invoking the `run` script follows this general process:

  ```
  mkdir /var/tmp/inst_hook
  cd /var/tmp/inst_hook
  zcat hook.tar.Z | tar xf -
  /sbin/sh run <stage>
  ```

- The script cannot cause the installation client to drop its network connection (for example, do not use the script to reboot the installation client or reconfigure the active network interface). If the installation client drops its network connection, the OS provisioning process will fail.
- The `run` script must exit normally. If the script exits with a non-zero value, the OS provisioning process will end. However, the Jumpstart process will continue when a pre-installation hook fails (exits with a non-zero value). When creating the run script,

you should ensure that the JumpStart process does not continue when a pre-installation hook fails.

- The `run` script should not take an exceptionally long time to complete, otherwise the OS provisioning process might time out.

Sample Solaris Build Customization Script

```
#!/sbin/sh
pre_jumpstart() {
    #
    # strip any partitioning information out of profile, and
    # replace it with keywords to use default partitioning, but
    # to size swap equal to the amount of physical RAM
    #
cat $SI_PROFILE | grep -v partitioning | grep -v filesys > /tmp/
profile.$$
    echo "partitioning default" >> /tmp/profile.$$
    echo "filesys any $SI_MEMSIZE swap" >> /tmp/profile.$$
    cp /tmp/profile.$$ $SI_PROFILE
    rm -f /tmp/profile.$$
}
post_jumpstart() {
    #
    # set local-mac-address eeprom setting
    #
    eeprom 'local-mac-address?=true'
}
pre_agent() {
    : # do nothing
}
post_agent() {
    : # do nothing
}
case "$1" in
          Pre-JumpStart)  pre_jumpstart ;;
          Post-JumpStart) post_jumpstart ;;
          Pre-Agent)      pre_agent ;;
          Post-Agent)     post_agent ;;
esac
```

Linux Build Process

The following table describes in detail the exact steps that occur in the OS Provisioning Subsystem to provision an installation client with Red Hat or SUSE Linux.

A user initiates the build process with Steps 1 and 6. The rest of the build process steps happen automatically in the OS Provisioning Subsystem.

It is important to understand the Linux build process before you include a build customization script in a Linux OS definition. See Table 3-6 for the Linux build process.

Table 3-6: Linux Build Process

PHASE	BUILD PROCESS STEPS
Pre-installation	1 A user boots the installation client from PXE or the Linux Boot CD ROM. 2 The installation client loads a custom Red Hat AS 3.0 boot image and mounts the second stage image specified by the kernel parameters. 3 Anaconda is replaced by a custom Opsware script that is used to invoke the OS Build Agent. 4 The OS Build Agent registers with the Opsware Build Manager. 5 The Linux build script probes the hardware configuration of the installation client and registers it with the Opsware System, causing the installation client to appear in the Server Pool list in the Opsware Command Center.
Phase One	6 In the Opsware Command Center, a user selects the target version of Linux to install on the installation client. 7 The Linux build script creates a 10 cylinder partition at the beginning of the disk and copies the target boot image from the Boot Server to this partition. 8 The Linux build script copies GRUB onto the partition and installs it into the MBR. 9 The Linux build script configures GRUB to boot this partition and kernel arguments are set to do an NFS installation on the location indicated by the MRL. 10 If the Custom Attribute `kernel_arguments` is set for the OS definition, these kernel arguments are appended. 11 The OS Build Agent exits and the server reboots.

Table 3-6: Linux Build Process

PHASE	BUILD PROCESS STEPS
Phase Two	**12** The target boot image loads and runs the OS Build Agent. **13** The Linux build script verifies that the media indicated by the MRL is the same version as the boot image under which it is running. **14** The Linux build script writes the configuration file defined by the MRL to the disk. **15** If it exists, the Linux build script runs the build customization script. **16** The Linux build script runs in the background. The OS Build Agent and Anaconda starts. The Linux installation starts normally by using the configuration file written to the disk. Concurrently, the OS Build Agent monitors the installation process providing feedback, which is displayed in the Opsware Command Center. **17** After all packages have been installed, as part of the post install, the OS Build Agent copies the Opsware Agent Installer and the OS Build Agent to the server and sets up an `init` script to start the OS Build Agent after the reboot. **18** When the OS installation completes, Anaconda reboots the installation client, which will boot from the newly installed OS.
Phase Three	**19** On entering multi-user mode, the OS Build Agent is invoked and contacts the OS Build Manager. **20** The Linux build script installs the Opsware Agent. **21** The Linux build script exits. The OS installation section of provisioning is complete.

About Linux Build Customization Scripts

The Linux build script runs a single installation hook that gives you the ability to customize the Linux build process before Anaconda loads.

The installation hook is run in a RAM disk right before the installation program runs but after the network has been brought up.

Requirements for Linux Build Customization Scripts

To use a build customization script for Linux, you must meet the following requirements:

- Create an executable script and name it `run`.
- Include the `run` script in an archive file in `tar.gz` format. Include the script at the top level of the archive. During OS provisioning, the `tar.gz` archive is unpacked on the installation client and the script is executed.
- The `run` script is unpacked in its own directory with the other files in the archive. This directory serves as the current working directory when the `run` script is invoked. Based on this fact, correctly refer to the other files in the archive. For example, unpacking and invoking the `run` script follows this general process:

```
mkdir /tmp/installhook
cd /tmp/installhook
tar -xzf hook.tgz
./run 2>&1
```

- The `run` script should not take an exceptionally long time to complete, otherwise the OS provisioning process might time out.
- The `run` script must exit normally. If the script exits with a non-zero value, the OS provisioning process will end.
- The `run` script must have execute permissions to function properly.

Microsoft Windows Build Process

Table 3-7 describes in detail the exact steps that occur in the OS Provisioning Subsystem to provision an installation client with Windows.

A user initiates the build process with Steps 1 and 6. The rest of the build process steps happens automatically in the OS Provisioning Subsystem.

Table 3-7: Microsoft Windows Build Process

PHASE	BUILD PROCESS STEPS
Pre-installation	1 A user boots an installation client over the network by using a PXE network bootstrap program or by using the Windows Boot Floppy. 2 The user selects `Windows` from the boot menu on the console for the PXE network bootstrap program. 3 PXE boots the Windows Opsware OS Build Agent over the network. 4 The Opsware OS Build Agent prompts the user to create a FAT boot partition on which to install Windows. 5 The Opsware OS Build Agent collects pertinent hardware information and registers the information with the Opsware System. The server is ready to be provisioned and is available for selection from the Server Pool in the Opsware Command Center.
Phase One	6 The user selects an DOS server from the Server Pool list in the Opsware Command Center and assigns a Windows OS definition or a Windows template to the server. 7 The Windows build script mounts the Windows installation media as indicated by the Media Resource Location (MRL). 8 The Windows build script initiates Windows unattended setup. 9 The Windows build script waits for Windows unattended setup to complete and Windows to boot for the first time.
Phase Two	10 Windows boots for the first time. 11 If a build customization script was specified in the OS definition, it is executed by the Windows build script. 12 The Windows build script installs the Opsware Agent. The Windows build script exits and Phase Two is complete.

About Windows Build Customization Scripts

The Windows build script includes one installation hook that runs after the Windows OS is installed but before the Opsware Agent is installed on the server.

The installation hook must be packaged as a Microsoft cabinet file. During the provisioning process, the cabinet file is downloaded to the server being provisioned and extracted into a private temporary directory.

The OS Provisioning Subsystem expects to find a file named `run.bat` in the top level directory of the cabinet archive. If the file is found, the OS Provisioning Subsystem executes the `run.bat` file in a command shell and returns the output of the command to the Opsware Command Center.

If running the `run.bat` file returns a non-zero exit code, the OS Provisioning Subsystem detects the failure and ends the build process for that server.

A customer can use the hook as an opportunity to perform common post OS-installation tasks for Windows, such as modifying the Windows registry or applying security templates.

Working with OS Definitions

This section provides information on OS definitions within the Opsware System and contains the following topics:

- About Conditional Packages for Solaris
- Overview of Installation Order for Solaris and Linux
- About Hardware Signature Files for Windows
- Defining an Operating System
- About Editing OS Definitions
- Changing the Properties for an OS Definition
- Modifying the Way an OS Is Installed on Servers
- Modifying Which Packages an OS Definition Installs
- Viewing the History of Changes for an OS Definition
- Deleting an OS Definition

About Conditional Packages for Solaris

A metacluster specified in a JumpStart profile can include conditional packages. Conditional packages are packages that the Solaris installation program might (or might not) install during JumpStart. The Solaris installation program determines which packages to install based on the hardware attributes of the server being provisioned. For example, the presence of a specific graphics card would cause the drivers for that card to be installed.

When you upload a Solaris profile in the Prepare Operating System Wizard, the OS Provisioning Subsystem extracts the list of packages specified in the profile and displays them on the Review Packages page. The Review Packages page does not display Solaris conditional packages because the Opsware System cannot determine, at that time, whether the conditional packages will be installed.

You can specify to always install conditional packages by adding them to the Packages List. Adding packages to the Packages List does not change the Solaris profile. The Opsware System installs the packages even if the JumpStart Installer does not install them.

See “Defining an Operating System” on page 212 in this chapter for more information.
See “About Editing OS Definitions” on page 216 in this chapter for information about adding packages to or removing packages from the Packages List.

Overview of Installation Order for Solaris and Linux

For Sun Solaris, Red Hat Linux, and SUSE Linux, installation order is defined by the vendor. These dependencies control the order that software packages are installed during JumpStart, Kickstart, and YaST2.

However, the Opsware Command Center provides the ability to specify additional OS packages to install on servers after JumpStart or Kickstart or YaST2 completes. You can specify installation order for these additional packages. You set the package installation order when you define the OS.

See “Defining an Operating System” on page 212 in this chapter for information about how to specify the installation order while reviewing packages in the Prepare Operating System Wizard.

About Hardware Signature Files for Windows

A Windows response file contains information that is applicable to any hardware make and model. The remaining part of the configuration file is hardware-specific, taking into account differences between specific models of servers.

The generic part of the response file specifies how to install and configure the Windows OS. Typically, the hardware-specific part specifies hardware dependent configuration for devices such as mass storage.

Based on the hardware you expect to provision, you can upload hardware-specific files for each Windows OS definition. You map a signature for that hardware to the correct hardware-specific profile. The OS Provisioning Subsystem selects the correct Hardware Signature file at build time based on the hardware signature of the server that is about to be provisioned.

Certain x86-processor-based hardware requires pre-installation configuration of the OS. You usually perform this configuration by running vendor-supplied utilities with certain parameters. Because the utilities are hardware specific, you can script these configuration steps by using a Hardware Signature file.

Utilities referenced by the Hardware Signature file must be accessible through the network during build time.

Using Hardware Signatures is not required for Sun Solaris or Red Hat Linux operating systems because Solaris and Linux distributions do not need to be tailored for particular hardware models.

Defining an Operating System

The Prepare Operating System Wizard helps you define an OS installation that you can use during the OS provisioning process.

Perform the following steps to define an operating system:

1. From the Opsware Command Center home page, click the Prepare OS link in the Tasks panel.

   Or

   From the navigation panel, click Software ➤ Operating Systems. The Operating Systems page appears. Click the Prepare OS button.

The Prepare Operating System Wizard appears, as Figure 3-3 shows.

Figure 3-3: Overview Page in the Prepare Operating System Wizard

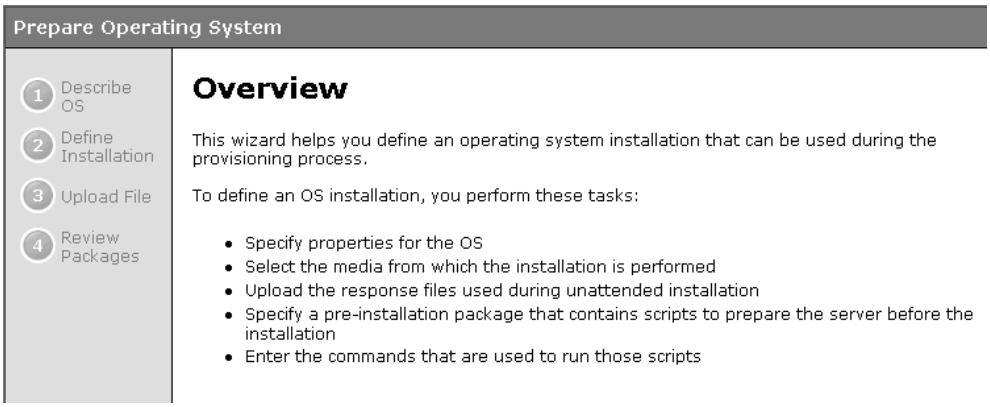

2 Describe the OS by specifying the following information:

- **(Required)** Name – sets the display name for the OS.
- **(Required)** Customer – associates the OS with a specific customer; to set up the OS for use by all customers, select Customer Independent.
- **(Required)** OS Version – sets the version of the OS (selected from the pre-populated list of the operating systems that the Opsware System supports).
- (Optional) Description – provides a long text description; using the description to identify the platform and hardware support is recommended.

3 Click the Next button. The Define Installation page appears, as Figure 3-4 shows.

Figure 3-4: Define Installation Page in the Prepare Operating System Wizard

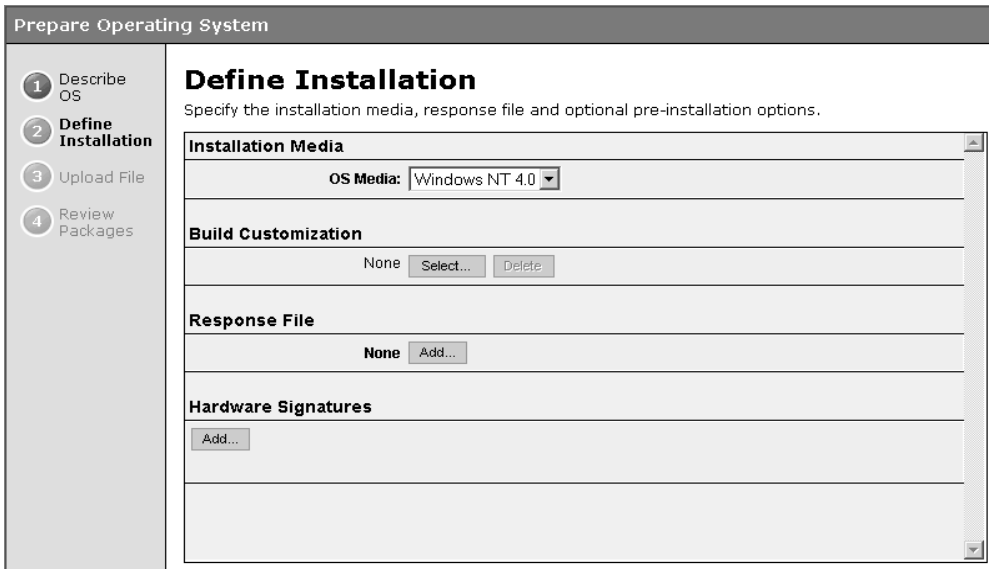

4 Define the installation by specifying the following information:

- (Required) OS Media – sets the MRL for the OS (select one MRL from the pre-populated drop-down list).

  See "OS Media Management" on page 185 in this chapter for information about how to set up OS media so that the MRL for the media appears in the drop-down list.

- (Required) A configuration file – indicates a JumpStart profile, Kickstart configuration file, YaST2 autoinst.xml, or Windows response file to upload into the OS Provisioning Subsystem.

  The file that you upload can have any filename. However, the OS Provisioning Subsystem renames the file with the correct filename for use by the vendor installation program.

- (Optional – Windows only) Hardware Signatures – defines the list of hardware that the OS supports.

  Click the Add button to open the Add Hardware Signature Setting window. The Applies To field is pre-populated with the hardware makes and models that have been successfully built so that they appear in the Managed Server list.

You can add multiple Hardware Signature files to a Windows OS definition.

- (Optional) Build Customization Script – customizes the way the build process operates for that OS (select a file from the popup window).

  The way you can customize the build process is specific to each build script. You must follow the requirements for build customization scripts to use this feature. Scripts appear in the popup window after you upload them through the Opsware Command Center.

  See “Build Customization Scripts” on page 199 in this chapter for more information.

5 Click the Upload button.

The Opsware System creates the OS definition and uploads the configuration file (and parses packages for Sun Solaris and Red Hat and SUSE Linux). A progress bar appears that shows the progress of the OS preparation process.

6 Click the Next button to review the packages. A page appears that shows the list of packages, as Figure 3-5 shows.

Figure 3-5: Review Packages Page in the Prepare Operating System Wizard

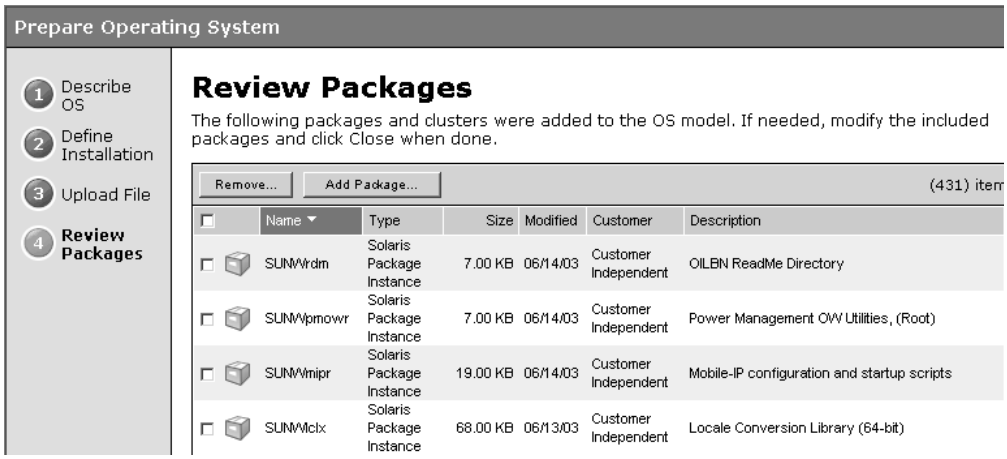

For Solaris and Linux, the list shows the vendor packages that were specified in the Solaris profile or Linux configuration file.

For Windows, the list is empty because you cannot specify specific packages in the Windows response file. You can add packages to the Windows OS definition by clicking the Add Package button.

7 (Optional) Click the Remove or Add Package buttons to modify the list of software that the OS definition installs or to change the installation order.

See "About Conditional Packages for Solaris" on page 211 in this chapter for information about how to ensure Solaris conditional packages are always installed.

8 Click the Close button to end the Wizard.

About Editing OS Definitions

You can edit an OS definition in the following ways:

- By changing the properties for the OS, such as which customer can use the OS definition to provision servers
- By modifying the way that the OS is installed on servers by changing the configuration file or customizing the way the build process works for that OS definition
- By adding custom attributes to the OS definition to override default values in the build process

  See "Default Values for the OS Build Process" on page 221 in this chapter for more information.

  See "Managing Nodes on the Software Tree" on page 315 in Chapter 6 for information about how to set custom attributes for software nodes.
- By modifying which packages are installed with the OS definition

  Modifying the list of packages in an OS definition does not change the configuration file uploaded for the OS definition. The Opsware System installs the packages after the OS installation technology (Sun Solaris JumpStart or Red Hat Linux Kickstart or SUSE Linux YaST2) installs the packages specified in the configuration file. For Microsoft Windows, the response file cannot specify specific packages to install; however, you can add Windows packages so that the Opsware System installs them with the OS.
- By setting up configuration tracking for an OS definition

  See "Configuration Tracking Policies" on page 495 in Chapter 11 for information about how to set a configuration tracking policy for the OS definition.

Changing the Properties for an OS Definition

Perform the following steps to change the properties for an OS definition:

1 From the navigation panel, click Software ➤ Operating Systems. The Operating Systems page appears.

2 Click the display name of the OS that you want to edit. The Edit Operating System page appears.

3 Click the Properties tab (see Figure 3-6) and modify the following settings:

- Name – sets the display name for the OS.
- Description – provides a long text description of the OS.
- Customer – associates the OS with a specific customer.

  If an OS definition is used (a server is provisioned by using the OS definition), you cannot change the customer association for that OS definition.

Figure 3-6: Properties Tab for an OS Definition in the Opsware Command Center

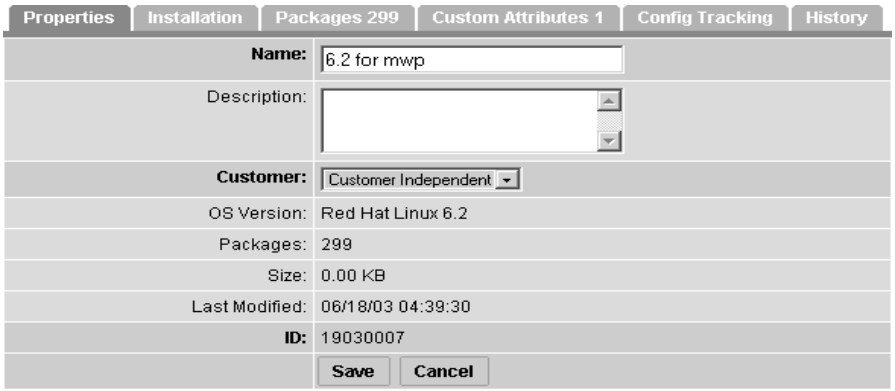

4 Click the Save button.

Modifying the Way an OS Is Installed on Servers

Perform the following steps to modify the way an OS is installed on servers:

1 From the navigation panel, click Software ➤ Operating Systems. The Operating Systems page appears.

2 Click the display name of the OS that you want to edit. The Edit Operating System page appears.

3 Click the Installation tab. The installation resources defined for the OS definition appear, as Figure 3-7 shows.

Figure 3-7: Installation Tab for an OS Definition in the Opsware Command Center

4 Modify the following settings:

- Installation Media – sets the MRL for the OS. Click the Select button and select an OS media from the list in the popup window.
- Build Customization Script – customizes the way the build process operates for that OS. Click the Select button and select a build customization package from the list in the popup window.

  Scripts appear in the popup window after you upload them through the Opsware Command Center.
- Configuration file – indicates a JumpStart profile, Kickstart configuration file, YaST2 configuration file, or Windows response file to upload into the OS Provisioning Subsystem. Click the Upload button and enter the filename or browse to the file.

  The file that you upload can have any filename. However, the OS Provisioning Subsystem renames the file with the correct filename for use by the vendor installation program.
- Hardware Signatures for Windows *only* – defines the list of hardware that the OS supports. Click the Add button and select the hardware signature that you want to include in the OS definition.

  Hardware signatures appear in the list box after a server with that make and model are successfully built so that it appears in the Managed Server list.

5 Click the Save button.

Modifying Which Packages an OS Definition Installs

Perform the following steps to modify which packages an OS definition installs:

1. From the navigation panel, click Software ➤ Operating Systems. The Operating Systems page appears.
2. Click the display name of the OS that you want to edit. The Edit Operating System page appears.
3. Click the Packages tab. The list of packages that the OS definition installs appears, as Figure 3-8 shows.

Figure 3-8: Packages Tab for an OS Definition in the Opsware Command Center

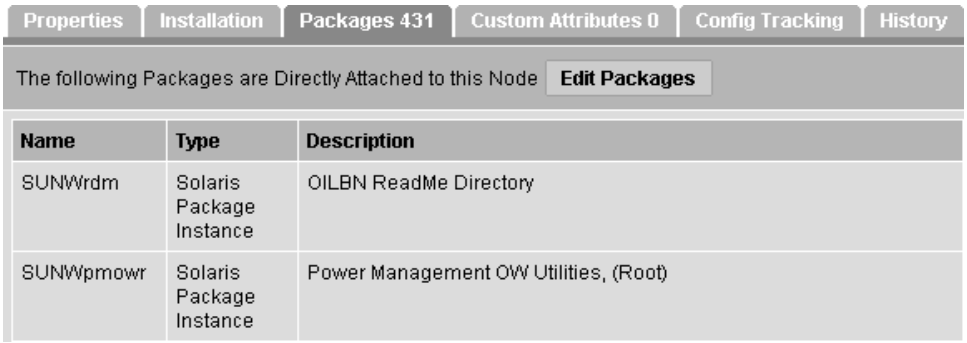

4. Click the Edit Packages button. The Software Directly Attached page appears.
5. To add a package for installation, click the Add Software button and specify or search for the software package that you want to add to the list.
6. To remove software packages, select them in the list and click the Remove Software button. The packages are deleted from the list in the page but are not actually removed from the OS definition until you click the Save Edits button.
7. To change the order in which the packages are installed on servers, select the package that you want installed in a different order and click the Up or Down arrows.
8. Click the Save Edits button.

Viewing the History of Changes for an OS Definition

By default, the OS Provisioning Subsystem maintains information about the changes to OS definitions for 180 days.

The following actions create an entry in the History tab for an OS definition:

- The customer association is changed for the OS definition.
- A server uses the OS definition to install an OS.
- Packages are added to or removed from the Package List in the OS definition.

Perform the following steps to view the history of changes for an OS definition:

1. From the navigation panel, click Software ➤ Operating Systems. The Operating Systems page appears.
2. Click the display name of the OS for which you want to review the history of changes. The Edit Operating System window appears.
3. Click the History tab. The list of events and changes appears, as Figure 3-9 shows.

Figure 3-9: History Tab for an OS Definition in the Opsware Command Center

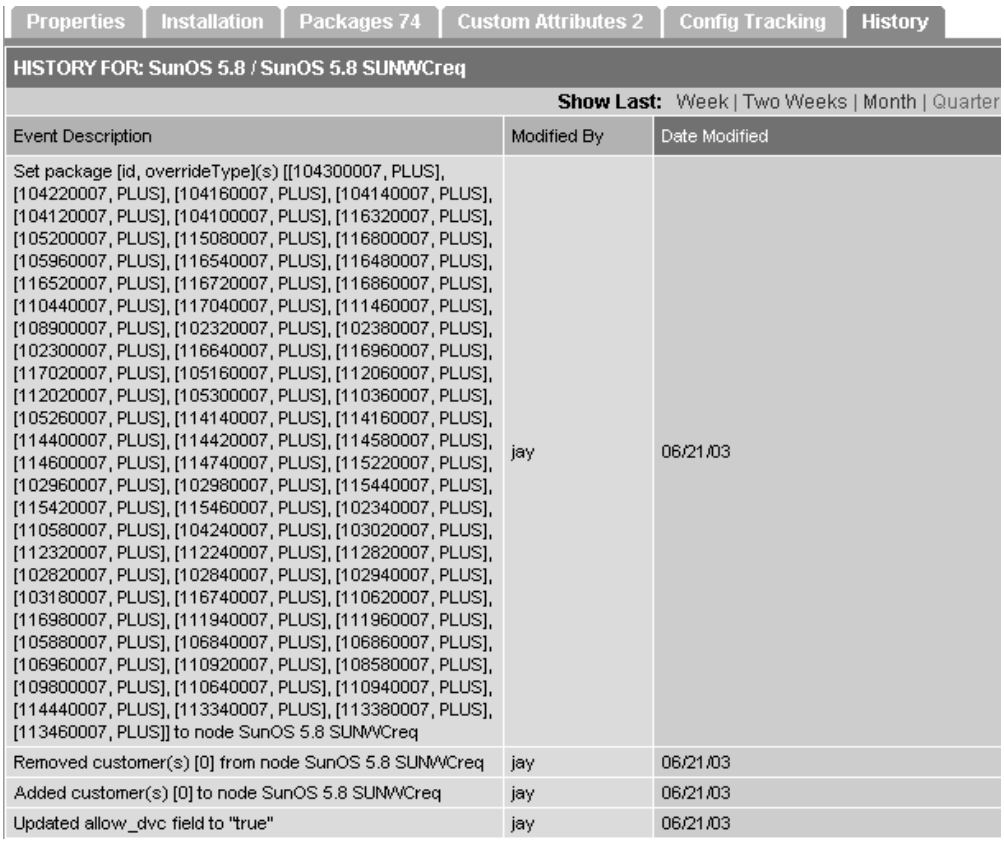

Deleting an OS Definition

If a server is using the OS definition or the OS definition is included in a template, you cannot delete it.

Perform the following steps to delete an OS definition:

1. From the navigation panel, click Software ➤ Operating Systems. The Operating Systems page appears.
2. Select the OS that you want to delete.
3. Click the Delete button. (If a server has used the OS definition or the OS definition is included in a template, a warning message appears.)

   The list of OS definitions re-appears.

Default Values for the OS Build Process

This section provides information on default values for the OS build process within the Opsware System and contains the following topics:

- Default Values for the OS Build Process Overview
- Custom Attributes for Sun Solaris
- Custom Attributes for Linux
- Custom Attribute for Microsoft Windows
- Adding Custom Attributes to an OS Definition

Default Values for the OS Build Process Overview

In addition to the customization provided by using build customization scripts, each build script uses custom attributes.

Opsware Command Center provides a data management function by allowing users to set custom attributes for servers. These custom attributes include setting miscellaneous parameters and named data values. Users can write scripts that use these parameters and data values when performing a variety of functions, including network and server configuration, notifications, and CRON script configuration.

See "Custom Attributes Set for the Environment" on page 349 in Chapter 6 for information about custom attributes.

For OS provisioning, the Opsware System uses custom attributes to pass specific information to each build script to configure aspects of the installation process.

You can edit an OS definition to override the default values used by the build process. You override these default values by setting custom attributes for the OS definition.

See "Adding Custom Attributes to an OS Definition" on page 224 in this chapter for information about the steps to set custom attributes for an OS definition.

Custom Attributes for Sun Solaris

The build script for Solaris OS provisioning uses a number of custom attributes. Several of these custom attributes correlate with an equivalent setting that would be defined normally by a Solaris `sysidcfg` file.

You cannot modify the `sysidcfg` file that the OS Provisioning Subsystem used. However, you can override specific values specified in the default `sysidcfg` file. You can set custom attributes for a Solaris OS definition in the Opsware Command Center.

The custom attributes correspond to the equivalent keywords in the `sysidcfg` file. See Table 3-8.

Table 3-8: Sun Solaris Custom Attributes

KEYWORD	DESCRIPTION
`root_password`	Sets the encrypted value for the password on an installation client. One way to obtain an encrypted value is by using `/etc/shadow`. If a value is not set, the system will not have a root password.
`timezone`	Sets the time zone in which to configure the installation client (sets `TZ` in `/etc/default/init`). The directories and files in the directory `/usr/share/lib/zoneinfo` provide the valid time zone values. By default, the timezone value is UTC. For example, the time zone value for Pacific Standard Time in the United States is `US/Pacific`. You can also specify any valid Olson time zone.

Table 3-8: Sun Solaris Custom Attributes

KEYWORD	DESCRIPTION
`system_locale`	Sets the language in which to configure the installation client (sets `LANG` in `/etc/default/init`). Valid locale values are installed in `/usr/lib/locale`. If you set this attribute, you should also use the `locale` keyword in the operating system profile so that the appropriate locale is installed. By default, the value for this keyword is `system_local=C`.
`required_patches`	This keyword is reserved by the Solaris build script. Using it might cause the installation process to fail. To specify required patches, include them with the OS definition in a template. See "Opsware Patch Management" on page 403 in Chapter 8 for information about how patch management works in the Opsware System.

Custom Attributes for Linux

You can use custom attributes to specify additional arguments to the kernel under which the install is running. By specifying these arguments, you can accomplish tasks such as pinning interfaces. The OS Provisioning Subsystem appends the contents of the custom attribute to the kernel arguments for the kernel that is installing the OS.

Set a custom attribute for the OS definition (edit the OS definition and click the Custom Attributes tab). The custom attribute must have the name `kernel_arguments`.

The kernel arguments are separated by spaces (like they are when you type them after the boot prompt for the CD-ROM or DVD). For example:

```
name=value jones=barbi
```

To have the kernel arguments persist after the base OS is installed, you must set them in the uploaded configuration file. Setting kernel arguments by using custom attributes only allows you to create a completely automated installation (as if you were installing the OS from CD-ROM or DVD).

Custom Attribute for Microsoft Windows

For a Windows OS definition, you can set a value for the `timeout` custom attribute. Setting this value controls the timeout value after an error.

Set this value to the amount of time (in minutes) it takes Windows setup to complete.

If Windows setup does not complete in the specified amount of time, the OS installation will fail with a timeout error. By default, this value is set to 60 minutes.

Adding Custom Attributes to an OS Definition

Perform the following steps to add custom attributes to an OS definition:

1. From the navigation panel, click Software ➤ Operating Systems. The Operating Systems page appears.
2. Click the display name of the OS that you want to edit. The Edit Operating System page appears.
3. Click the Custom Attributes tab. The list of custom attributes specified for the OS definition appears, as Figure 3-10 shows.

Figure 3-10: Custom Attributes Tab for an OS Definition in the Opsware Command Center

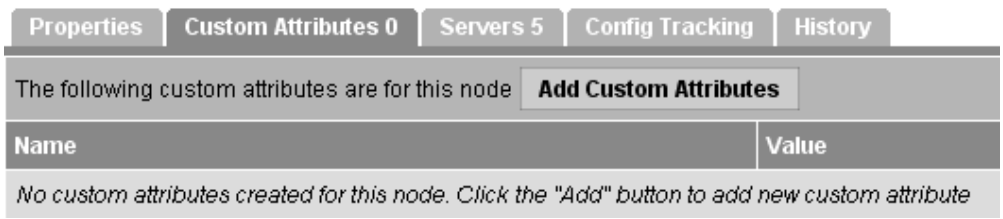

If the OS Definition contains custom attributes, the Edit Custom Attributes button appears in the page. Click the Edit Custom Attributes button to add new attributes and edit existing ones.

4. Click the Add Custom Attribute button. A page appears in which you can enter the names and values for custom attributes.
5. Enter a name and a value for the custom attribute.
6. Click the Save button. The list of custom attributes set for the OS definition reappears. The new custom attribute is added to the list.

Including OS Definitions in Templates Overview

In the Opsware System, you can create templates to automate building complete server baselines.

By using Opsware templates, system administrators can define and provision servers with standard configurations, sometimes called *server baselines*. For example, system administrators can define a Windows baseline for databases, a Windows baseline for Web and application servers, and a different Windows baseline for messaging servers. Each baseline can include a different variant of the following items:

- A base operating system
- The latest operating system patches
- System utilities such as SSH or PC Anywhere
- Security tools such as TripWire or anti-virus software
- Widely shared system software such as the latest Java Virtual Machine

Using templates, which are pre-packaged collections of installable software, users can provision an entire software stack, including the base operating system, the latest set of operating system patches, system utilities such as SSH and the latest JVMs, middleware including databases, Web servers, and application servers, and so on, up to the custom business applications that the server ultimately runs.

See “Working with Templates” on page 353 in Chapter 6 for information about how to create and edit templates to include OS definitions.

Hardware Support in OS Provisioning

This section provides information on hardware support in OS provisioning within the Opsware System and contains the following topics:

- Hardware Support in OS Provisioning Overview
- PXE Images for Windows and Linux Overview
- Windows and Linux Boot Floppies Overview
- About NIC Support in Windows Floppy Images
- Adding NIC Support to a Windows Floppy Image
- Sample Mapfile

- Sample Mapfile for an Intel 8255x-based PCI Ethernet Adapter
- Prerequisites for Creating Windows Floppy Images
- Creating a Windows Boot Floppy
- Updating the PXE Image for Windows
- Adding Hardware Support to a Linux Build Image
- Creating a Linux Boot Image

Hardware Support in OS Provisioning Overview

The OS Provisioning Subsystem ships with support for a broad range of hardware platforms out of the box. Additionally, customers can add support to the OS Provisioning Subsystem for hardware models not initially supported.

Preparing the OS Provisioning Subsystem to provision new hardware is straightforward. The process involves packaging and uploading system utilities that the server manufacturer provided into the Opsware System.

At a minimum, the boot processes for Windows and Linux (Opsware Boot Floppies or CDs and the PXE boot system) must be updated to support the new hardware.

Additionally, the Linux build images themselves might need to be updated to support the new hardware.

See "Adding NIC Support to a Windows Floppy Image" on page 229 in this chapter for more information.

See "Adding Hardware Support to a Linux Build Image" on page 233 in this chapter for more information.

PXE Images for Windows and Linux Overview

The OS Provisioning Subsystem supports booting new x86-processor-based servers with the Preboot Execution Environment protocol (PXE).

When the Opsware System was installed with the Opsware Installer, a default boot image was added to the PXE system for Windows and for Linux so that new servers can be booted for the first time over the network. The boot image is used by the Opsware System as the second stage PXE image for PXE network bootstrap programs such as PXELinux.

For Linux, the Opsware System includes a boot image that contains the `bootnet.img` CD for Red Hat Linux AS 3.0. The image has changes to the `syslinux.cfg` and `boot.msg` files; however, the kernel and `initrd.img` are identical to the files on the Linux OS media.

The default boot images include common NIC drivers for many hardware makes and models. The Opsware System uses these NIC drivers to boot new x86-processor-based servers for the first time.

The Linux PXE image and Linux Boot CD contain the same NIC drivers included on the Red Hat Linux AS 3.0 installation CD-ROM or DVD.

Table 3-9 shows the Windows boot floppy image, which includes the following set of common NIC drivers.

Table 3-9: NIC Drivers Included with the Windows Boot Image

DRIVER NAME	DESCRIPTION
B57	Broadcom NetXtreme Gigabit Ethernet NDIS2 Driver v5.20 (021025)
DC21X4	Digital 2104x/2114x 10/100 mbps Ethernet Controller v3.00
E1000	Intel 8254X Based Adapter (pro/1000 gigabit) v1.28 040302
E100B	Intel(R) PRO PCI Driver v4.35 042902
EL59X	3Com DOS NDIS driver for 3C59X Family Adapters v1.2f
EL90X	3Com Etherlink PCI DOS NDIS driver v5.2.2
ELNK3	3Com DOS EtherLink 10 ISA (3C509b) Network Driver v3.1
ELPC3	3Com Megahertz Ethernet PC Card 589E DOS Netw. Driver v1.9.002
ELPC575	3Com Megahertz 10/100 LAN CardBus PC Card DOS NDIS driver v3.4b
FA31X	Netgear FA310TX Fast Ethernet PCI Adapter
FETND	VIA Rhine Family Fast Ethernet Adapter Driver v4.05
N100	Compaq Fast Ethernet and Gigabit NDIS 2 NIC Drivers 7.0a (25Jan02)
NE2000	Microsoft NE2000 NDIS Driver
NETFLX3	Compaq NetFlex-3 DOS NDIS 2.02 driver
PCNTND	AMD PCNet Family Ethernet Adapter NDIS v2.0.1 MAC Driver v3.12
RTSND	Realtek RTL8139/810X Family PCI Fast Ethernet v3.23 07/28/99

Table 3-9: NIC Drivers Included with the Windows Boot Image

DRIVER NAME	DESCRIPTION
SMC9432	SMC EtherPower II 10/100 (9432TX) v1.02c (970605)

If the NIC drivers that you need for your environment are not included in the default set, you must perform the following tasks:

- Add them to the boot image for Windows, Linux, or both.
- Update the Windows or Linux boot image in the PXE system with the new boot images.

Windows and Linux Boot Floppies Overview

For environments with servers that do not support network boot technology, Opsware supports floppy- or CD-based booting.

You can create a Windows Boot Floppy from the default boot image for Windows. The Opsware System includes the Opsware Build Image Administrator, a tool for creating a Boot Floppy for Windows.

For Linux, you can download the boot image for Linux from the Opsware Command Center. Search for the package name `bootfloppy` and package type `Unknown` in the Packages section of the Opsware Command Center. Download and create a Linux Boot CD from this image.

See "Creating a Linux Boot Image" on page 234 in this chapter for more information.

About NIC Support in Windows Floppy Images

The Opsware System includes a default set of common NIC drivers for many hardware makes and models. If the NIC drivers you need for your environment are not included in the default set, you must add them to the boot floppy image for Windows.

The Opsware Build Image Administrator has the ability to dynamically detect your server's PCI network adapter. It does this by scanning the PCI bus for PCI information and comparing the information against each entry in a driver catalog until it finds a match. The driver catalog is constructed each time you create a build image with the Opsware Build Image Administrator.

Each properly formatted cabinet file in the directory `\content\drivers\ndis` under the Opsware Build Image Administrator directory is included as an entry in the driver catalog.

Adding NIC Support to a Windows Floppy Image

Before you perform this procedure, you must obtain the appropriate NDIS2 network drivers and `protocol.ini` file for the card from the manufacturer of the card.

Perform the following steps to add NIC support to a Windows floppy image:

1. Create a temporary working directory for accumulating files that becomes part of the cabinet file.
2. Place NIC drivers and the `protocol.ini` files in the temporary directory.
3. Create a text file called `ndis.pci` in the temporary working directory.
4. Using a PCI bus scanner, determine the PCI vendor ID and device ID of the NIC card.

   For example, the 8255x-based PCI Ethernet Adapter from Intel has vendor ID 8086 and device ID 1229.
5. Using the vendor ID and device ID you obtained for the NIC, construct the mapfile `ndis.pci`.

   In the mapfile, lines that begin with a semicolon (;) are treated as comments and ignored.

   The sample mapfile in this chapter contains comment lines so that you can use it as a header for your mapfile.
6. Create a file named `ndis.txt` in the temporary directory that contains the following single line of text:

   ```
   [basename of cabinet file] “[Driver description string]”
   ```

   The information in this file is used to make up a selection list if the PCI adapter cannot be automatically detected.

   Example `ndis.txt` file for the `E100B.CAB`:

   ```
   E100B   “Intel(R) PRO PCI Driver v4.35 042902”
   ```
7. Create the cabinet file by using cabarc and copy the cabinet file to the directory `.\content\drivers\ndis` under the Opsware Build Image Administrator directory. (Cabarc is a Microsoft utility that creates, extracts, and lists the contents of cabinet files.)

   ```
   E:\temp\temp_cab>cabarc N e100b.cab *
   Microsoft (R) Cabinet Tool - Version 5.2.3718.0
   Copyright (c) Microsoft Corporation. All rights reserved.
   Creating new cabinet 'e100b.cab' with compression 'MSZIP':
   ```

```
-- adding e100b.dos
-- adding e100b.ini
-- adding ndis.pci
-- adding ndis.txt
Completed successfully
```

Sample Mapfile

Modify the contents of this sample mapfile. This sample mapfile contains comment lines so that you can use it as a header in the mapfile that you create.

```
; Mapfile for PCISCAN "PCI PnP for DOS"
;
; Syntax:
;   ret="string_to_return"
;   ven=<vendorID> ["Vendor description"]
;   dev=<deviceID> ["Device description"]
;
; Example:
;   ret="aspi8dos.sys"
;   ven= 9004 "Adaptec"
;   dev= 7078 "Adaptec AIC-7870 PCI SCSI Controller"
;      7178 "Adaptec AHA-294X/AIC-78XX PCI SCSI Controller"
;      7278 "SCSI Channel on Adaptec AHA-3940/3940W PCI SCSI
;      Controller"
;      7478 "Adaptec AHA-2944 PCI SCSI Controller"
;      7578 "SCSI Channel on Adaptec AHA-3944 PCI SCSI
;      Controller"
;      7678 "Adaptec AIC-7870 based PCI SCSI Controller"
```

Sample Mapfile for an Intel 8255x-based PCI Ethernet Adapter

```
ret="E100B"
ven=8086 "Intel"
dev=1002 "PRO 100 Mobile Adapters"
   1031 "PRO/100 VE Network Connection"
   1032 "PRO/100 VE Network Connection"
   1035 "PRO/100 VM Network Connection"
   1036 "82562EH based Phoneline Network Connection"
   1038 "PRO/100 VM Adapter"
   1039 "PRO/100 VE Network Connection"
   103b "PRO/100 VM Network Connection"
   103c "PRO/100 VM Network Connection"
   103d "PRO/100 VE Network Connection"
   103e "PRO/100 VM Network Connection"
   1059 "PRO 100 Mobile Adapters"
   1229 "8255x-based PCI Ethernet Adapter (10/100)"
```



```
2449 "PRO/100 VE Desktop Adapter"
2459 "82562 based Fast Ethernet Connection"
245d "82562 based Fast Ethernet Connection"
```

Prerequisites for Creating Windows Floppy Images

The Opsware Build Image Administrator is used to create a Windows Boot Floppy that installs the OS Build Agent on servers. The Opsware Build Image Administrator is packaged with MSI.

You must meet the following requirements to use the Opsware Build Image Administrator:

- The machine on which the Opsware Build Image Administrator is installed must have a Python interpreter installed. You can obtain a Python interpreter from ActiveState.
- The Opsware System includes a default set of common NIC drivers for many hardware makes and models. If the NIC drivers that you need for your environment are not included in the default set, you must add them to the floppy image.

  See “Adding NIC Support to a Windows Floppy Image” on page 229 in this chapter for more information.

Creating a Windows Boot Floppy

Perform the following steps to create a Windows boot floppy:

1. Download the MSI package that contains the Opsware Build Image Administrator by downloading the file `opswbia-<version>-0.msi` from the Opsware Command Center.

   Where `<version>` is the latest version of the Opsware Build Image Administrator tool for the release of the Opsware System installed at your facility. Only one version of the Opsware Build Image Administrator tool is available on the Software Repository.

   See “Downloading a Package” on page 306 in this chapter for more information.

2. Install the MSI package that contains the Opsware Build Image Administrator tool on a Windows server that has a Python interpreter.

   By default, the Opsware Build Image Administrator is installed in the following directory:

   ```
   %SystemDrive%\Program Files\OPSWBIA
   ```

3. Change directories to the Opsware Build Image Administrator installation directory:

   ```
   \Program Files\OPSWBIA >
   ```

4. Insert a floppy into drive A.

5. Run the python script `mkimage.pyc`.

   ```
   \Program Files\OPSWBIA > python mkimage.pyc <options>
   ```

   If you do not enter any options, the Opsware Build Image Administrator creates a PXE build image file `dosopsw.1` in the current working directory. Enter the `-w` option to write the file `dosopsw.1` to a floppy.

Details: Options for the Opsware Build Image Administrator

You can use the options that Table 3-10 shows when you run the Opsware Build Image Administrator from the command line.

Table 3-10: Opsware Build Image Administrator Command Line Options

OPTION	DESCRIPTION
`-a` <*drive*>	Writes the boot floppy image to this drive (Default drive: `A`)
-c	Makes an el-torito bootable CD image in addition to the boot floppy image
`-d`	Enables debugging of the OS Build Agent in the generated image
`-f`	Formats the floppy first when writing to a floppy
`-h` <*host*>	Specifies the build hostname for the generated OS Build Agent (Default hostname: `buildmgr`)
`-i` <*file*>	Specifies the filename for the generated floppy image (Default filename: `dosopsw.1`)
`-n` <*directory*>	Specifies the directory where NDIS driver packages are located (Default directory: `./content/ndis`)
`-o` <*OS*>	Sets the OS for the boot floppy image (Default OS: `dos622`)
`-p` <*port*>	Sets the port for the OS Build Agent to use to contact the build host (Default port: `1017`)
`-t`	Performs a test image generation and does not execute any commands
`-w`	Writes the generated image to a floppy in the drive specified by the option `-a`

Updating the PXE Image for Windows

When the Opsware System was installed with the Opsware Installer, a image was added to the PXE system by default. You only need to update the PXE image when you have added support for additional NIC drivers to the image.

See "Adding NIC Support to a Windows Floppy Image" on page 229 in this chapter for more information.

After adding NIC support to the boot floppy image, install the floppy image by using `scp` to copy the floppy image file into the `/opt/OPSWboot/tftpboot` directory on the Opsware Build Server.

Adding Hardware Support to a Linux Build Image

You can modify the OS Provisioning Subsystem to add new hardware support to a Linux build image. To provision servers with a Linux OS, the Opsware System uses two types of Linux build images:

- A Linux Boot Image – the Opsware System uses a modified version of Red Hat Linux AS 3.0 as a bootstrap image. The Linux Boot Image is loaded on servers when they are booted up for the first time by using the Linux Boot CD or by using PXE. The server appears in the Server Pool list and is ready to be provisioned with an OS.
- A Linux Build Image that installs the target OS – the Opsware System uses this type of Linux Build Image to install the target Linux OS on servers.

To add new hardware support to a Linux Build Image, you must recompile the kernel and modules, and insert the modules into the `initrd.img` file and replace the kernel if it changed.

The Linux Build Images are located on the OS Build Manager host in the following directories:

```
/cust/buildscripts/linux/bi-<version>
```

Where `<version>` is the version of Linux.

When you modify the Linux Boot image, include the following options in the kernel:

```
CONFIG_PACKET=y
CONFIG_FILTER=y
```

Setting these options is required if you want to retrieve the Build Manager parameters from DHCP. The existing Linux Boot image is compiled with these options.

See the Red Hat Linux or SUSE Linux documentation for information about how to add hardware support.

Creating a Linux Boot Image

The Opsware System includes a command line utility, OPSWlinuxdbootiso, that you can use to create a Linux Boot Image on CD.

1. In the Opsware Command Center, search for the package name `OPSWlinuxdbootiso*` and operating system Red Hat Enterprise Linux AS 3.0. See "Searching for Packages" on page 290.
2. Download the package to a server or desktop running Linux. "Downloading a Package" on page 306 for information.
3. On the server or desktop where you downloaded the OPSWlinuxdbootiso utility, verify that version 1.10-4 of the mkisofs utility is installed.
4. From the following directory, run the mkcdrom.sh script:

   ```
   /opt/OPSWlinuxbootiso>./mkcdrom.sh
   ```

   ```
   Usage: ./mkcdrom.sh <outputiso> sin:/opt/OPSWlinuxbootiso>./
   mkcdrom.sh /tmp/boot.iso
   ```

5. At the prompts, enter the following information:
   - The IP address or hostname of Build Manager (default hostname: `buildmgr`).
   - The port for the OS Build Agent to use to contact the Build Manager [default: 1017].
   - The IP address or hostname of the Boot Server (default hostname: `buildmgr`).
   - The path to the media for the OS Build Agent (default path: `/opt/OPSWboot/kickstart`).
   - The server from which to run Linux Kickstart.

Running the OPSWlinuxdbootiso utility creates an iso file that you can write to CD-ROM.

Example: Usage of the OPSWlinuxdbootiso Utility

```
sin:/opt/OPSWlinuxbootiso>./mkcdrom.sh /tmp/boot.iso
Please enter IP or hostname of Build Manager[buildmgr]:
buildmgr.c07
Please enter bootagent port of Build Manager[1017]:
Please enter IP or hostname of Boot Server[buildmgr.c07]:
Please enter path to bootagent media[/opt/OPSWboot/kickstart]:
Please enter device you want to kickstart from
```



```
just press enter for default[]: eth0
buildmgr.c07
1017
buildmgr.c07
/opt/OPSWboot/kickstart
*****************************************
* Rewritting isolinux.cfg                *
*****************************************
*****************************************
* Building iso... /tmp/boot.iso
*****************************************
Size of boot image is 4 sectors -> No emulation
Total translation table size: 2048
Total rockridge attributes bytes: 0
Total directory bytes: 2048
Path table size(bytes): 26
Max brk space used 4000
1552 extents written (3 Mb)
sin:/opt/OPSWlinuxbootiso>
```


Chapter 4: Operating System Provisioning

IN THIS CHAPTER

This chapter provides the following information about installing operating systems on servers by using the OS Provisioning Subsystem:

- Supported Environments for OS Provisioning
- OS Provisioning
- OS Provisioning Process
- Hardware Preparation Overview
- Booting New Servers
- OS Installation with Opsware Command Center

Before users can install operating systems on servers with the OS Provisioning Subsystem, the operating systems must be defined in the Opsware System. The OS media must be made available in the Opsware System. Additionally, OS definitions must be created in the Opsware System.

Supported Environments for OS Provisioning

The OS Provisioning Subsystem supports installation of the following versions of Red Hat Linux, Sun Solaris, and Microsoft Windows operating systems:

- Red Hat Linux 7.1
- Red Hat Linux 7.2
- Red Hat Linux Advanced Server 2.1
- Sun Solaris 2.6
- Sun Solaris 7

- Sun Solaris 8
- Sun Solaris 9
- SUSE Linux Standard Server 8.0
- SUSE Linux Enterprise Server 8.0
- SUSE Linux Enterprise Server 9.0
- Windows NT 4.0
- Windows 2000
- Windows 2003

The OS Provisioning Subsystem does not provision HP-UX or AIX operating systems. However, you can integrate the Opsware System with Network Installation Management (NIM) to provision AIX and Ignite-UX to provision HP-UX. See Appendix C, "OS Installation Integration" for more information about how to integrate the Opsware System with HP-UX and AIX OS provisioning systems.

The OS Provisioning Subsystem supports a large variety of hardware models from different manufacturers out of the box. For hardware not supported out of the box by the OS Provisioning Subsystem, you can update the OS Provisioning Subsystem to provision new hardware.

See "Hardware Support in OS Provisioning" on page 225 in this chapter for information about how to extend the OS Provisioning Subsystem to support new hardware.

The OS Provisioning Subsystem works in floppy or CD environments or network-boot environments.

OS Provisioning

This section provides information on OS Provisioning within the Opsware System and contains the following topics:

- OS Provisioning Overview
- Permissions Required for OS Provisioning
- The Server Lifecycle for OS Provisioning

OS Provisioning Overview

In the OS Provisioning Subsystem, server provisioning is installation-based instead of image-based. The OS Provisioning Subsystem uses Red Hat Linux Kickstart, Sun Solaris JumpStart, and Microsoft Windows unattended installation to install operating systems on servers.

The OS Provisioning Subsystem is fully integrated with the Opsware System; users can install an OS on the following types of servers:

- A bare metal server that does not have an OS installed
- A server that the Opsware System already manages
- A server that is running in the environment but the Opsware System does not manage it

The OS Provisioning Subsystem facilitates installing operating systems on servers in the following ways:

- Each OS definition in the OS Provisioning Subsystem contains all the information necessary to build and maintain a server with that OS.
- When installing an OS on a server, the OS Provisioning Subsystem displays information about server hardware and which operating systems are compatible with that hardware architecture.

Permissions Required for OS Provisioning

In the Opsware Command Center, users access only the areas of functionality relevant to their responsibilities in the managed server environment. If access is allowed to a functional area in the Opsware Command Center, the link for that function displays in the navigation panel and on the home page.

To provision servers with operating systems in the Opsware Command Center, you must have the following permissions to perform the tasks that this chapter describes:

- Access to the Server Pool so that you can view which servers are available and waiting to be provisioned with an OS
- Access to templates so that you can view and edit templates to include OS definitions
- Access to the Install Operating System Wizard so that you can use this Wizard to install operating systems on servers

- Access to specific customers if you want to associate the servers you provision with specific customers

To obtain the required permissions to perform OS provisioning, contact your Opsware administrator.

The Server Lifecycle for OS Provisioning

The Opsware System is designed to enable multiple teams to work together to provision servers. The OS Provisioning Subsystem allows IT teams to separate the tasks of readying servers for provisioning (such as racking servers, connecting them to power and a network) from provisioning the servers with operating systems.

Someone mounts a new server in a rack and connects it to the Opsware build network. Then they boot the server for the first time by using an Opsware Boot Floppy or CD or by using the network. At a later time, a different system administrator can select the available server from the Server Pool list and provision it with an OS. In the *available* state, servers do not have an OS installed and might not have access to disk resources.

See Table 4-1 for the lifecycle values for servers. During OS provisioning, servers progress through the following Opsware lifecycle state changes:

Unprovisioned ➤ Available ➤ Installing OS ➤ Managed

Table 4-1: Opsware System Lifecycle Values for Servers

OPSWARE LIFECYCLE VALUE	DESCRIPTION
Server Pool Values	
Available	Indicates a server on which the OS Build Agent was installed and is running, but an OS has not been installed on the server. The OS Build Agent is a small agent that can run in the memory of the bare metal server. See "About Installation of OS Build Agents" on page 251 in this chapter for more information.
Installing OS	Indicates that a user is installing an OS on the server. The server stays in the Server Pool list until the installation process finishes successfully; then, the server moves to the Managed Server list.

Table 4-1: Opsware System Lifecycle Values for Servers

OPSWARE LIFECYCLE VALUE	DESCRIPTION
Build Failed	Indicates a server on which the OS Build Agent was installed and is running, but the installation of an OS failed. The server will remain in the Server Pool list with this status for 7 days before the Opsware System deletes the entry. See "Recovering When an OS Installation Fails" on page 260 in this chapter for more information.
Managed Server Values	
Managed	Indicates a server that the Opsware System is managing. The Opsware System performs reachability checks for managed servers. After a server reaches this lifecycle state, the entry for the server moves from the Server Pool list to the Managed Servers list.
Deactivated	Indicates an Opsware-managed server that was removed from management. However, the server's history still exists in the Opsware System. Deactivated servers are not reachable.

OS Provisioning Process

This section provides information about the OS provisioning process within the Opsware System and contains the following topics:

- OS Provisioning Process Overview
- Solaris OS Provisioning
- Linux OS Provisioning
- Windows OS Provisioning

OS Provisioning Process Overview

The process for provisioning new servers of all supported operating systems includes the following steps in the OS Provisioning Subsystem:

1. A system administrator unpacks a server, mounts it in a rack, and attaches the server to power and a network that can reach the Opsware System.

2. The system administrator prepares the hardware for OS provisioning.

   See "Hardware Preparation Overview" on page 245 in this chapter for more information.

3. If necessary, the system administrator inserts a bootable floppy or CD provided with the Opsware System. Using a bootable floppy is not necessary for Intel-based servers that support PXE or Unix servers that support DHCP because these types of servers are capable of booting over a network.

   See "Booting New Servers" on page 246 in this chapter for more information.

4. The system administrator turns the server on.

   For servers capable of booting over the network, powering the server on causes the server to initiate its network boot process. For example, the server sends a boot request to a PXE server.

   The Opsware OS Build Manager responds to this network boot request by delivering the Opsware OS Build Agent, a small agent that can run in the memory of the bare metal server. (For servers not capable of booting over the network, the Opsware OS Build Agent is on the bootable floppy or CD)

   The Opsware OS Build Agent constructs an inventory of the server (including server manufacturer, server model, MAC address, available memory, and available storage) and delivers that information to the Opsware OS Build Manager.

5. In the Opsware Command Center, the system administrator sees this server and its hardware inventory in a list of available servers ready to be provisioned.

   See "Verifying Installation of an OS Build Agent" on page 251 in this chapter for more information.

6. The system administrator selects the OS or a complete server baseline (which can include a base OS, a set of OS patches, system utilities, and middleware software) to provision.

   The system administrator selects to install the OS or complete server baseline on the server at that time or schedule the installation for some time in the future.

   See "Installing an OS by Using a Template" on page 255 in this chapter for more information.

   See "Installing an OS by Using a Custom Installation" on page 258 in this chapter for more information.

The OS Provisioning Subsystem installs the selected software onto the server.

7 The system administrator uses the Opsware System to configure networking for the newly provisioned server.

See “Configuring Networking for an Opsware Managed Server” on page 128 in Chapter 2 for more information.

Additionally, the system administrator might choose to reprovision servers running Red Hat Linux or Sun Solaris operating systems by using the OS Provisioning Subsystem.

See “Reprovisioning a Solaris or Linux Server” on page 263 in this chapter for more information.

Solaris OS Provisioning

The OS Provisioning Subsystem includes a DHCP-based JumpStart configuration that hides the complexity of JumpStart from the end user. Unlike typical JumpStart systems, the OS Provisioning Subsystem does not require configuration updates to the JumpStart server for each installation that you provision.

Instead, an OS definition exists in the OS Provisioning Subsystem for each version of the Solaris OS that you want to install on servers in your environment.

The process for Solaris OS provisioning follows the general OS provisioning process that the OS Provisioning Subsystem established.

See “Sun Solaris Build Process” on page 200 in Chapter 3 for information about the detailed steps that occur during the Solaris build process.

Linux OS Provisioning

The OS Provisioning Subsystem includes a Kickstart or YaST2 system that hides the complexity of Kickstart or YAST2 from the end user.

Mapping a specific installation client to a particular Kickstart or YaST2 configuration is a simple procedure in the OS Provisioning Subsystem. The OS Provisioning Subsystem allows users to easily choose a particular Kickstart or YaST2 configuration through the Opsware Command Center at installation time.

The process for Linux OS provisioning follows the general OS provisioning process that the OS Provisioning Subsystem established.

See “About Linux Build Customization Scripts” on page 207 in Chapter 3 for information about the detailed steps that occur during the Linux build process.

Windows OS Provisioning

In the OS Provisioning Subsystem, system administrators can perform unattended, scripted installations of Windows NT and Windows 2000 on bare metal servers.

The installation-based approach allows system administrators to adapt to variations in hardware. The OS Provisioning Subsystem can be set up to install Windows operating systems on the known hardware makes and models in the managed environment. At build time, the OS Provisioning Subsystem provisions the server with the correct hardware-specific software and drivers based on the hardware signature of the server about to be provisioned.

The process for Windows OS provisioning follows the general OS provisioning process that the OS Provisioning Subsystem established.

See "Microsoft Windows Build Process" on page 208 in Chapter 3 for information about the steps that occur during the Windows build process.

Hardware Preparation Overview

Before you use the OS Provisioning Subsystem to install an OS on a server, the server must meet certain requirements, or the hardware must be prepared in certain ways, as Table 4-2 shows.

Table 4-2: Required Hardware Preparation for Opsware-managed Servers

OPERATING SYSTEM	HARDWARE REQUIREMENTS
Microsoft Windows	Before you install Windows on a server, you need to prepare the hardware by performing the following tasks: • If the hardware has a SCSI RAID controller, you must extend the Windows OS media distribution based on vendor specific requirements. The Windows OS media from Microsoft Corporation does not include the necessary drivers for these SCSI-RAID controllers. • Depending on the version of Windows, create a FAT16 partition or FAT32 on which to install the Windows OS You can create this required partition when using the Windows Boot Floppy or PXE to boot a server the first time. The boot image contains the functionality to create the required partition. See "Booting a Windows or Linux Server with PXE" on page 247 in this chapter for more information. See "Booting a Windows or Linux Server" on page 249 in this chapter for more information.
Sun Solaris	To install Solaris on a server, the hardware must meet the following requirements: • Have a DHCP-capable PROM (older servers can be upgraded to DHCP-capable PROM) • Be part of the sun4u system architecture (platform group) You do not need to perform any Opsware-specific preparation of the hardware before you install Solaris on a server.
Linux	Before you install Linux on a server, you need to prepare the hardware by configuring valid, logical drives for RAID.

Booting New Servers

This section provides information on booting new servers with the Opsware System and contains the following topics:

- Booting New Servers Overview
- About the OS Build Agent
- Booting a Windows or Linux Server with PXE
- Booting a Windows or Linux Server
- Booting a Solaris Server Over the Network
- About Installation of OS Build Agents
- Verifying Installation of an OS Build Agent
- Recovering When an OS Build Agent Fails to Install

Booting New Servers Overview

On Intel-based servers, you can boot a new server over a network in a hands-off fashion by using PXE. For environments with servers that do not support network boot technology, Opsware supports floppy- or CD-based booting.

For Windows and Linux servers, the Opsware Boot Floppy and CD respectively contains a small operating system, network drivers, software required to mount a network drive, and the Opsware OS Build Agent. The Opsware Boot Floppy or CD has the software that is otherwise delivered over the network as part of the network boot process.

For Solaris servers, you can provision an OS over the network by using DHCP.

To boot servers over the network, the installation client must be able to reach the Opsware DHCP server on the Opsware core network. If the installation client is running on a different network than the Opsware core network, your environment must have a DHCP proxy (IP helper). Alternatively, for Linux and Windows installation clients, you can boot the servers by using an Opsware Boot CD or Floppy instead of booting the servers over the network.

About the OS Build Agent

The OS Provisioning Subsystem de-couples the task of readying a server for provisioning from provisioning the server with an OS. This de-coupling of tasks is possible because of the OS Build Agent.

Booting a new server for the first time installs an OS Build Agent on the server; however, the server does not have the target OS installed and might not have access to disk resources. The Opsware System can still communicate with the server and perform commands on it remotely because the OS Build Agent is running an OS installed in memory.

The OS Build Agent performs the following functions:

- Registers the server with the Opsware System when the OS Build Agent starts
- Listens for command requests from the Opsware System and performs them

The OS Build Agent can perform commands even though the target OS is not installed.

Booting a Windows or Linux Server with PXE

Perform the following steps to boot a Windows or Linux Server with PXE:

1. After you mount the new server in a rack and connect it to the Opsware build network, set up the server to boot by using PXE.

   See the hardware vendor's documentation on how to prepare a server to boot by using PXE.

2. Power on the server and select the option to boot the server with PXE.

   The Opsware System menu appears and prompts you to select the type of Opsware Build Agent to install on the server.

```
windows   - Windows Build Agent (DOS 6.22)
undi      - Windows Build Agent (DOS 6.22 + UNDI)
win98     - Windows Build Agent (DOS 7.01)
undi98    - Windows Build Agent (DOS 7.01 + UNDI))
linux - Linux Build Agent
localdisk - Normal boot from localdisk (default after 10 sec)
```

Which version of the Windows Build Agent you should select depends on the type of x86 hardware being provisioned. The images for the Windows Build Agents vary in terms of the memory management software, disk partitioning capabilities, and network drivers – DOS or universal network device interface (UNDI) – that they contain.

For example, if you are provisioning a server that has more than 2GB of RAM, you should select the `Win98` or `Undi98` Boot Image. If an incompatible Boot Image is selected for the hardware, an error message appears at the console. The error message can appear at any point during the provisioning process; for example, it might appear when the Windows Build Agent is booting and DOS is loading or it might appear later in the process when the Windows Installer is loading. See Table 4-3 for the differences between images for the Windows Build Agents.

Table 4-3: Differences Between Images for the Windows Build Agents

BOOT IMAGE	NETWORK DRIVERS	MEMORY MANAGEMENT SOFTWARE	DISK PARTITIONING CAPABILITIES
`windows`	DOS	DOS 6.22	FAT16
`undi`	UNDI	DOS 6.22	FAT16
`win98`	DOS	Windows 98	FAT32
`undi98`	UNDI	Windows 98	FAT32

If you do not select an option after 10 seconds, the server defaults to booting from local disk.

If you select Windows as the option for booting the server, an additional set of Opsware menus appears on the console so that you can partition the hardware disk.

3. For Windows servers only, select the menu choices to partition the disk based on your specifications.

   After the booting process finishes successfully, a message appears on the console that indicates that the server is ready for OS provisioning. An OS Build Agent was installed on the server and the server appears in the Server Pool list in the Opsware Command Center.

4. (Optional) Record the MAC address of the server so that you can locate the server in the Server Pool list in the Opsware Command Center.

You should verify that the newly racked server shows up in the Opsware Command Center and is ready to hand off for OS installation.

See “Verifying Installation of an OS Build Agent” on page 251 in this chapter for more information.

When booting a Linux or Windows server by using PXE, the DHCP relay must be running on the router of the build network for PXE to function properly.

Booting a Windows or Linux Server

You can boot different types of x86 hardware by using an Opsware Boot Floppy (Windows, Windows 98) or by using an Opsware Boot CD (Red Hat Linux or SUSE Linux) because a Boot Floppy or CD can contain multiple NIC drivers.

When you boot a Windows server with a boot floppy, select the Windows or Windows 98 boot floppy based on the server’s memory and disk partitioning requirements.

See Table 4-3 on page 248 for the differences between the Windows and Windows 98 OS build images.

Perform the following steps to boot a Windows or Linux server:

1. After you mount the new server in a rack and connect it to the Opsware build network, insert the Windows Boot Floppy or Linux Boot CD (depending on which OS you want to install on the server).
2. Power on the server. A hardware-vendor specific message appears on the console.

   If you selected Windows as the option for booting the server, Opsware menus appear on the console so that you can partition the hardware disk.
3. For Windows servers only, select the menu choices to partition the disk based on your specifications.

   After the booting process finishes successfully, a message appears on the console that indicates that the server is ready for OS provisioning. An OS Build Agent was installed on the server and the server appears in the Server Pool list in the Opsware Command Center.
4. (Optional) Record the MAC address of the server so that you can locate the server in the Server Pool list in the Opsware Command Center.

You should verify that the newly racked server shows up in the Opsware Command Center and is ready to hand off for OS installation.

See "Verifying Installation of an OS Build Agent" on page 251 in this chapter for more information.

Booting a Solaris Server Over the Network

When the Opsware System was installed in your facility, the OS Provisioning Subsystem was set up so that the Opsware Boot Server listens for broadcast requests from new servers and it responds by using DHCP.

Perform the following steps to boot a Solaris server over the network:

1. Mount the new Solaris server in a rack and connect it to the network.

   The installation client on this network must be able to reach the Opsware DHCP server on the Opsware core network. If the installation client is running on a different network than the Opsware core network, your environment must have a DHCP proxy (`IP helper`).

2. Enter one of the following commands at the prompt:

   ```
   ok boot net:dhcp - install
   ```

   Or

   ```
   ok boot net:dhcp - install <interface_setting>
   <buildmgr=hostname|IP_address>
   ```

   Where `<interface_`*`setting`*`>` is one of the following options:

   ```
   autoneg, 100fdx, 100hdx, 10fdx, 10hdx
   ```

You can include an interface setting with the boot command to set the network interface to a specific speed and duplex during OS provisioning. When the Opsware System was installed in the local facility, a default value was provided for this interface setting. Specifying this boot argument allows you to override the default interface setting.

To continue setting the network interface with a specific speed and duplex, you can use a variety of methods, including using a Solaris build customization script or specifying the values in a Solaris Package or RPM in the OS media.

See "About the Solaris Build Customization Script" on page 203 in Chapter 3 for more information.

Details: How the OS Build Agent Locates the Opsware Build Manager

For Solaris OS Provisioning, the JumpStart build script runs the OS Build Agent, which contacts the Opsware Build Manager. The Solaris `begin` script attempts to locate the Opsware Build Manager in the following ways:

- By using information that the Opsware DHCP server provided
- By looking for the hostname `buildmgr` in DNS as configured by the DHCP server

You can override the way that the OS Build Agent contacts the Opsware Build Manager by specifying a boot argument at the prompt when you boot a new Solaris server:

```
ok boot net:dhcp - install [buildmgr=hostname|IP_address]
```

About Installation of OS Build Agents

After you install an OS Build Agent on a server by booting the server with PXE or an Opsware Boot Floppy (Windows and Linux) or by using the network (Solaris), the server appears in the Server Pool list.

You should verify that the newly racked server shows up in the Opsware Command Center and is ready to hand off for OS installation.

The Server Pool list displays the servers that have registered their presence with the Opsware System but do not have the target OS installed on them. From here, you can install an OS by selecting the server and clicking the Install OS button.

Verifying Installation of an OS Build Agent

Perform the following steps to verify the installation of an OS build agent:

1. Log in to the Opsware Command Center.
2. From the navigation panel, click Servers ➤ Server Pool. The Server Pool page appears, as Figure 4-1 shows.

Figure 4-1: Server Pool List in the Opsware Command Center

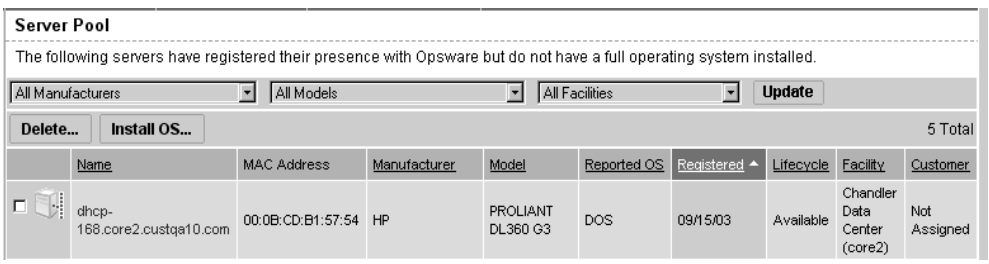

3 (Optional) From the drop-down lists, select the manufacturer, model, or facility of the server that you want to verify and click the Update button.

4 For Intel x86 processor-based servers, locate the MAC address of the server that you just booted.

Or

For Sun SPARC processor servers, locate the chassis ID of the server that you just booted.

The chassis ID for Sun SPARC processor servers appears in the MAC Address column in the Server Pool list.

The Lifecycle column indicates the progress or success of the OS Build Agent installation. If the OS Build Agent was successfully installed, the Lifecycle column indicates that the server is available for OS provisioning.

See "The Server Lifecycle for OS Provisioning" on page 240 in this chapter for more information.

Recovering When an OS Build Agent Fails to Install

When an OS Build Agent fails to install on a server, the server does not appear in the Server Pool list.

You can check the server console for error messages and try to boot the server again with PXE or by using the Opsware Boot Floppy or CD.

If all errors were successfully resolved, the initial boot occurs, the OS Build Agent is installed on the server, the server appears in the Server Pool list, and the Lifecycle column indicates that the server is available.

If you are unable to resolve the error condition and install the OS Build Agent on the server so that it appears in the Server Pool list, contact your Opsware administrator for troubleshooting assistance.

OS Installation with Opsware Command Center

This section provides information on OS installation with the Opsware Command Center and contains the following topics:

- OS Installation with Opsware Command Center Overview
- Ways to Install Operating Systems on Servers

- Installing an OS by Using a Template
- Installing an OS by Using a Custom Installation
- Recovering When an OS Installation Fails
- Network Configuration for Servers after OS Provisioning
- Requirements for Reprovisioning Solaris and Linux Servers
- Reprovisioning a Solaris or Linux Server

OS Installation with Opsware Command Center Overview

You begin the OS provisioning process by reviewing the servers in the Server Pool list. Servers in the Server Pool have registered their presence with the Opsware System but do not have the target OS installed. From there, you can install an OS by clicking the Install OS button. See Figure 4-2.

Figure 4-2: Server Pool List in the Opsware Command Center

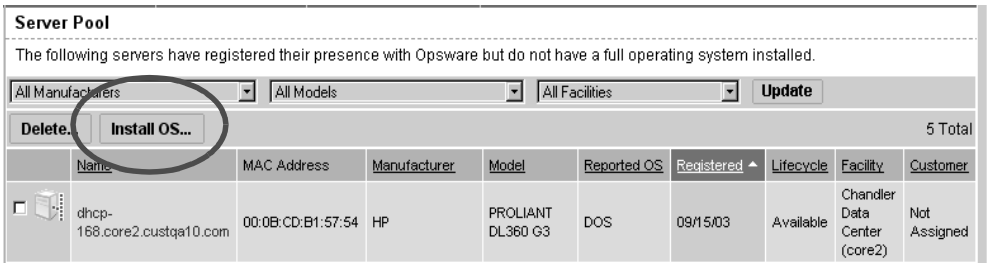

The Server Pool provides the following information about each server waiting to be provisioned with the target OS:

- The hostname set by booting the server the first time over the network or by using an Opsware Boot Floppy or CD
- The MAC address or chassis ID
- The manufacture and model of the server
- The OS of the OS Build Agent (Windows, Red Hat Linux, or Solaris)

  You use this information to select the target OS for servers.
- The customer association

- Additional hardware information (clicking the server name opens a window that displays specific hardware information, as Figure 4-3 shows.

Figure 4-3: Information Displayed in the Edit Server Page for a Server in the Server Pool

HARDWARE INFORMATION	
MAC Address:	00:50:8B:E2:4A:69
Serial Number:	6J0CFCX2J0Y3
Manufacturer:	COMPAQ
Model:	PROLIANT DL360
Memory:	4.00 MB SYSTEM MEMORY
Processors:	800. INTEL () 800. INTEL ()
Storage:	(not set)
ADDITIONAL INFORMATION	
Server ID:	1540001
MID:	1540001

Ways to Install Operating Systems on Servers

You can install an OS on a server by using one of the following methods:

- Selecting a pre-defined template, which is a pre-packaged collection of installable software

  The template includes the base operating system, and can include software to provision the entire software stack, such as the latest set of operating system patches, system utilities (SSH or the latest JVMs), middleware including databases, Web servers, and application servers, and so on, up to the custom business applications the server ultimately runs.

- Performing a custom installation, which includes defining the installation on-the-fly by selecting an OS definition, patches, and other applications to install

  After performing a custom installation, you can then save the selections in a new template for later use on other servers.

After the installation has begun on a set of servers, you can view progress or results either in the Opsware wizard itself or through the My Jobs interface.

See "Templates and Server Management" on page 38 in Chapter 2 for information about how the Opsware System works with templates.

See “Working with Templates” on page 353 in Chapter 6 for information about how to create and manage templates.

Installing an OS by Using a Template

During OS installation, you cannot select a server from the Server Pool list that has the status Installing OS.

Perform the following steps to install an OS by using a template:

1. From the Opsware Command Center home page, click the Install OS link in the Tasks panel.

   Or

   From the navigation panel, click Servers ➤ Server Pool. The Server Pool page appears. Select the servers that you want to provision and click the Install OS button.

   When you select multiple servers to provision, you must select servers with similar hardware architecture – x86-processor-based hardware or SPARC. If you select servers that have different hardware architecture, an error message appears.

   The Install Operating System Wizard appears. See Figure 4-4.

Figure 4-4: Overview Page in the Install Operating System Wizard

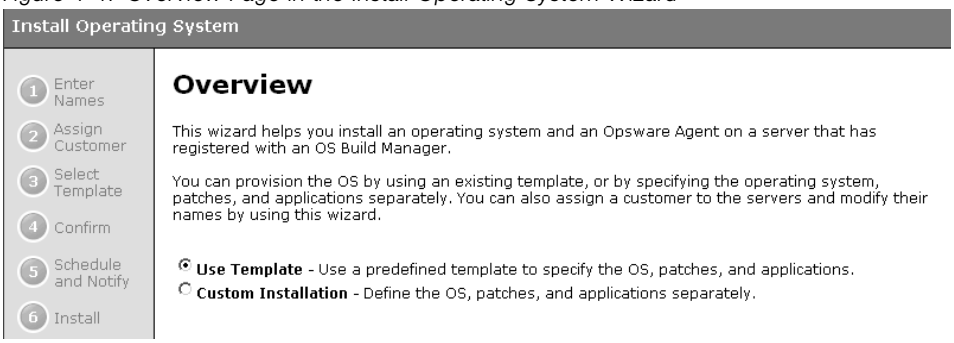

2. If necessary, select the Use Template option and click the Start button. If you did not select servers in Step 1, the Select Servers page appears.

3. If prompted, select the servers that you want to provision and click the Next button. You can find the servers that you want to provision by browsing the list or by searching.

   See “Searching with Advanced Search” on page 48 in Chapter 2 for information about how to use the advanced search features in the Opsware System.

Servers in the Server Pool list waiting to be provisioned are identified by MAC address or chassis ID because they are still using DHCP addresses.

You must select servers with similar hardware architecture (based on the value in the Reported OS column in the Server Pool list) or an error message appears.

4 Enter a name for each server and click the Next button. The Assign Customer page appears.

By default, the OS Provisioning Subsystem entered the server hostname in this field. You can enter new names that adequately describe each server.

The name that you enter appears as the display name for the servers in the Opsware Command Center UI.

5 Select a customer for the servers and click the Next button. The Select Template page appears.

In the Assign Customers page, you only see customers listed that you have the permission to access with your user account. Additionally, the customer that you select controls which templates you can use to install an OS and software on the servers. Depending on the customer that you select, you see only the templates associated with that customer, that are customer independent, or are not assigned to a customer.

See "Customer Accounts in the Opsware System" on page 58 in Chapter 2 for information about the distinction between the customers *Customer Independent* and *Not Assigned.*

6 Select the template to use to provision the servers and click the Next button. Templates appear in the list only if they meet these requirements:

- The templates were created for the same OS as the servers that are being provisioned.
- The templates contain OS software to install on servers.

  See "Templates and Server Management" on page 38 in Chapter 2 for information about how the Opsware System works with templates.

  See "Working with Templates" on page 353 in Chapter 6 for information about how to create and manage templates.

To view the details about the templates before you select one, click the template name. A window appears that displays general information about the template, as Figure 4-5 shows.

Figure 4-5: Template Details Page in the Install Operating System Wizard

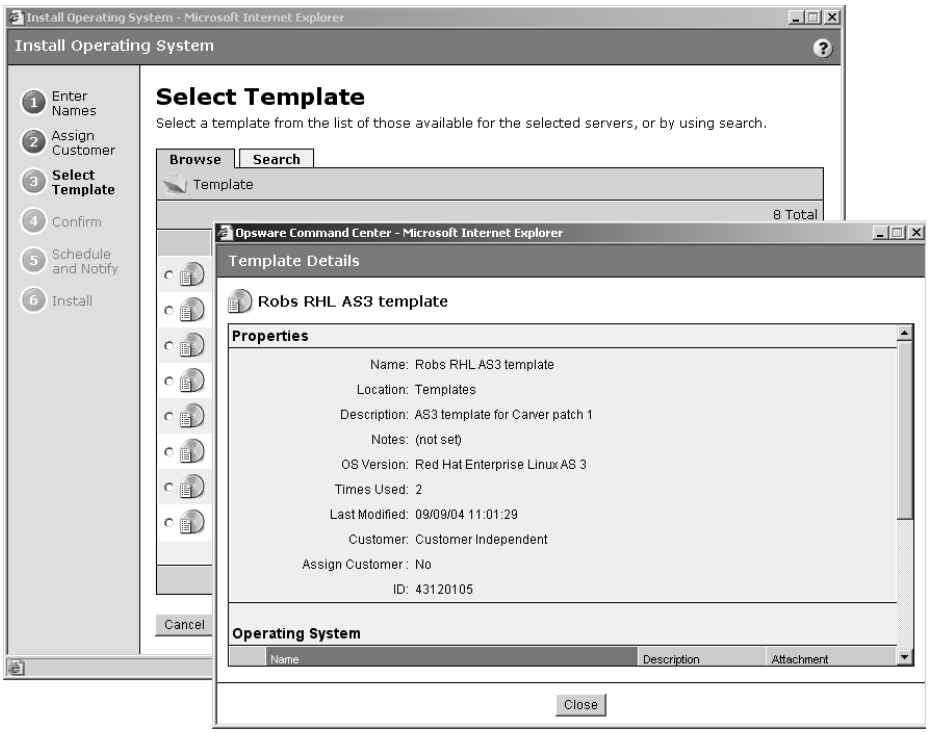

The Select Template page only displays the templates available for the OS that you chose to install.

To view detailed information about the OS that the template installs, such as the settings in the configuration file, click Software ➤ Operating Systems from the navigation panel, then click the Installation tab. The installation resources for the OS appear. You can view information about the OS installation, such as the contents of the configuration file.

After you click the Next button, a confirmation page appears that shows details about the template that will be installed on the servers.

7 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.

- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.
- If you select to run the job at that time, a progress bar appears that shows the progress of the OS installation.

8 (Optional) When the installation finishes, click the View Details button to see progress or the results of the installation.

9 Click the Close button to end the wizard.

Closing the wizard does not stop the installation if it is still running. After you close the wizard, you can view the progress of the installation by viewing My Jobs.

See "Scheduling and Notifying Server Management Tasks" on page 106 in Chapter 2 for information about the My Jobs feature.

Installing an OS by Using a Custom Installation

During OS installation, you cannot select a server from the Server Pool list that has the status Installing OS.

Perform the following steps to install an OS by using a custom installation:

1 From the Opsware Command Center home page, click the Install OS link in the Tasks panel.

Or

From the navigation panel, click Servers ➤ Server Pool. The Server Pool page appears. Select the servers that you want to provision and click the Install OS button.

You must select servers with similar hardware architecture (x86-processor-based hardware or SPARC) or an error message appears.

The Install Operating System Wizard appears, as Figure 4-6 shows.

Figure 4-6: Overview Page in the Install Operating System Wizard

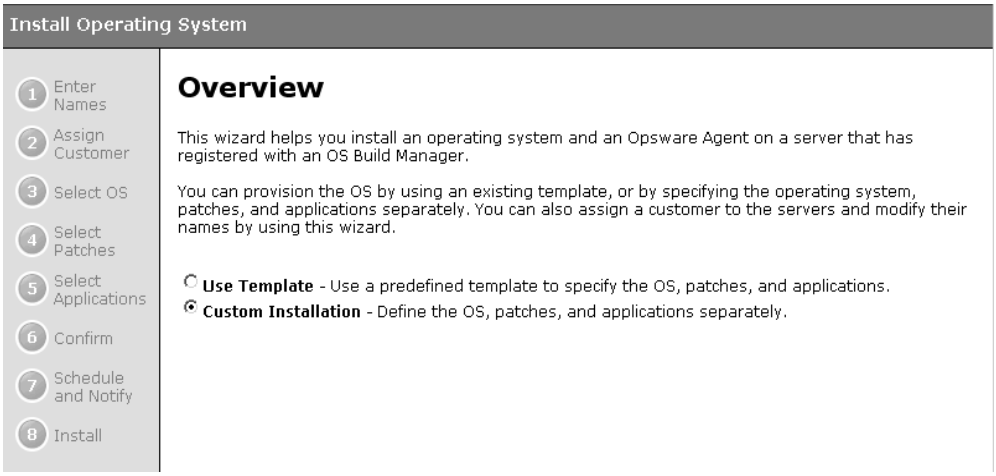

2. Select the Custom Installation option and click the Start button. If you did not select servers in Step 1, the Select Servers page appears.
3. If necessary, select the servers that you want to provision and click the Next button. The Select Operating System page appears.

   You must select servers with similar hardware architecture (based on the value in the Reported OS column in the Server Pool list) or an error message appears.
4. Enter a name for each server and click the Next button. The Assign Customer page appears.

   By default, the OS Provisioning Subsystem entered the server hostname in this field. You can enter new names that adequately describe each server.

   The name that you enter appears as the display name for the servers in the Opsware Command Center UI.
5. Select a customer for the servers and click the Next button.

   The customer that you select controls which operating systems and applications you can select to install on the servers. Depending on the customer that you select, you see only the operating systems and applications associated with that customer, that are customer independent, or are not assigned to a customer. Patches are always customer independent.
6. Select the OS for the servers and click the Next button.

7 Select the OS patches that you want to apply to the servers and click the Next button. The Select Applications page appears.

See "Opsware Patch Management" on page 403 in Chapter 8 for information about how the Opsware System performs patch management.

8 Select any applications that you want to install on the servers and click the Next button.

9 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.
- If you select to run the job at that time, a progress bar appears that shows the progress of the OS installation.

10 (Optional) When the installation finishes, click the View Details button to see progress or the results of the installation.

11 Click the Close button to end the Wizard.

Closing the Wizard does not stop the installation if it is still running. After you close the Wizard, you can view the progress of the installation by viewing My Jobs.

See "Scheduling and Notifying Server Management Tasks" on page 106 in Chapter 2 for information about the My Jobs feature.

Recovering When an OS Installation Fails

Servers waiting to be provisioned appear in the Server Pool list with the status available. When an OS installation fails for an unprovisioned server, the server status changes to Build Failed.

Perform the following steps to recover when an OS installation fails:

1 From the Opsware Command Center home page, click the link for the failed installation in the My Jobs panel.

Or

From the navigation panel, click My Jobs. The My Jobs page appears and displays the operations that you performed with the Opsware System, as Figure 4-7 shows.

Figure 4-7: My Jobs Panel in the Opsware Command Center Home Page

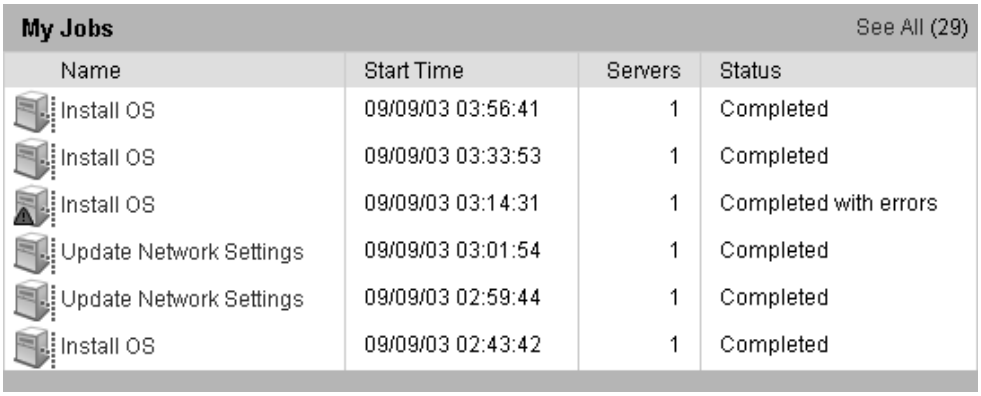

2 Locate the failed OS installation that you initiated and click the link for the job to open a window for the job. An error message appears in the window, as Figure 4-8 shows.

Figure 4-8: Details Page for a Job in the My Jobs Panel

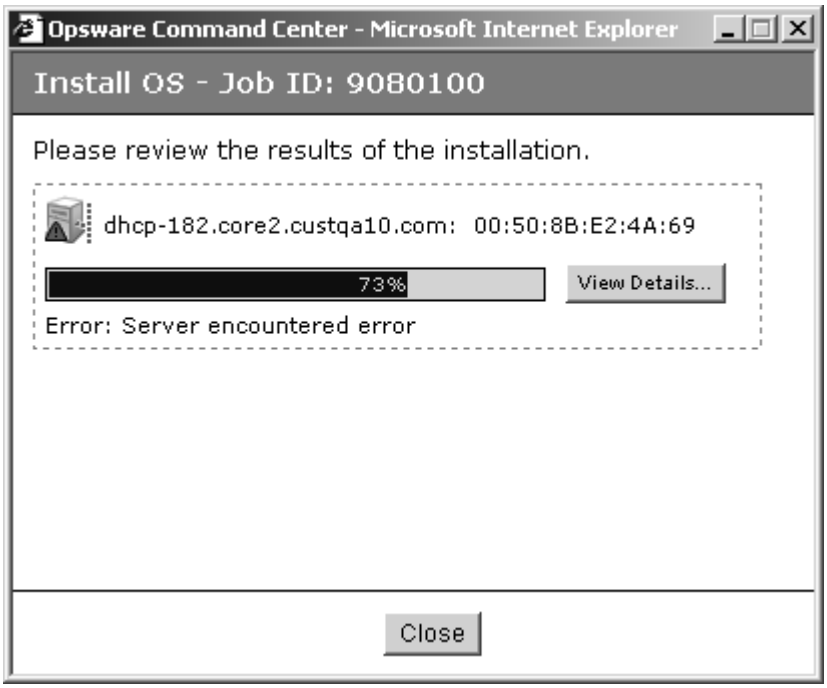

3 Click the View Details button to see detailed information about the job. The My Jobs information contains a build log for the OS installation. This build log contains any error messages that the OS Provisioning Subsystem generated, as Figure 4-9 shows.

Figure 4-9: Details Page in the My Job Window

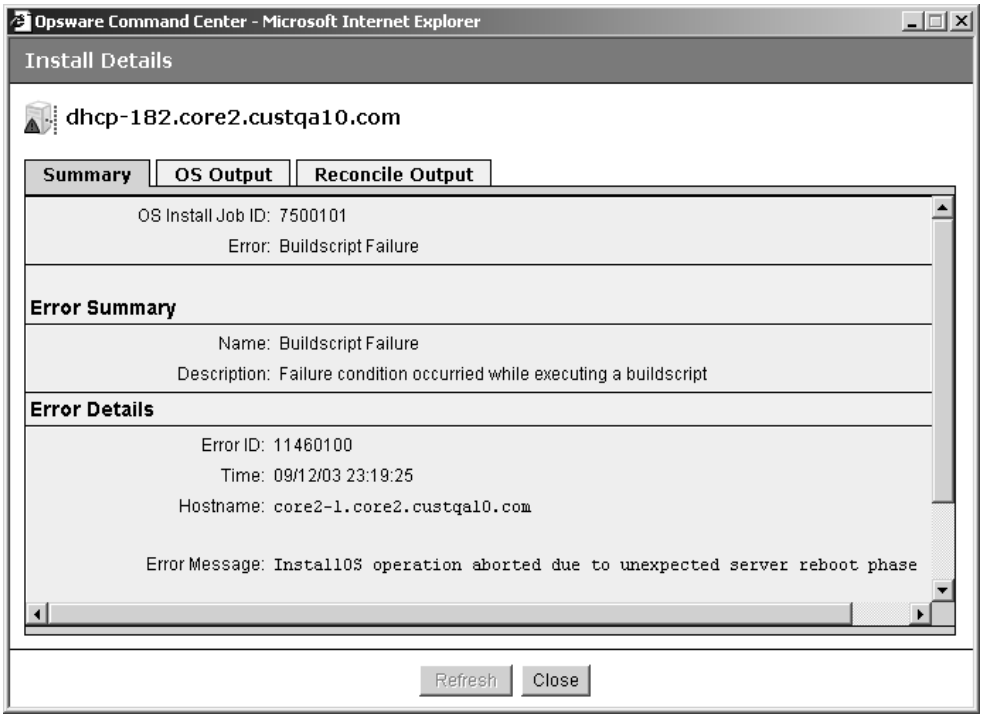

4 Review and fix any errors that occurred during the build process.

5 (Optional) Delete the failed OS installation from the Server Pool list by selecting the server and clicking the Delete button.

You can leave the entry for the failed OS installation in the Server Pool list because the Opsware System automatically replaces the entry when you reboot the server.

Alternatively, the Opsware System removes the entry from the list after 7 days and all information about the server is removed from the Opsware System.

6 Reboot the server by using the network (PXE for Windows and Linux servers or DHCP for Solaris servers) or by using an Opsware Boot Floppy (for Windows) or by using an Opsware Boot CD (for Linux servers). If successful, the server appears in the Server Pool list with the status available.

7 Install the OS on the server by using a template or custom installation.

See “Installing an OS by Using a Template” on page 255 in this chapter for more information.

See “Installing an OS by Using a Custom Installation” on page 258 in this chapter for more information.

If all errors were successfully resolved, the OS is installed on the server and the server moves from the Server Pool list to the Managed Server list.

See “Scheduling and Notifying Server Management Tasks” on page 106 in Chapter 2 for information about using My Jobs to obtain a history of your operations.

Network Configuration for Servers after OS Provisioning

The OS Provisioning Subsystem provisions servers with an OS by using DHCP addresses. By using the Opsware Command Center, users can configure network settings, including a static IP address, hostname, default gateway, DNS server addresses, subnet masks, and so on, after the base operating system is installed on servers.

Because DHCP servers often assign temporary IP addresses to servers that boot over a network, system administrators typically need to assign static IP addresses (and other network properties) before the servers can be put into service. The Opsware System enables system administrators to do this through the Opsware Command Center rather than logging onto the server manually after OS provisioning is complete.

See “Configuring Networking for an Opsware Managed Server” on page 128 in Chapter 2 for information about how the Opsware System configures networking for servers.

Requirements for Reprovisioning Solaris and Linux Servers

When you choose to preserver network configuration for a server you are about to reprovision, the following requirements apply:

- The server must be in a DHCP-enabled network when the reprovisioning begins.
- The server must be in its original network after the OS installation is complete.

Because of these requirements, re-provisioning Solaris and Linux servers functions best when the build and production networks are overlaid (namely, the production network is running a DHCP server).

Reprovisioning a Solaris or Linux Server

You can reprovision Solaris and Linux servers so that they are running another version of the same OS so long as the hardware supports that new version of the OS.

You can reprovision servers built by the Opsware System and Opsware assimilated servers by using this feature.

You cannot reprovision a Linux server so that it runs a Windows OS.

Perform the following steps to reprovision a Solaris or Linux server:

1. From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers page appears.
2. Select the servers that you want to reprovision from the list.

   You must select servers with similar hardware architecture (x86-processor-based hardware or SPARC) or an error message appears.

   The check boxes for servers running the Windows OS are unavailable.
3. From the Software menu, choose Install ➤ Operating System.

   The Install Operating System Wizard appears. The Wizard contains a warning that you are about to reprovision the servers. See Figure 4-10.

Figure 4-10: Overview Page for the Install Operating System Wizard When Re-Provisioning

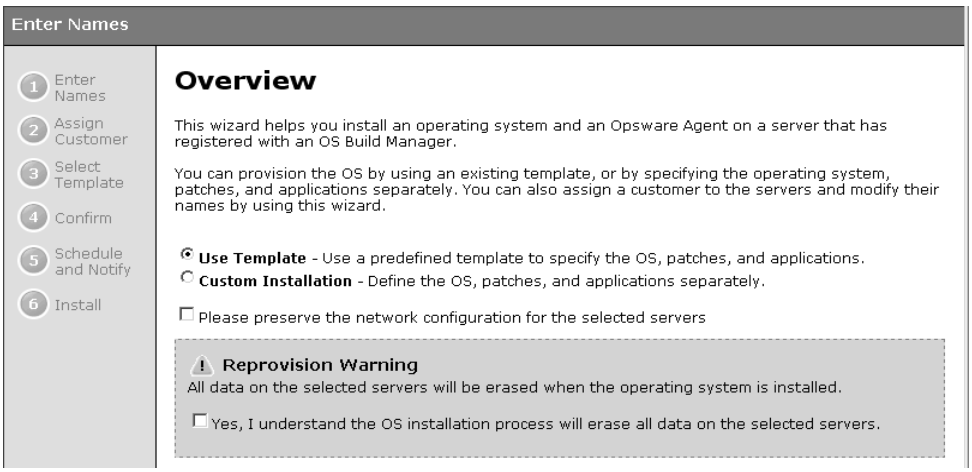

4. Select the check box in the Warning (indicating that you understand that the Opsware System erases all data from the servers).

5 (Optional) Select the check box to preserve the network settings for the server.

See “Configuring Networking for an Opsware Managed Server” on page 128 in Chapter 2 for information about how the Opsware System configures networking for servers.

6 Select the Use Template option or the Custom Installation option and click the Start button.

The installation proceeds normally through the Install Operating System Wizard.

7 Complete the installation by using the Install Operating System Wizard.

See “Installing an OS by Using a Template” on page 255 in this chapter for more information.

See “Installing an OS by Using a Custom Installation” on page 258 in this chapter for more information.

Details: Reprovisioning Solaris Servers

When you reprovision a Solaris server, the OS Provisioning Subsystem reboots the server automatically. Therefore, you cannot provide boot arguments at the prompt (for example, to specify the interface setting or the location of the Opsware Build Manager).

Alternatively, you can provide a boot argument by including a custom attribute for the server named `reboot_command` that has the value for the interface setting that you want to use.

See “Booting a Solaris Server Over the Network” on page 250 in this chapter for information about how to use interface settings when you boot a Solaris server.

Chapter 5: Package Management

IN THIS CHAPTER

This chapter discusses how to manage packages in the Opsware System and covers the following package management topics:

- Package Management Overview
- Supported Operating Systems and Package Types
- AIX Package Management
- HP-UX Package Management
- Linux Package Management
- Solaris Package Management
- Windows Package Management
- ZIP Package Management
- Package Management Tasks Overview
- Searching for Packages
- Viewing Packages Assigned to Nodes
- Uploading a Package
- Overwriting a Package
- Editing Package Properties
- Deleting a Package
- Deprecating a Package
- Downloading a Package

You must have specific permissions to manage packages by using the Opsware Command Center. Contact your Opsware administrator to obtain the necessary access rights.

This chapter assumes that you have already packaged your software applications.

Package Management Overview

Software applications are made available in the Opsware System by uploading the packages to the Software Repository with the Opsware Command Center or by using the Opsware Command Line Interface. See Appendix A, "Opsware Command Line Interface" for more information.

The Opsware Command Center provides options to perform the following tasks:

- Upload packages to the Software Repository.

  The Software Repository provides a data store for all software that Opsware technology manages. It contains packages for operating systems, applications (for example, BEA WebLogic or IBM WebSphere), databases, customer code, and software configuration information.

- Perform other package setup and management functions.

After you upload packages to the Software Repository, you install packages by adding them to nodes, assigning the nodes to servers, and reconciling the servers. The Opsware System reconciles the differences between how a server is configured by assigned nodes and what software is actually installed. The Opsware Agent installed on each Opsware-managed server coordinates installation of software packages that are missing or need to be upgraded based on updates to nodes, and removes the packages that should no longer be installed.

See "Application Provisioning Setup" on page 307 in Chapter 6 for information about how to add software to nodes.

The Opsware System wizards automatically assign servers to nodes and reconcile them.

Container Packages and Installable Packages

For some operating systems, the distribution or container package might contain more than one installable package. For example, Solaris packages can contain multiple installable packages, called *instances*.

When a container package is uploaded in the Opsware System, package entries are automatically created for all the installable packages it contains.

Container packages cannot be directly attached to nodes, only installable packages. Container packages cannot be deleted directly; you can only delete the container package, and only if none of the installable packages it contains are attached to nodes.

Solaris patch clusters are an exception to this general rule. They are container packages, but can also be attached to nodes themselves.

Supported Operating Systems and Package Types

Each operating system that the Opsware System supports has a list of package types that you can upload. The Opsware System supports these package types on the supported operating systems, as Table 5-1 shows.* For certain package types, the

Table 5-1: Supported Operating Systems and Package Types

OPERATING SYSTEM	PACKAGE TYPE	FILE FORMATS	ADDITIONAL METADATA*
AIX	LPP (contains an update fileset or base filesets)	.bff	N/A
	RPM	.rpm	N/A
HP-UX	Depot (contains products and filesets)	.tar	N/A
Linux	RPM	.rpm	N/A
Solaris	Patch	.jar, .tar, tar.gz, .tar.Z, t.gz, .zip	N/A
	Patch Cluster (contains patches)	.tar, .tar.gz, tar.Z, .t.gz, .zip	N/A
	Solaris package (contains package instances)	Datastream File	N/A
	RPM	.rpm	N/A

Table 5-1: Supported Operating Systems and Package Types

OPERATING SYSTEM	PACKAGE TYPE	FILE FORMATS	ADDITIONAL METADATA*
Windows	Hotfix	.exe	N/A
	Security Patch	.exe	N/A
	MSI	.msi	Product version and name
	OS Service Pack	.exe	Service Pack Level
	Windows Utility (Microsoft Security Baseline Analyzer and qchain)	.exe	N/A
	Microsoft Patch Database (contains a description of available patches) See "About the Microsoft Patch Database" on page 412 in Chapter 8 for more information.	.xml, .cab	N/A
	ZIP	.zip	N/A
OS Independent	Unknown	All	N/A

Opsware Command Center requires that you provide additional metadata for the package.

The Build Customization Script feature is available for Linux, Solaris, and Windows. See "Build Customization Scripts" on page 199 in Chapter 3 for more information.

The Opsware System verifies that RPM files, Solaris patch clusters, AIX LPPs, Solaris packages, and HP-UX depots that are uploaded are the correct package type. You can upload packages that are designated OS Independent, but you cannot attach them to nodes.

AIX Package Management

This section provides information on AIX package management within the Opsware System and contains the following topics:

- AIX Package Management Overview
- LPP Metadata

AIX Package Management Overview

LPPs are the container packages for AIX. LPPs have the following characteristics:

- An LPP contains either one or more base filesets or an update fileset.
- When an LPP contains multiple filesets, frequently only a subset of those filesets is installed because users might want to install only certain filesets.

The basic unit of AIX packages is the fileset. Filesets have the following characteristics:

- Filesets are versioned.
- The two types of filesets are base and update.
- Users add filesets to nodes. Therefore, the Opsware System adds filesets to and removes filesets from servers through reconcile.

Filesets are delivered as part of an LPP file, which users upload to the Software Repository. The Opsware System automatically creates package entries for all the filesets that the LPPs contain. When viewing an LPP in the Opsware Command Center, users see which filesets it contains.

The Opsware Agent reports which filesets and Authorized Program Analysis Reports (APARs) are installed on servers because servers only report filesets and APARs (and cannot report LPPs). The Opsware Command Center shows filesets and APARs in the Installed Packages list for a server.

See "About AIX Patches" on page 418 in Chapter 8 for information about how the Opsware System manages AIX APARs.

LPP Metadata

The Opsware System uses the metadata contained in LPPs when creating the package entries in the list of packages. An LPP contains the following metadata:

- The name of the LPP

- The name, version, and description of each fileset in the LPP
- For an update fileset, a list of APARs addressed by the fileset
- For each APAR listed, the list of filesets that make up that APAR

The Opsware System does not support bundles (which are abstract sets of filesets, drawn from multiple LPPs) or Program Temporary Fix (PTFs), which are similar to APARs without the metadata. However, users can still model a bundle or PTF by creating a node and attaching the filesets included in the bundle or PTF to that node.

When a user uploads an LPP, the Opsware System performs the following actions:

- Opens the LPP and parses its metadata
- Automatically creates entries in the list of packages for the filesets in the LPP and registers them as installable
- Automatically creates entries in the list of packages for the APARs defined by the update filesets in the LPP (if any)
- Registers the LPP as a non-installable package

HP-UX Package Management

This section provides information on server management jobs within the Opsware System and contains the following topics:

- HP-UX Package Management Overview
- Depot Metadata
- Prerequisites to HP-UX Package Management
- Example Commands: Converting a Depot
- Example File: Script to Split a Depot by Product
- Example File: Script to Split a Depot by Bundle

HP-UX Package Management Overview

Depots are the container packages for HP-UX. Depots have the following characteristics:

- A depot either contains products that contain filesets, or it contains patch products that contain patch filesets.
- When a depot contains multiple products and filesets, frequently only a subset of them are installed because users might want to install only certain products or filesets.
- A depot is a special type of directory formatted for use by HP Software Distributor (SD-UX) commands. SD-UX, a software management system, is the distribution mechanism for all HP software for HP-UX.
- A depot can be a local directory, a CD-ROM, tape, or it can reside on a server on the network.
- Multiple depots can be created for different applications or purposes.
- Users upload depots to the Software Repository in TAR format.
- Users can upload depots as HP-UX 11.00 or 11.11 depots. However, HP-UX software can be compatible with both 11.00 and 11.11. When the software in a depot is compatible with both 11.00 and 11.11, upload the depot to the Software Repository for both 11.00 and 11.11.
- Depots cannot be differentiated by hardware platform, such as s700 or s800.
- HP-UX depots have two basic formats:
  - Directory – the format for depots saved on a server or CD-ROM
  - Tape – the format for standalone depot files and the format required for uploading HP-UX packages into the Opsware System

HP-UX depots that contain both products and patch products cannot be uploaded to a specific customer. They can only be uploaded to Customer Independent.

Products and filesets are the installable packages for HP-UX. They have the following characteristics:

- Products and filesets are versioned.
- Filesets are the smallest installable unit. A fileset can belong to only one product, but can be included in multiple subproducts or bundles.
- Subproducts are logically related filesets and are not versioned; for example, X11.Manuals.

- Products are supersets of filesets.
- Bundles are logical groups of filesets; for example, HP-UX Support Tools Bundle.

The Opsware System supports products, filesets, and patch products as installable software.

The Opsware System does not support bundles (which are abstract sets of filesets, drawn from depots) or subproducts by automatically creating nodes for bundles and subproducts when users upload depots. However, users can still model bundles and subproducts by creating nodes for them and attaching the filesets for the bundles and subproducts. The Opsware System does not support using HP-UX codewords.

When a user uploads a depot, the Opsware System performs the following actions:

- Opens the depot and parses its metadata
- Automatically creates entries in the list of packages for the products and filesets in the depot and registers them as installable
- Registers the depot as a non-installable package

If a depot contains different software for HP-UX 11.00 and 11.11, create OS-specific depots for each HP-UX version and upload the depots to the Software Repository. The Opsware Command Center does not check the OS compatibility of the products and filesets in a depot when a user uploads the depot. When attaching products or filesets to a node, the products and filesets are attachable only when the associated OS of their depot matches the OS specified for the node.

The format of HP-UX version information can be inconsistent, making it difficult to determine whether one version is older than another when installing a package that has another version already installed. The Opsware System attempts to install it anyway. An error results if a newer version is already installed.

The Opsware System does not provide alternate root support for HP-UX. Do not include commands that require alternate root support in the Install Flags text box of the Packages: Edit Properties page. See "Editing Package Properties" on page 300 in this chapter for more information. By default, the HP-UX `swinstall` command does *not* replace a newer version of a fileset or product with an older version. However, the Opsware System does overwrite newer versions of filesets and products with older versions. The Opsware System does not support relocating packages for HP-UX.

Depot Metadata

The Opsware System uses the metadata contained in depots when creating the package entries in the list of packages. A depot contains the following metadata:

- The name, version, and description of each product in the depot
- The list of filesets in each product in the depot
- The name, version, and description of each fileset in the depot

Prerequisites to HP-UX Package Management

Before you upload a depot to the Software Repository, perform the following tasks:

1. Convert the depot on the installation media (CD-ROM) from directory format to tape format by using the `swpackage` command:

   ```
   swpackage -x media_type=tape -s <directory depot> <software selection> @ <file depot>
   ```

2. Split the depot into depots for each product.

You can perform this step manually by using NIM utilities or you can run a script to automate this step. See "Example File: Script to Split a Depot by Product" on page 277 in this chapter for more information. See "Example File: Script to Split a Depot by Bundle" on page 277 in this chapter for more information.

Example Commands: Converting a Depot

The following example shows the commands used to create a Quality Pack file depot from the Support Plus CD-ROM for HP-UX 11.00:

1. Mount the directory on the CD-ROM that contains the Quality Pack file depot:


```
mount -F cdfs /dev/dsk/c2t1d0 /cdrom
```

2. Convert the depot on the CD-ROM from directory format to tape format by using the swpackage command:

```
swpackage -x media_type=tape -s /cdrom/QPK1100 QPK1100 @ \
  /var/tmp/QPK1100.depot
```

Entering this command copies the QPK1100 bundle contained in the depot to a file that can be uploaded into the Opsware System.

Example File: Script to Split a Depot by Product

```
# This is an example script that splits a depot into individual
# product depots that can then be uploaded to the Opsware
# Software Repository

for product in `swlist -l product -s <location of depot> | \
  cut -f1 | grep -v ^# | grep '[A-z]'`
do
swpackage -x media_type=tape -s <location of depot> $product \
  @ /var/tmp/$product.depot
done
```

Example File: Script to Split a Depot by Bundle

```
# This splits a depot into individual bundle depots that can
# then be uploaded to the Opsware Software Repository

for bundle in `swlist -l bundle -s <location of depot> | \
  cut -f1 | grep -v ^# | grep '[A-z]'`
do
swpackage -x media_type=tape -s <location of depot> $bundle \
  @ /var/tmp/$bundle.depot
done
```

Linux Package Management

Linux packages are RPMs, which have the following characteristics:

- RPMs are both uploaded and installed as a unit so there is no distinction between container and installable packages.
- RPMs are versioned.

RPM Metadata

The Opsware System uses the metadata contained in RPMs when creating the package entries in the list of packages. An RPM contains the following metadata - the name, version, and release of the RPM

When a user uploads an RPM, the Opsware System performs the following actions:

- Opens the RPM and parses its metadata
- Registers the RPM as an installable package

Solaris Package Management

Solaris packages are the container packages for Solaris. Solaris packages have the following characteristics:

- A Solaris package contains one or more package instances.
- When a Solaris package contains multiple instances, frequently only a subset of those instances will be installed because users might want to install only certain instances.
- Solaris packages have two basic formats:
  - Filesystem format – the format for packages stored in a directory structure
  - Datastream format – the format for standalone package files. This format is required for uploading Solaris packages into the Opsware System.

The basic unit of Solaris packages is the package instance. Package instances have the following characteristics:

- Package instances are versioned.
- Users add package instances to nodes. The Opsware System adds package instances to, and removes package instances from, servers by using the reconcile function. See "Reconcile" on page 441 in Chapter 9 for more information.

In the Opsware Command Center, you can upload, view, download, and delete Solaris packages, and you can view, deprecate, and attach to nodes the instances that they contain.

The Opsware System supports Solaris packages in the following ways:

- Users upload Solaris packages in the uncompressed datastream file format.

- The Opsware System can install interactive and non-interactive Solaris package instances. Interactive Solaris package instances require response files.
- The Opsware System displays the name and version number for Solaris packages in the following way:

```
SUNW125f-1.0,REV=2001.03.21.17.00
SUNW1394h-11.9.0,REV=2002.04.06.15.27
```

- The Solaris utilities (such as `pkgadd`) use an admin file. The admin file stores settings regarding how the utilities should work. Each Opsware Agent on managed servers includes its own admin file that it uses when installing Solaris package instances. The admin file that the Opsware Agent uses is *only* used by the Opsware System and does *not* set defaults for other applications using `pkgadd`.
- In some instances, a Solaris package might only get partially installed. A partial installation generally occurs when a package contains an installation script (other than the checkinstall script - for example, a preinstall or postinstall script) and that script exits non-zero during package installation. A partially installed Solaris package can be removed as if it were installed as a full package by removing it, or by overwriting it with a new package.
- For more information on pkginfo, pkgadd, and pkgrm, see the man pages.

Response files are text files. The entries in a response file occur as name = value pairs; for example, `BASEDIR="/opt/SUNWexplorer"` is a valid entry.

The Opsware System supports response files in the following ways:

- Users create response files outside of the Opsware System by using the `pkgask` Solaris utility.
- By using the Package Properties page in the Opsware Command Center, users upload, overwrite, view, and delete response files that are associated with Solaris package instances.
- Each response file is accessible *only* in the context of the Solaris package instance to which it belongs.
- Each Solaris package instance can have zero or one response file. Response files are not shared by different Solaris package instances.
- Attaching an interactive package to a node includes the response file because the Opsware System stores the response file with the package. You do not need to attach the response file to the node.

- After a Solaris package instance has a response file, the Opsware System uses that response file whenever the Solaris package instance is installed.
- If a Solaris package instance requires a response file and that file is missing in the Opsware Command Center, the Opsware System might report an error when any server is reconciled with that Solaris package instance.

When a user uploads a Solaris package, the Opsware System performs the following actions:

- Opens the package and parses its metadata
- Automatically creates entries in the list of packages for the package instances in the package and registers them as installable
- Registers the Solaris package as uninstallable

Solaris Package Metadata

The Opsware System uses the metadata contained in Solaris packages when creating the package entries in the list of packages. A Solaris package contains the following metadata - the name, version, and description of each package instance in the package.

Prerequisites to Solaris Package Management

The Solaris package must be in datastream format before you can upload it to the Opsware Software Repository. If it is in file system format, you can convert it by using the `pkgtrans` command:

```
pkgtrans -s <location of package> <new package> all
```

Windows Package Management

The Opsware System supports the following Windows packages:

- Microsoft Installer Packages
- Microsoft Hotfixes, Security Patches, and Service Packs

Microsoft Installer Packages

Microsoft Installer packages (MSI) have the following characteristics:

- Contain all the information that the Microsoft Installer requires to install an application or product

- Contain information that the installer requires to run the setup user interface

MSI packages contain:

- An installation database
- A summary information stream
- Data streams for various parts of the installation

The Opsware System supports .msi files as installable software.

MSI Package Metadata

The Opsware System catalogs each MSI package by its ProductName and ProductVersion. These properties are defined in the Properties table of the MSI installation database. When you upload an MSI package to the Opsware System, you are required to provide ProductName and ProductVersion exactly as they appear in the Properties table.

To discover the ProductName and ProductVersion, use the Orca tool that Microsoft provides as part of its MSI SDK, available for download from

www.microsoft.com

Perform the following steps to discover the ProductName and ProductVersion of an MSI package:

1. Launch the Orca application.
2. Select File ➤ Open to open the target MSI package file.
3. In the Tables Column, select Property.
4. Note the exact value for the ProductName and ProductVersion properties.

Prerequisites to MSI Package Management

The Opsware System supports the Microsoft Windows Installer versions 1.1 and 2.0. Version 1.1 is included with Windows 2000, and version 2.0 is included with Windows 2003.

Windows NT does not include a version of the Windows Installer, but the Microsoft Windows redistributable can be obtained for download at

www.microsoft.com

or by including the `-withmsi` option on the Opsware Agent Installer command line.

See “Server Assimilation” on page 139 in Chapter 2 for information about the steps to install an Opsware Agent on a server.

Microsoft Hotfixes, Security Patches, and Service Packs

These packages include:

- Hotfixes
- Service Packs
- Security Patches

Hotfixes are issue specific and should only be applied if you experience the exact issue addressed by the hotfix, and only if you are using the current operating system version that has had the latest service pack applied.

Service packs are groups of hotfixes. They are more thoroughly tested than individually-released hotfixes, and are available to all customers, not just those with the specific problem.

Security patches are similar to hotfixes, but are mandatory if you are experiencing the specific problem they are created to address, and they need to be deployed as soon as they are made available.

Microsoft Patch Metadata

When you upload a Service Pack, Opsware requires the user to provide the version of the service pack.

When you upload Hotfixes and Security Patches, Opsware requires the user to provide the operating system version and the patch type.

Microsoft Patch Management Prerequisites

You must have Internet Explorer 5.0.1 or later installed on the server in order to use its native Microsoft Baseline Security Analyzer (MBSA) tool. The Opsware System uses the MBSA tool for patch management.

You must also have an XML parser such as MSXML version 3.0 SP2 installed in order for the tool to function correctly.

Windows NT Service Pack 6a must be installed in order to add Microsoft Installer support to Windows NT.

ZIP Package Management

This section provides information on ZIP package management within the Opsware System and contains the following topics:

- ZIP Package Management Support
- ZIP Packaging
- Creating ZIP Packages
- Uploading ZIP Packages
- Defining Package Installation and Remove Scripts
- Editing Properties for ZIP Packages
- Info-Zip Compatible Zip Packages
- Windows Performance for Uploading Packages

ZIP Package Management Support

The Opsware Command Center adds support for ZIP packages on the following platforms:

- Windows NT4
- Windows 2000

When you specify the installation directory, post-installation, and pre-uninstallation script names for a ZIP package, do not use quotation marks to enclose the entire directory path and the script names.

ZIP Packaging

Use ZIP packages primarily to deliver code that can be run on a server. You can also use them to deliver application files for installing applications.

When a user installs a ZIP package on a server, the files are automatically extracted and saved to a directory that the user selects; otherwise, a default directory is used. The Opsware System keeps track of all ZIP packages that it has installed, which prevents you from installing a ZIP package with the same name twice.

A ZIP package has no limits or restrictions on the size, format, or number of files that it contains.

The Opsware System supports ZIP encapsulation for application package files that were built using other standalone installation programs, for example, InstallShield.

The Opsware System requires *silent install* operation for programs designed for interactive installation. When you package these program files to upload to the Opsware System, use the silent install options to play back automatic responses to provide unattended installation.

For information on how to construct ZIP files that use silent install features for unattended installation operations, refer to the documentation provided with the archive program that you are using.

Creating ZIP Packages

The Opsware System supports the ZIP file format for application package files that are built using non-MSI standalone installation programs, for example, InstallShield. Programs such as InstallShield were originally designed to provide for interactive installation. However, using the silent install feature, InstallShield users can play back a recording of a previous application installation that creates an unattended installation file with a suffix ISS.

The interactive installation recording is saved in the form of a setup.iss file that contains the responses to the interactive dialog boxes and popup menus that typically display during an interactive installation. After the response file is recorded, you can pass the setup.iss file as an argument to setup.exe executed from the command line to perform an unattended installation.

Similarly (using InstallShield), uninstalling an application can be set up to run unattended using the UnInst.exe command invoked with the -a and -y options to instruct the installer to run uninstall in silent mode.

See the documentation provided with your specific installer software for more information on silent install features and options.

Uploading ZIP Packages

You can upload, re-upload, download, delete, deprecate, and attach ZIP packages to Roles with the OCC. You can also use the Opsware Command Line Interface (OCLI) to upload and download ZIP packages.

Enter Windows ZIP file as the value for the `--pkgtype` argument to upload a ZIP package by using the OCLI.

See Appendix A, "Opsware Command Line Interface" for more information.

Defining Package Installation and Remove Scripts

Because you can invoke silent installation and uninstallation from the command line, you can create scripts that perform silent installation and uninstallation. You can then include the scripts as part of the package file properties that you specify when you upload a ZIP file to the Package Repository.

Editing Properties for ZIP Packages

After you upload a ZIP package to the Package Repository, the Opsware Command Center displays a page in which you can enter additional package properties. On the Edit Properties page, you can define scripts to run when you install or uninstall ZIP packages.

Perform the following steps to edit the properties for a ZIP package as defined in the Opsware System:

1. Click the ZIP package link for the package that you want to edit. The Manage Packages: Edit Properties page appears.

   The Manage Packages: Edit Properties page for ZIP packages does not include the Install and Remove Flags text boxes.

2. Enter an installation directory in the Installation Directory text box. The default installation directory, if nothing is entered, is `%SystemDrive%\Program Files\[basename of ZIP file]`.

When you specify the installation directory, post-installation, and pre-uninstallation script names for a ZIP package, do not use quotation marks to enclose the entire directory path and the script names.

When you uninstall a ZIP package, the extracted ZIP files that were installed on the server are not removed from the server. To uninstall those files, you must run an uninstallation script by specifying that script in the Pre-Uninstall Script Filename text box.

3. Configure the location of the post-installation or pre-uninstallation scripts and enter the script names in the Post-Install Script Filename or Pre-Uninstall Script Filename text boxes. The installation and uninstallation scripts must be included in the ZIP archive.

4. Click the *If non-zero return status, halt operation* check box for either the installation or uninstallation so that the script stops when it fails.

Extracting a ZIP package fails if the top-level directory in the file's pathname does not exist.

Info-Zip Compatible Zip Packages

The Opsware System offers package management support for Info-Zip compatible .zip packages. The files that are archived within Info-Zip are installable files on the Opsware System. You can download the .zip package creation tool from

www.info-zip.org

Info-Zip Compatible Package Metadata

The Opsware System uses the ZIP package filename to uniquely identify a ZIP package.

Prerequisites of Info-Zip Compatible Package Management

Full support for managing ZIP packages on a server is included with the Windows Opsware Agent.

Windows Performance for Uploading Packages

When you upload packages from a Windows computer, users can improve the performance of the computer used to upload by changing TCP stack registry settings that affect upload speeds. The recommended change to the Windows registry file increases the default tcp-send buffer size from 8 KB to 16 KB.

Consult your system administrator before you make this change.

Perform the following steps to change the tcp-send buffer setting:

1. Using regedit, navigate to the following registry key:

```
HKEY_LOCAL_MACHINE
  SYSTEM
    CurrentControlSet
      Services
        Afd
          Parameters (Create this key if it does not already
exist)
```

2. Set the following value for the key:

```
Name: DefaultSendWindow
Value Type: REG_DWORD
Value: 16384 (decimal)
```

After you set the value, reboot the machine for the changes to take effect.

Package Management Tasks Overview

The Package Management feature of the Opsware System provides tools for defining package types and uploading packages onto managed servers. The tasks that you can perform when you use this feature include:

- Displaying Packages
- Searching for Packages
- Viewing Packages Assigned to Nodes
- Uploading a Package
- Overwriting a Package
- Editing Package Properties
- Deleting a Package
- Deprecating a Package
- Downloading a Package

Displaying Packages

By default, the Opsware Command Center displays results (up to 10 package listings per page) in alphabetical order by name. Links to consecutive pages (if there are more than 10 packages to display) are provided at the bottom of the page. In addition, you can click the Show All link to display up to 500 packages at a time in a scrollable list.

You can also sort the results of the packages display by clicking any column heading, for example, Name, Description, or OS Version, to sort the list by that column and toggle the display of packages between ascending and descending order.

Perform the following steps to display results:

1 From the navigation panel, click Software ➤ Packages.

You can also use the Search option in the navigation panel to search for specific package names or locate packages for specific customers.

The Packages: Browse Packages page appears that lists packages available in the Software Repository. By default, the packages for the customer with which the user is associated display, as Figure 5-1 shows.

Figure 5-1: Packages: Browse Packages Page

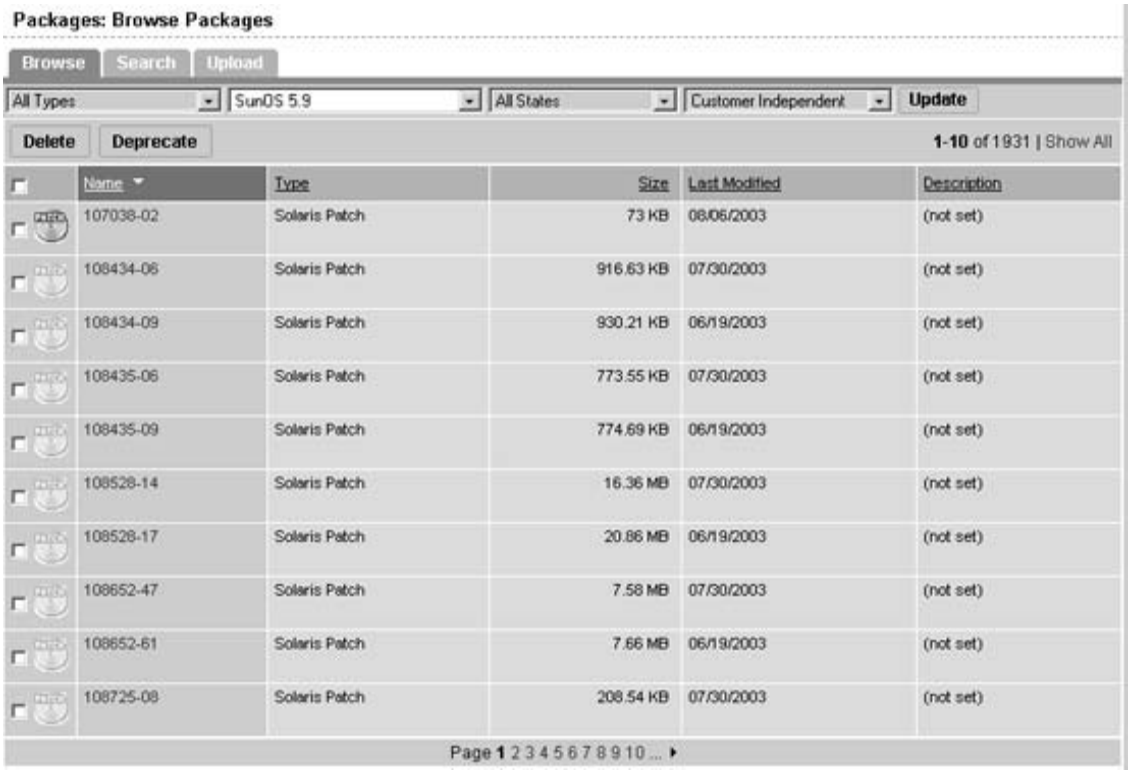

For each package available for a selected customer, operating system, type, and state, the Opsware Command Center displays the package's name, size, last modified date, and description, in addition to the selected operating system and customer.

2 To display the properties of the package, click the package name.

See "Editing Package Properties" on page 300 in this chapter for more information.

Details: Filtering the Packages to Display

The top panel of the page displays four filters that you can specify to qualify the packages that the Opsware Command Center displays, as Figure 5-2 shows.

Figure 5-2: Filters for Selecting Packages to Display

- Package Type – specifies the package type for the operating system that you select.
- Operating System – specifies the operating system for which packages are available. You can select All Operating Systems or choose specific operating systems that your Opsware administrator sets up.
- Package State – specifies whether the display lists all packages or only those that are available or deprecated. By default, the Opsware Command Center shows all packages. Selecting *Deprecated* from the list displays only those packages that are deprecated in the Software Repository. Selecting *Available* displays only those packages that are not deprecated.
- Customer – specifies the customer for which packages are available. Options for customer selections are those that your Opsware administrator defines by using the Opsware Administration channel.

Searching for Packages

Knowing how package names are defined helps you perform more precise searches. For example, the package name:

```
iPlanet_Web_Server-4.1sp9-LC~0.sparc64.rpm
```

specifies a package that contains iPlanet Web Server software, version 4.1sp9, running on the Sun Solaris operating system.

Perform the following steps to search for packages:

1 From the navigation panel, click Software ➤ Packages.

You can also use the Search options in the navigation panel to search for specific package names or locate packages for specific customers.

The Packages: Browse Packages page appears that lists the packages in the Software Repository that match the criteria set by the filters at the top of the page.

By default, the packages display for the customer with whom the user is associated.

2 Click the Search tab. A search form appears in which you can enter more specific criteria to locate packages, as Figure 5-3 shows.

Figure 5-3: Searching for Packages

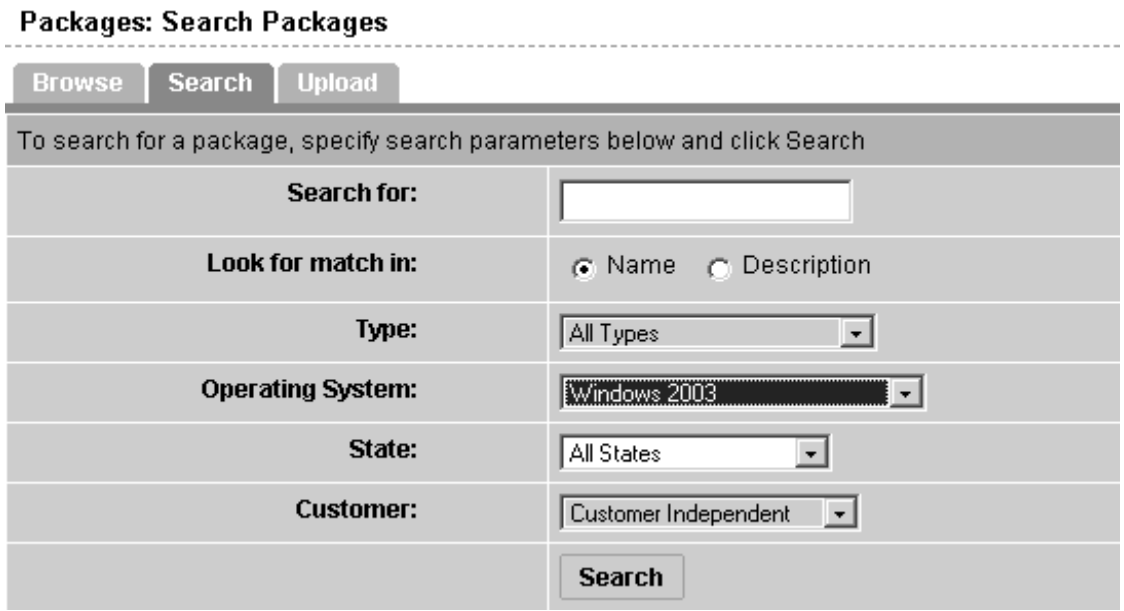

3 Enter your search criteria:

- In the Search for field, enter a simple text string to match the name or description of specific packages. You can also specify wildcard characters to match multiple characters as part of the search string. (The Opsware Command Center automatically adds a "*" wildcard character after any text string you specify, if you do not specify any wildcards in your search string.)
- Select a Look for match in option button to specify whether to match a text search string in a package's name or description (that was entered in the Package Properties page).
- Restrict your search by package type, operating system, file state, and customer:
  - Type – specifies the package type for the operating system selected. By default, all package types display.

- Operating System – specifies the operating system for which packages are available. You can select All Operating Systems or choose specific operating system selections that your Opsware administrator sets.
- State – specifies whether the display lists all packages or only those that are available or deprecated. By default, the Opsware Command Center shows all packages. Selecting Deprecated from the list displays only those packages that are deprecated in the Software Repository. Selecting Available displays only those packages that are not deprecated.
- Customer – specifies the customer for which packages are available. Options for customer selections are those that your Opsware administrator defines by using the Opsware Administration channel.

4. Click the Search button. The Opsware Command Center returns a list of packages that match the search criteria that you specified.
5. Click a package name to view the package properties.

   OR

   Click the Return to Search Packages hyperlink.

Viewing Packages Assigned to Nodes

Perform the following steps to view packages that are assigned to nodes:

1. Locate the package whose attachment to nodes you want to view. Locate the package either by browsing (See "Displaying Packages" on page 288 in this chapter for more information) or by searching (See "Searching for Packages" on page 290 in this chapter for more information).
2. Click the package name link for the package whose attachment to nodes you want to view. The Packages: Edit Properties page appears.
3. Click the Nodes tab. The Package: View Nodes page appears. The page displays the nodes to which the package is already attached, as Figure 5-4 shows.

If you are viewing the properties for an AIX LPP, Solaris package, Unknown package, Windows utility, Microsoft Patch Database, Build Customization Script, or HP-UX depot, the Nodes tab does not appear because these packages cannot be attached to nodes.

Figure 5-4: Viewing Packages Attached to Nodes Page

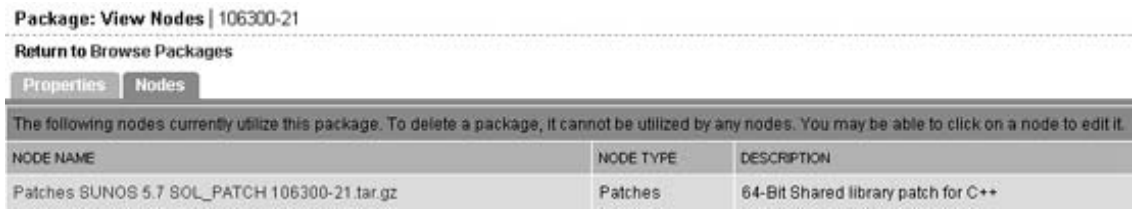

4 (Optional) Click the node name to edit the node. For example, you can edit the node to remove the package.

You cannot add deprecated packages to nodes.

Uploading a Package

Each operating system supported by the Opsware System supports certain package types. See “Supported Operating Systems and Package Types” in Chapter 3 for more information.

If a package that is being uploaded already exists in the Software Repository, the Opsware System overwrites the package. See “Overwriting a Package” in Chapter 3 for more information.

If you upload Solaris patch clusters that contain patches that already exist in the Software Repository, the patches are overwritten. However, the Opsware System preserves any reboot options or flags set for the patches in the Opsware Command Center.

If an update fileset is part of many different APARs, the Opsware System can take a long time to upload the LPP that contains it, because it must create a large number of APARs in the list of packages. The package upload can appear to time-out in the Opsware Command Center. However, the Opsware System continues to upload the package.

When you upload a Windows MSI package, the information that you enter about the package must match the internal data (ProductName and ProductVersion) stored in the file or the Opsware System reports errors when the package is installed on servers.

To obtain the ProductName and ProductVersion from an MSI package, use the Orca tool from Microsoft Corporation. You can download this tool as part of the MSI software developers' kit (SDK) v1.2. See "MSI Package Metadata" on page 281 in this chapter for information about how to obtain and use Orca.

Perform the following steps to upload a package:

1. From the navigation panel, click Software ➤ Packages. The Packages: Browse Packages page appears.
2. Click the Upload tab. The Packages: Upload Package page appears, as Figure 5-5 shows.

Figure 5-5: Packages: Upload Package, Specify Basic Properties Page

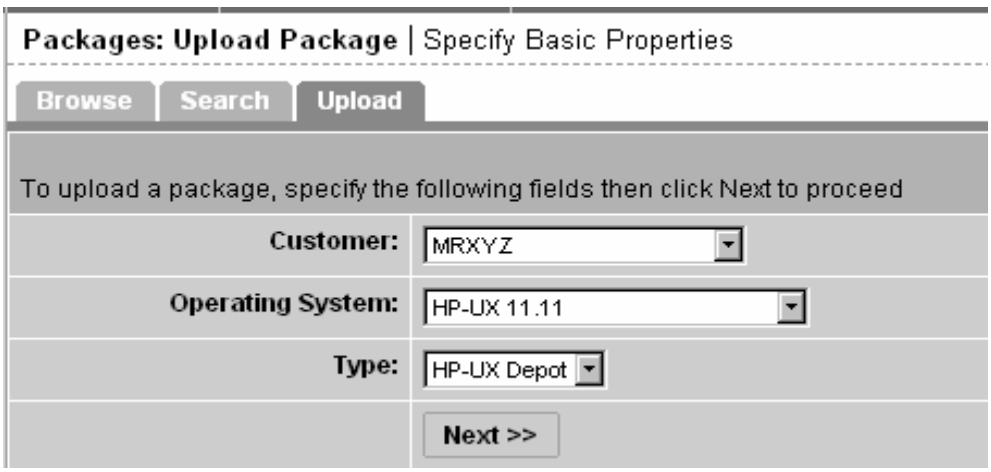

3. Specify the operating system and customer for which the package is to be used:

   Customer - specifies the customer for which packages will be used. Options for customer selections are those that your Opsware administrator defines by using the Opsware Administration channel.

When you upload any patch, Windows utility, or Microsoft Patch database, the file is automatically associated with Customer Independent regardless of the customer that you select from the drop-down list. HP-UX depots that contain both products and patch products cannot be uploaded to a specific customer. They can only be uploaded to Customer Independent.

Operating System - specifies the operating system for which the package is to be used. You can select from specific operating systems that your Opsware administrator set up.

When you select a Windows utility or Microsoft database, the operating system is automatically set to Windows 2000.

Type - specifies the package type for the operating system selected. Certain package types, for example, patches or Windows utilities, force the customer to be Customer Independent.

The Opsware System verifies that RPMs, Solaris patch clusters, AIX LPPs, Solaris packages, and HP-UX depots uploaded are the correct package type.

4 Click the Next button. If you select Windows MSI or OS Service Pack as the package type, a second form appears that prompts you to enter additional metadata for the package, as Figure 5-6 shows.

Figure 5-6: Packages: Upload Package Specify Location Page

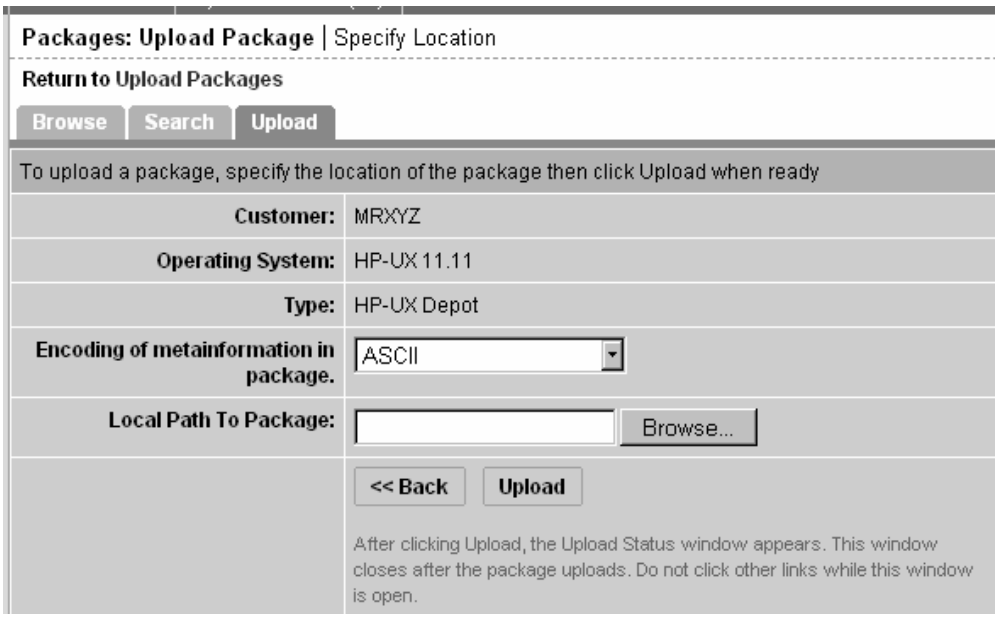

5 If necessary, enter additional required information for the package and click the Next button. A page appears that prompts you to specify the location of the package file.

- In the **Local Path to Package** field, enter the name and directory location of the package that you want to upload or click the Browse button to locate and select the package to upload.
- In the **Encoding of metainformation in package** field, select the encoding scheme to be used by the package.

You need to specify an encoding scheme so that the Opsware System can extract the metadata contained in the package and correctly display the information in non-ASCII characters in the Opsware Command Center (for example, in the Package Properties pages).

6 Click the Upload button.

If a package (excluding AIX LPPs, HP-UX depots, and Solaris packages) with the same name already exists in the Software Repository for the same customer and operating system, the Opsware Command Center overwrites the existing file after it prompts you to confirm your choice. See “Overwriting a Package” on page 299 in this chapter for more information.

After you upload the package, the Opsware Command Center displays a page where you can enter additional package file property information.

If you are uploading a Solaris package, you can upload a response file for each of its instances by editing the response file's package properties for those instances.

See “Editing Package Properties” on page 300 in this chapter for information about additional properties that you can specify and the steps to upload a response file for a Solaris package instance.

Encoding Schemes for Package Metadata and Scripts

In Opsware System 4.7, you can specify encoding schemes for package metadata and scripts in the following ways:

- Specify the encoding scheme for package metadata when uploading packages in the Opsware Command Center (in the Packages channel and Software Install Wizard) or by using the Opsware Command Line Interface (OCLI).

  When specified, the Opsware Command Center correctly displays in non-ASCII any package metadata, description fields, and error and status message returned by the operating system of the managed servers.

- Specify the encoding scheme for scripts when uploading them in the Opsware Command Center (in the Run Distributed Script Wizard and Scripts channel).

  A user must specify an encoding scheme for an uploaded script so that the Opsware System can convert the bytes inside the script into UTF-8 format by using the encoding scheme with which the script was created.

  After scripts run, users can download a zip file that contains the results encoded in UTF-8 format.

  For example, on Unix operating systems, you can use iconv (the code set conversion function) to interpret the downloaded results of the script execution.

The Opsware Command Center includes the following selections for encoding schemes:

- ASCII
- BIG5
- BIG5-HKSCS
- CP850, CP862, CP866, CP874, CP932, CP949, CP950, CP1133, CP1250, CP1251, CP1252, CP1253, CP1254, CP1255, CP1257, CP1258, CP1266
- EUC-CN, EUC-JP, EUC-KR, EUC-TW
- GB18030, GBK
- GEORGIAN-ACADEMY
- GEORGIA
- N-PS
- HZ
- ISO-2020-CN, ISO-2020-CN-EXT, ISO-20202-KR, ISO-2022-JP, ISO-8859-1, ISO-8859-2, ISO-8859-3, ISO-8859-4, ISO-8859-5, ISO-8859-6, ISO-8859-7, ISO-8859-8, ISO-8859-9, ISO-8859-10, ISO-8859-13, ISO-8859-14, ISO-8859-15, ISO-8859-16
- JOHAB
- KOI8-R, KOI8-RU, KOI8-T, KOI8-U
- MULELAO-1
- SHIFT_JIS
- TCVN
- TIS-620
- UCS-2, UCS-4
- UTF-8
- VISCII

Overwriting a Package

You can overwrite an uploaded package with a new package that has the same name (the contents of the package might be different.) The new package keeps the same property information as the current uploaded package and the same nodes to which the current package is already assigned or attached.

You *cannot* overwrite container packages (LPPs, HP-UX depots, or Solaris packages). To update a container package in the Opsware Command Center, delete the existing package and upload the file again. You can only delete the container package if none of its installable packages are attached to nodes.

When you overwrite a package, the Opsware System deletes the existing package and uploads the new package. If you have manually changed any configuration in the package on a server, those changes are removed when you reconcile the server. The Opsware System removes manual changes because the package was upgraded with a new version.

If you upload a Solaris patch cluster that contains patches that already exist in the Software Repository, the patches are overwritten. However, any reboot options or flags set for the patches in the Opsware Command Center are not affected.

See “Editing Package Properties” on page 300 in this chapter for more information.

Perform the following steps to overwrite a package:

1. Locate the package that you want to replace with a newer version. Locate the package by browsing (See “Displaying Packages” on page 288 in this chapter for more information) or by searching (See “Searching for Packages” on page 290 in this chapter for more information).
2. Click the package name link for the package that you want to overwrite. The Packages: Edit Properties page appears.
3. Click the Replace button at the bottom of the page. The Packages: Upload Package page appears.
4. In the Local Path to Package field, enter the name and directory location of the newer package you want to upload.

   OR

Click the Browse button to locate and select the package to upload.

5 Click the Upload button.

Editing Package Properties

After you upload a new file or select an existing package in the Software Repository, the Opsware Command Center displays a page that you can use to add or update additional package properties.

You *cannot* change the operating system or customer association of a package by editing the package properties.

Perform the following steps to edit package properties:

1 Locate the package whose properties you want to edit. Locate the package by browsing (See "Displaying Packages" on page 288 in this chapter for more information) or by searching (See "Searching for Packages" on page 290 in this chapter for more information).

2 Click the package name link for the package that you want to edit. The Packages: Edit Properties page appears.

Some of the properties fields discussed in Step 3 might not appear on the page because the fields and information that appear are based on package type.

3 Edit the following properties for the package:

- Description – specifies a short description that is used to indicate the package's contents.
- Upgradable: Allow package to be installed with -U option - an optional argument used to remove the old version of RPMs and install the new version in a single step. By default, this flag is selected.
- Install Flags – optional arguments that you can specify to run when this package is installed on servers (a node that includes this package is assigned to a server and the server is reconciled).

- Pre-Install Script – enter the script required to run before you install the package.
- End this and subsequent installs if this script fails – select this option if you want this and all other installations to stop if this script fails.
- Post-Install Script – enter the script required to run after you install the package.
- End subsequent installs if this script fails – select this option if you want to end all installations if this script fails.
- Reboot on Successful Install – click this option if you want the system to reboot when the package is successfully installed.
- Uninstall Flags – optional arguments that you can specify to run when this package is uninstalled (a node including this package is removed from a server and the server is reconciled).
- Pre-Uninstall Script – enter the script required to run before the package is uninstalled.
- Post-Uninstall Script – enter the script required to run after the package is uninstalled.
- Reboot on Successful Uninstall – click this option if you want the system to reboot when the package is successfully uninstalled.

You must specify valid command line options in the Install and Uninstall Flags text boxes. Specifying invalid command line options in these text boxes can cause package installations and uninstallations to fail when Opsware users reconcile servers. Scripts must likewise be entered with care for the same reason.

Do not use quotation marks to enclose directory paths or script names or the installation will fail.

- Deprecated – optional selection that you can use to deprecate a package uploaded in the Software Repository. After you deprecate a package, it is no longer available to be added to new nodes.
- Notes - Enter any notes about the installation or uninstallation.

For Windows Zip Files:

Unlike the previous fields, the name of the script file is entered here, not the script itself. Also, the Install Flags and Uninstall Flags fields do not appear for this type of file.

- Installation Directory – enter the path where you will install the zip file. If you do not enter a path, the default directory is `%SystemDrive%\Program Files\[basename of zip file].`
- Post-Install Script Filename – enter the name of the post-installation script included in the archive being installed.
- Pre-Uninstall Script Filename - enter the name of the pre-uninstallation script included in the archive being uninstalled.

Through the Edit Properties page, you can upload, delete, and overwrite response files for Solaris package instances. The buttons to manage response files only appear when you are viewing the properties for a Solaris package instance.

4 To upload a response file for a Solaris package instance, follow these steps:

1. Click the Upload button in the Response File field. This button and field only appear for Solaris package instances. A popup window appears that prompts you to specify the path for the response file.
2. Specify the file path by entering it in the text box or browsing.
3. Click the Upload button. The popup window closes and the response filename appears in the page.

See "Solaris Package Management" on page 278 in this chapter for information about how to use response files with Solaris package instances.

5 To delete a response file for a Solaris package instance, click the Delete button in the Response File field. This button and field only appear for Solaris package instances.

6 To overwrite a response file for a Solaris package instance, click the Replace button and specify the path to the new file. The file can have the same name or a different filename as the existing file. The Opsware System replaces the response file with the new file.

7 Click the Save button.

You can also replace (overwrite), delete, or download a package directly from the Edit Properties page.

Deleting a Package

You can delete packages so that they are no longer available to add to nodes. The ability to delete a package depends on the package type that you select.

Deleting a Solaris patch cluster does *not* delete the patches contained in the cluster from the Software Repository.

When you delete a Solaris package instance, any response file associated with the package is also deleted.

Restrictions on Deleting Packages

You cannot delete packages if the following conditions are true:

- The package is attached to a node.
- The package is attached to a patch node that has a server also attached.
- The package is attached to a patch node that is in a template.
- Any of the items contained in a package are attached to a node.
- The user does not have permission to access the resources for that customer.
- The package type is not a physical package type.
- The package type is not deletable because the system needs it.
- The package contains a package that is attached to a patch node that has a server also attached.
- The package contains a package that is attached to a patch node that is in a template.

Do *not* delete a Solaris patch when the patch is contained in a Solaris patch cluster. Deleting a Solaris patch contained in a Solaris patch cluster can cause problems when you install the patch cluster on a server.

Perform the following steps to delete a package:

1. From the navigation panel, click Software ➤ Packages. The Packages: Browse Packages page appears.
2. Select the check boxes next to the packages that you want to delete.

You can also click the check box at the beginning of the list to select or clear package selections.

3 Click the Delete button. A confirmation page appears that prompts you to confirm the deletion.

The Opsware Command Center also indicates whether packages selected for deletion are still assigned to nodes. To delete packages, you must remove the packages from assigned nodes before you delete them, as Figure 5-7 shows.

Figure 5-7: Attempting to Delete a Package Attached to a Node

4 Click the Delete button to remove the packages from the Software Repository.

Deprecating a Package

You can deprecate packages so that they are no longer available to add to nodes. The ability to deprecate a package depends on the package type that you select.

Deprecating an APAR does not deprecate the update filesets within it. Conversely, deprecating an update fileset does not deprecate the APARs of which the update fileset is a part. If a user reconciles a node for an APAR that contains a deprecated update fileset, the Opsware System installs the deprecated fileset. Additionally, users can install available update filesets for deprecated APARs. See "AIX Package Management" on page 272 in this chapter for more information.See "About AIX Patches" on page 418 in Chapter 8 for information about how the Opsware System manages APARs.

In most cases, a new package replaces the deprecated version. You can record this information in the Notes field for the package that you are deprecating.

You can also deprecate a package while you view the package properties. See “Editing Package Properties” on page 300 in this chapter for more information.

Perform the following steps to deprecate a package:

1. From the navigation panel, click Software ➤ Packages. The Packages: Browse Packages page appears.
2. Select the check boxes next to the packages that you want to deprecate.
3. Click the Deprecate button. The Packages: Deprecate page appears, as Figure 5-8 shows.

Figure 5-8: Packages: Deprecate Page

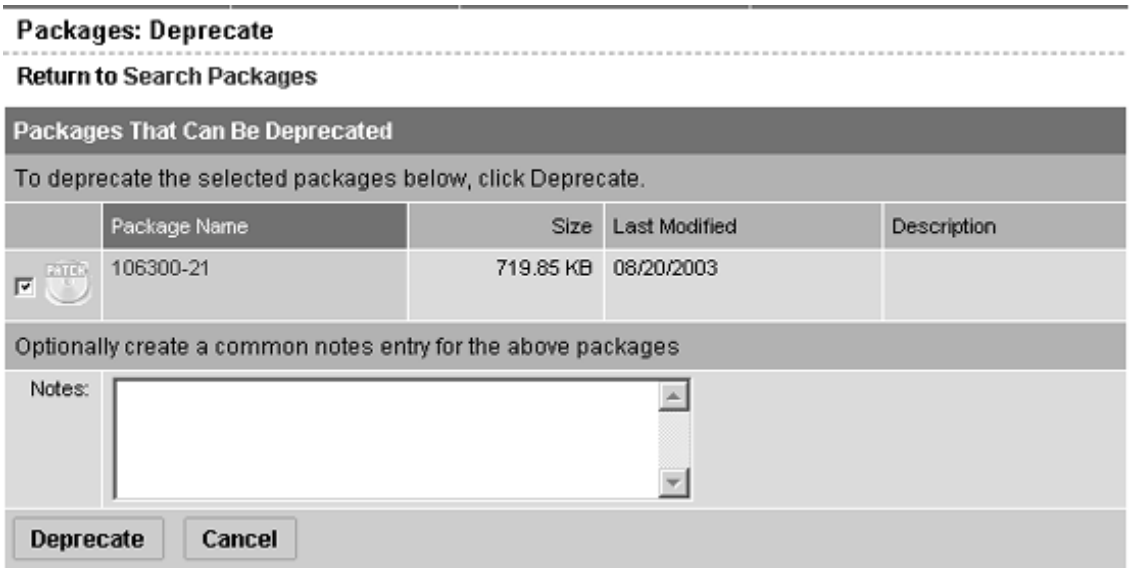

4. To record the version of a new package that replaces the deprecated version, enter the information in the Notes field.
5. Click the Deprecate button.

Restrictions on Deprecating Packages

Deprecation is not allowed if the following conditions are true:

- The package is not attachable to nodes.
- The user does not have permission to access resources for that customer.
- The package is already deprecated.

Downloading a Package

You can download a package to your local computer so that you can check the installation of the package on a test or staging machine.

Package types that are not physical files – like APARs – cannot be downloaded so no download button appears for those file types.

Perform the following steps to download a package:

1. Locate the package that you want to download. Locate the package by browsing. (See "Displaying Packages" on page 288 in this chapter for more information) or by searching (See "Searching for Packages" on page 290 in this chapter for more information).
2. Click the package name link for the package that you want to download. The Packages: Edit Properties page appears.
3. Click the Download button at the bottom of the page.

Chapter 6: Application Provisioning Setup

IN THIS CHAPTER

This chapter discusses how to prepare to provision applications by setting up nodes in the Software Tree and creating templates. Topics in this chapter include:

- Software Tree
- Managing Nodes on the Software Tree
- Software Attached to Nodes
- Custom Attributes Set for the Environment
- Working with Templates
- Folders and Templates

A user must have specific permissions to manage applications that use the Opsware Command Center. Contact your Opsware administrator to obtain the necessary access rights.

This chapter assumes that users have already packaged software applications and uploaded the packages to the Software Repository. Chapter 5, “Package Management” on page 267 of this guide for more information about how to upload packages to the Software Repository.

Software Tree

This section provides information on the software tree within the Opsware System and contains the following topics:

- Software Tree Overview
- Example of a Software Tree
- Guidelines for Setting Up the Software Tree

- How to Use the Software Tree
- Understanding How Software Is Reconciled onto Servers
- When to Reconcile
- How to Reconcile

Software Tree Overview

Many IT organizations define policies for creating systems in their operational environment by setting standards for configuration and processes. Defining policies and standardizing repeated processes minimizes custom work and creates efficiency and reliability because the same policies are used throughout the installation.

Opsware technology enables IT organizations to implement their policies by using Opsware's model-based approach to manage an operational environment. The model can:

- Encapsulate policies for creating systems in an operational environment
- Contain a comprehensive picture of an operational environment

See "Opsware System's Model-Based Approach" on page 3 in Chapter 1 for more information.

The Opsware Command Center visualizes the model as a tree called the *Software Tree*. The Software Tree is made up of nodes and subnodes, which model the interrelationships and dependencies among the software and customer accounts in an operational environment. Users navigate the Software Tree to perform specific operations.

The top level of the Software Tree for Applications has six default main categories. These categories display on the Applications page as Figure 6-1 shows. You can access the Applications page from the navigation panel by clicking Software ➤ Applications.

Figure 6-1: Top Level of Applications Page

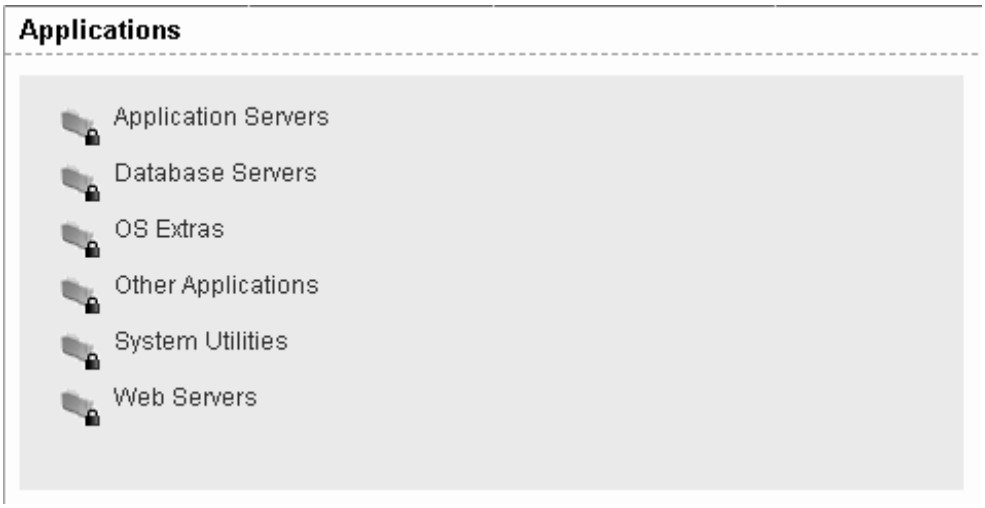

These six default main categories are defined as follows:

- Application Servers, such as WebLogic, connect the database information with the client program, such as the Web browser.
- Database Servers, such as Oracle or Microsoft SQL Server, contain and manage the database.
- OS Extras, which are optional, contain applications that need to be installed directly after an operating system, for example, VCS or special utility software.
- Other Applications is a catch-all category for software that does not fit into the other categories. Try to add software to one of the other categories before you add a node in this category.
- System Utilities are such as PC Anywhere, JDK, Open SSL, or Winzip.
- Web Servers, such as Apache, Microsoft Internet Information Server, or iPlanet are server processes running at a Web site that deliver Web pages in response to browser requests.

No additional categories can be added, and categories cannot be deleted or modified.

The Software Tree has the following characteristics:

- Each point in the Software Tree is called a *node*.
- A node inherits properties or software from the nodes above it.
- Software or a server might be attached to a node depending on its location in the Software Tree. See "Software Attached to Nodes" on page 334 in this chapter for more information.
- Attaching a server to a node determines what software is installed on that server and how it is configured.
- Users can add nodes and subnodes within each category.
- The node information at the top of the Software Tree is more general, and as you travel farther down the tree, each successive subnode contains more details and specific information that relates to the node above it.

Example of a Software Tree

Figure 6-2 shows how three categories and the nodes within them might be defined at a site.

Figure 6-2: Example of How Nodes and Subnodes Are Defined Within the Software Tree

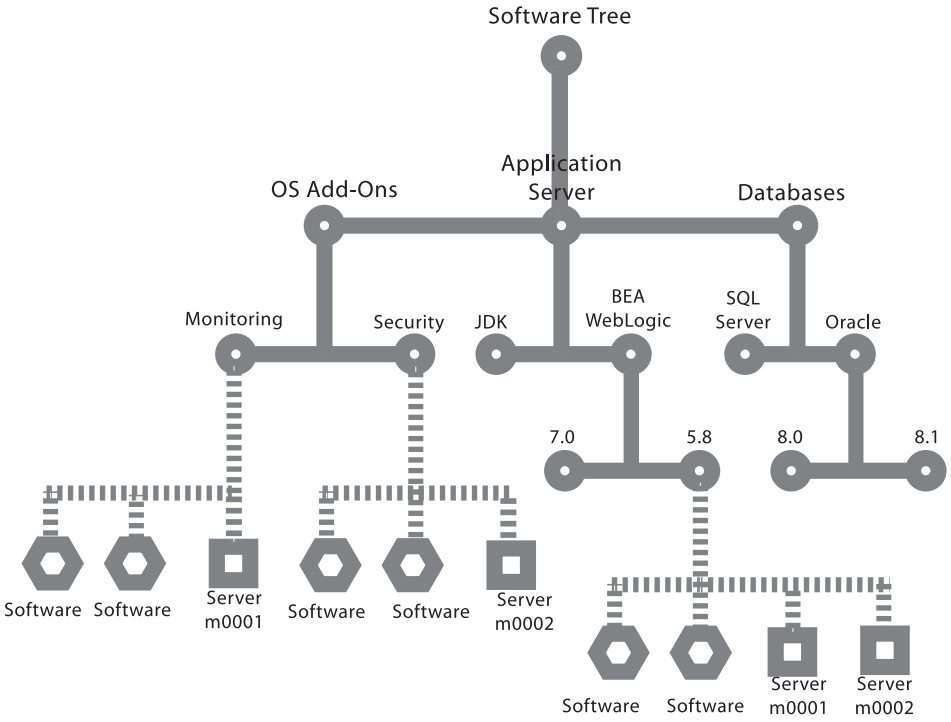

In the figure, servers and software are attached to the Software Tree in the following locations:

- Server hostname m0001 is attached to the Application Servers BEA WebLogic 5.8 node and the Monitoring node.
- Server hostname m0002 is attached to the Application Servers BEA WebLogic 5.8 node and the Security node.

Users can attach servers to multiple nodes within the same category, as well as nodes in different categories.

Guidelines for Setting Up the Software Tree

When you build the model of an operational environment, the Software Tree should:

- Describe all facets of the managed environment.
- Be customizable.
- Be extensible—when something new is learned, the model is updated.

To keep the information consistent, Opsware Inc. recommends that users follow this structure when they create new nodes within a category:

Category Name ➤ Product Name ➤ Operating System ➤ Product Version ➤ Node Version

For example:

Application Servers ➤ Vignette ➤ SunOS 5.7 ➤ 5.05 ➤ v1

- Some nodes do not have software or attributes attached to them because either they inherit the software or attributes from their parent nodes, or they are used for organizational purposes.
- In most cases, you attach servers to nodes at the end of the Software Tree path because servers need to be attached to nodes at the point where the proper operating systems, software, versions, and so forth that belong on that particular server have been identified.
- Opsware Inc. recommends organizing your Software Tree with the operating system preceding the product version because all software versions might not be available for each operating system.
- The best general rule is to keep the structure consistent, organized, and simple.

How to Use the Software Tree

Figure 6-3 demonstrates the node hierarchy as it goes from the general to the specific—System Utilities is the category, Winzip is a node in that category, Windows 2000 is a subnode of Winzip, version 8.0 is a subnode of Windows 2000, and Winzip Patch and Addons are subnodes of version 8.0. The tabs used for all functions in Applications are also shown.

Figure 6-3: Application Hierarchy

Locked nodes appear in the Software Tree with an icon to show that the set of software contained in the node is locked. Only users with the correct permissions can lock a node, unlock a node, and edit a locked node.

When you navigate through the Software Tree, you can perform actions on a node by using the tabs. Table 6-1 gives descriptions of the available tabs.

Table 6-1: Tabs Available for Applications

TAB PAGE	DESCRIPTION
Properties	View general information about the node, including restrictions for using the node (such as whether servers can be attached, and the operating systems and customers of the packages and servers that can be attached).
Packages	Manage all software associated with a node, including adding and deleting packages, changing the order of installation, and overriding software attachments that would be inherited from a parent node.

Table 6-1: Tabs Available for Applications

TAB PAGE	DESCRIPTION
Custom Attributes	Set custom attributes for servers that apply to all servers attached to a node. These custom attributes include setting miscellaneous parameters and named data values. Users can write scripts that use these parameters and data values when performing a variety of functions, including network and server configuration, notifications, and CRON script configuration.
Install Order	Create and delete dependencies that specify the order of software installation on servers.
Servers	Manage all servers associated with a node. Manage servers by reconciling servers between nodes, installing applications, installing patches, and so forth. See "Server Management Overview" on page 30 in Chapter 2 for information about how to use these functions.
Config Tracking	View and add entries in the configuration tracking policy for a node. See "Automated Configuration Tracking" on page 493 in Chapter 11 for more information.
History	View the audit trail for changes made to a node. View events such as when servers were attached to nodes, subordinate nodes were added, or customers or operating systems were associated with nodes.

Understanding How Software Is Reconciled onto Servers

A user installs software on a server by using the packages that are stored in the Software Repository. In the Opsware Command Center, the user adds a package to a node, and then attaches a server to that node. When the user reconciles the server, the Opsware System installs the software on the server.

When your node hierarchy is completely defined in the Software Tree, and software and servers are attached, you are ready to perform the reconcile function to install the software on those servers. The Opsware System retrieves the information about the software currently installed on the server, compares it to the list of software about to be installed, and determines any differences. The reconcile function installs the specified software and possibly also removes software, depending upon the type of reconcile being

performed. Before you install the software, the Opsware Command Center displays a preview of the actions about to be performed so that you can make any necessary changes before you proceed.

When to Reconcile

Whenever you make a change to a node that deals with attached software or with custom attributes, you have to reconcile any servers that are attached to that node in order for those changes to be applied. Whenever an application is installed or uninstalled using the Software Install or Software Uninstall Wizard, the Wizard automatically starts a reconcile.

How to Reconcile

See “Reconcile” on page 441 in Chapter 9 for information about how to reconcile.

Managing Nodes on the Software Tree

This section provides information on how to manage nodes on the Software Tree and contains the following topics:

- Application Provisioning Setup Tasks Overview
- Adding a Node to the Software Tree
- Editing a Node in the Software Tree
- Deleting a Node in the Software Tree
- Copying a Node in the Software Tree
- Moving a Node in the Software Tree
- Creating a Software Tree Using the Add Many Function
- Managing a Node’s Configuration Tracking Policy
- Viewing Node History

Application Provisioning Setup Tasks Overview

It is important to manage the nodes on the Software Tree by adding them appropriately and attaching packages appropriately in order to make sure that your servers are properly provisioned with applications.

Users perform the following tasks when setting up and managing the Software Tree and templates:

- Create, edit, and delete nodes.
- Add to, remove from, and define the order of software packages on nodes.
- View and manage software inherited from another node.
- Add, remove, and edit custom attributes.
- Define package installation order so that packages are installed on servers in a specific order.
- Create and edit a node's configuration tracking policies. See "Automated Configuration Tracking" on page 493 in Chapter 11 for more information.
- View the changes to nodes.
- Create, edit, and delete templates.

Adding a Node to the Software Tree

Within each of the categories, users can create subordinate nodes that extend a node, inheriting software packages or other attributes from their parent nodes and adding new software, updates, or other information.

Reasons for Adding a Node

Normally, you do not modify an existing node that has attached servers. The next time a user reconciles the servers, the users inadvertently upgrade the packages to the latest version or remove packages from the servers. So, if you need to apply a change to only a subset of servers attached to a node, consider creating a subordinate node. If you want to add a later version of a package, consider creating a sibling node with the new version number and moving the servers that you plan to reconcile to it. See "Reconcile" on page 441 in Chapter 9 for more information.

For example, you might want to keep package version 1 separate from version 1.1 because you want to upgrade only one customer to version 1.1. In this example, nine servers for three different customers are attached to the node for version 1. You do not want to upgrade all the customers' servers. Create a new node for version 1.1 and move the three servers that you want to upgrade to the version 1.1 node.

Restrictions When Adding Nodes

Nodes can be added at any level on the Software Tree by users with the proper permissions. The only restriction is that no nodes can be added at the category level at the top of the tree.

How to Add a Node

Perform the following steps to add a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Click the appropriate category.
3. If you are adding a new node to the Software Tree, click the Add button.
4. If you are adding a node to an existing branch on the Software Tree, first navigate to the point in the Software Tree where you want to add a node (for example, System Utilities ➤ Winzip ➤ Windows 2000) and then click the Add button.

   The Add Sub-Node form appears, as Figure 6-4 shows. (If a lot of nodes display, you might need to scroll down to see the form.)

Figure 6-4: Adding a Sub-Node to an Existing Node or Category

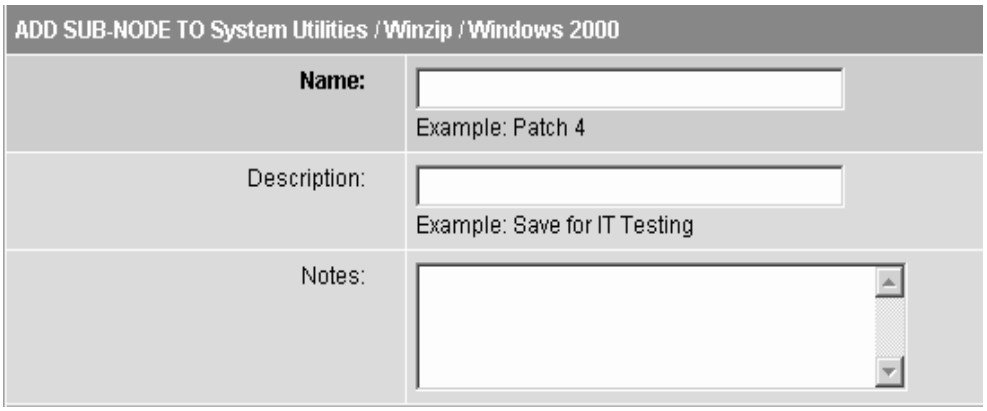

5. Complete the entries that define the node:

   - Enter a name (required).

     The name is limited to 50 characters; names for nodes under the same parent node must be unique. Names are used in the display of the node navigation path, for example:

Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1

- (Optional) In the Description text box, enter a one-line description (up to 500 characters) and indicate why the node was created and what is special about it.
- (Optional) In the Notes text box, enter any other more detailed notes (up to 4000 characters) that you want others to know about the node.

The bottom section of the page shows information about what attributes the new node inherits from its parent, as well as what customers and operating systems the user can edit after creating the node.

6 Click the Save button at the bottom of the page. The Edit Attributes form appears.

7 Follow the instructions in the How to Edit a Node section in this chapter to complete this form.

In general, the most common attributes that need to be changed after you add a new node are Allow Servers, Change Customer, and Change Operating Systems.

Editing a Node in the Software Tree

After you create a node, you might need to make changes to the properties of the node instead of deleting it and creating a new node.

Reasons for Editing a Node

You can edit nodes to change the following properties:

- The name of the node in order to be consistent with your company's naming structure
- The description or notes about the node to clarify its function and purpose
- The Locked attribute to limit, or to remove limits on, access to the node
- The Allow Servers attribute that allows or disallows servers to be assigned to the node
- The Customer and Operating System attributes that define what software and templates can be attached

Restrictions When Editing Nodes

You cannot edit nodes in the following cases:

- The node is a category—Application Servers, Database Servers, OS Extras, Other Applications, System Utilities, or Web Servers—that appears on the top-level Applications page.

- The node is locked and you do not have the correct permissions to edit locked nodes for that category.
- The node is reserved (system -level attribute; user cannot use or modify).
- The node is derived (system-level attribute; user cannot use or modify).
- In a multiple facility environment, the node's data in one Model Repository conflicts with the data of the same node in another Model Repository.
- You do not have the correct read/write permissions for the customer of that node.

Each software package and each node has customer and operating system attributes associated with it that must match the package that is uploaded to the Software Repository. These attributes must match in order for the software to attach to a node. For example, software with a Solaris operating system (SunOS node) cannot be attached to a node associated with a Linux operating system.

How to Edit a Node

Perform the following steps to edit a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the point in the node hierarchy where you want to edit a node (for example, System Utilities ➤ Winzip ➤ Windows 2000).

3 Click the Edit button. The Edit Attributes for [*node*] page appears, as Figure 6-5 shows.

Figure 6-5: Edit Attributes Page

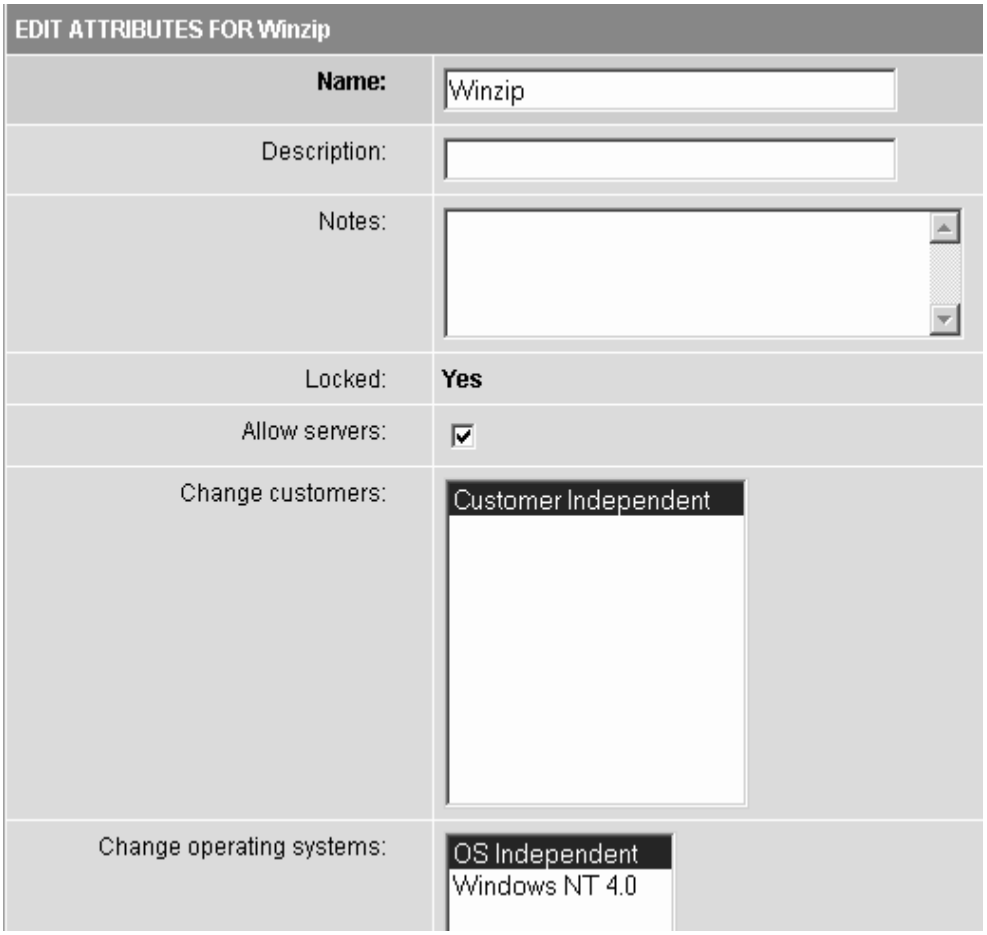

4 Table 6-2 shows the fields that you can edit to update the node. Click the Save button when you are done.

Table 6-2: Field Data to Update for the Node

FIELD	DESCRIPTION
Name	Name, limited to 25 characters, that is used in the display of the node navigation path. For example: Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1 Constraint: Names for nodes under the same parent node must be unique.
Description	One-line description that indicates why the node was created and what is special about it. Limited to 500 characters.
Notes	More detailed information that you want others to know about the node. Limited to 4,000 characters.
Locked	Locked nodes indicate that the node can only be edited by users with Opsware Locking Permissions for that category. If the user has the correct permission, the user is able to edit the attributes of the node, add and remove software, and add and remove custom attributes. Constraints: • Only privileged users—those with Opsware Locking Permissions—can lock and unlock nodes. • Nodes cannot be locked if the parent node is not locked. • Nodes cannot be unlocked if one or more subnodes is locked.
Allow Servers	Whether a server can be assigned to this node. If this option is selected, users have the node listed as an option in the Install Applications Wizard, the Server Assign dialog box, and the Server Reassign dialog box when managing servers (Manage Servers, My Servers, Server Search, and the Node Server tab). In addition, only nodes with this attribute selected can be assigned to templates. Constraint: Must be checked if servers are already assigned to the node or templates are assigned to the node.

Table 6-2: Field Data to Update for the Node

FIELD	DESCRIPTION
Change Customers	The customer that the node is associated with. This attribute determines who can view and edit the node (with the appropriate Opsware customer permissions), what software and servers can be attached, and what templates this node can be assigned to. Each node can only have one customer associated with it. Only customers that are valid for the node display. Sometimes, only the current customer is on the list, meaning that it cannot be changed. The list of customers that displays in the selection box varies. Only valid choices display. The following conditions affect what displays in the list: • If there are subnodes and what customers are associated with those nodes • If packages are attached and what customers are associated with those packages • If servers are attached to the node • Templates that the node is added to • What the parent node's customer is

Table 6-2: Field Data to Update for the Node

FIELD	DESCRIPTION
Change Operating Systems	Specifies the operating systems that the node is associated with. This attribute determines what software and servers can be attached and what templates this node can be assigned to. Each node can only have one operating system associated with it. Only operating systems that are valid for the node display. Sometimes only the current operating system is on the list, meaning that it cannot be changed. Nodes with the operating system OS Independent cannot have software attached. The list of operating systems that display in the selection box vary. Only valid choices display. The following conditions affect what the list displays: • If there are subnodes and what operating systems are associated with those nodes • If packages are attached and what operating systems are associated with those packages • If servers are attached to the node • Templates that the node is added to • What the parent node's operating system is

Deleting a Node in the Software Tree

You should only delete nodes if you no longer need to retain that intellectual property. All history, knowledge of that software in the model, custom attributes, and configuration tracking policies are removed when you delete a node. You cannot undo a deleted node.

Reasons for Deleting a Node

You might want to delete a node after you create it to test a new software installation. To keep the display of your node hierarchy well-organized, delete nodes that are no longer needed and whose information is no longer needed.

Restrictions When Deleting Nodes

You cannot delete a node when the node:

- Has servers attached
- Contains subordinate nodes

- Is a category (a top-level node)
- Belongs to a template
- Is locked (you must have special permission to delete these nodes)
- Is reserved (system level attribute; user cannot use or modify)
- Is derived (system level attribute; user cannot use or modify)
- The node's data in one Model Repository conflicts with the data of the same node in another Model Repository (multiple facility environment)

How to Delete a Node

Perform the following steps to delete a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the node that you want to delete (for example, System Utilities ➤ Winzip ➤ Windows 2000).
3. Click the Delete button.

The Delete button does not appear if one or more of the above restrictions is true.

The Opsware Command Center displays a page that shows how many servers, subordinate nodes, and templates the node has. The number should be zero for all in order to delete the node. See Figure 6-6.

Figure 6-6: Applications: Delete Node Page

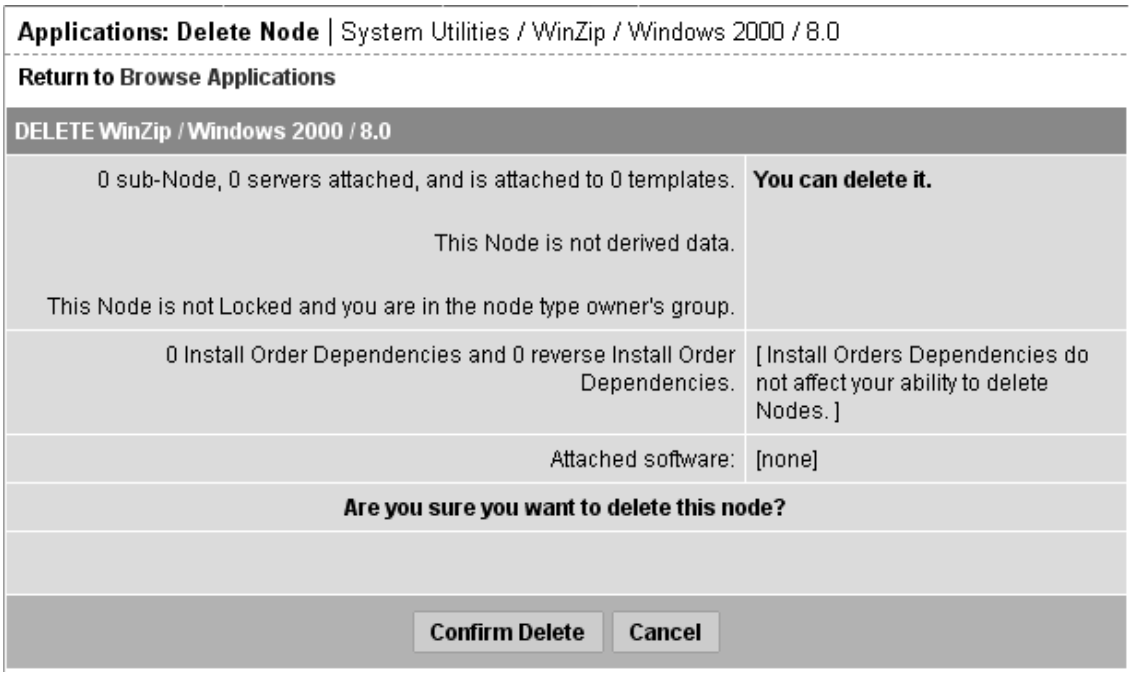

The display confirms how many software packages and installation order dependencies are associated with the node. The presence of software does not prevent a user from deleting a node.

You are prompted to confirm the deletion.

4 To delete the node, click the Confirm Delete button.

The Opsware Command Center returns you to the Properties page for the parent node.

Copying a Node in the Software Tree

The copy function is a convenient way to duplicate information that is contained in an existing node without having to enter the information all over again.

Reasons for Copying a Node

Use this function when you want to copy the structure and attributes from one node to a new node at the same level in the Software Tree.

Restrictions When Copying Nodes

You can copy a node at any level, except the top-level categories.

How to Copy a Node

Perform the following steps to copy a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the node that you want to copy (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Copy button. A page appears, called Copy [*node path*], as Figure 6-7 shows.

Figure 6-7: Copy Node Page

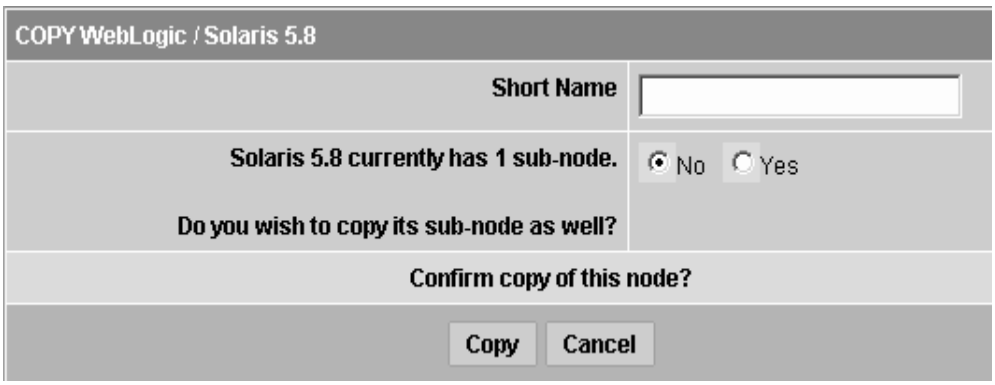

4. Enter a short name for this node. Remember that this node name must be unique. The limit is 25 characters.
5. Click Yes to copy any subnodes. Click No if you only want to copy this node and no subnodes.
6. Click the Copy button. The Applications [*node path*] page appears that displays the information for the newly-copied node, which is now available for editing, deleting, copying, or creating subnodes.

Moving a Node in the Software Tree

The Move Node Wizard allows you to move a node to nearly any other location in the Software Tree. When you move a node to a new destination node, that node becomes a child of the destination node

When you move a node, all of that node's dependencies and all servers attached to that node remain intact. All other attributes of the node are inherited from the new parent and all direct attachments remain intact.

For example, if you move a node that has devices on it, it is likely that the new location will have different packages or custom attributes inherited. That means that the software model has changed and the devices attached to the node need to be reconciled. The Move Node Wizard allows you to perform a reconcile in order to resolve any discrepancies that the move caused.

Restrictions for Moving a Node in the Software Tree

The following lists of restrictions for moving a node:

- You can only move a node to a location on the Software Tree that has a compatible OS. For example, you cannot move a node that has OS "A" to a node that belongs to OS "B." However, you can move a node with a specific OS node to a node that is OS independent.
- You can only move a node to a location on the Software Tree that has a compatible Customer. For example, you cannot move a node that belongs to Customer A to a node that belongs to Customer B. However, you can move a customer dependent node to a customer independent node.
- You cannot move a node to itself or to a child of itself.
- You cannot move a locked node to a destination node that is not locked, but you can move an unlocked node to a locked node. If you want to move a locked node to another locked node, you need locking permissions to do so.
- You cannot move a top-level node (such as Application Servers, OS Extras, Web Servers, and so on).
- You cannot move nodes within the Hardware and Opsware channels.

How to Move a Node

Perform the following steps to move a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the node that you want to move (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).

3 Click the Move button. The Move Node Wizard appears.

Figure 6-8: Move Node Wizard, Select Destination page

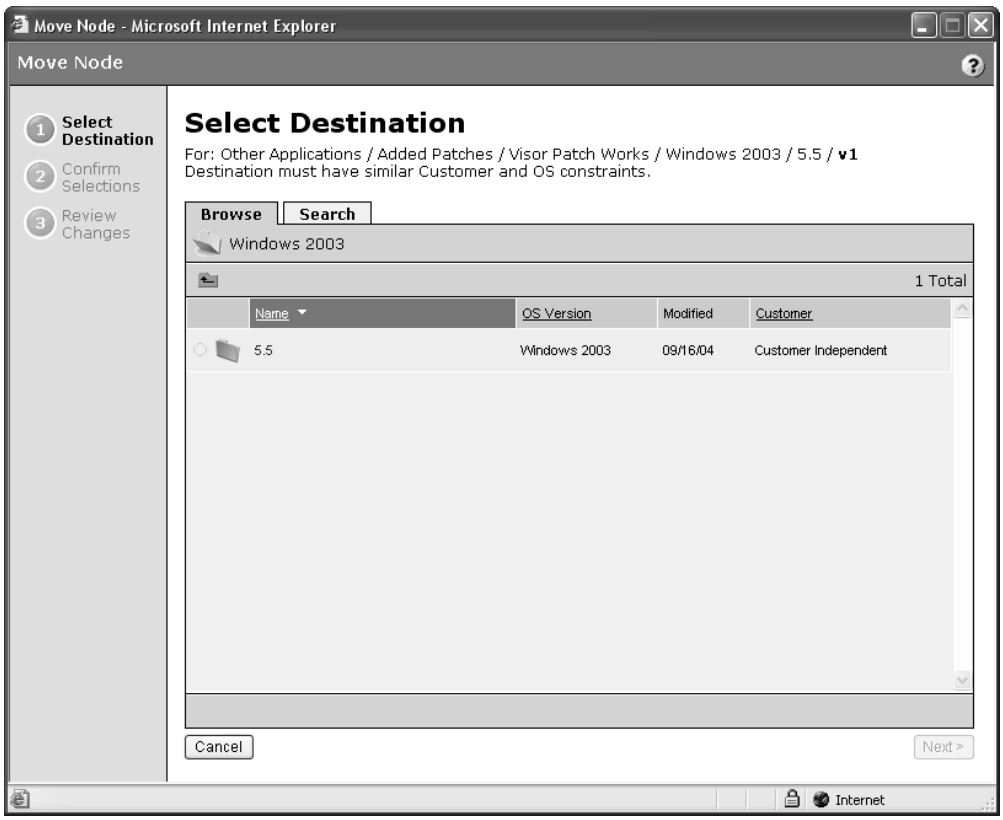

4 In the Select Destination page of the Wizard, choose a destination for the node. To choose a destination, you have two choices:

- Browse for a Node—From the Browse tab you can browse to a location in the software tree by clicking the link of a node's name. When you locate the destination node that you want to move to, select the radio button next to the node name to choose that node as the destination node.
- Search for a Node—If you want to search for a node by name, then select the Search tab. You will notice that the OS and Customer criteria are already selected and reflect the properties of the node that you are going to move. To search by name, choose a search criteria (is or contains) from the drop-down list, then enter a name for the node that you would like to search for. In the results page, select a node.

5 Click Next.

6 In the Confirm Selection page of the Wizard, check to make sure that you are moving the correct node to the desired destination, and then click Move to move the node.

7 The Review Changes page lists the details of the move node operation. If any devices (servers) are attached to the node and need to be reconciled, then you will see the Reconcile Servers button. The Review Changes page lists the details of the move node operation. If any servers are attached to the node then you will see the Reconcile Servers button. Opsware recommends that you reconcile the servers so that the moved node conforms to its changed software model.

8 Click Close when you finish.

Creating a Software Tree Using the Add Many Function

Use the Add Many function to create an entire node hierarchy by specifying the information for all the nodes in one page, as opposed to creating a node hierarchy a single node at a time by using the Add function.

Reasons for Using the Add Many Function

From the top-level categories, you can create all the nodes necessary to define a software product line to be deployed in an operational environment by using the Add Many function. The node hierarchy is created as described in “Guidelines for Setting Up the Software Tree” on page 312.

Restrictions When Using the Add Many Function

You can use the Add Many function only at the top-level categories of the Software Tree. It is not available at other levels in the tree.

How to Use the Add Many Function

In a category, you can create all the nodes and subnodes necessary to define a software product line that is deployed in an operational environment.

When you create a Software Tree by using the Add Many function, by default each node in the new hierarchy is associated with the customer Customer Independent and the Allow Servers attribute is checked for every node in the hierarchy.

Perform the following steps to use the Add Many function:

1 From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.

2 Click the link for the category in which you want to create a node hierarchy.

3 Click the Add Many button. The Create Model page appears, as Figure 6-9 shows.

Figure 6-9: Create Model Page

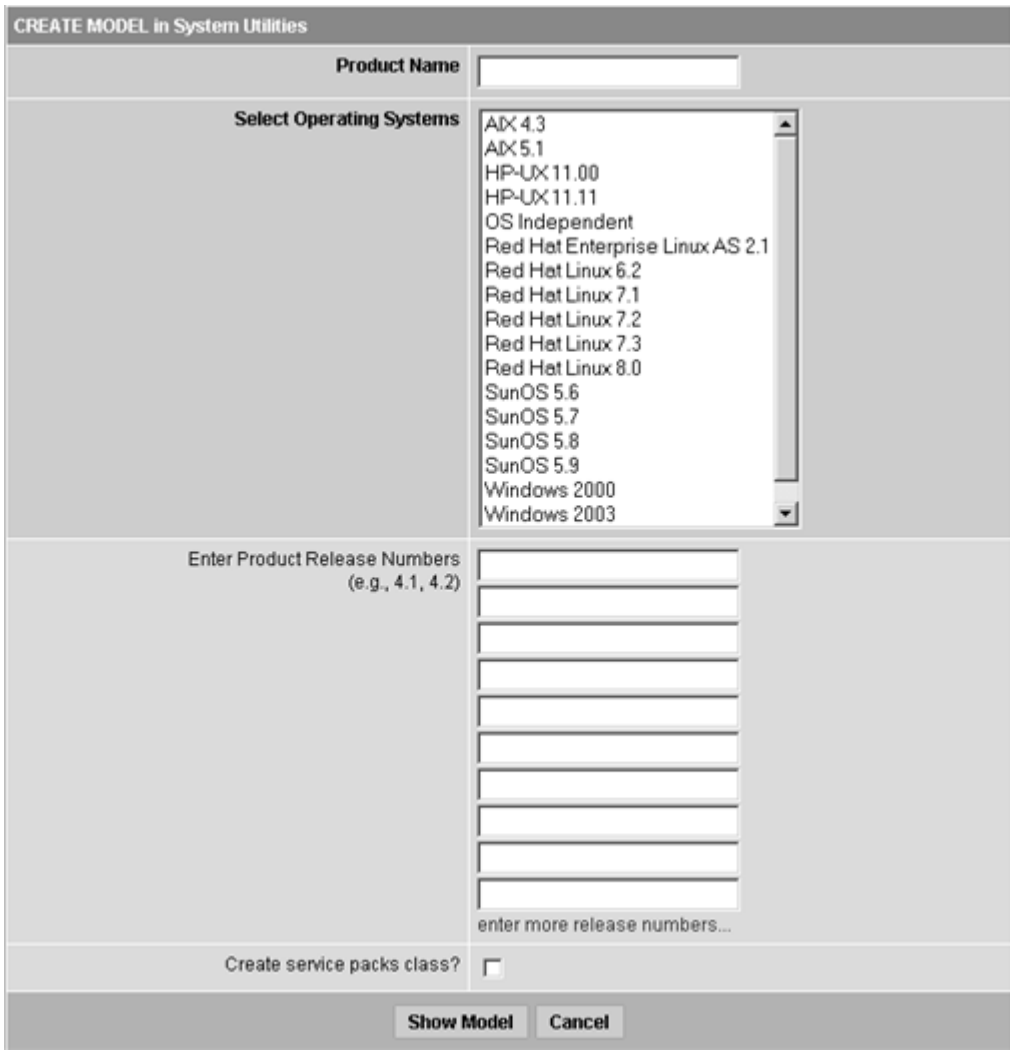

4 Complete the following entries to define the node hierarchy:

- Product Name (Required)

  Product names appear at the top of the hierarchy for the node navigation path. The product name is limited to 25 characters. Product names for nodes under the same category must be unique.

- Select Operating Systems (Required)

Creates branches within the hierarchy for the operating systems selected. Also, specifies the operating system that each node, by default, is associated with. You can later edit this field for individual nodes in the Opsware Command Center.

- Enter Product Release Numbers

  Creates nodes within the hierarchy under each operating system for each product version that is specified. You can later edit the nodes for versions in the Opsware Command Center.

- Create service packs class?

  Select this option when you create a product line that contains software for Windows operating systems.

5 Click the Show Model button at the bottom of the page. The Display of Model area appears at the bottom of the page, as Figure 6-10 shows.

Figure 6-10: Display of Model Area

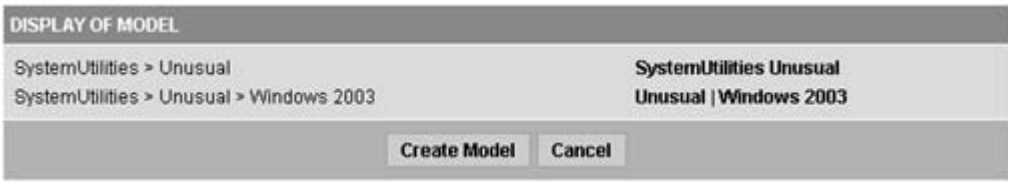

6 Verify that the node hierarchy has the structure that you want.

7 (Optional) Make changes to the entries that define the node hierarchy and click the Show Model button again.

8 Click the Create Model button. A message appears at the bottom of the page that the model was created.

9 Click the Refresh button. The page appears for the category in which you created the hierarchy. The top level of the hierarchy you just created appears as a branch in the category.

You can later edit the attributes for individual nodes in the Opsware Command Center. See “Editing a Node in the Software Tree” on page 318 in this chapter for more information.

Managing a Node's Configuration Tracking Policy

The Opsware System's configuration tracking feature allows you to monitor selected configuration files and configuration databases for change and to take certain actions when change is detected.

To use the configuration tracking feature, you create configuration tracking policies, which specify which files to monitor and what actions to take when a change is detected.

The preferred method of creating configuration tracking policies is to use Opsware nodes. Using nodes allows to you take advantage of the node's model-based architecture. Through Opsware nodes, you can create configuration policies and deploy them to the appropriate servers based on what nodes the servers are attached to.

See "Automated Configuration Tracking" on page 493 in Chapter 11 for information about how to use the Opsware configuration tracking feature.

Viewing Node History

Using the Opsware Command Center, you can view the audit trail for changes made to a node. For example, you can see who has modified the node that you are browsing.

Most actions performed on the node through the Opsware Command Center are recorded in the history. The following list describes the actions that create an entry in the history for a node:

- Editing any of the node's attributes such as name, description, notes, locked, allow servers, customers, or operating systems
- Adding software to or removing software from the node
- Moving a node
- Changing the software inheritance for the node
- Creating or deleting subordinate nodes
- Moving servers from one node to another
- Adding servers to or removing servers from a node

Each history entry contains three pieces of information, as Table 6-3 shows.

Table 6-3: History Entry Information

HISTORY ENTRY	DESCRIPTION
Event Description	Description of operation performed, for example: Updated lc_certified field to “true”
Modified By	User name of individual who made the change
Date Modified	Date change was made, for example: 18-Aug-2003 08:54:37 PM

The history is read-only.

Perform the following steps to view the History for Node page:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the level in the node hierarchy where you want to view the history (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the History tab. The History for Node page appears, as Figure 6-11 shows.

Figure 6-11: History for Node Page

By default, the view shows changes made within the past week. To see changes earlier than the past week, select Two Weeks, Month, or Quarter.

The Opsware Command Center only maintains the history of changes for the last three months.

Software Attached to Nodes

This section provides information about software that is attached to nodes within the Opsware System and contains the following topics:

- Software Attached to Nodes Overview
- Modeling Software in Nodes
- Software Configuration Settings
- Viewing Software Attached to a Node
- Adding a Software Package to a Node
- Removing a Software Package from a Node
- Changing the Installation Order of Software
- How Software Is Inherited from Other Nodes
- Changing the Override Values of Inherited Software
- Dependencies Between Nodes for Software Installation
- Viewing Software Installation Dependencies
- Adding Software Installation Dependencies
- Removing Software Installation Dependencies

Software Attached to Nodes Overview

Software packages are attached to nodes. Using nodes to install software packages provides an automated and consistent mechanism for software installation based on the best practices that your organization defines.

See "Package Management" on page 267 in Chapter 5 for information about how to manage (upload and download) packages through the Opsware Command Center.

A user installs software on a server by using packages stored in the Software Repository. In the Opsware Command Center, the user adds a package to a node and then attaches a server to that node. When the user uses the Install Software Wizard, the Opsware System reconciles the server and installs the software on the server. See “Reconcile” on page 441 in Chapter 9 for more information.

After you define the software directly attached to a node, the packages might not be listed in the order in which they need to be installed on a server. This condition can be true if you recently added new packages to the list. Using the Up and Down arrows next to the package list, you can change the installation order of software packages.

When you modify software that belongs to a node, remember to click the Save Edits button or the changes are not saved.

You can save software edits made to a particular node and commit them all at once. For example, you can add packages, remove packages, move packages up and down in the installation order, reverse the override value on inherited software, and then click the Save Edits button to commit all the changes at the same time. See “How Software Is Inherited from Other Nodes” on page 344 in this chapter for more information.

Modeling Software in Nodes

The Opsware System supports Advanced IBM Unix (AIX), Hewlett-Packard Unix (HP-UX), Linux, Solaris, and Windows operating systems. Each software package and each node has operating system attributes associated with it. The types of packages users attach to nodes vary by operating system. See “Application Provisioning Setup” on page 307 in Chapter 6 for information about the full list of package types that the Opsware System manages.

Table 6-4 shows only the package types that you can attach to nodes.

Table 6-4: Package Types Organized by Operating System

PACKAGE TYPE	FILE FORMATS	NOTES
AIX		
AIX Update Fileset	N/A	None
AIX Base Fileset	N/A	None

Table 6-4: Package Types Organized by Operating System

PACKAGE TYPE	FILE FORMATS	NOTES
AIX APAR	N/A	Users can attach complete and incomplete APARs to nodes. An APAR is incomplete when the Opsware System does not have all required update filesets for the APAR. Users can install incomplete APARs on servers when the Software Repository contains the update filesets corresponding to the installed base filesets. Possibly, the missing update filesets are not applicable because their base filesets have not been installed. The Opsware System warns users when they try to reconcile servers when necessary update filesets have not been uploaded.
RPM	.rpm	None
HP-UX		
HP-UX Product	N/A	When products are installed, all filesets that make up that product are installed automatically.
HP-UX Fileset	N/A	None
HP-UX Patch Product	N/A	None
Linux		
RPM	.rpm	None
Solaris		
Solaris Patch	.tar, .tar.Z, .zip, .tar.gz, .tgz, .jar	None
Solaris Patch Cluster	.tar, .tar.Z, .zip, .tar.gz, .tgz, .jar	None
Solaris Package	Datastream File	The Opsware System supports interactive (requires a response file) and non-interactive Solaris packages.

Table 6-4: Package Types Organized by Operating System

PACKAGE TYPE	FILE FORMATS	NOTES
Solaris Package Instance	N/A	If a Solaris package requires a response file, add it to the package after you upload the package. Attaching an interactive package to a node includes the response file. You do not need to attach the response file to the node.
RPM	.rpm	None
Windows Operating System		
Windows Hotfix	.exe	None
Windows MSI	.msi	None
Windows OS Service Pack	.exe	None
Windows ZIP File	.zip	None

Software Configuration Settings

Each software package should contain the configuration settings necessary to run the application on a server. If a configuration change is required in an operational environment, Opsware Inc. recommends that you repackage the software and add the updated package to the Software Repository.

See “Package Management” on page 267 in Chapter 5 for information about how to manage (upload and download) packages with the Opsware Command Center.

System administrators or operations personnel might need to adjust the configuration settings on a server after it is provisioned with an application. Operations personnel might have to further modify configurations after a server is deployed to the environment. As a result, the Opsware System might be unable to rebuild a server to its latest configuration. Be sure to backup any manual configuration changes. See “Automated Configuration Tracking” on page 493 in Chapter 11 for more information.

Even though the Opsware Command Center allows users to edit nodes, use caution when you modify existing nodes that have servers attached, so that you preserve the ability to rebuild server configurations.

Viewing Software Attached to a Node

Perform the following steps to view software that is attached to a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to a node (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Packages tab. The Packages Attached to a Specific Node page appears, as Figure 6-12 shows.

Figure 6-12: Packages Attached to a Specific Node

The following Packages are Directly Attached to this Node **Edit Packages**

Name	Type	Description
SUNWbzip	Solaris Package Instance	The bzip compression utility
SUNWbzipx	Solaris Package Instance	The bzip compression library (64-bit)

The following Packages are Inherited to this Node

No Inherited Packages

This page has the following two sections:

- Package Directly Attached–Directly attached packages are software that is attached specifically at this particular point in the node hierarchy.
- Package Inherited–Inherited Packages are software that is attached at some point higher in the node hierarchy. Inherited software can come from a parent node or from any ancestral node and also can be a combination of packages from several levels of nodes.

Adding a Software Package to a Node

Each software package and each node has customer and operating system attributes associated with it. These attributes must match in order for the software to attach to a node. For example, software with a Solaris operating system (SunOS node) cannot be attached to a node associated with a Linux operating system. You cannot attach software to a node that is associated with the operating system OS Independent.

Perform the following steps to add a software package to a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to a node (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Packages tab. A page that lists the software attached to the node appears.
4. Click the Edit Packages button. A page appears that shows the directly attached packages and the inherited packages.

When a node contains a large number of software packages, the page might take a few seconds to load.

5 Click the Add Packages button. The Search for Packages page appears that allows you to search for the software that you want to add, as Figure 6-13 shows.

Figure 6-13: Search for Packages Form

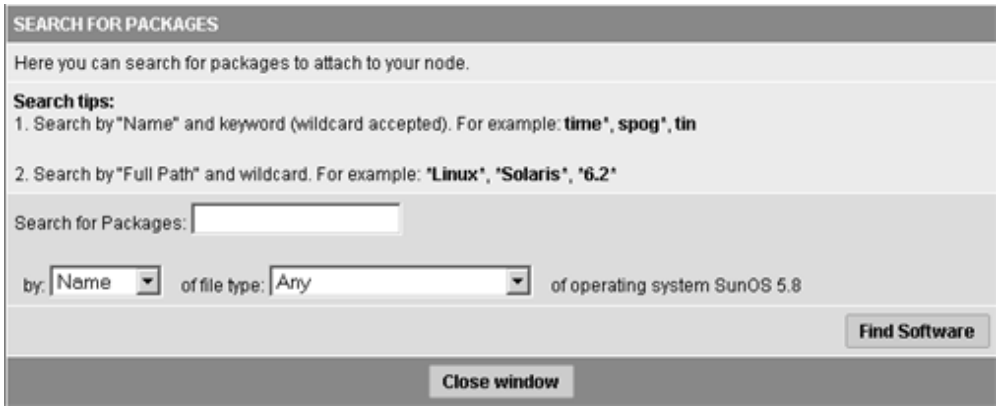

6 Search by the package name (if you know it), or search by the path (if you know the directory of the package in the Software Repository).

7 After you specify your search criteria, click the Find Software button.

The time it takes the Opsware Command Center to return search results depends on the number of items a search query needs to process. The more general the query, the slower the results will be. If the query results in more than 500 packages found, the following message appears, "Your search returned too many results. Please try to narrow your query."

8 From the search results list, select the packages that you want to add to a node by selecting the check box next to each package. (Software packages already attached to the node do not have a check box.)

9 Click the Add to Node button. A dialog box appears, which prompts you to search for additional packages.

10 Click Cancel to end your search or click OK to continue searching for additional packages to add to the node.

The packages are added to the end of the list of packages already attached. If you need to specify a different installation order for the packages, use the Move functions. See "Changing the Installation Order of Software" on page 343 in this chapter for more information.

11 Click the Save Edits button.

Additions and changes are not stored in the Opsware System until you click the Save Edits button.

Details: Searching for Software Packages

You can search by package name or full path, use wildcards in your search terms, and narrow search results by specifying the file type (RPM, Windows Hotfix, and so forth) of a software package. Only file types that are valid for the current node's operating system are available as choices.

- If you know the exact path of the software package that you want to attach, you can enter that path in the search box. Change the first drop-down menu to Full Path, and then click the Find Software button. For example:

  ```
  /osimage/Linux/6.2/LC-1.0/cdrom/RedHat/RPMS/vixie-cron-
  3.0.1-40.i386.rpm
  ```

- If you know the name of the software package, but not its location, set the first drop-down menu selection to Name and search for just the name of the software (without version information or file extension). For example:

  ```
  vixie-cron
  ```

- If you want to find all software packages that start with a certain string, set the first drop-down menu selection to Name, type a wildcard query string, and click Find Software. For example:

  ```
  time
  ```

  (You do not need to type the wildcard character "*" in this case.)

- If you want to find all the software packages that are located in a particular directory, set the first drop-down menu selection to Full Path and search for a wildcard query string that specifies the directory that you want. For example:

  ```
  *RedHat*
  ```

Removing a Software Package from a Node

Users can delete one, many, or all of the software packages attached to a node. Be careful because deleting software from nodes might cause problems for other users of the same nodes. Check with other users who might be using the nodes before you delete any associated software.

Perform the following steps to remove a software package from a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to a node (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Packages tab. A page appears that shows the list of packages attached to the node.
4. Click the Edit Packages button. A page appears that shows a form with the packages currently attached displayed inside a box that allows additional packages to be added, the order of the packages to be changed, or packages to be removed, as Figure 6-14 shows.

Figure 6-14: Edit Packages Page

5. Click the package name in the Packages Directly Attached box to highlight the package to be deleted.

You can highlight all the packages at one time by clicking the Select All button located below the list. (You can also reset the selection by clicking the Deselect All button.) You can select multiple packages to delete by holding down the CTRL key while selecting packages.

6 Click the Remove Packages button.

The selected packages disappear from the select box but the packages are not yet detached from the node.

7 Click the Save Edits button at the bottom of the page to complete detachment of the software.

Additions and changes are not stored in the Opsware System until you click the Save Edits button.

To apply changes that you made to a node to servers that already have the node attached, you need to reconcile those servers. See “Reconcile” on page 441 in Chapter 9 for more information.

Changing the Installation Order of Software

Perform the following steps to change the installation order of the software:

1 From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.

2 Navigate to a node (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).

3 Click the Packages tab. A page appears that shows the list of packages attached to the node.

4 Click the Edit Packages button. A page appears that shows a form with the packages currently attached displayed inside a box that allows additional packages to be added, the order of the packages to be changed, or packages to be removed.

5 Highlight the single package or multiple packages (Ctrl-Click) that you want to move in the installation order.

6 Click the up arrow button or the down arrow button.

Highlighted packages are moved up or down one position in the installation order each time you click the button.

Click the move to top arrow button or move to bottom arrow button.

Highlighted packages are moved to the top or bottom of the installation order.

7 When you finish making changes to the installation order, click the Save Edits button at the bottom of the page.

Installation order changes are not stored in the Opsware System until you click the Save Edits button.

How Software Is Inherited from Other Nodes

If a node inherits software from its parent (or from a node farther up the node hierarchy), the software displayed for the node includes directly attached software as well as software inherited from parent nodes. For the inherited software, you can override whether or not software is installed by the current node, but you cannot change the order of the inherited packages.

When you reconcile, inherited software is installed first, then directly attached software is installed.

Examples: Ways to Use Inheritance for Software

A subordinate node can inherit the properties of a parent node. For example, IIS only runs on Windows NT. The parent node indicates that the operating system for the node is Windows NT only. The subordinate nodes created under that node can only be used for Windows NT servers.

In a second example, you create a node for a version of IIS (version 1). Then, you create a subordinate node and attach IIS version 1.1, which has the same list of software except that it adds a few more packages. The subordinate node of IIS inherits the list of software that is attached to the node for version 1.

A third way that you can use inheritance is for installing and testing new versions of software. For example, you might create a subordinate node (Q1) of the node (Q) that contains software to be updated or patched. You could then move a group of test servers from node Q to the subordinate node Q1. In the inherited software list of Q1, you could

override packages that contain changed software, so as not to inherit similar packages from the parent node. Then, you can directly attach the new packaged version of the software to Q1. When you reconcile, the overridden package is removed from the test servers and the new directly attached package is installed instead.

Changing the Override Values of Inherited Software

In the inherited software list, most of the packages have a plus symbol (+) to the left. These packages are referred to as *plus override*. Any software package with a plus symbol (+) by its name is inherited and installed on servers attached to this node.

If a software package has a minus symbol (–) to the left of its name, it is referred to as *minus override*. This means that the node is aware that its parent has a particular package, but the node does not inherit this package and servers attached to the node do not have this package installed.

Perform the following steps to change the override values of inherited software:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to a node (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Packages tab. A page appears that shows the list of packages attached to the node.
4. Click the Edit Packages button. A page appears that shows a form with the packages currently attached displayed inside a box that allows additional packages to be added, the order of the packages to be changed, or packages to be removed.
5. In the Inherited Packages box, highlight one package at a time or select multiple packages (Ctrl-Click) for which you want to change override settings.
6. When you finish your selection, click the Save Edits button.

   Selected packages that were prefaced with a plus symbol (+) are prefaced with a minus symbol (–), and packages that were prefaced with a minus symbol (–) are prefaced with a plus symbol (+).
7. Click the Save Edits button.

The changes made to override values are not committed to the Opsware System until you click the Save Edits button.

After the edits are saved, the page refreshes, updated to show your changes.

Dependencies Between Nodes for Software Installation

In addition to specifying installation order for directly attached and inherited software within nodes, users can define special dependencies between nodes to specify the order of software installation on servers. Spelling out any special dependencies that exist for specific software packages allows those software packages to be installed in the correct order when a server is reconciled.

The Opsware System reconciles and installs software on servers in the following default order (from first to last):

- Operating System
- OS Extras
- System Utilities
- Database Server
- Application Server
- Web Server
- Other Applications
- Patches

Sometimes the default order needs to be overridden so that software packages in one category are installed before another package in a different category. You can prescribe a change from the default package installation order by clicking the Dependency tab and specifying any special dependencies of different nodes. In doing so, you can make one node dependent on another so that the software is installed in the correct order.

For example, if a Web server package and a corresponding application server plug-in package are both installed on the same server, the Web server package must be installed before the application server plug-in (to ensure correct configuration of the plug-in). In that case, you can specify an installation order dependency for plug-ins that Web server packages must be installed before the plug-in software packages.

These dependencies are for installation ordering only, not to check for the presence of a particular piece of software.

Viewing Software Installation Dependencies

If the current node is locked (a lock icon displays beside the node name at the top of the page), then only authorized staff can add or remove dependencies. If you have permission to modify node dependencies, you might see two orange buttons at the top of the screen: Remove Dependency and Add Dependency. The Remove Dependency button does not appear if there are no dependencies.

The Opsware Command Center displays a list of nodes that are dependent on the current node. You see a list of nodes that should precede and follow the current node when determining installation order for software attached to those nodes.

For convenience, each node in each list is a hyperlink. To see details about each node, click the hyperlinked node name.

You cannot remove the Install Order for nodes that are to be installed before the current node. If a dependency listed in this section needs to be removed, you need to go to that node and remove the dependency from that node's Install Order tab.

Perform the following steps to view software installation dependencies:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the node for which you want to view an installation order dependency (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Install Order tab. The Install Current Node After list and Install Current Node Before list display, as Figure 6-15 shows.

Figure 6-15: Install Order Tab

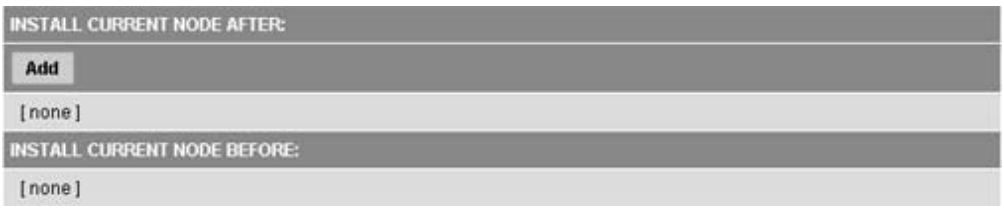

The Opsware Command Center displays [none] when there are no installation order dependencies.

Adding Software Installation Dependencies

When you create software installation order dependencies, Opsware Inc. recommends that you follow these rules:

- A node cannot be made dependent on itself.
- Do not create a circular dependency between nodes; for example, Node A is dependent on Node B and Node B is dependent on Node A. If a server is attached to a node with a circular dependency and you attempt to reconcile the server, the Opsware Command Center displays an error message and does not complete the reconcile.

Perform the following steps to add software installation dependencies:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the node for which you want to add an order dependency (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Install Order tab. The Install After list and Install Before list display.
4. Click the Add button. A window appears that shows the categories.
5. Navigate to the node for which you want to specify an installation order before the current node (for example, System Utilities ➤ Java JDK ➤ SunOS 5.8 ➤ 1.4).

6 Click the Add button. The Opsware Command Center closes the window to display the new software installation order dependency, as Figure 6-16 shows.

Figure 6-16: New Installation Order Dependencies

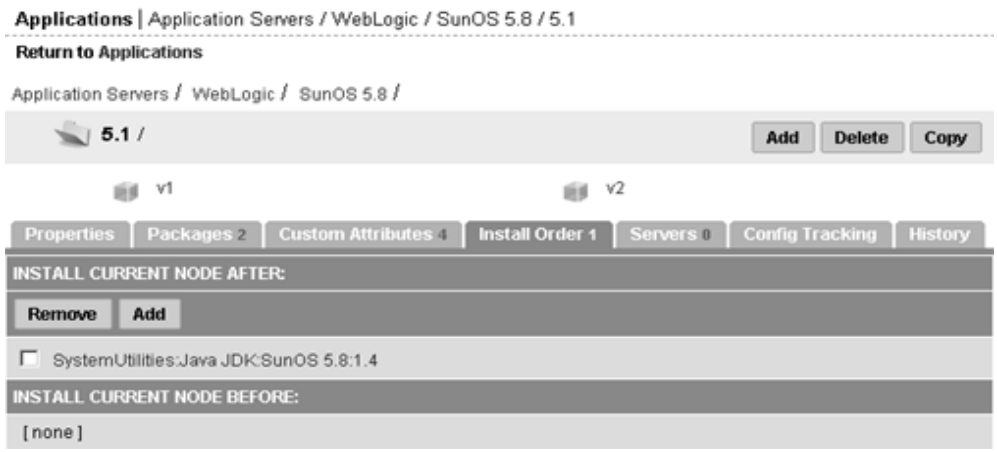

Removing Software Installation Dependencies

You can delete multiple dependencies at the same time by selecting multiple check boxes before you click Remove Dependency.

Perform the following steps to remove software installation dependencies:

1 From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.

2 Navigate to the node for which you want to remove an install order dependency (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).

3 Click the Install Order tab.

4 Click the check box next to the specific dependency that you want to delete.

5 Click the Remove Dependency button.

Custom Attributes Set for the Environment

This section provides information on custom attributes set for the environment within the Opsware System and contains the following topics:

- Custom Attributes Set for the Environment Overview

- Managing Custom Attributes
- Adding Custom Attributes for a Node
- Editing Custom Attributes for a Node
- Deleting Custom Attributes for a Node

Custom Attributes Set for the Environment Overview

Users might need to store specific miscellaneous information in the Opsware Model Repository, for example, to facilitate server or application installation and configuration or scripting.

The Opsware Command Center provides a data management function that allows users to set custom attributes for servers. These custom attributes include 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.

For information about how to set custom attributes required by the software running on a specific server, contact the group responsible for packaging your applications.

Managing Custom Attributes

Even though the Opsware Command Center allows users to edit nodes, use caution when you modify existing nodes, especially those with servers attached, so that you preserve the ability to rebuild server configurations.

To set custom attributes that affect a specific server, attach the custom attributes directly to the specific servers.

See "Server Management" on page 29 in Chapter 2 for information about custom attributes and servers.

To set custom attributes that affect every server for a customer, you can set custom attributes for that customer.

Adding Custom Attributes for a Node

When you add custom attributes to a node, the attributes and values affect every server attached to that node.

Perform the following steps to add custom attributes for a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the node for which you want to add a custom attribute (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Custom Attributes tab.

   If there are no custom attributes, the button is called Add Custom Attributes. If there are custom attributes and you want to add more, click the Edit Custom Attributes button, and then the Add Custom Attributes button.
4. Enter the name and value for the custom attribute that you want to add. You can enter up to 10 custom attributes.

If you want to enter a value that is longer than the space available in the field, click the "..." button to open a window with a larger text box.

5. Click the Save button. The View Custom Attributes page appears.
6. Click the Edit Custom Attributes button.
7. Click the Save Edits button.

Editing Custom Attributes for a Node

To change the name of a custom attribute entry, create a new custom attribute with the new name and delete the old custom attribute.

Perform the following steps to edit custom attributes for a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the node for which you want to add a custom attribute (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Custom Attributes tab.

4 Click the Edit Custom Attributes button. A new form appears, which shows three columns of information: Name, Inherited Value, and Local Value. See Figure 6-17.

Figure 6-17: Edit Custom Attributes Page

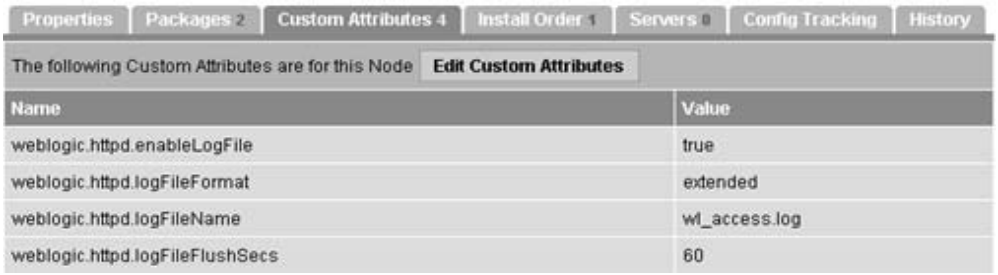

- Name is the name of the attribute.
- Inherited Value is the value that was entered in an ancestral node for the name that displays.
- Local Value is the value for the current node, which will override the inherited value if one displays.

5 If there are inherited values, the label next to the check box is labeled Inherit. Check it if you want to use the inherited value. If it is not checked, the local value is used.

If there are no inherited values, the label next to the check box is Delete. Check it if you want the attribute to be deleted when the Save Edits button is clicked.

6 If there is an inherited value, and the Inherit check box is checked, the inherited value is the one being used. To override that value, deselect the check box. The Local Value text box is enabled to allow text to be entered.

7 If there is {more...} in the local value text box, the value is too large to put in the text box. To edit this value, select the "..." button to open a window with a larger text entry box.

8 Click the Save Edits button to commit the changes.

The changes made to custom attributes are not committed to the Opsware System until you click the Save Edits button.

Deleting Custom Attributes for a Node

Perform the following steps to delete custom attributes for a node:

1. From the navigation panel, click Software ➤ Applications. The Applications page appears that shows the categories.
2. Navigate to the node for which you want to delete a custom attribute (for example, Application Servers ➤ WebLogic ➤ SunOS 5.8 ➤ 5.1).
3. Click the Custom Attributes tab. A page appears that shows the custom attributes for the selected node.
4. Click the Edit Custom Attributes button. The three-column form appears with a check box labeled Delete available next to the Local Value field, as Figure 6-18 shows.

Figure 6-18: Delete Custom Attributes

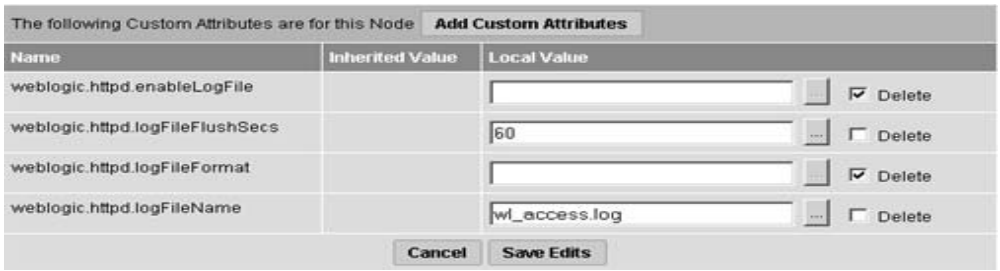

5. Select the check box next to the attributes that you want to delete.

   If the label for the attribute is Inherit, and you check the box, only the local value is deleted. To completely remove an inherited attribute, delete it from the node higher in the hierarchy from which it was inherited.
6. Click Save Edits to commit the change and delete the attributes.

Working with Templates

This section provides information on how to work with templates within the Opsware System and contains the following topics:

- Templates Overview
- Templates, Folders, and Inheritance
- Template Inheritance
- Attachments: Local, Inherited, and Blocked
- Blocking and Reattaching Inherited Attachments

Templates Overview

Opsware templates allow you to create groups of logical relationships between nodes and a server. With templates, you can arrange related sets of software packages together so that you can apply them to a server in a single operation by using the Opsware template wizard.

Templates represent collections of objects such as software, service levels, operating systems, and patches that are simultaneously attached to a server when the template is applied. These collections of nodes are referred to as *attachments* in the Templates channel of the Opsware System. Unlike attaching software to nodes in the software tree, however, templates are applied but do not stay permanently associated with the server.

The two basic types of Opsware templates are:

- Templates that include the installation of an operating system
- Templates that do not include an operating system installation

For example, you can use an Opsware template to quickly bring new servers into production. Such a template might include an operating system for a new server, the latest security patches for the operating system, plus all the applications that are required to run a full-fledged Web service. In another case, you could create a template that consists of a set of applications and patches to install a new service on servers that are already in production.

Templates, Folders, and Inheritance

Using template folders, you can create a deep hierarchy of templates to take advantage of commonalities between server configurations. You can add as many subfolders and other templates inside folders as you would like, grouping together any number of templates to be applied simultaneously to a server.

Folders are more than empty containers into which you can organize your templates. They can also have attachments. Importantly, folders operate under the principle of inheritance, which means that any nested folders or templates inherit all the attachments of the parent folder they belong to. Inheritance minimizes the amount of maintenance required to manage templates in a hierarchy.

A unique feature of templates is the ability to apply either a template or a folder to a server. Because a folder can take attachments, the process of selecting and applying the folder is the same as applying a template. It is important to know that templates do not

have a persistent connection to the servers they are used to configure. Consequently, changes made to a template do not affect servers that are already configured with that template.

Template Inheritance

Template inheritance works in such a way that templates and folders inherit all attachments of the folder they reside in. Inheritance is propagated from parent (folder) to child (template or folder) – and all children of children. It is important to know that child folders and templates always inherit the attachments of their ancestor folders, unless you specifically block the inheritance of an attachment inside a template or folder.

You can attach the following nodes to a template:

- Operating System (only one per template)
- Applications
- Patches
- Service Levels

You can attach nodes to a folder or template in the following three ways:

- At the current or local level
- At the current level by inheritance from an ancestor folder
- Blocked at the current level

  Blocking at the current level prevents an attachment from being inherited by child folders.

Take care in organizing your folders so that you can take maximum advantage of inheritance when you create child folders. Organize your folders so that the most common attachments that servers use are contained in the parent folders. That means no matter how many child folders are created from that parent, if an attachment needs to change in all folders in a family, you only need to make that change once, in the parent folder, and all child folders and child templates will inherit the change.

Attachments: Local, Inherited, and Blocked

Templates take advantage of commonalities among server configurations. You can use the power of inheritance to further minimize maintenance of your templates as configurations need to change. This section discusses how to define attachments to achieve this goal.

All attachments in a folder are automatically propagated to any child folders or templates. The three possible states that an attachment can be in are *local*, *inherited*, or *blocked*.

Local

If a node is attached directly to a folder or a template, the attachment is considered *local*, which means that the attachment is not inherited from an ancestor folder. It is directly attached to the folder. Figure 6-19 shows that Solaris 5.8 and Basic Veritas VFS Software v1 are attached locally to Folder 1.

Figure 6-19: Folder with Nodes Attached Locally

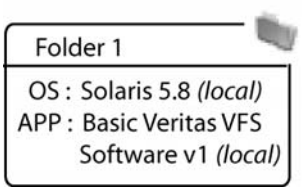

This form of attachment is the simplest. Understanding local attachments is important, as more complex examples are shown.

Inherited

If an attachment is not attached locally to a folder or a template, but is inherited from an ancestor folder, the attachment is considered *inherited*.

Figure 6-20 shows that Folder 1 from the previous diagram now has a child Folder 2 and a child Template 1. The major difference between the two children–the folder and the template–is that you can further customize the folder with more levels of folders and templates, while you cannot do that to the template.

Both the template and the folder inherit Solaris 5.8 and Basic Veritas VFS Software v1 from the parent folder, plus Template 1 has the Patch 105181-34 attached locally that is inherited by its children, as Figure 6-20 shows.

Figure 6-20: Folder with a Child Folder and a Template Showing Attachments Inherited from the Parent Folder

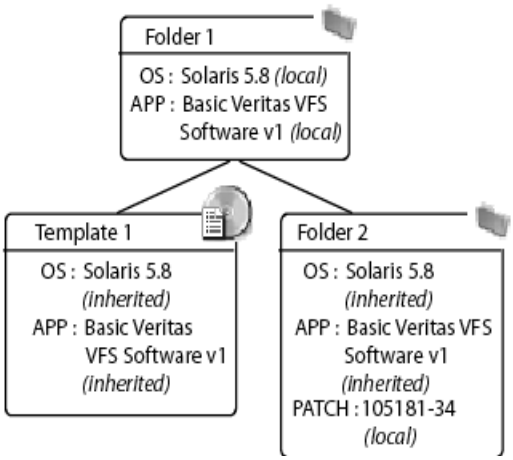

Applying Folder 2 to a group of servers configures them with the inherited OS (Solaris 5.8) and the Basic Veritas VFS Software v1, in addition to the one patch. At this point, you could use either Folder 1 or Template 1 with the same results.

Blocked

You can also block an attachment, which means that it is not installed when that template is applied. Also, blocked attachments do not appear in child templates or folders.

Figure 6-21 shows Template 1 and Folders 1 and 2 from the previous example, but now Basic Veritas VFS Software v1 has been blocked. Consequently, the children created under it, Template 2 and Folder 3, do not inherit that application.

You can only block inherited attachments. You can delete local attachments and their inheritance down the tree is removed. If the parent has an attachment, and you do not want the children to have it, use blocking to prevent it from propagating to its children.

Figure 6-21: Folders with Child Folders and Templates That Show Attachments Blocked from Inheritance

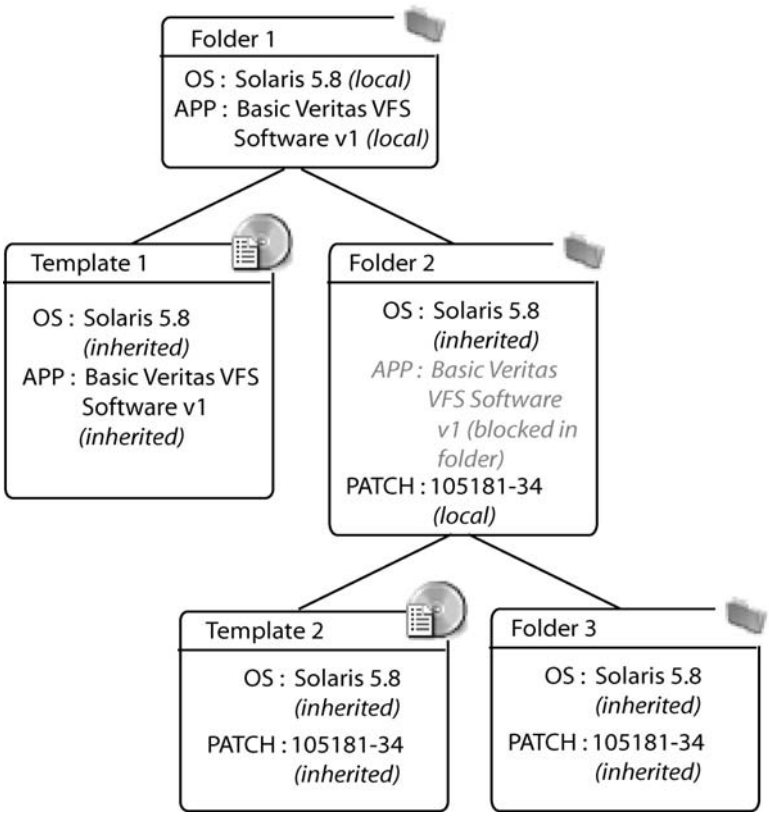

Blocking and Reattaching Inherited Attachments

Blocking and then reattaching an attachment lower in the tree is possible, but this action can be complex and should be done with care.

An attachment that you block in an ancestor folder is not inherited by the children of that folder. You can add the blocked attachment again to a folder or template lower down the hierarchy, and that attachment is again inherited by children still farther down the hierarchy, until it is blocked.

In Figure 6-22, you can see in Step A a new version of Basic Veritas VFS Software that is attached to Folder 2, called v2. Template 2 inherits this new version. But the block on Folder 3 would not apply the v2 software to Folder 3, and subsequently, all children of Folder 3 do not inherit this new version of the Veritas software.

In order for other children in the hierarchy to inherit the Veritas v2 software, you would need to reintroduce (reattach) the software into a child folder. The diagram in Figure 6-22 shows how Folder 5 has had the Veritas software reintroduced, and thus all children of Folder 5, Template 5 and Folder 6, now inherit the new version of the Veritas software.

Figure 6-22: Example of Blocked-Unblocked Attachments in a Template Folder Hierarchy

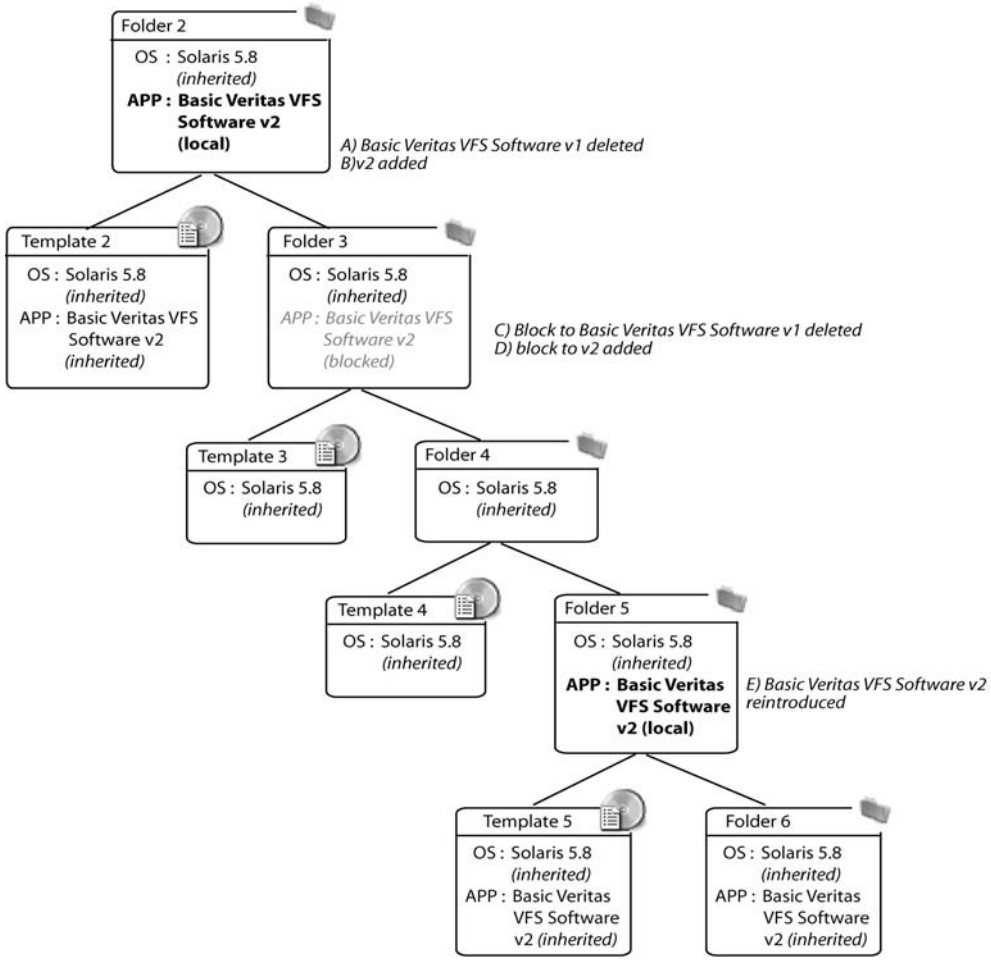

Delete old attachments before you attach the new version.

Blocks do not disappear automatically when you delete the related local attachment. An *orphan* block (one with no local attachments of the same kind) does not affect the system.

Folders and Templates

This section provides information about folders and templates within the Opsware System and contains the following topics:

- Folders and Templates Overview
- Applying Templates and Folders
- Creating Templates
- Creating Folders
- Operating System in Templates or Folders
- Copying Templates and Folders
- Deleting Templates or Folders
- Blocking Folders and Templates from Inheriting
- Blocking vs. Removing Attachments

Folders and Templates Overview

The Opsware System allows you to create both folders and templates. Create folders when you plan to develop a hierarchy that uses inheritance to define generalized parent folders with increasingly specific child folders below. Create templates when you need only a single level of attachments to be defined and you do not need those attachments to be inherited. If later you decide that you want a template to become a folder so that you can build a hierarchy and use inheritance, you can copy the template as a folder and make that copy a folder.

Consider creating folders instead of templates to allow for the future growth of your folder hierarchy.

You can specify an operating system version and customer for the folder or template, and the Opsware System only allows the combination of operating system and customer for the child templates and folders.

Folders and templates require little setup and you can create and deploy them quickly. When you create a folder or template, you select applications, patches, operating systems, or service levels that are already configured and tested for installation, and add them to the folder or template. If, for example, a new patch is released, you can edit the folder or template and add the new patch. Any application of the template to a server would then include the patch.

Applying Templates and Folders

If the template that you create includes an operating system, the template is applied through the Install OS Wizard. If the template you create does not include an operating system (OS Independent), the template is applied through the Install Template Wizard.

If you edit a template, the changes only affect new applications of the template. Servers that already had the template applied are not automatically modified to match the changed template.

Creating Templates

Perform the following steps to create a template:

1. In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.

   You can either create a new template here or navigate to the location in the folder hierarchy where you would like to create the new template.

2 Click the New Template button. The Templates: Create Template page appears, as Figure 6-23 shows.

Figure 6-23: Templates: Create Template Page

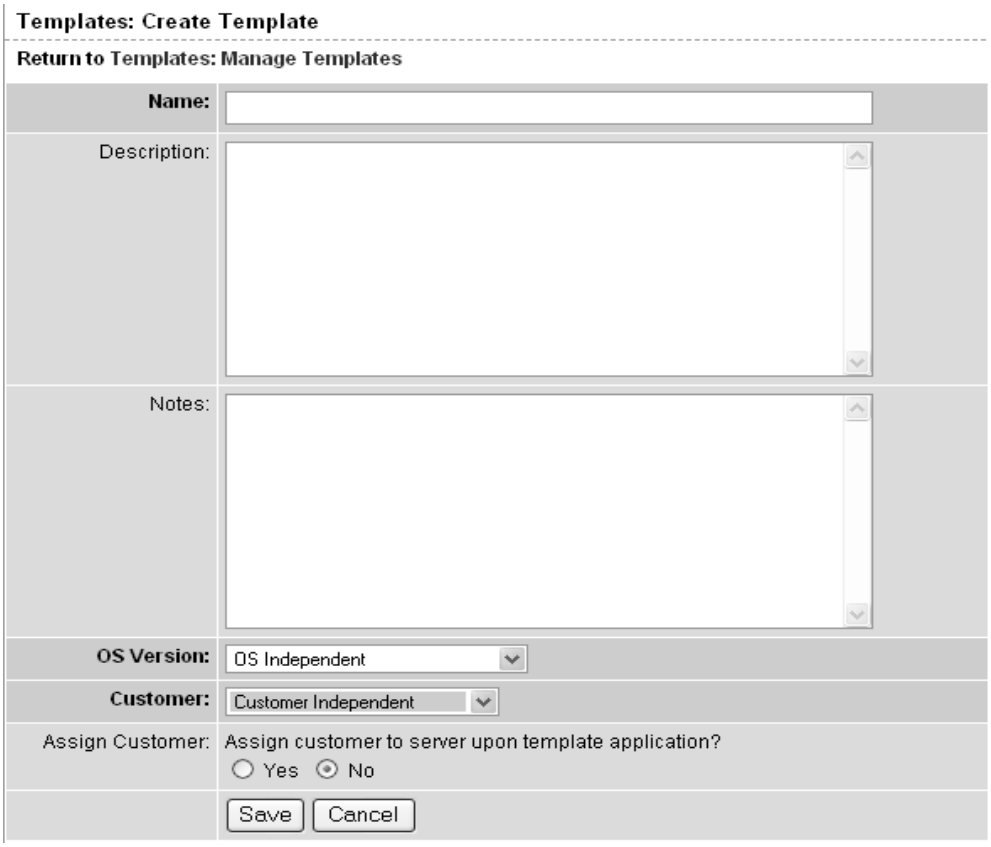

3 Enter a name for the template and then select an operating system version and a customer (all required.) You can also provide a description of the template.

4 If you want the server to which you apply the template to be automatically assigned to the customer associated with the template, select the Yes option in the Assign Customer field.

If the server is already assigned to a customer, selecting this option causes the server to be reassigned when the template is installed.

5 Click Save to create the template. The Templates: Edit Folder | <My New Template> page appears, as Figure 6-24 shows.

Figure 6-24: Templates: Edit Folder | My New Template Page

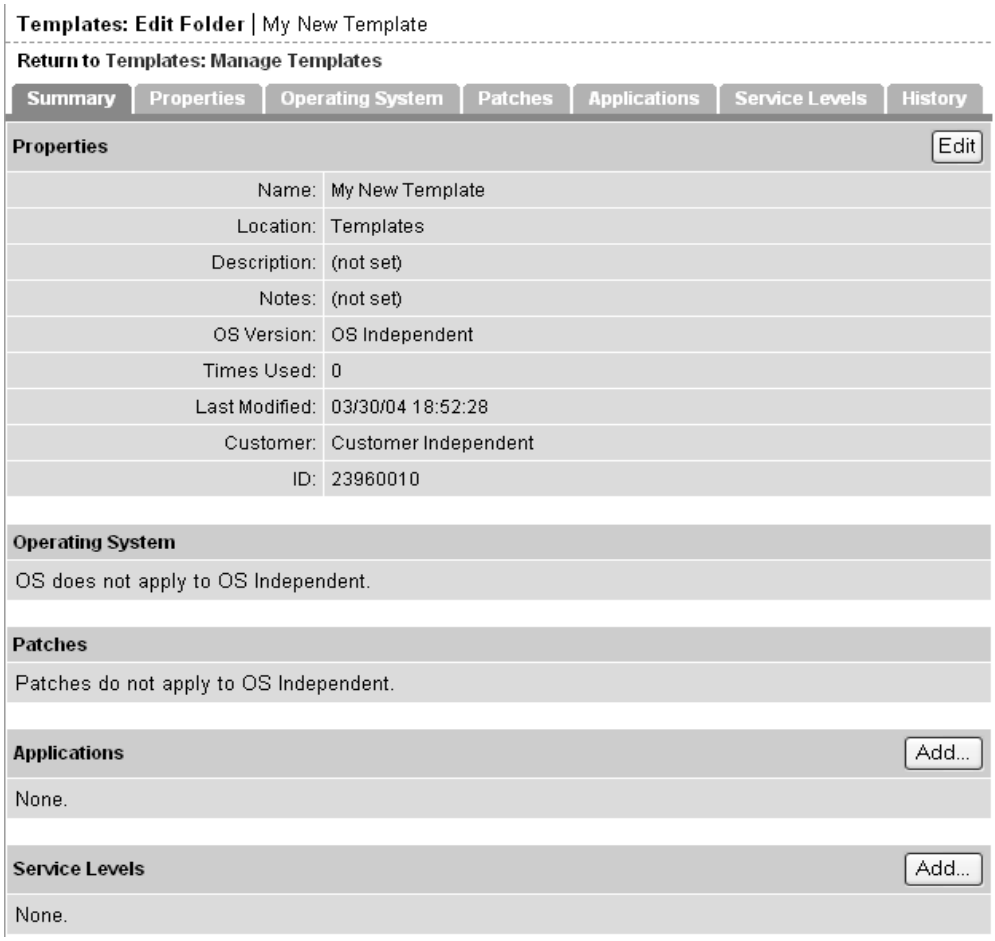

6 You can now select the operating system, patches, applications, and service levels that you want to add to the template.

Creating Folders

Perform the following steps to create a folder:

1 In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.

You can either create a new folder here or navigate to the location in the hierarchy where you would like to create the new folder.

2 Click the New Folder button. The Templates: Create Folder page appears, as Figure 6-25 shows.

Figure 6-25: Templates: Create Folder Page

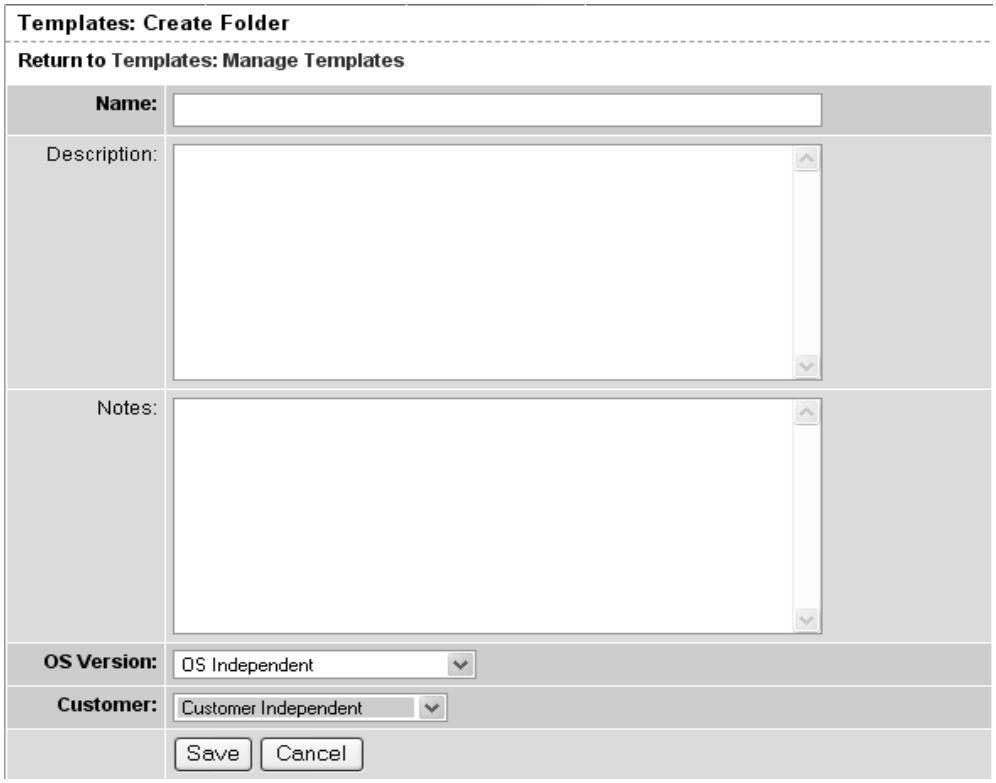

3 Enter a name for the folder, select an operating system version and select a customer (all required). Optionally, you can add a description and notes.

4 Click Save. The new folder is created and the Templates: Edit Folder | <My New Folder> page appears, as Figure 6-26 shows.

Figure 6-26: Templates: Edit Folder | My *New Folder Page*

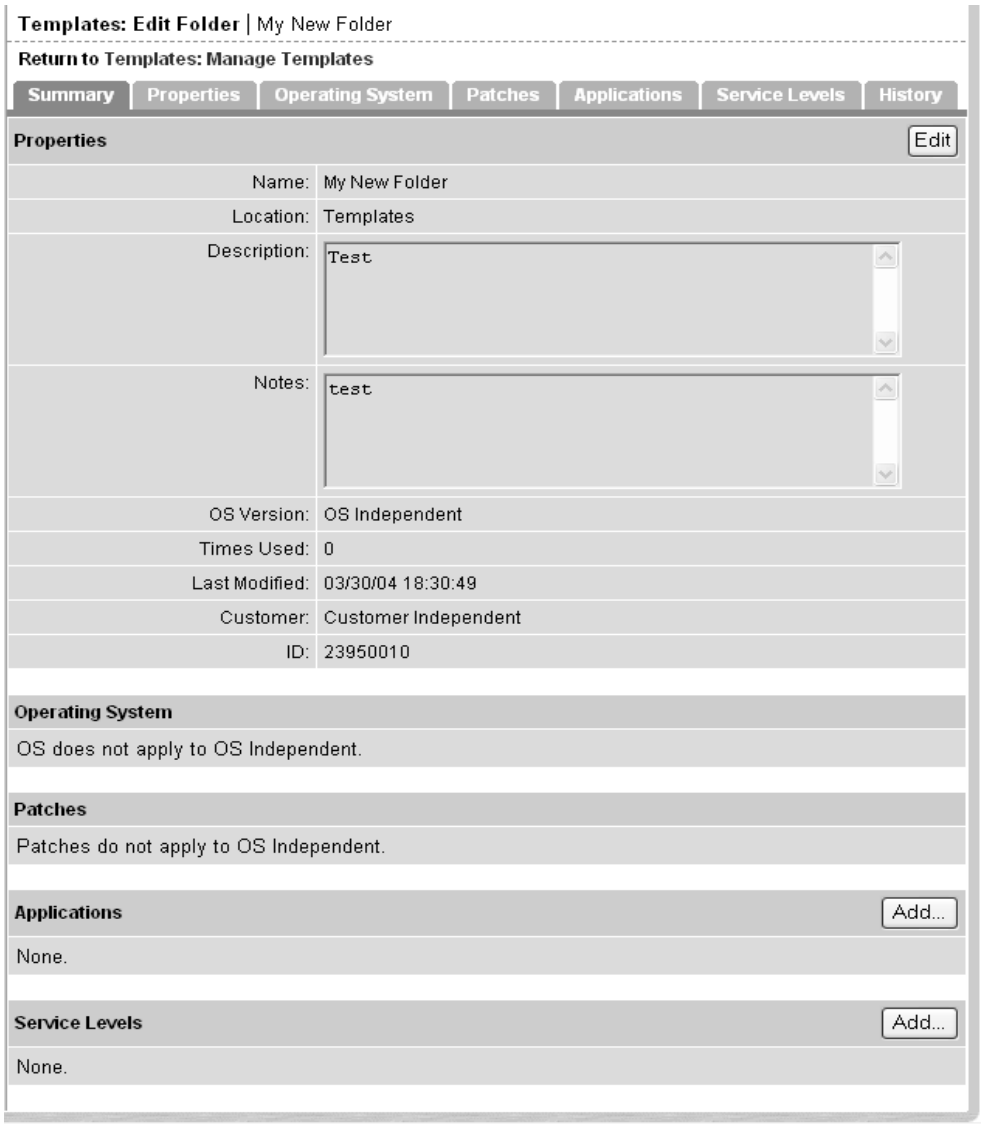

Operating System in Templates or Folders

The following tasks show you how to add, change, and remove an operating system for templates or folders:

- Adding an Operating System to a Template

- Removing an Operating System from a Template
- Adding an Operating System to a Folder
- Removing an Operating System from a Folder

You cannot change or remove an operating system from a template or folder that is OS Independent.

Adding an Operating System to a Template

You cannot add an operating system to a template if the template inherits an operating system from a parent folder.

Perform the following steps to add an operating system to a template:

1. In the navigation panel of the Opsware Command Center, click Software ➤Templates. The Templates: Manage Templates page appears.
2. Select the template where you want to change the operating system. (If a template is inside another folder, you need to navigate to its location before you can select it.) The Templates: Edit Template page appears, as Figure 6-27 shows.

Figure 6-27: Templates: Edit Template Page for Adding an Operating System to a Template

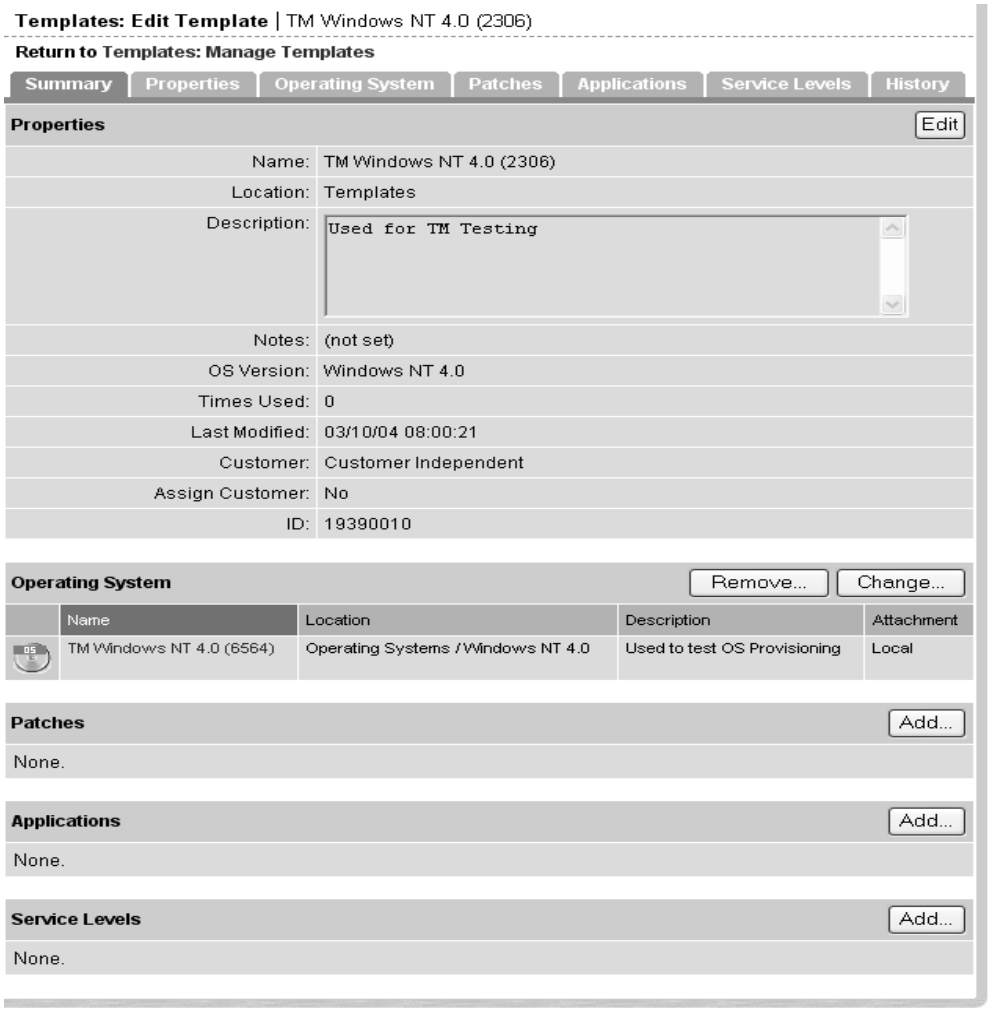

3 Click the Select button in the Operating System field. The Select OS window appears, as Figure 6-28 shows. (If the Template already has an operating system, click the Change button to change the operating system.)

Figure 6-28: Select OS Window

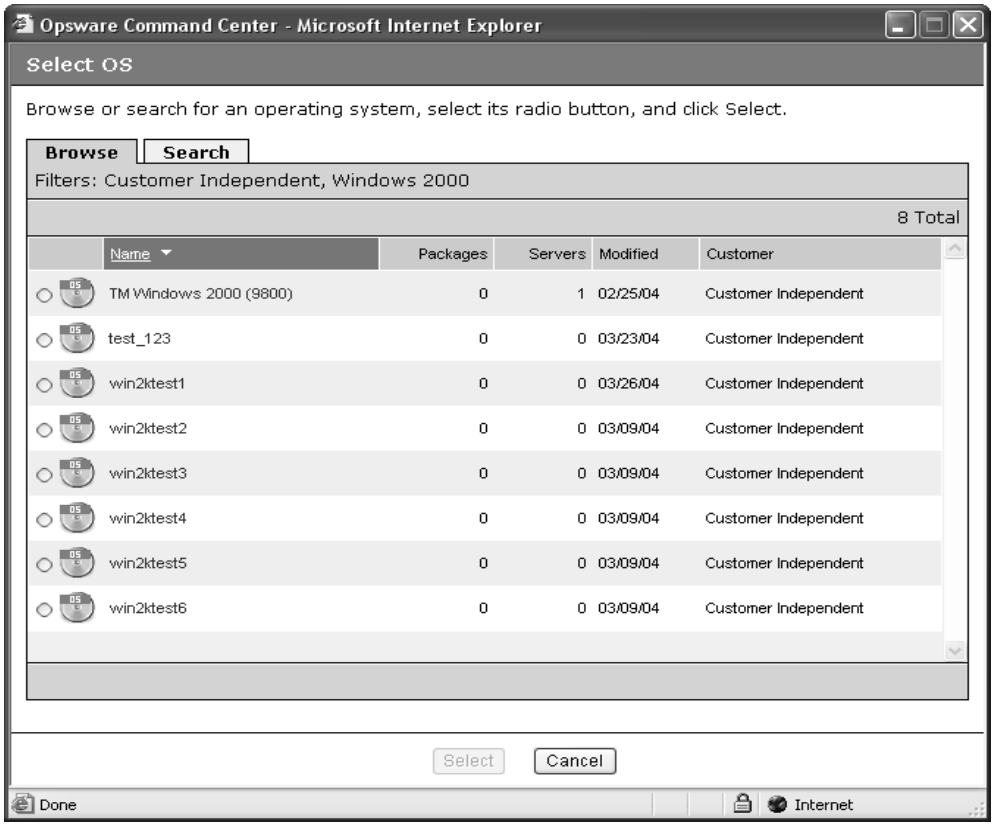

4. Browse or search for the operating system that you want to add to the template and select it by clicking its radio button.

5. Click the Select button to add or change the operating system.

Removing an Operating System from a Template

If the template inherits an operating system from a parent folder, then you cannot remove an operating system from the template.

Perform the following steps to remove an operating system from a template:

1. In the navigation panel of the Opsware Command Center, click Software and then click Templates. The Templates: Manage Templates page appears.

2 Select the template where you want to change the operating system. (If a template is inside another folder, you need to navigate to its location.) The Edit Templates page appears, as Figure 6-29 shows.

Figure 6-29: Templates: Edit Template Page for Removing an Operating System from a Template

3 In the Operating System field, click the Remove button.

4 You are asked to confirm that you want to remove the operating system from the template. Click Yes to remove the operating system from the template.

Adding an Operating System to a Folder

You cannot add an operating system to a folder that has any children that already have an operating system, or that inherits an operating system from a parent folder.

Perform the following steps to add an operating system to a folder:

1. In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.
2. Select the folder where you want to change the operating system. (If the folder is inside another folder, you need to navigate to its location before you can select it.) The folder contents appear, as Figure 6-30 shows.

Figure 6-30: Folder Contents

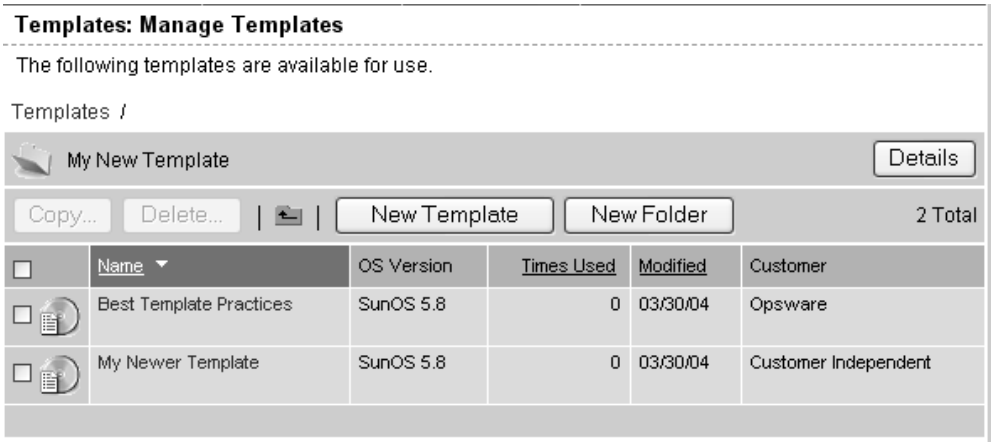

3. Click the Details button. The folder's property summary page appears, as Figure 6-31 shows.

The Details button is only available when the folder has no children that have an operating system or does not inherit an operating system from a parent folder.

Figure 6-31: Folder Properties

4. To add an operating system, from the Operating System section, click the Select button.
5. In the Select OS Window, browse or search for the operating system that you want to install and select it by clicking its radio button.
6. Click the Select button to add the operating system to the folder.

Removing an Operating System from a Folder

You cannot remove an operating system from a folder if the operating system is inherited from the parent.

When you remove an operating system from a folder, the folder no longer appears in the Install OS Wizard when you use a template to install an operating system on a server.

Perform the following steps to remove an operating system from a folder:

1. In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.
2. Select the folder where you want to remove an operating system. (If the folder is inside another folder, you need to navigate to its location before you can select it.)
3. In the Folder Contents page, click the Details button.
4. In the Summary page for the selected folder, in the Operating System field, click the Remove button, as Figure 6-32 shows.

Figure 6-32: Remove Operating System from a Folder

5 You are asked to confirm that you want to remove the operating system from the folder. Click Yes to remove the operating system from the folder.

If you remove an operating system from a folder that has child folders, you are not able to add another operating system at the same level. Instead, use the Change button in the Operating System field.

Adding Applications to Templates or Folders

Perform the following steps to add an application to a template or to a folder:

1. From the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.
2. Select the template or folder where you want to add an application. (If a template or folder is inside another folder, you need to navigate to its location before you can select it.)
3. If you are adding an application to a folder, click the Details button. If you are adding an application to a template, go to the next step.
4. To include applications in your template, click the Add button in the applications field. The Add Applications page appears, as Figure 6-33 shows.

Figure 6-33: Add Applications Page

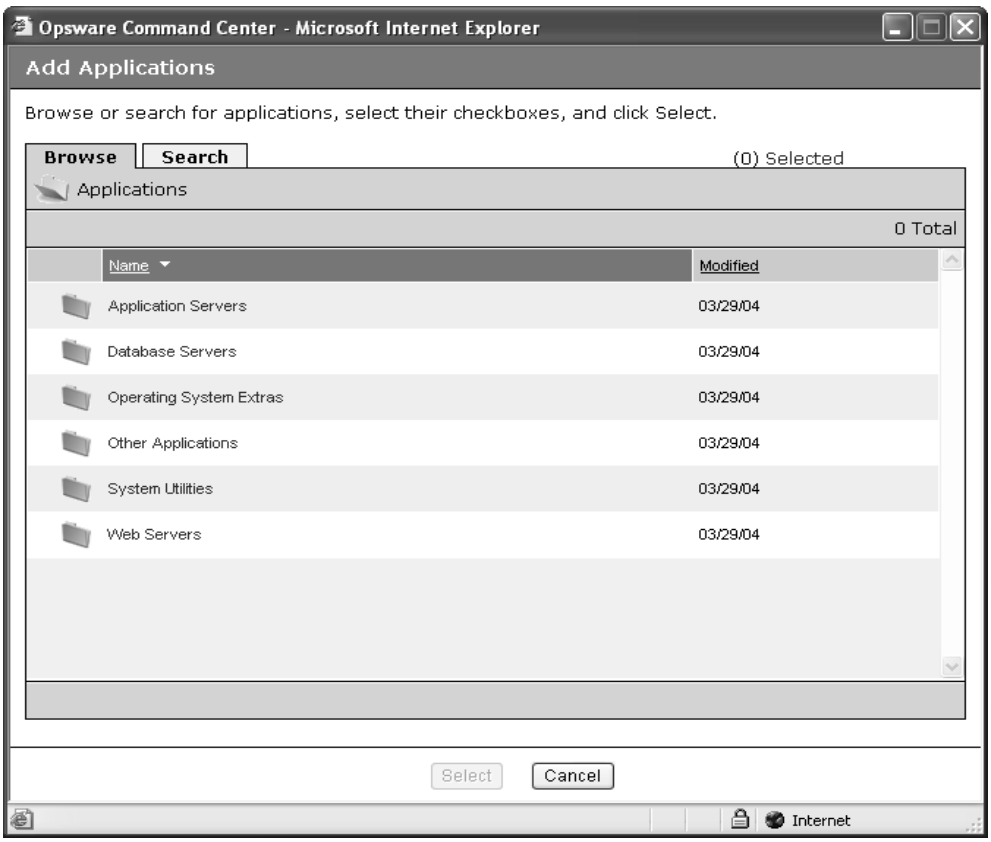

5 Browse or search for the applications that you want to add. Select the applications by clicking their checkboxes, and then click the Select button to add them to your template or folder, as Figure 6-34 shows.

Figure 6-34: Adding Applications to a Template

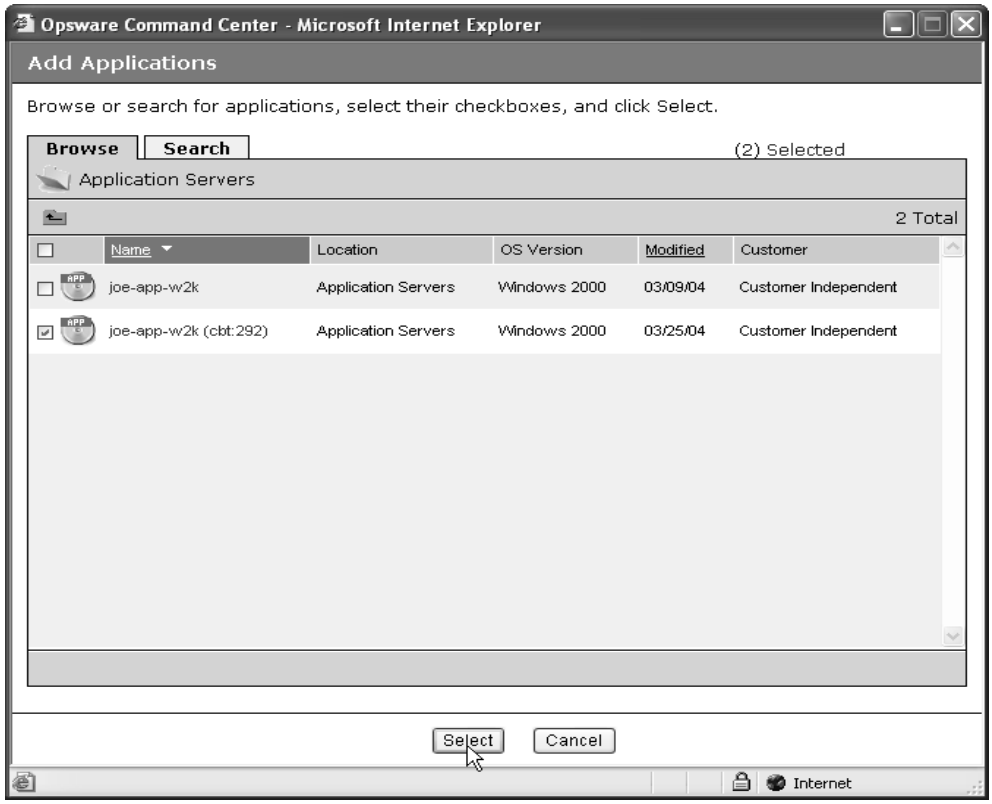

6 When you add the application to the template or folder, it shows in the Applications field of the Edit Template page, as Figure 6-35 shows.

Figure 6-35: Applications Added to a Template

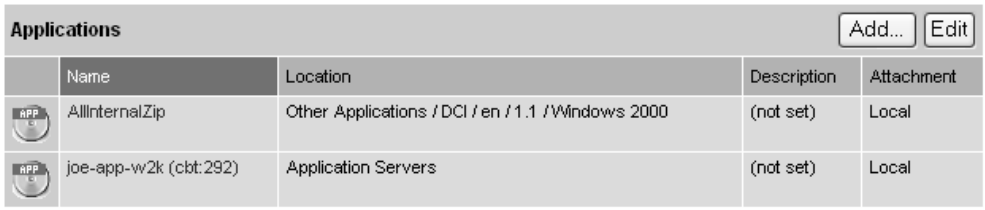

Editing or Removing Applications for Templates or Folders

Perform the following steps to edit or to remove applications:

1. In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.
2. Select the folder or template where you want to edit an application. (If a template or folder is inside another folder, you need to navigate to its location before you can select it.)
3. If you are adding an application to a folder, click the Details button. If you are adding an application to a template, go to the next step.
4. Click the Edit button in the Applications field. The Templates: Edit Template page appears with the Applications tab active, as Figure 6-36 shows.

Figure 6-36: Edit Applications Page

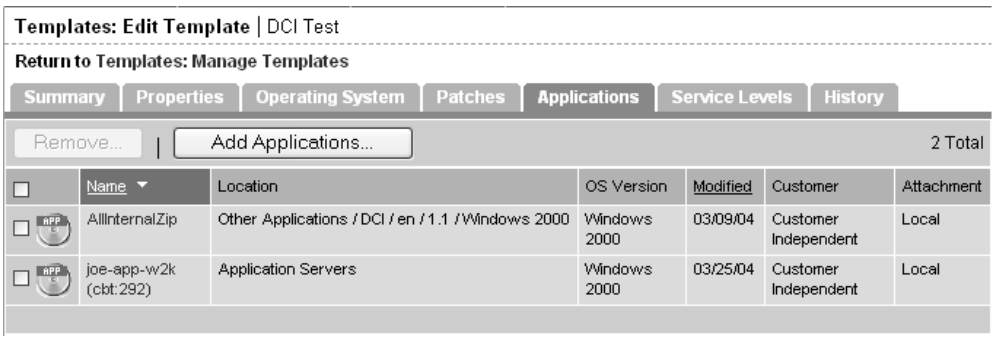

5. To remove applications, click the checkbox next to the name of the application that you want to remove.

Inherited and blocked applications do not have a selectable checkbox so you cannot remove them here. You have to remove them from the folder where they are attached locally.

6. Click the Remove button.
7. The Remove Applications confirmation page appears. Click the Remove button to remove the applications.

8 You are asked to confirm that you want to remove the selected application. Click Yes to remove the application.

Adding Patches to Templates or Folders

In the Patches area of the Edit Template page, you have the option of either adding or editing patches. If you choose to edit patch attachments, you can both remove existing patches and add new patches.

Perform the following steps to add patches to a template or a folder:

1 In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.

2 Select the folder or template where you want to add a patch. (If a template or folder is inside another folder, you need to navigate to its location before you can select it.)

3 If you are adding a patch to a folder, click the Details button. If you are adding a patch to a template, go to the next step.

4 Click the Add button in the Patches field. The Add Patches window appears, as Figure 6-37 shows.

Figure 6-37: Add Patches Window

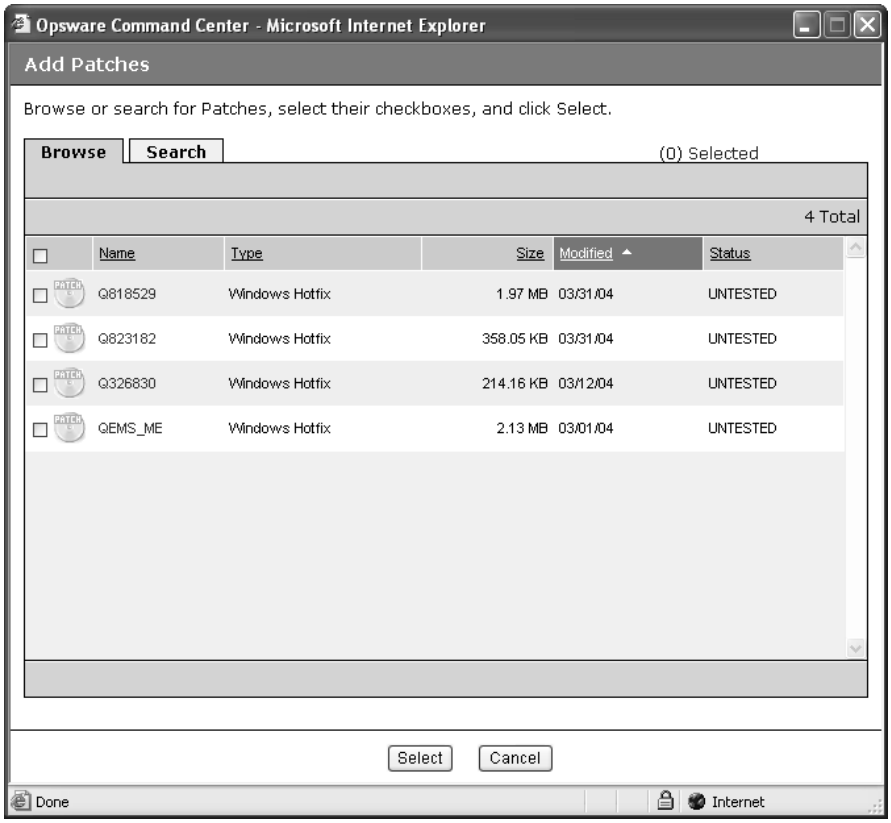

5 Browse or search for the patches that you want to add to the template or folder. Click the checkboxes for the desired patches and then click the Select button to add the patches.

Editing or Removing Patches for Templates or Folders

In the Patches area of the Edit Template page, you have the option of either adding or editing patch attachments. If you choose to edit patch attachments, you can remove existing patches *and* add new patches.

Perform the following steps to edit or remove patches:

1 In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.

2. Select the folder or template where you want to edit a patch. (If a template or folder is inside another folder, you need to navigate to its location before you can select it.)
3. If you are editing the patch of a folder, click the Details button. If you are adding a patch to a template or folder, go to the next step.
4. Click the Edit button in the Patches field. The Templates: Edit Template page appears with the Add Patches tab active, as Figure 6-38 shows.

Figure 6-38: Templates: Edit Template Page - Edit Patches

5. To remove patches, click the checkbox next to the name of the patches that you want to remove.

Inherited and blocked patches do not have a selectable checkbox so you cannot remove them here. You must remove them from the folder where they are attached locally.

6. Click the Remove button
7. The Remove Patches confirmation page appears. Click the Remove button to remove the patches from the template or folder.
8. To add patches to the template or folder, click the Add button. The Add Patches page appears.
9. Browse or search for the patches that you want to add. Select them by clicking their checkboxes and then click the Select button.

Adding Service Levels in Templates or Folders

Perform the following steps to add service levels in templates or folders:

1. In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.
2. Select the template or folder where you want to add a service level. (If a template or folder is inside another folder, you need to navigate to its location before you can select it.)
3. If you are adding a service level to a folder, click the Details button. If you are adding a service level to a template, move to the next step.
4. Click the Add button in the Service Levels field. The Add Service Levels page appears, as Figure 6-39 shows.

Figure 6-39: Add Service Levels Page

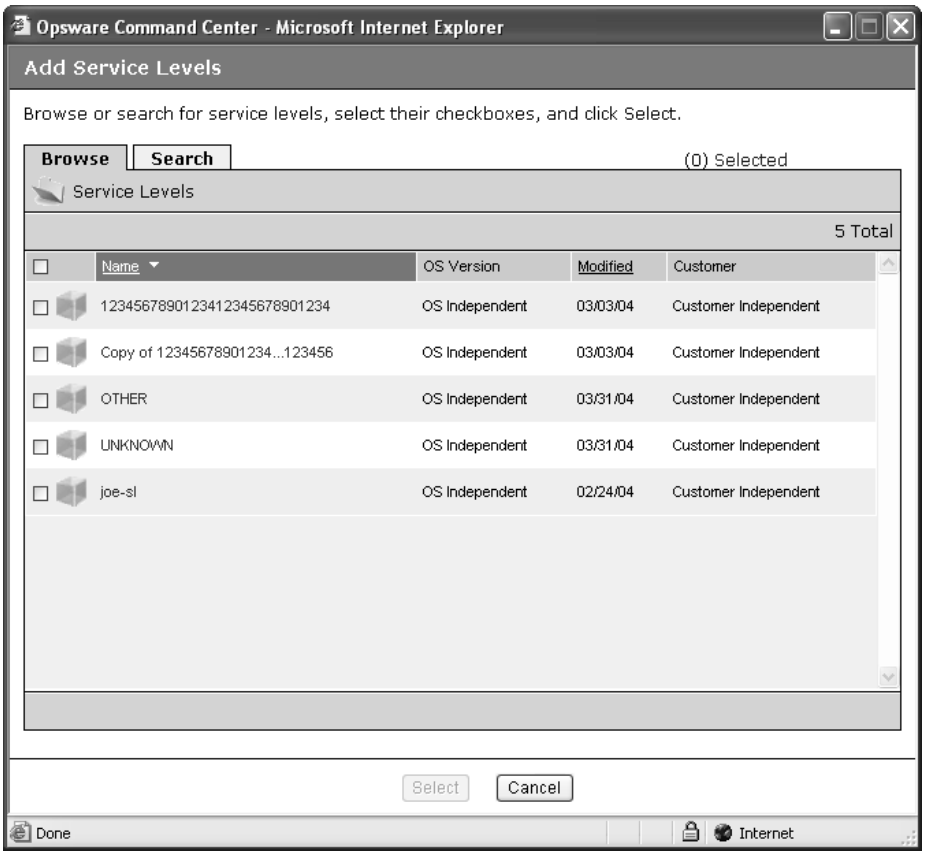

5. Browse or search for the service levels that you want to include in your template. Select the service levels by clicking the checkboxes, and then click the Select button to add them to your template or folder.

See “Service Levels” on page 165 in Chapter 2 for information about Service Levels.

Editing or Removing Service Levels in Templates or Folders

Perform the following steps to edit or remove service levels in templates or folders:

1. Click the Edit button in the Service Levels field. The Edit Template page appears with the Service Levels tab active.
2. To remove service levels, click the checkbox next to the name of the service levels that you want to remove.

Inherited and blocked service levels do not have a selectable checkbox so you cannot removed them from here. You have to remove them from the folder where they are attached locally.

3. Click the Remove button.
4. The Remove Service Levels confirmation page appears. Click the Yes button to remove the service levels from the template or folder.

Copying Templates and Folders

You can copy a template or a folder as a convenient way to use existing folders or templates without having to re-create a template or an entire hierarchy of folders.

When you copy a template, the copy has the same attachments as the original template. If the original template has attachments inherited from an ancestor folder, the copy also has all the ancestor folders and the local attachments that were inherited by the template.

When you copy a folder, the copy has the same attachments and the same children with all of the attachments as the original folder.

Perform the following steps to copy a template or a folder:

1. In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.
2. Select the folder or template that you want to copy. (If a template or folder is inside another folder, you need to navigate to its location before you can select it.)
3. If you are copying a folder, click the Details button. If you are copying a template, move to the next step.

4 Click the Copy button.

5 In the Copy Template/Folders window, enter the name that you want to apply to the copied folder or template or accept the default, as Figure 6-40 shows.

Figure 6-40: Copy Template Window

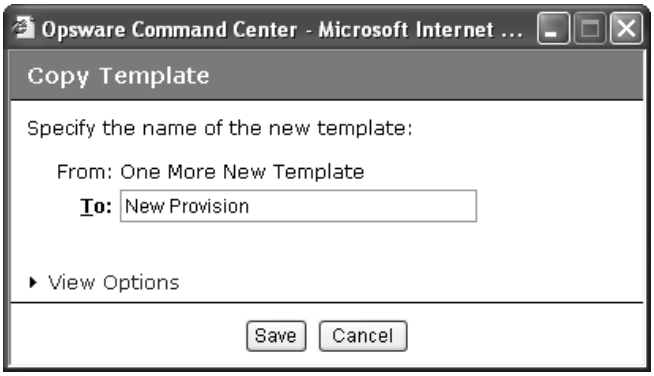

The default option is to copy a folder as another folder and a template as a template.

6 You have the option of clicking the View Options link to show the Copy as Folder or Copy as Template options.

If you copy a folder that already has children, and you would like to make a copy of the folder and all its children, make sure you copy as a folder. If you copy a template or a folder as a template, it is unable to have children.

7 Click the Save button. The new folder or template now appears on the list of templates.

Deleting Templates or Folders

Perform the following steps to delete a template or a folder:

1 In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.

2 Select the folder or template that you want to delete. (If a template or folder is inside another folder, you need to navigate to its location before you can select it.) Selecting a folder or template activates the Copy and Delete buttons.

3 Click the Delete button.

4 You are asked to confirm that you want to delete the folder. Click Yes to complete the deletion, or No to cancel. The deleted folder or template is removed from the list of templates.

Blocking Folders and Templates from Inheriting

Making a change to an attachment at the top level of a template hierarchy results in a change to all the children of that top-level folder, unless a child folder has already had the attachment blocked.

You can only block inherited attachments. The equivalent for a local attachment is to delete it.

Blocking vs. Removing Attachments

The following list describes some issues to think about when you are deciding to block or remove an attachment from a folder or template:

- You have the option of blocking attachments and preventing inheritance from a parent folder.
- You have the option of removing local attachments.
- Block an attachment when you do not want it to be inherited at the current level or at any subsequent levels of the hierarchy.
- Remove a local attachment when you do not want it to be in your hierarchy at all, or if you plan to attach a different version of the node. If you plan to attach a different version, remove the old version first to avoid conflicts.
- If you remove an attachment, any blocks to that attachment remain in the hierarchy until you remove them. These blocks continue to appear, but when you move your mouse over them you see that they are inactive. If you click an inactive block, a window appears asking if you want to delete the block. Click Yes to remove it.

To Block an Attachment from Being Inherited

Perform the following steps to block an attachment from being inherited:

1 In the navigation panel of the Opsware Command Center, click Software ➤ Templates. The Templates: Manage Templates page appears.

2 Navigate to the folder whose attachment you want to block.

3 Locate the operating system, patch, application, or service level attachment that you want to block.

4 Click the link that says Inherited. It changes to Blocked. Click it again to change it to Inherited if you want to remove the block, as Figure 6-41 shows.

Figure 6-41: Blocking Inherited Attachment

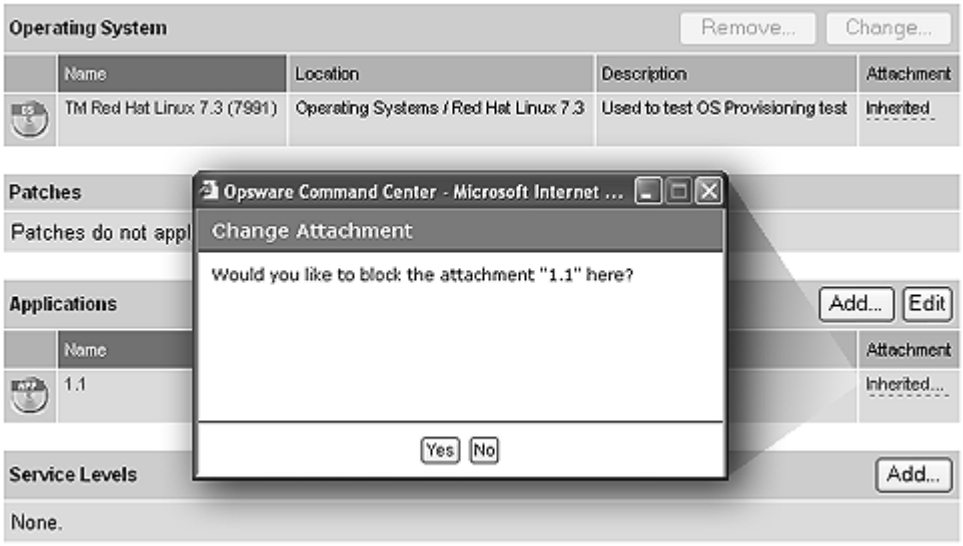

5 To remove the block, click the link named Blocked. A dialog box allows you to change it back to Inherited.

Chapter 7: Application Provisioning

IN THIS CHAPTER

This chapter provides information about how to use the Opsware System wizards to install and uninstall applications on Opsware-managed servers.

This chapter contains the following sections:

- Installing and Uninstalling Software Overview
- Software Installation and Uninstallation Issues
- Installing Templates

Installing and Uninstalling Software

This section provides information on how to install and uninstall software by using the Opsware System and contains the following topics:

- Installing and Uninstalling Software Overview
- Software in the Opsware System
- Ways to Install Software
- Types of Software
- Summary of Application Provisioning Features
- Platform-Specific Reconcile Overview

Installing and Uninstalling Software Overview

The Opsware System automates the time-consuming process of installing and uninstalling software on servers. Using the Opsware System, you can quickly deploy applications across a large number of servers with a minimum amount of downtime. The Opsware System allows you to select software packages from the Software Repository, select the servers that you want to install the software on, preview the results of the installation, and install the software all in a single operation.

The Opsware System provides detailed feedback about which packages are installed, what events occurred (such as a server reboot), the output of scripts, and any errors that occurred. The Opsware System also provides the same degree of control for uninstalling applications.

You can uninstall any application that you installed through the Opsware System. You can also preview the results of an installation and then schedule the installation for some later time. If, for example, the Opsware System reports that the software installation requires multiple reboots, you can schedule the installation for a time when the reboots cause the least disruption to your services.

Software in the Opsware System

In the Opsware System, software packages reside in a central Software Repository. Opsware users upload the software and also specify options that ensure that the software is correctly installed. The users add pre- and post-install and uninstall scripts to software packages that help control the way that the software is installed.

See "ZIP Package Management" on page 283 in Chapter 5 for information about how to upload packages to and manage packages in the Software Repository.

The Opsware System maintains detailed information about the state of every server under management in a central database called the Model Repository.

Ways to Install Software

The Opsware System provides two primary ways to install software. You can select a single application package and install it, or you can select a template that includes a number of different (and usually related) software packages that you can install in a single operation. Opsware users create and test the templates, and they provide the same level of control and consistency that is available when you install individual software applications. Templates, for example, allow you to install a set of applications, such as a Web server, an Oracle database, and related applications that allow you to quickly deploy a server in a fully operational state.

Types of Software

The Model Repository usually stores information about hundreds of different types of software packages. To make the selection of software easier, the Opsware System organizes software into the following categories:

- Application Servers

- Database Servers
- Operating System Extras
- Other Applications
- System Utilities
- Web Servers

Each of these categories can contain software for any of the operating systems that the Opsware System supports. You can browse through the list of software, or you can also use the Opsware System's advanced search capabilities to locate the packages or templates that you want to apply.

Summary of Application Provisioning Features

The Application Provisioning Subsystem has the following features:

- A central location (the Software Repository) for software package storage
- Consistency of installation with user-created scripts that specifies options to ensure that software is installed in a uniform way
- The ability to preview the installation and uninstallation processes to see what packages will be installed and what server operations are required (such as rebooting)
- The ability to install individual software packages quickly across a large number of servers
- The ability to quickly apply software templates, which includes bundles of applications, across a large number of servers
- Auditing of all changes made to managed servers and quick uninstallation in case problems result

Platform-Specific Reconcile Overview

The Reconcile step is the last step that you perform when installing software. During that step, the Opsware System provides you with information about software that is currently installed on a server, along with information about the software that is about to be installed, and software that is about to be removed. Chapter 9 describes this process in detail. Table 7-1 gives details about platform-specific reconciles.

Table 7-1: Platform-Specific Reconcile

OPERATING SYSTEM	SEE THIS PAGE
AIX	See "AIX Reconcile" on page 449 in Chapter 9 for more information.
HP-UX	See "HP-UX Reconcile" on page 449 in Chapter 9 for more information.
Linux	See "Linux Reconcile" on page 450 in Chapter 9 for more information
Solaris	See "Solaris Reconcile" on page 450 in Chapter 9 for more information.

See "Utilities Used During Reconcile Sessions" on page 448 in Chapter 9 for more information. This table shows the utilities used during the reconcile process for all supported platforms.

Software Installation and Uninstallation Issues

When you upload software to the Software Repository, a user can specify a number of options that affect what happens when you install software by using the Install Software Wizard. The following sections explain the consequences of these options:

- Script Error Conditions
- Inheritance and the Install and Uninstall Software Wizards
- Installing Software with the Install Software Wizard
- Installing Packages on Servers with Low Disk Space
- Uninstalling Software with the Uninstall Software Wizard

Script Error Conditions

A software package that you select in the Install Software Wizard can include installation scripts. A user can specify what happens if the pre- or post-install script encounters an error condition. The user can specify an option that allows the software installation or uninstallation to proceed in spite of an error, or the user can select an option that prevents the software installation if the install script returns an error (for example, exits with a non-zero return code.) You can view the error message when you preview the installation or uninstallation.

Inheritance and the Install and Uninstall Software Wizards

Software in the Opsware System is organized into a Software Tree. Each branch of the tree supports inheritance. A node lower in the branch inherits all the software and attributes of the nodes above it in the same branch.

See “How Software Is Inherited from Other Nodes” on page 344 in Chapter 6 for information about how to set up inheritance for software nodes.

Ordinarily, a branch in the tree is set up for one software package. In some cases, however, the branch might include further nodes. For example, the child node of an application node might contain patches for the application.

This same inheritance tree displays in the Install and Uninstall Software Wizards. The branches of a software tree display as a familiar folder hierarchy.

In most cases, the last level of the folder hierarchy contains a particular application that you want to install or uninstall. Keep in mind, however, that if the folder hierarchy continues with additional software (such as patches), all software above the node that you select is also installed or uninstalled. For example, if you are using the Install Software Wizard and you select a subfolder that contains a patch for an Oracle database – and the parent folder contains the Oracle database itself – selecting the patch for the database also causes the database itself to be installed.

Installing Software with the Install Software Wizard

Perform the following steps to install software with the Install Software Wizard:

1. Launch the Install Software Wizard from the Software Provisioning pane on the Opsware Command Center home page. You are presented with an overview of the steps for installing software.
2. Click Start to continue. The Select Software page appears, as Figure 7-1 shows.

3 Select the category of software that you want to install (Application Server, Database Server, and so forth), or select the Search tab and search for the software package. Select one or more software packages to install. (After you make your first selection, you are only able to select additional software that belongs to the same operating system, version, and customer.) Figure 7-1 shows an example of possible software selections.

Figure 7-1: Select Software Page

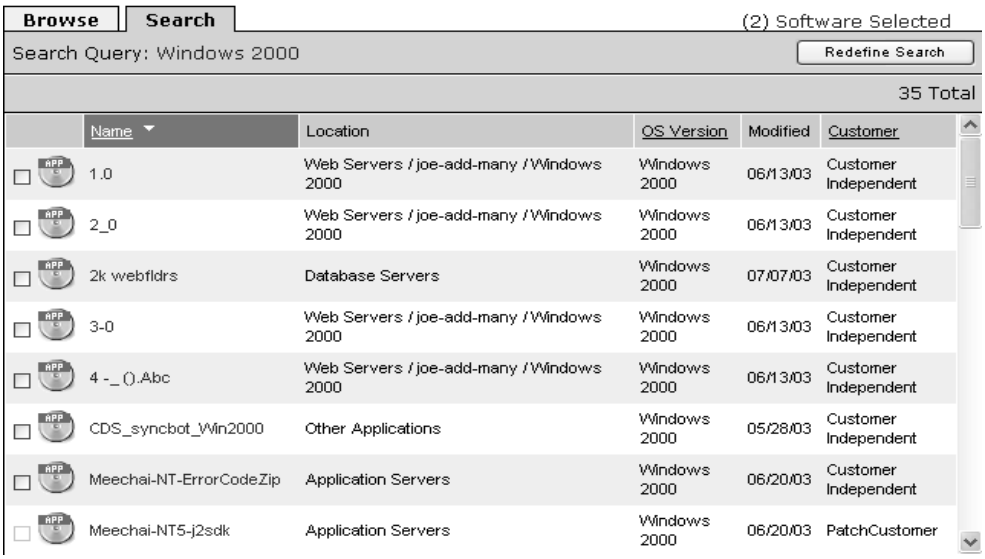

4 (Optional) Click the name of any of the software packages to display additional information about the package, as Figure 7-2 shows.

Figure 7-2: Software Details Page

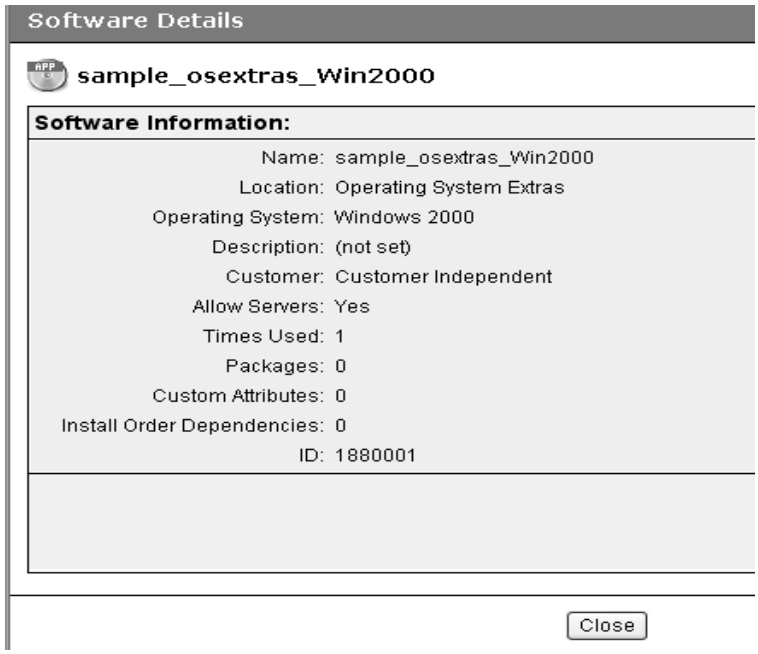

5 Click Next to continue. The Select Servers page appears, as Figure 7-3 shows.

Figure 7-3: Select Servers Page

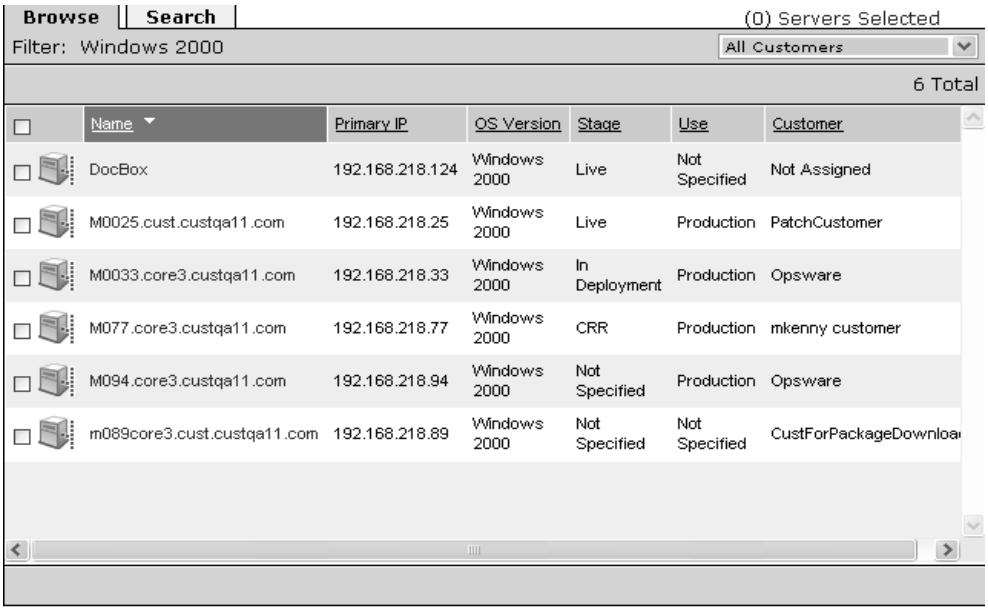

The selection of servers is limited to those that match the operating system and customer of the software that you selected.

6 Select the servers where you want to install the software by browsing or searching. Click Next to continue. The Confirm Selection page appears that displays the servers and software that you selected in previous steps.

7 Review your selections and click Preview to continue. The Preview Page appears. The Preview Page initially displays status bars for each server that indicates the progress of the preview process.

8 (Optional) When the process is completed, click View Details to see what happens when you actually install the software, as Figure 7-4 shows.

Figure 7-4: View Details Page

Packages that Opsware determined should not be installed:

	Name	Reboot	Messages
	qa_tagMSI_Unit20	No	qa_tagMSI_Unit20 1.0 is already installed
	qa_tagMSI_Unit21	No	qa_tagMSI_Unit21 1.0 is already installed
	qa_tagMSI_Unit22	No	qa_tagMSI_Unit22 1.0 is already installed
	qa_tagMSI_Unit23	No	qa_tagMSI_Unit23 1.0 is already installed
	qa_tagMSI_Unit24	No	qa_tagMSI_Unit24 1.0 is already installed
	qa_tagMSI_Unit25	No	qa_tagMSI_Unit25 1.0 is already installed
	qa_tagMSI_Unit1	No	qa_tagMSI_Unit1 1.0 is already installed

9 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

10 (Optional) When the process is completed, click View Details to see what packages were installed or removed, and the output, if any, from the installation scripts.

Installing Packages on Servers with Low Disk Space

If a managed server does not have enough disk space to download packages from the Software Repository to be installed, you can specify locations such as shared network drives, or a CD-ROM, where the Opsware System should look for packages to install.

Specifying Paths for Package Installation

Perform the following steps to specify paths for package installation:

1 Log in to the Opsware Command Center.

2. From the navigation panel, select Servers ➤ Managed Servers.
3. Navigate to the servers where you are defining a new custom attribute, and click the check box next to the name of each server that you are defining a new custom attribute for.
4. Click the Custom Attributes tab. The Managed Servers: Custom Attributes | [*server name*] page appears.

   ```
   /shared/hpux_depots
   ```

5. Click the New button to open a new Custom Attribute form.
6. Enter OPSWpackage_paths in the name field. Be sure to use this exact spelling and case for each server for which you specify paths for package installation.
7. In the Value field, enter each path where the Opsware System should look for the package to install. For example:

   ```
   /mnt/cdrom
   ```

   or

   ```
   /networkshare/packages/SunOS/5.6/
   ```

   ```
   /mnt/cd0
   ```

   You can enter any number of paths, which the Opsware System attempts to use to find the package. The Opsware System tries each path, one after the other, until the package is located. If it is not found in any of the specified locations, the Opsware System looks in the Software Repository. In that case, the Opsware System attempts to download the package if there is enough disk space. If the Opsware System calculates that there is not enough disk space, an error message appears, and the packages are not downloaded.

   Check that permissions on files are set so that the Opsware System has read access.

You can specify the Software Repository as a path by using opsware_repository as one of the values. This is useful in cases where you enter a number of pathnames and want to disable the feature temporarily without having to re-enter the pathnames when you are ready to enable the feature again. Enter this value at the top of the list of values.

Table 7-2 shows the operating systems and package types that you can use with the Low Disk Space feature.

Table 7-2: Supported Operating Systems for the Low Disk Space Feature

OPERATING SYSTEM	FILE TYPE
Sun Solaris	RPM only
Linux	RPM
IBM AIX	RPM, APAR, Base Fileset, Update Fileset, Maintenance Level,
HP-UX	Depots, disk format only

Special Operating System Requirements

Some operating systems have additional requirements when you use this low disk space feature:

- HP-UX patches must be in disk format.
- When you add new install packages to an AIX installation directory, always run the `inutoc` command to ensure that the installation subsystem recognizes the new packages. This command creates a new .toc file.

  Run the `inutoc` command in the AIX installation directory or with the installation directory as the only parameter.

Write access as root to the installation directory is required for NFS-mounted installation directories. For example, if the installp packages are in /tmp/sys/inst.images, run the command as follows:

```
cd /tmp/sys/inst.images ; inutoc
```

or

```
inutoc /tmp/sys/inst.images
```

Uninstalling Software with the Uninstall Software Wizard

The process for uninstalling software is nearly identical to the process of installing software. When you use the Uninstall Software Wizard, you can select any number of servers running the same version of the same operating system to remove software from. (The servers must, however, be assigned to the same customer as the customer associated with the software.)

Perform the following steps to uninstall software with the Uninstall Software Wizard:

1. Launch the Uninstall Software Wizard from the Software Provisioning pane on the Opsware Command Center home page. You are presented with an overview of the steps for uninstalling software.
2. Click Start to continue. The Select Servers Page appears.
3. Select the OS version of the servers that you want to uninstall software from by browsing or searching. The type of server is defined by its operating system type and operating system version, as Figure 7-5 shows.

Figure 7-5: OS Versions Page

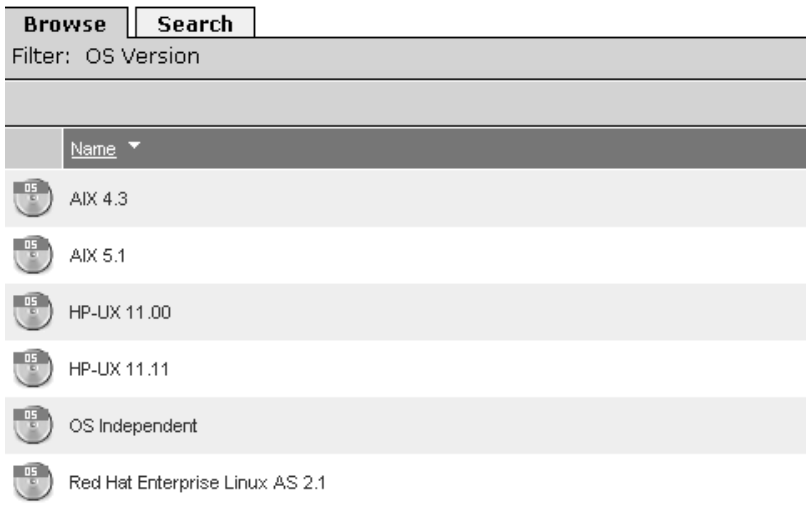

After you select the OS version, the Select Servers page appears. The Select Servers page presents a list of all servers of the type that you selected, as Figure 7-6 shows.

Figure 7-6: Select Servers Page

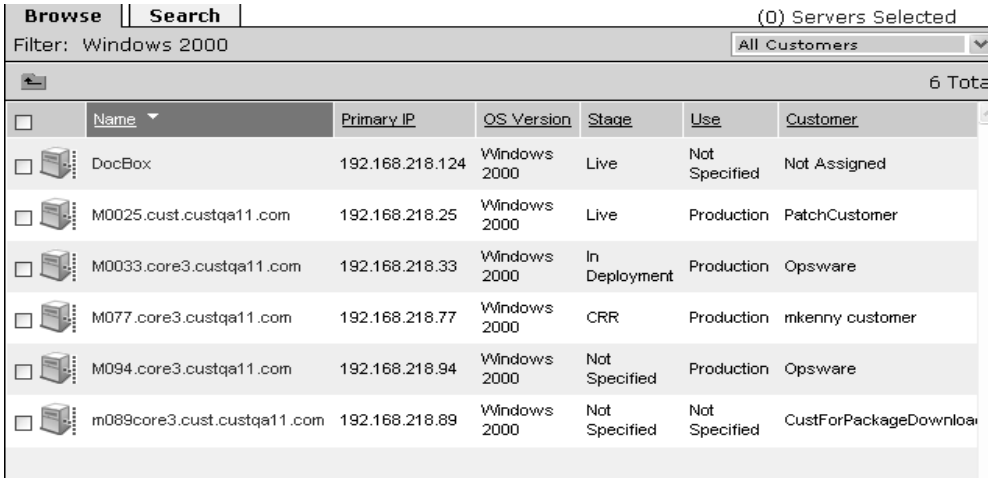

4 Select one server or multiple servers and click Next. The servers that you select, however, must all have at least one common software node assigned to them in order for you to select software to uninstall.

Alternatively, you can use the search function to find servers that all have a particular package installed on them, as Figure 7-7 shows.

Figure 7-7: Search Servers Page

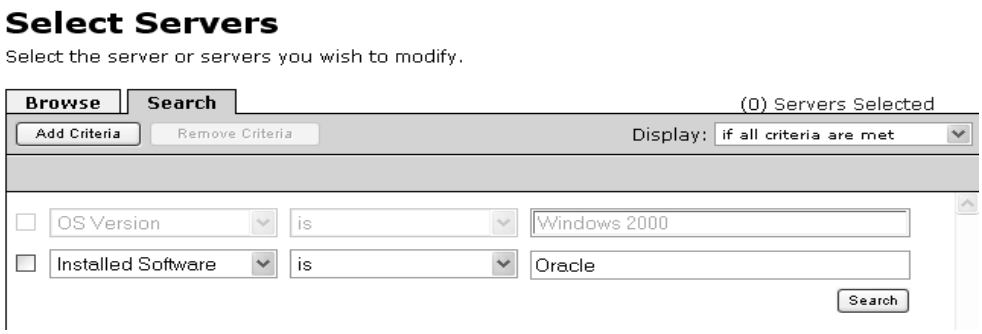

When searching, the first search condition, OS Version, is preselected and cannot be changed. This is because you already selected the operating system version in the previous step.

After you click Next, the Select Software page appears. The Select Software page shows a list of software packages common to all the servers that you selected, as Figure 7-8 shows.

Figure 7-8: Software Common to Selected Servers Page

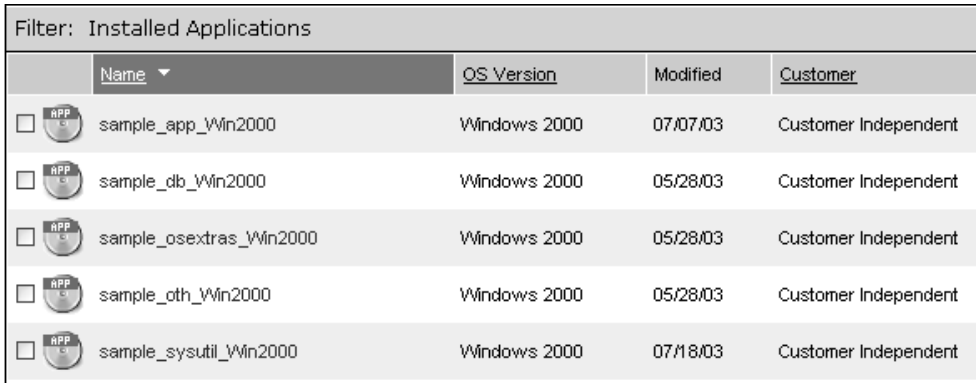

If you selected a group of servers that has no software nodes in common, this list is empty.

5. Select the software that you want to uninstall and click Next to continue. The Confirm Selection page appears. The Confirm Selection page displays the servers and software that you selected to uninstall.

6. Review your selection and click Preview to continue. The Preview page appears. The Preview page displays a progress bar for each server, which shows the status of the preview process.

7. When the preview process completes, click View Details to see what occurs when you uninstall software.

   From the summary, you can see what packages will be uninstalled, including packages that you have not selected but that will be uninstalled as a consequence of uninstalling the package that you selected. You can also see if any reboots are required.

8. On the Schedule and Notify page, you have the following options:

   - Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.

- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

If you chose to uninstall immediately, progress bars appear that indicate the status of the installation process. You can watch the progress bars to see when the software is uninstalled.

9 When the process is complete, click View Details to see what packages were installed or removed, and the output, if any, for the installation scripts.

Installing Templates

This section provides information on how to install templates within the Opsware System and contains the following topics:

- Installing Templates Overview
- Installing Templates with the Install Templates Wizard

Installing Templates Overview

Templates are predefined suites of software that you can install on a large number of servers running the same version of the same operating system in a single operation. Templates can include both applications and patches. Using templates, you can quickly deploy a suite of applications that together provide a particular service, such as a Web server and the applications that run on top of it that provide a Web service.

See “Working with Templates” on page 353 in Chapter 6 for information about how to create templates in the Opsware System.

You cannot uninstall the packages in a template as a group cannot though an Opsware wizard. You can, however, use the Opsware Uninstall Software Wizard to uninstall the individual packages that a template contains.

Installing Templates with the Install Templates Wizard

Perform the following steps to install a template with the Install Templates Wizard:

1 Launch the Install Template Wizard from the Software Provisioning pane on the Opsware Command Center home page. You are presented with an overview of the steps for installing a template.

2 Click Start to continue. The Install Templates Wizard page appears, as Figure 7-9 shows.

Figure 7-9: Install Templates Wizard Page

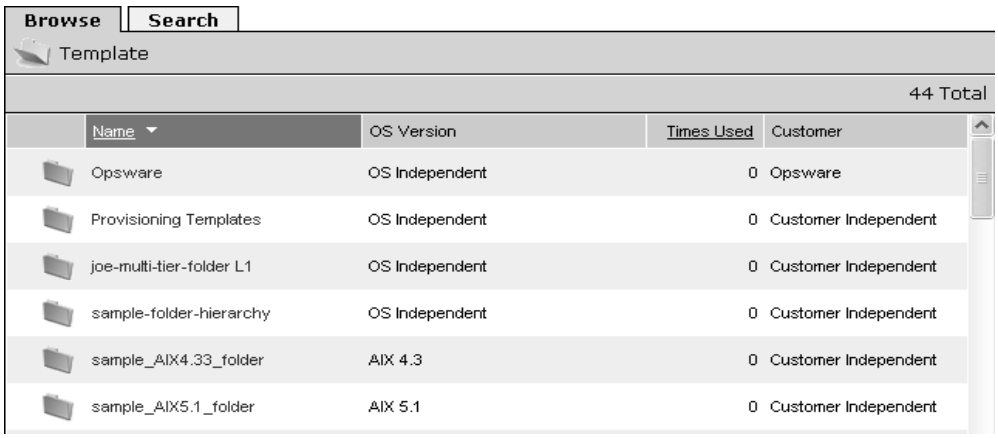

3 Select the template that you want to install by browsing or searching. Templates are organized by operating system type and operating system version and by customer. You can select only one template at a time to install. Select the template and click Next to continue.

The Select Servers page appears. The Select Servers page displays a list of servers limited to the operating system type and version and the customer type of the template that you selected.

4 Select the servers that you want to apply the template to by browsing or searching, and click Next to continue. The Confirm Selection page appears. The Confirm Selection page displays the servers and template that you selected to install.

5 Review your selection and click Preview to continue. The Preview page appears. The Preview page displays one progress bar per server that shows the status of the preview process.

6 When the preview process completes, click View Details to see the templates that will be installed. You can also see if any reboots are required.

7 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.

- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

If you choose to install immediately, progress bars appear that indicate the status of the installation process on each server that you selected.

8 When the process is complete, click View Details to see what templates were installed or removed, and the output, if any, from the installation scripts.

Chapter 8: Patch Management Subsystem

IN THIS CHAPTER

This chapter discusses how you can use the Patch Management Subsystem to automate patch management on your managed servers.

By using the Opsware Command Center, you can perform the following patch management tasks:

- Opsware Patch Management
- Patch Management Roles
- Setting Up for the Patch Management Subsystem
- Uploading Patches
- Patch Administration Using the Opsware Command Center
- Overview of Installing and Uninstalling Patches
- Overview of the Microsoft Patch Update Wizard

Opsware Patch Management

This section provides information about patch management within the Opsware System and contains the following topics:

- Patch Management Overview
- Support for Patch Testing and Installation Standardization
- Special Support for Windows Servers
- Summary of Features
- Supported Operating Systems and Supported Patch Types
- Supporting Technologies for Patch Management
- How the Opsware System Supports Patch Management

Patch Management Overview

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 server failures.

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.

Patch management is a fully integrated component of the Opsware System, and leverages the Opsware System's complement of server automation features. The Opsware System, for example, maintains a central database (called the Model Repository) that has detailed information about every server under management, the patches and software installed on the servers, and the patches and software available for installation. You can use this data to determine the severity of your exposure to a newly discovered threat, and to help you assess the benefits of rolling out a patch versus the costs in downtime and testing requirements.

By automating the patching procedure, the Patch Management Subsystem can reduce the amount of downtime required for patching. The Opsware System also allows you to schedule patch activity, so that patching occurs during off-peak hours.

After the patch is integrated into your environment, you can make it part of your standard builds with Opsware templates.

Support for Patch Testing and Installation Standardization

The Opsware System offers features to minimize the risk of rolling out patches. First, when a patch is uploaded into the Opsware System, its status is marked as *untested* and only administrators with special privileges can install it.

The patch administrator then defines patch installation and uninstallation options and tests the patch. Only after the patch is tested and the patch administrator marks it as *available* for use can other administrators install the patch.

The Patch Management Subsystem allows you to standardize the way that patches are installed and uninstalled, thereby preventing ad-hoc installation procedures. Patch administrators standardize patch installation by providing pre- and post-install scripts, install and uninstall flags, and instructions on when to reboot and how to handle error codes from the pre- and post-install scripts.

Special Support for Windows Servers

The Patch Management Subsystem offers an even higher degree of patch automation for Windows servers. The Opsware System takes advantage of the Microsoft Patch Database, which contains information about what patches are available and how the patches should be applied. The Opsware System compares all Windows servers to this database, which allows the patch administrator to determine which patches need to be applied. When the Microsoft Patch Database is updated, the new patches that have not yet been uploaded display in the Opsware Command Center, with links to the patches that allow the patch administrator to immediately download the patches.

Summary of Features

The Opsware System automates patch management by providing the following features:

- A central repository where patches are stored and organized in their native formats
- A database that includes information on every patch that has been applied
- Customized scripts that can be run before and after a patch is installed
- Microsoft servers are analyzed against a Microsoft database to determine which servers need which patches
- Advanced search abilities to identify servers that require patching
- Auditing abilities so that security personnel can track the deployment of important patches

Supported Operating Systems and Supported Patch Types

The Patch Management Subsystem supports all of the operating system versions that the Opsware System supports, except for Linux.

Linux does not support patches in the ordinary sense. The packages are not patchable. Instead, new versions of the RPM are delivered. Linux systems that the Opsware System manages are therefore not viewable through the Patch Management Subsystem interfaces. New Linux packages and updates should be managed and applied though the software channel.

Table 8-1 shows the operating system versions and the patch types that the Patch Management Subsystem supports.

Table 8-1: Supported Operating System Versions and Patch Types

OS VERSIONS	PATCH TYPES
AIX 4.3	AIX Update Fileset APARs
AIX 5.1	AIX Update Fileset APARs
AIX 5.3	AIX Update Fileset APARs
HP-UX 11.00	HP-UX Patch Filseset HP-UX Patch Products
Solaris 5.6	Solaris Patch Solaris Patch Cluster
Solaris 5.7	Solaris Patch Solaris Patch Cluster
Solaris 5.8	Solaris Patch Solaris Patch Cluster
Solaris 5.9	Solaris Patch Solaris Patch Cluster
Windows NT 4.0	Windows Hotfix Windows OS Service Pack

Table 8-1: Supported Operating System Versions and Patch Types

OS VERSIONS	PATCH TYPES
Windows 2000	Windows Hotfix Windows OS Service Pack
Windows 2003	Windows Hotfix Windows OS Service Pack

Special Requirement for Windows NT Systems

If you have any Windows NT 4.0 servers under management, you must install Windows NT 4.0 SP6a, with Internet Explorer 6.0 or later.

Internet Explorer 6.0 or later is required because it provides the ability to parse XML, and parsing XML is required in order to compare the server against the Microsoft Patch database.

For information about how to create a silent installable version of Internet Explorer 6.0 or later, see the *Opsware System Installation Guide*.

Supporting Technologies for Patch Management

The Patch Management Subsystem uses patching utilities and technologies for each supported operating system. The Opsware System uses these tools behind the scenes, which allows you to perform patch management though a single interface, without having to worry about invoking a number of different patching utilities.

The Opsware System models the way it treats patches on the way the underlying utility treats a patch. For example, if the Solaris patchadd utility is not able to install one patch contained in a patch cluster, the Solaris utility continues to install the remaining patches in the patch cluster. The Opsware System respects this behavior and allows that patch installation operation to continue. Any patches that are not installed are reported at the end of the installation operation.

Table 8-2 shows the patch management and installation tools that are used for each of the supported operating systems.

Table 8-2: Supporting Technologies for Patch Management

WINDOWS	SOLARIS	AIX	HU-UX
Qchain enables single reboot when installing more than one Hotfix	Patchadd installs Solaris patches	Installp installs and uninstalls filesets	Swlist lists patch products, files, products, and filesets
mbsacli lists and verifies installed Hotfixes and Service Packs	Patchrm uninstalls Solaris patches	Lslpp lists installed LPPs	Swinstall installs a depot
	Showrev lists installed Solaris patches	Instfix lists installed APARs	Swremove removes a depot
	Pkgadd installs Solaris packages		
	Pkginfo lists installed Solaris packages		

How the Opsware System Supports Patch Management

When a server is brought under management by the Opsware System, the Opsware Agent installed on the server registers the server's hardware and software configuration with the Opsware System. (The Opsware Agent repeats this registration every twenty-four hours.) This information, which includes data about the exact OS version, hardware type,

installed software and patches, is immediately recorded in the Model Repository. Also, when you first provision a server with the Opsware System, the same data is immediately recorded.

When a new patch is issued, you can use the Opsware Command Center to immediately identify which servers require patching. The Opsware System provides a Software Repository where you upload patches and other software. Users access this software from the Opsware Command Center to install patches on the appropriate servers.

After a server is brought under management, you should install all patches through the Patch Management Subsystem. If you install a patch manually, the Opsware System does not have data about that patch until the next software registration. If you install a patch manually, it can take as long as twenty-four hours until the data about that server in the Model Repository is up-to-date.

Whenever you install or uninstall software or patches with the Opsware System, however, the Opsware Agent immediately updates the information about the server in the Model Repository.

Patch Management Roles

This section provides information on patch management roles and contains the following topics:

- Patch Management Roles Overview
- About the Patch Administrator
- About the System Administrator
- Opsware System Permissions for Patch Management

Patch Management Roles Overview

The Opsware System provides support for rigorous change management by assigning the functions of patch management to two different types of administrators:

- The patch administrator (often referred to as the security administrator), who has the authority to upload and test, and edit patch options
- The system administrator, who applies the patches (that have been approved for use) uniformly and automatically according to the options that the patch administrator specifies

About the Patch Administrator

In most organizations, patch administrators are responsible for reviewing the latest security threats and the patches that vendors have released to address these problems. The patch administrators are generally experts in the operating systems and applications that they manage, and are able to assess the necessity of applying patches issued by vendors. They are able to diagnose common problems that arise after patches are installed, allowing them to thoroughly test the patch application process.

In the Opsware System, patch administrators are granted specific permissions that allow them to upload patches into the Opsware System, test the patches, and then mark them as *available* for use. Basic users can upload patches, but they cannot install them or mark them as available. They are also able to edit patch options (such as installation scripts) though patch management. Other types of users are not allowed to upload or edit patches.

Typically, the patch administrator uploads patches and then tests them on non-production reference hardware. After testing the patches and determining that the patches are safe to apply to production systems, they mark the patches available in the Opsware Command Center, and then advise the system administrators that they must apply the approved patches.

About the System Administrator

The Opsware users are system administrators who are responsible for the day-to-day maintenance of the servers in a deployment. These users are not required to have the same level of expertise in low-level system details as the patch administrator.

Because the patch administrator has set up the patch installation, the system administrators can apply the patches to a large number of servers with a few mouse clicks. They are responsible for searching for the servers that require the approved patch, running the Patch Installation Wizard, and verifying that the patches are installed successfully.

Opsware System Permissions for Patch Management

The patch administrator must be assigned to the advanced user role. The system administrator must be assigned to the basic user role. (For more information about user roles, see the *Opsware System 4.7 Administration Guide*.)

The patch administrator has the following privileges:

- Uploading patches through the patches channel or by using the Upload Patch Wizard

- Installing patches that are in *untested* as well as the *available* state
- Uninstalling patches
- Editing patch options though the Opsware Command Center
- Performing Microsoft patch updates
- Uploading the Microsoft Patch Database

The system administrator cannot use the Patches link on the Opsware Command Center home page navigation panel. (This link gives access to advanced features not available through the Patch Management Wizards.) The system administrator has the following privileges:

- Uploading patches through the Upload Patch Wizard (but cannot install a patch until the patch administrator has marked it as available)
- Installing patches that have are *available* for use
- Uninstalling patches
- Performing Microsoft Patch Updates

Setting Up for the Patch Management Subsystem

This section provides information about the set up for the patch management subsystem and contains the following topics:

- Setting Up for the Patch Management Subsystem Overview
- About the Microsoft Patch Database
- If a user performs a patch analysis of a Windows server immediately after uploading a new version of the Microsoft database, the analysis does not yet include the data from the new Patch Database. Instead, the Opsware System reports the data from the last time that the Opsware Agent recorded the results of its comparison.
- Uploading the Microsoft Patch Database
- Products Tracked in the Microsoft Patch Database
- Selecting Which Microsoft Products to Track

Setting Up for the Patch Management Subsystem Overview

Before you upload any patches, you should first upload the latest version of the mssecure.xml Microsoft Patch Database file if you have any Windows servers in your facility. The Patch Management Subsystem also relies on two additional utilities, Qchain.exe and mbsacli.exe. These two utilities are uploaded when the Opsware System is first installed. For additional requirements for Windows Servers to support patch management, see the *Opsware System 4.7 Installation Guide*.

About the Microsoft Patch Database

Once every twenty-four hours, the Opsware Agent on a Windows server compares the server's current state against the Microsoft Patch Database that has been uploaded into the Opsware System by the patch administrator. The Opsware Agent reports the results of that comparison, and the data is stored in the Model Repository. When a user requests an analysis of a Windows server (for example, by using the Microsoft Patch Update Wizard), the data is retrieved from the Model Repository and displayed in the Opsware Command Center. By storing the data in the Model Repository, rather than performing an actual comparison on the server itself when a user requests an analysis, the data can be quickly retrieved and displayed.

If a user performs a patch analysis of a Windows server immediately after uploading a new version of the Microsoft database, the analysis does not yet include the data from the new Patch Database. Instead, the Opsware System reports the data from the last time that the Opsware Agent recorded the results of its comparison.

About Uploading the Microsoft Patch Database

To upload the Microsoft Patch Database, you have two options. If your Opsware Command Center server has access to the Internet, you can specify the URL of the Microsoft Patch Database.

You must still re-upload from that URL when a new version of the database is released. The Microsoft Patch Database is generally updated once a month and you must re-upload the new version on a monthly basis.

If your Opsware Command Center is isolated from the Internet, you must periodically download the Microsoft XML database to a location on your network accessible to the Opsware Command Center and then upload it.

You can either upload the Microsoft Patch Database as a CAB archive, which contains an XML file, or you can upload the XML file directly. You can upload the patch database either with the Opsware Command Center or the Opsware Command Line Interface OCLI). For instructions on how to use the OCLI, see the “Opsware Command Line Interface” on page 597.

Uploading the Microsoft Patch Database

Perform the following steps to upload the Microsoft Patch Database:

1. From the navigation panel in the Opsware Command Center, click Software ➤ Patches. The Patches page appears.
2. Click the Patch Preferences tab. See Figure 8-1.

Figure 8-1: Patch Preferences Tab

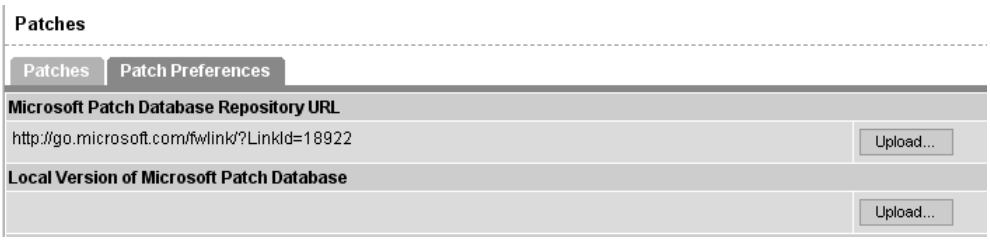

3. If you want to upload the mssecure.cab file from the Internet, click the Upload button under the Microsoft Patch Database Repository URL. The Upload URL page appears.
4. Specify the URL of the Microsoft Patch Database Repository URL and then click Upload.
5. If you do not want to upload the Microsoft Patch Database from the Internet, first copy the file to a location accessible to your Opsware Command Center server.
6. Specify the fully qualified path of the Microsoft Patch Database or click the Browse button and navigate to the database.
7. Click the Upload button. See Figure 8-2.

Figure 8-2: Upload Local Version of the Microsoft Patch Database

Products Tracked in the Microsoft Patch Database

The Microsoft Patch Database contains information about a wide range of Microsoft products. After you upload the database, you must select the products that you want to track. The Opsware System tracks the data for the products that you select, and ignores information about the products that you do not use. If you do select products that you do not actually use, the patches for those products show up in the Opsware Command Center.

Selecting Which Microsoft Products to Track

Perform the following steps to select which Microsoft products to track:

1. From the navigation panel in the Opsware Command Center, click Patches ➤ Patch Preferences.
2. Under the Patch Options, click the Select button. The Patch Options page appears, as Figure 8-3 shows.

Figure 8-3: Microsoft Patch Options Page

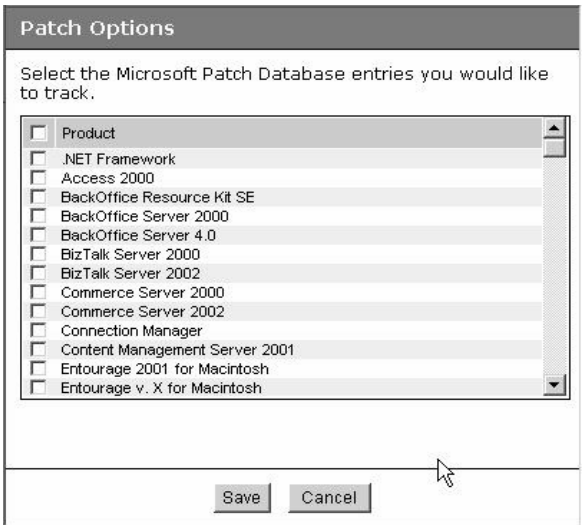

3. Click the check boxes for all of the Microsoft products whose Hotfixes you want to track.
4. Click Save. The next time the patch database is uploaded, Hotfixes that affect the set of selected products are modelled in the Opsware System.

Uploading Patches

This section provides information about how to upload patches with the Opsware System and contains the following topics:

- Uploading Patches Overview
- Uploading Patches with the Opsware Command Center or the Opsware CLI
- About Windows Patches
- About AIX Patches
- About Solaris Patches
- About HP-UX Patches
- Installation Scripts Overview
- Installation and Uninstallation Flags Overview
- Preparing to Upload Patches
- Uploading a Patch with the Upload Patch Wizard
- About Testing Patches

Uploading Patches Overview

When a patch is uploaded, you associate the patch with a specific version of an operating system. When you upload a Solaris patch, for example, you must select the version of the Solaris operating system that this patch applies to, such as Solaris 5.6 or 5.9. You can only install this patch on servers that are running that version of the operating system.

If, for any reason, you need to install a given patch across servers running different versions of the same operating system, you need to upload the patch multiple times and associate the patch with each of the operating system versions that the patch applies to.

For example, if the same Solaris patch needs to be installed on servers running Solaris 2.7 and 2.8, you must upload the patch two times. The first time that you upload the patch, you associate it with the Solaris 2.7. You then repeat the procedure and associate the patch with Solaris 2.8. (This procedure also allows you to specify different installation options. The different versions of the same operating system can sometimes require different installation scripts, installation flags, and so forth.)

In the case of application patches, it is even more common that you need to upload a patch multiple times. A Solaris patch for Oracle, for example, often needs to be applied to instances of Oracle running on slightly different versions of the Solaris operating system.

Uploading Patches with the Opsware Command Center or the Opsware CLI

If you upload a patch though the Opsware Command Center, the Upload Patch Wizard guides you through the process. The Upload Patch Wizard allows you to specify a number of options for the patch, including install and post-install scripts, install and uninstall flags, and other options.

Because the Opsware Command Center is a browser-based interface, you can only upload one patch or patch container (such as a Solaris patch cluster or an HP-UX depot) at a time. If you want to upload multiple patches at the same time, such as a large set of AIX LPPs, you can do so more quickly through the OCLI.

If you upload patches through the OCLI, however, you are not able to specify installation options during the upload process. Instead, you specify these options by editing the patches through the Opsware Command Center.

About Windows Patches

You can download most Windows patches from Microsoft directly though the Opsware Command Center. All of the Windows patches that affect any of the Microsoft products Opsware System is configured to track appear in the Patch Management pages. If the patch is new and has not yet been uploaded, a link to the patch on the Microsoft support site is provided in the Opsware Command Center that you can use to immediately download the patch.

After you download the Windows patch, however, you must still upload the patch through the Upload Patch Wizard or the OCLI.

If a Windows patch is already uploaded and installed on a server, for example, as part of an Opsware template, you receive a warning that the patch is already uploaded:

This patch already exists within the Opsware system and is utilized by node(s). The following entries are the optional settings for patch installation.

As with all patch types, you must be careful to specify the correct type of patch you are uploading, such as a Windows Hotfix or Windows OS Service Pack. Misidentifying Windows patches can cause additional problems because the information you provide about the patch conflicts with the data in the Microsoft Patch database.

The Patch Management Subsystem passes in install or uninstall flags that cause Windows patches to install and uninstall in *silent* mode, meaning that no dialog boxes should appear on the server while installing or uninstalling a patch. Some Windows patches, however, still produce modal dialog boxes when they are being installed or uninstalled, even when directed not to do so. The Opsware System has a special utility that automatically closes these modal dialog boxes.

About AIX Patches

AIX periodically releases Authorized Program Analysis Reports (APARs), which specify what update filesets (contained in LPPs) are necessary to fix an identified problem. An APAR only specifies the minimum version of an update fileset required to fix a problem; an APAR can therefore be satisfied with later versions of the same filesets. To maintain compatibility, however, the Opsware System always adopts the fileset with the lowest version number that meets the minimum version that APAR specifies. If a later version of the update fileset is uploaded, the Opsware System still associates the earlier version of the fileset with the APAR.

When an LPP is uploaded into the Patch Management Subsystem, Opsware recognizes which APARs the filesets contained in the LPP belong to. An entry is created for the APAR in the Patch Management Subsystem when the first fileset associated with an APAR is uploaded. (In some cases, a fileset is associated with more than one APAR. An entry is created for each APAR the fileset is associated with, if the entry does not already exist.)

If you want to be able to install all LPPs that APAR specifies, you must make certain to upload all of the specified LPPs into the Patch Management Subsystem, either through the Upload Patch Wizard or through the OCLI.

If you do not upload all of the LPPs that APAR specifies, it is still possible for the system administrator to browse for an APAR and install the partial set of LPPs that are uploaded. In such cases, the administrator receives a warning that the filesets for the APAR are not all installed.

The Patch Administrator must first upload and test an LPP before it is generally available in the Opsware System. The new fileset is integrated into the APAR only after the LPP is tested and approved. Even though the APAR is updated automatically, you still maintain control over the exact filesets that are allowed to be installed on your managed servers.

APAR update filesets cannot be installed on a server if the server does not already have the base filesets for which the update filesets are intended.

If, however, a server has a partial set of the base filesets, the APAR can be applied and only the applicable filesets for the base filesets are installed. For example, if an APAR specifies four update filesets to update four base filesets, and you attempt to apply the APAR to a server that has only three of the base filesets, three of the four update filesets from the APAR are installed.

About Solaris Patches

A Solaris patch cluster contains a set of selected patches for a specific Solaris release level. Ordinarily, after a patch cluster is installed, it is not possible to search for a particular patch cluster. The patches do not contain any metadata that relate them to the patch cluster in which they were originally bundled. You can only search for the individual patches.

If you install a Solaris patch cluster though the Patch Management Subsystem, however, the Opsware System keeps track of the patch cluster in the Model Repository. You can therefore search for a patch cluster to determine if a full patch cluster is installed. You can also uninstall the patch cluster if you installed it with the Patch Management Subsystem.

About HP-UX Patches

HP-UX patches are delivered exclusively as depots, which are patch products that contain patch filesets. The depot is uploaded directly into the Patch Management Subsystem.

If a depot is already uploaded and attached to a node, it cannot be uploaded through the Patch Management Subsystem. If you want to upload the depot into the Patch Management Subsystem, you must detach a depot from any nodes that it is attached to, and then delete it from the Software Repository.

For HP-UX 10.20, you can only apply patches through the Install Software Wizard because the Opsware System recognizes them as software and not patches.

Installation Scripts Overview

When you upload a patch, you can specify the following types of scripts:

- Pre-installation scripts that are executed before a patch is installed

- Post-installation scripts that are executed after a patch is installed
- Pre-uninstallation scripts that are executed before a patch is uninstalled
- Post-uninstallation scripts that are executed after a patch is uninstalled

A typical use of a pre-installation script is to shut down a process before you apply a patch to the application that the process belongs to. The post-installation script then restarts the process after the patch is applied.

A user can execute any kind of script that the operating system supports of the server where the patch is to be installed. You must make certain, however, that the proper shells, binaries, and so forth, are installed on the servers where you plan to run the scripts. For example, you can specify Python scripts to be run when the patch is installed, but you must make certain that Python is installed on the servers where you want to run the scripts. You must also call Python yourself. The Opsware System does not call Python on your behalf.

Installation and Uninstallation Flags Overview

You can specify installation and uninstallation flags that are applied whenever a patch is installed or uninstalled.

The Opsware System, however, also uses default installation and uninstallation flags. The Opsware System requires that patches are installed and uninstalled with these flags. You must therefore be certain that you do not specify any installation or uninstallation flags that override or contradict the default flags passed in by the Opsware System.

Some Windows Hotfixes do not support the -z flag, some do not support the -q flag, and some do not support either. In such cases, you must use a special expression: /-z or /-q or /-z -q respectively, to prevent the Patch Management Subsystem from passing in the -z or -q or -z -q flag.

Default Installation and Uninstallation Flags

Table 8-3 lists the default installation flags that the Opsware System uses.

Table 8-3: Default Installation Flags

OPERATING SYSTEM/PATCH TYPES	FLAGS
Windows Hotfix	`-q -z`

Table 8-3: Default Installation Flags

OPERATING SYSTEM/PATCH TYPES	FLAGS
Windows Security Rollup Package (treated identically to a Hotfix by the Patch Management Subsystem)	`-q -z`
Windows OS Service Pack	`-u -n -o -q -z`
AIX	`-a -Q -g -X -w`
HP-UX	None

Table 8-4 lists the default uninstallation flags that the Opsware System uses.

Table 8-4: Default Uninstallation Flags

OPERATING SYSTEM/PATCH TYPES	FLAGS
Windows Hotfix	`-q -z`
Security Rollup Package	`-q -z`
Windows OS Service Pack	Not uninstallable
AIX	`-u -g -X`
AIX Reject Options	`-r -g -X`
HP-UX	None

Preparing to Upload Patches

Before you upload a patch, you must copy it to a location that is accessible to the browser that you are using or the OCLI. If you are using the Opsware Command Center, you specify the path of the patch in the upload wizard, either by entering it directly or by browsing for the patch.

In some cases, you need to install patches in a particular order. You can create an installation order dependency by using the Opsware Command Center.

Uploading a Patch with the Upload Patch Wizard

Perform the following steps to upload a patch by using the Upload Patch Wizard:

1. Launch the Upload Patch Wizard from the Patch Management pane on the Opsware home page. You are presented with an overview of the upload process in the first page. Click Start to continue. The Select Patch page appears, as Figure 8-4 shows.

Figure 8-4: Select Patch Page

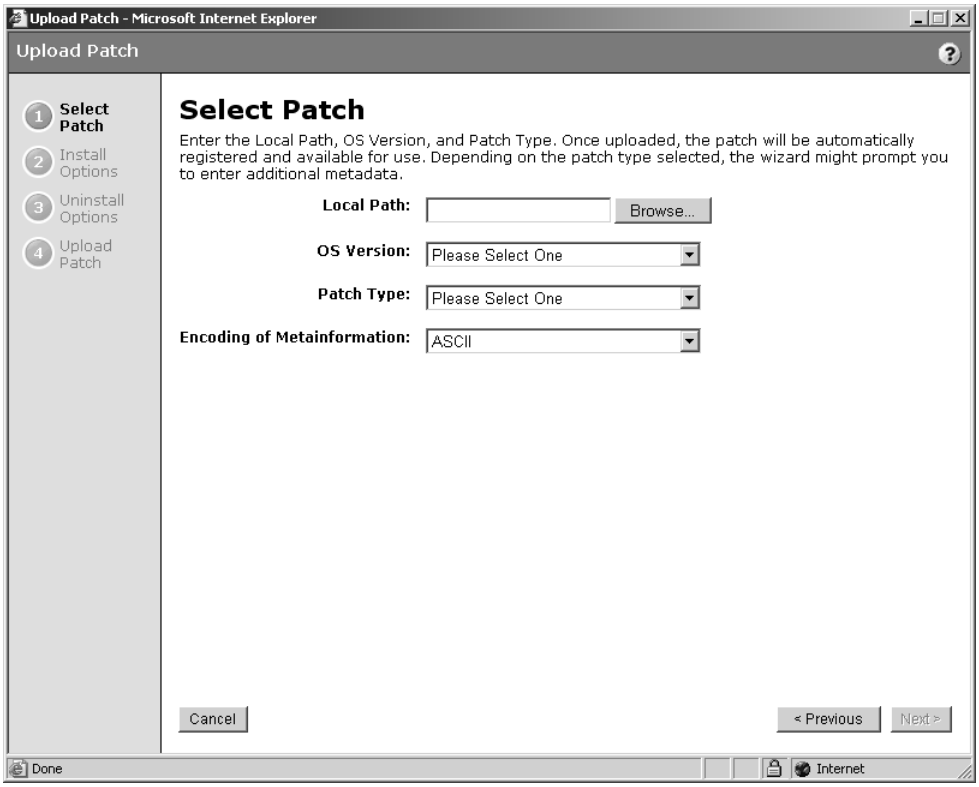

2. Either enter the fully qualified path of the patch that you want to upload, or click Browse and navigate to the patch that you want to upload.

3. Select the OS version of the patch that want to upload. You must be certain to select the correct operating version, or the patch will not be available for the correct operating system.

4. Select the type of patch that you are uploading. You must be careful to select the correct patch type or the patch will be misapplied, or uninstallable. The Opsware Command Center only allows you to select patches that are appropriate for the operating system that you select, but it is still possible to select the wrong kind of patch. (For example, selecting a Solaris patch when you intended to select a Solaris patch cluster.)

5 Click Next to continue to the Install Options page. In this page, you can specify a number of installation options:

- Install Flags passed directly to the patch installer. The Patch Management Subsystem also passes a number of default flags.
  - If you are installing a Windows Hotfix that does not support the -z flag, remember to use the /-z option here to prevent the Patch Management Subsystem from passing in the -z flag.
  - When installing an AIX Update fileset, the Patch Management Subsystem normally applies the fileset, which allows it to be rejected (uninstalled.) If you want to commit the fileset instead (so that it cannot be removed), use the -c option here.
- Pre-install Script. Enter the pre-install script into this box.
- If you want to terminate the installation of the patch if the script returns a non-zero return code, select the check box.
- Post-Install Script: Enter the post-install script into this box. If you want to terminate the installation of the patch if the script returns a non-zero return code, select the check box.
- Select the Reboot on install option if the patch you are removing requires a reboot. Keep in mind that other patches can be directly applied after this patch, so be sure to check this option if it is necessary.

6 Click Next to continue to the Uninstall Options page.

In this page, you can specify the following uninstallation options:

- Uninstall Flags passed directly to the installer. The Opsware Patch Management System passes a number of default uninstall flags to the installer.
- Pre-uninstall Script. Enter the Pre-uninstall script into this box.
- If you want to terminate the installation of the patch if the script returns a non-zero return code, select the check box.
- Post-uninstall Script: Enter the post-uninstall script into this box. If you want to terminate the installation of the patch if the script returns a non-zero return code, select the check box.
- Select the Reboot on Uninstall option if the patch that you are removing requires an immediate reboot. Keep in mind that other patches can be directly applied after this patch, so be sure to check this option if it is necessary.

7 Click Next to upload the patch. A progress bar appears.

8 After the patch is uploaded, you have the option to install the patch, as Figure 8-5 shows. Click Yes to install the patch, and then click Next. Otherwise, click No or click Close.

Figure 8-5: Upload Successful Message

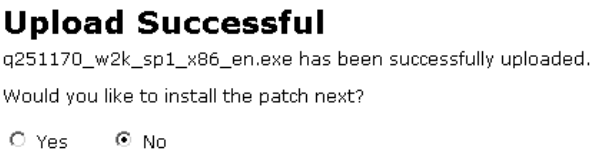

Remember that if you need to upload the same patch for multiple versions of the same operating system, you must repeat this process with the same patch.

About Testing Patches

After you upload the patch, you can install and test it using the Patch Install Wizard. As the patch administrator, you can install the patch, even though the patch is automatically set in the Untested state after you upload it the first time. When you finish testing the patch, use the Opsware Command Center to change the patch state to available so that system administrators can install the patch.

Patch Administration Using the Opsware Command Center

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

- Patch Administration Overview
- Patch Statuses Overview
- Editing Patch Options Overview
- About Patch Installation Order Dependencies
- Creating Patch Installation Order Dependencies

Patch Administration Overview

The Opsware Command Center allows you to search though all patches that have been uploaded. In addition, it lists patches from the Microsoft database that have not yet been uploaded. You can use the Opsware Command Center to edit patch options and create install order dependencies, and change the state of patches to *Available for Use* so that system administrators can install them. You can also view detailed information about individual patches, such as the number of times the patch has been installed.

Patch Statuses Overview

The patch administrator sets the status of a patch in the Opsware System. The status determines who can apply the patch, or if the patch can be applied at all. (One additional state, *Not Yet Uploaded*, is set automatically and applies only to Windows patches).

Table 8-5 describes the statuses that patches in the Opsware System can have.

Table 8-5: Patch Statuses in the Opsware System

STATUS	DESCRIPTION
Untested	Initial state of a patch after being uploaded. Only administrators with special privileges can install untested patches.
Available for Use	Has been uploaded and approved by the patch administrator and can be installed on servers.
Not yet uploaded (Windows only)	A patch is described in the Microsoft database for one of the products that you have selected to track. This patch, however, has not yet been uploaded and cannot be installed. (This status is set automatically.)
Deprecated	The patch is possibly still installed on some systems, but cannot be installed anymore, not even by a user who is a member of the Advanced User role.

Setting Patch Status

Perform the following steps to set the patch status:

1. From the navigation panel of the Opsware Command Center, click Software ➤ Patches. The Patches page appears.

2 Select the filter options from the drop-down menus to display the type of patch whose status you want to change. You select the operating system version, the patch type, and the patch state. For example, if you want to change the state of a Windows 2000 Hotfix from Untested to Available for Use, you would select Windows 2000, Hotfix, and Untested, as Figure 8-6 shows.

Figure 8-6: Search Example

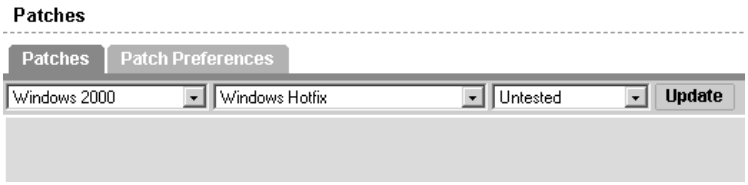

3 Click Update to display the list of patches that meet your selection criteria.

4 Locate the patch and click the name of the patch. The View Patch page appears.

5 In the patch summary section of the View Patch page, click Edit. The Edit Patch page appears.

6 Select the desired status from the Patch Status drop-down menu and then click Save.

Editing Patch Options Overview

You can edit any of the options that you specified for a patch that you uploaded using the Patch Upload Wizard. Additionally, if you uploaded a patch with the OCLI, you can specify the same options for the patch by editing the patch options.

Some patch option are not editable, due to the nature of the patch type. For example, you cannot change the Reboot on Install option of a Windows Service Pack from yes to no, because it is set to yes automatically. You cannot set patch status on an HP-UX patch fileset because you can only set the patch status on the parent HP-UX patch product. (After you change the status of the parent HP-UX patch product, the change is applied to the children filesets.) Other options cannot be set because they do not apply to the patch type that you are editing.

Editing Patch Options

Perform the following steps to edit patch options:

1. From the navigation panel of the Opsware Command Center, click Software ➤ Patches. The Patches page appears, as Figure 8-7 shows.

Figure 8-7: Patches Page

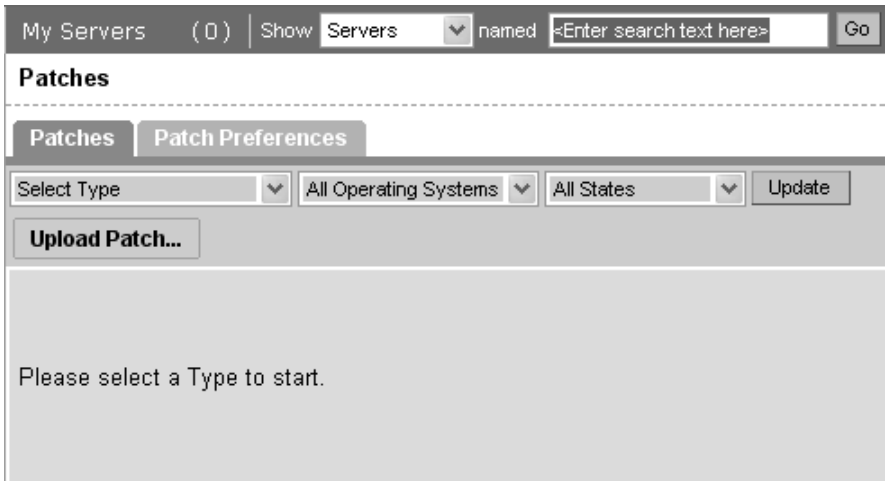

2. Select the options from the drop-down menus to display the type of patch that you want to edit. You select the operating system version, the patch type, and the patch state.
3. Click Update to display a list of patches that match your selected criteria.
4. Locate the patch that you want to edit and click the link for the patch name. The View Patch page appears.
5. Click the Edit button to edit the patch options. (Click the Edit button in the Install Options or in the Uninstall Options, as appropriate.)
6. Add or modify the patch install or uninstall options and click save.

If you are modifying the options of a patch that you already marked as Available, consider resetting the status of the patch back to Untested. Test the patch again with the new options, and set the status back to Available when you determine that it is safe to install the patch again.

About Patch Installation Order Dependencies

For some patch types, install order dependencies can be set. You create installation order dependencies through the Opsware Command Center.

To add patch dependencies to a patch, you must first upload the patch using the Patch Upload Wizard. You must then edit the patch through the Patches Channel.

Creating Patch Installation Order Dependencies

Perform the following steps to create patch installation order dependencies:

1. From the navigation panel of the Opsware Command Center, click Software ➤ Patches. The Patches page appears.
2. Select the options from the drop-down menus to display the type of patch that you want to edit. You select the operating system version, the patch type, and the patch state.
3. Click Update to display a list of patches that match your selected criteria.
4. Locate the patch that you want to edit and click the link for the patch name. The Patch Summary page appears.
5. Click the Edit button in the Install Order section. See Figure 8-8.

Figure 8-8: Install Order Section

6 Click the Add button to select the type of software that must be installed before the selected path. See Figure 8-9.

Figure 8-9: Software Type

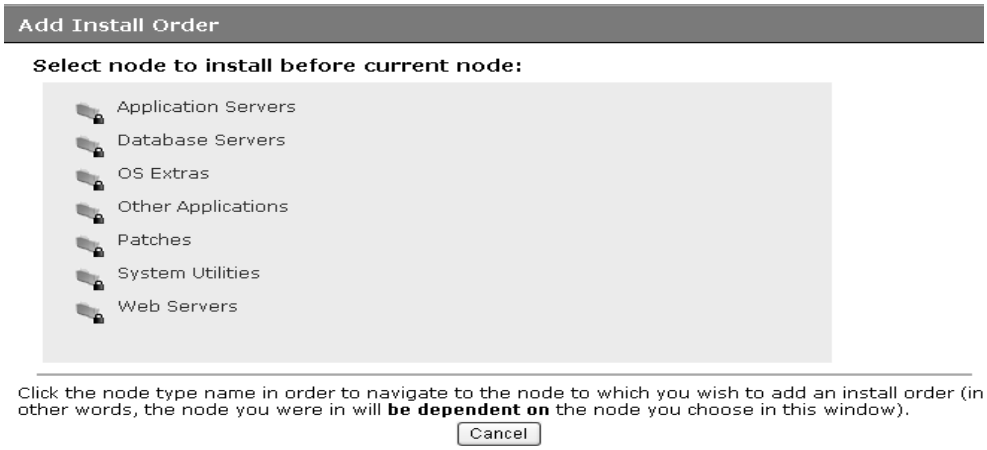

7 Browse for the software package or patch that must be installed before your selected patch.

8 Click the check box next to the desired software package and then click Add. See Figure 8-10.

Figure 8-10: Adding Install Order Dependency

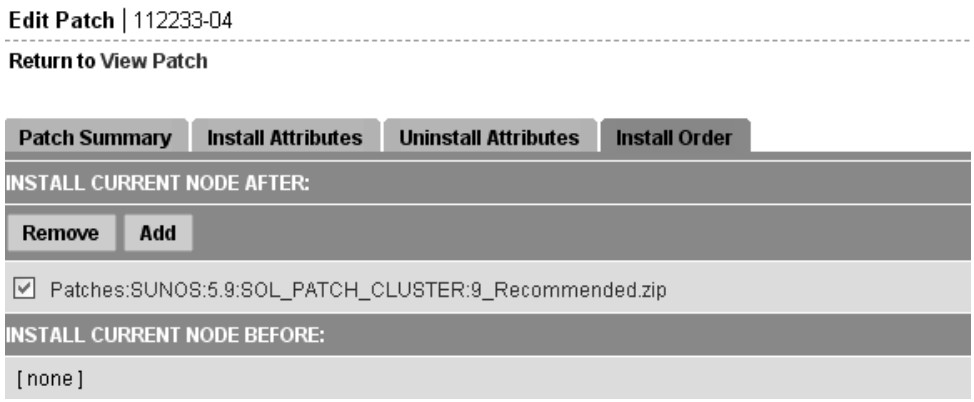

9 Confirm the dependency by clicking Add in the View Patch page. If you click Add again, the confirmation page does not appear. You will instead see the Add Install Order Dependency page, which allows you to add more packages.

10 Repeat the process if other dependencies must be expressed.

Overview of Installing and Uninstalling Patches

This section provides information about how to install and uninstall patches within the Opsware System and contains the following topics:

- Installing and Uninstalling Patches Overview
- About Installing and Uninstalling Application Patches
- About Patching Windows Servers
- Installing OS Patches with the Install Patch Wizard
- Installing Application Patches
- Uninstalling OS Patches with the Uninstall Patch Wizard
- Uninstalling Application Patches with the Uninstall Patch Wizard

Installing and Uninstalling Patches Overview

Installing and uninstalling patches are the primary patch management responsibilities of the systems administrator.

Typically, after the patch administrator has tested and approved a patch, the patch administrator notifies the system administrator that the patch is ready to be installed. The system administrator installs the patches with the Install Patch Wizard.

Patching and uninstalling patches frequently causes disruptions in service. Often, installing or uninstalling a patch will cause a server to reboot. In order to minimize the effects of service disruption, you can schedule the installation of patches.

The Patch Management Subsystem allows you to install both operating system patches and application patches. The procedure for installing application patches is slightly different than the process for installing operating system patches, because you must perform an additional search to find the servers where the application is installed.

The Install and Uninstall Patch Wizards allow you to select patches and servers by either browsing or searching.

About Installing and Uninstalling Application Patches

The Patch Management Subsystem does not allow you to apply a patch to an operating system for which the patch is not intended. When you are installing an application patch, however, the Patch Management Subsystem does not automatically filter out servers that

do not have the application installed that the patch is intended for. Though the Patch Management Subsystem does not prevent you from doing so, you must not attempt to apply application patches to servers that do not have the necessary application installed.

If an application patch is intended for an application that is running on more than one version of the same operating system, you cannot apply the patch to all of the servers at the same time. An application patch is associated with only one operating system version. You must first select the patch for one operating system, select the servers where the application is installed, and apply the patch. You must repeat this process for each version of the operating system where the application is installed.

Similarly, when uninstalling application patches that are installed on multiple versions of the same operating system, you cannot uninstall all of the patches at the same time. You must repeat the uninstallation process for each version of the operating system where the patch is installed.

About Patching Windows Servers

Every 24 hours, the Opsware Agent running on a managed server reports a manifest of data regarding the current state of software on the server; this data is stored in the Model Repository. The manifest includes information about all installed operating system software, application software and patches.

The Opsware Agent reports Windows software installed by MSI. It does not report all Windows software installed on the server (which is reported in the Add/Remove Programs list).

For Windows servers, the manifest also includes a list of patches that are recommend for installation on the server. This list is unique to the server and is generated based on the data contained in the Microsoft Patch Database. You can view the set of recommended patches for the server by using the Microsoft Patch Updates Wizard.

Installing OS Patches with the Install Patch Wizard

Perform the following steps to install operating system patches with the Install Patch Wizard:

1. From the Opsware Command Center homepage, click the link for the Install Patch Wizard. The Overview page appears. Click Start. The Select Patches page appears, as Figure 8-11 shows.

Figure 8-11: Select Patches Page

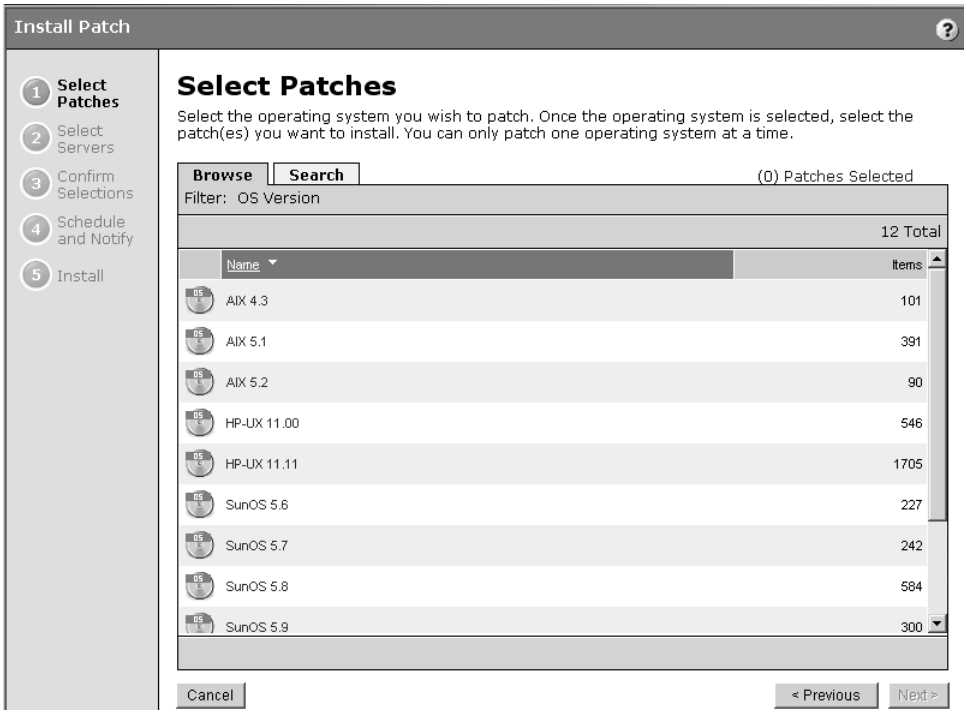

2. Select the operating system version for the patch that you want to apply. A list of all uploaded patches for that operating system appears, unless you selected a Windows operating system. The patch list for Windows servers can include patches that are in the Microsoft database but have not yet been uploaded.

3. Select one or more patches that you want to install and then click Next.

4 Select the servers that you want to patch and then click Next. The Confirm Selection page appears, as Figure 8-12 shows.

Figure 8-12: Confirm Selection Page

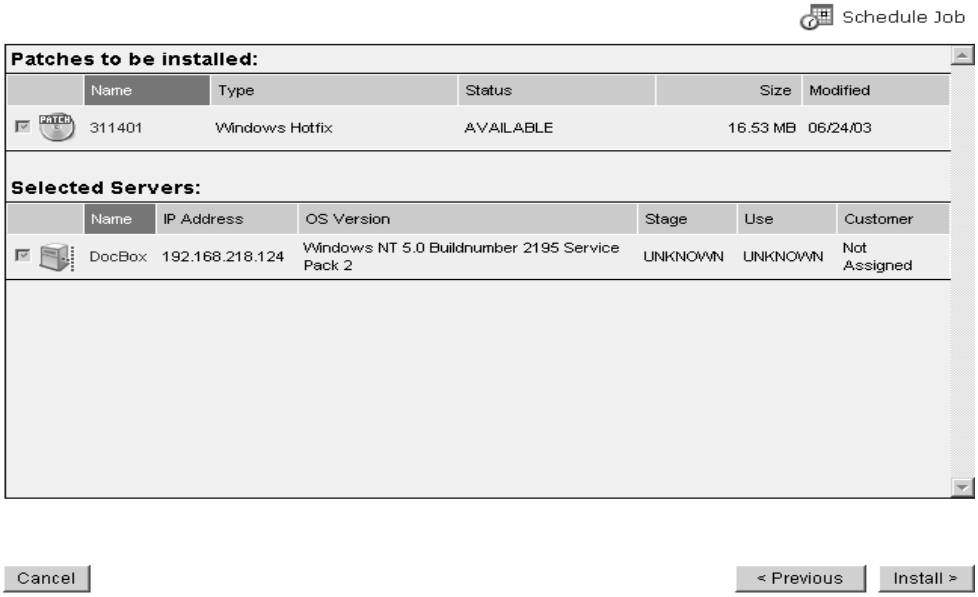

5 Review your selections.

6 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

If you decided to install the patches immediately by clicking Install, a progress bar appears, as Figure 8-13 shows.

Figure 8-13: Install Progress

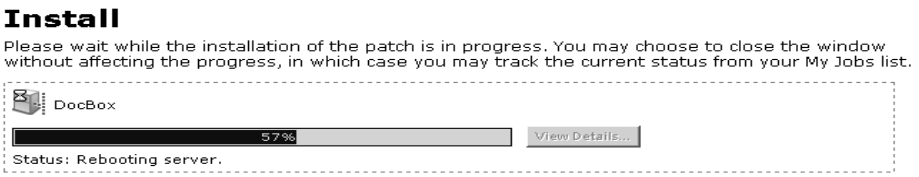

7 After the installation completes, click View Details for more information about the results of the installation operation.

Installing Application Patches

Perform the following steps to install application patches:

1 From the Opsware Command Center Patch Management pane, click the link for the Install Patch Wizard. The Overview page appears. Click Start. The Select Patch page appears.

2 Select the operating system version of the servers where the application is installed. A list of all uploaded patches for that operating system appears.

3 Select one or more application patches (for a given application) that you want to install and then click Next. The Select Servers page appears, as Figure 8-14 shows.

Figure 8-14: Select Servers

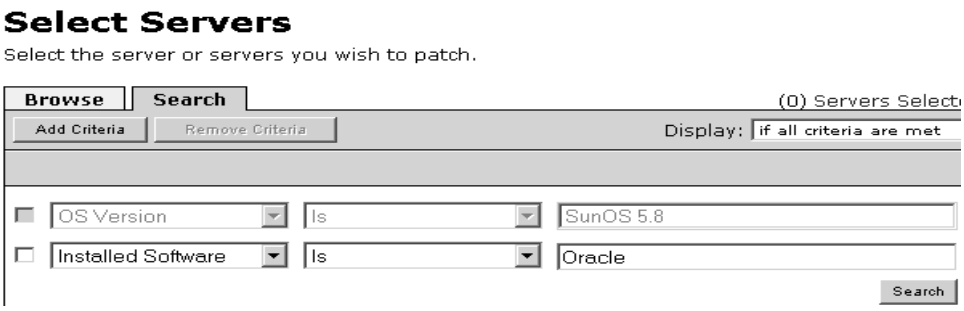

4 Click the Search tab and search for the servers where the application that you want to patch is installed and then click the Search button.

5 From the search results, select the servers running the application that you want to patch and then click Next. The Confirmation page appears, as Figure 8-15 shows.

Figure 8-15: Confirmation Page

6 Review your selections.

7 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

If you decided to install the patches immediately by clicking Install, a progress bar appears.

8 After the installation completes, click View Details for more information about the results of the installation operation.

Uninstalling OS Patches with the Uninstall Patch Wizard

Perform the following steps to uninstall operating system patches with the Uninstall Patch Wizard:

1. From the Opsware Command Center Patch Management pane, click the link for the Uninstall Patch Wizard. The Overview page appears. Click Start. The Select Patches page appears, as Figure 8-16 shows.

Figure 8-16: Select Patches Page

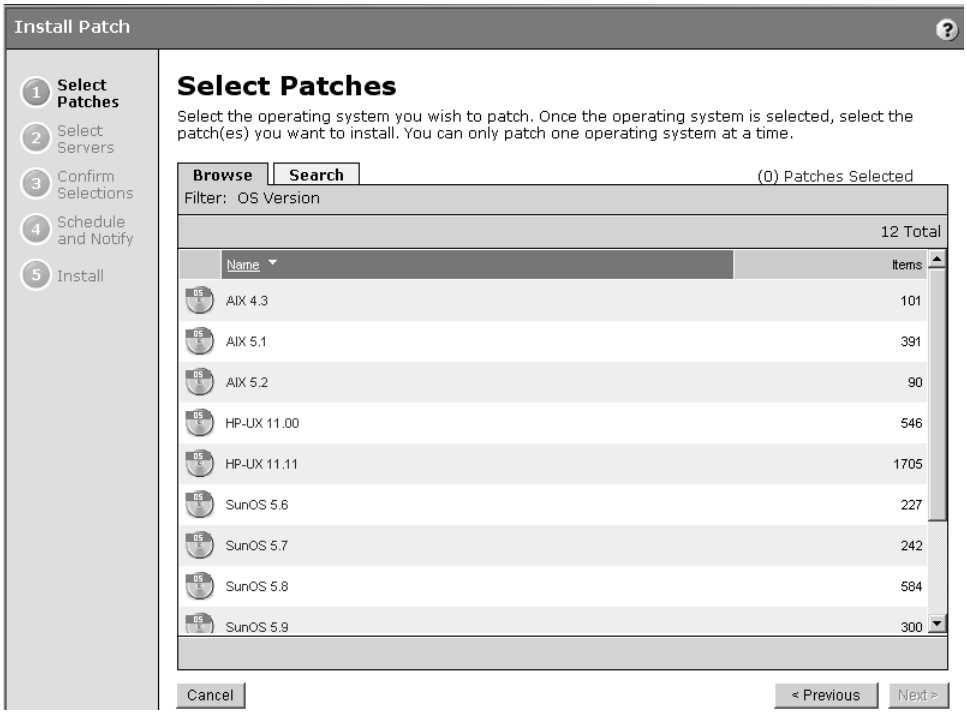

2. Select the operating system version for the patch to uninstall. A list of all installed patches for that operating system appears. You can also search for the patch. Some patches might display that you cannot select. These are patches that the Opsware System did not install.
3. Select the patch that you want to uninstall and then click Next. (You can only select one patch to uninstall.) A list of servers that have the selected patch installed appears.
4. Select the servers that you want to uninstall the patches from and then click Next. The Confirm Selection page appears.

5 Review your selections.

6 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

If you decide to uninstall the patches immediately by clicking Uninstall, a progress bar appears.

7 After the uninstallation completes, click View Details for more information about the results of the uninstallation operation.

Uninstalling Application Patches with the Uninstall Patch Wizard

Perform the following steps to uninstall application patches with the Uninstall Patch Wizard:

1 From the Opsware Command Center Patch Management pane, click the link for the Uninstall Patch Wizard. The Select Patch page appears.

2 Select the operating system version of the servers where the application is installed. A list of all uploaded patches for that operating system appears.

3 Select the application patches that you want to uninstall and then click Next. The Select Server page appears, as Figure 8-17 shows.

Figure 8-17: Select Servers Page

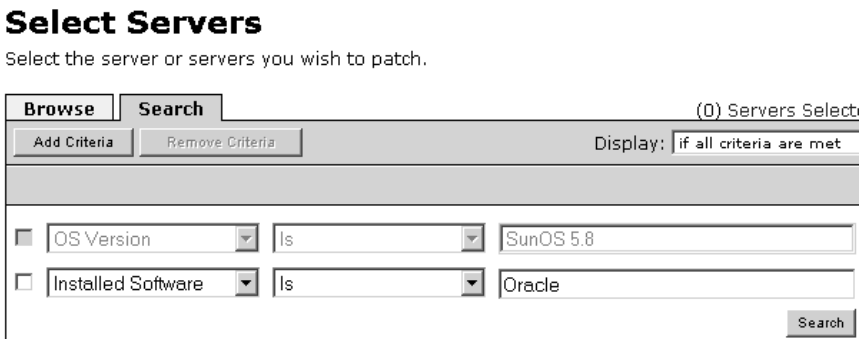

4. Click the search tab and search for the servers that have the patch that you want to remove and then click Search.
5. From the search results, select the servers whose application patch you want to remove and then click next. The Confirmation page appears.
6. Review your selections.
7. On the Schedule and Notify page, you have the following options:
   - Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
   - Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

   If you decide to uninstall the patches immediately by clicking Uninstall, a progress bar appears.
8. After the uninstallation completes, click View Details for more information about the results of the uninstallation operation.

Overview of the Microsoft Patch Update Wizard

The Microsoft Patch Update Wizard allows you to compare your Windows servers against the recommended list of patches contained in the Microsoft Patch Database that has been uploaded into the Opsware System.

If a new version of the Microsoft Patch database has recently been uploaded, a delay occurs before the servers are compared against the new Patch Database.

Using the Microsoft Patch Update Wizard

Perform the following steps to use the Microsoft Patch Update Wizard:

1. From the Opsware Command Center home page, click Microsoft Patch Update.
2. Select the customer whose servers you want to update or select All Customers to analyze all the Windows servers in your operational environment.

3. Select the version of the Windows operating system for the servers that you want to update. You can select Windows 2000, Windows 2003, NT 4.0, or you can select any operating system, which provides updates for all supported Windows operating systems.

4. Click Update to display the servers that match the customer and Windows OS version that you select.

5. Select the Windows servers that you want to compare against the Microsoft Patch Database. You can select multiple servers. Click Next to continue.

6. The recommended patches display for the servers that you select. Not all of the recommended patches are always available. Figure 8-18 shows that some of the patches have not been uploaded (as indicated in the status column), and others are unavailable because the patch administrator has not yet approved them.

Figure 8-18: Select Patches

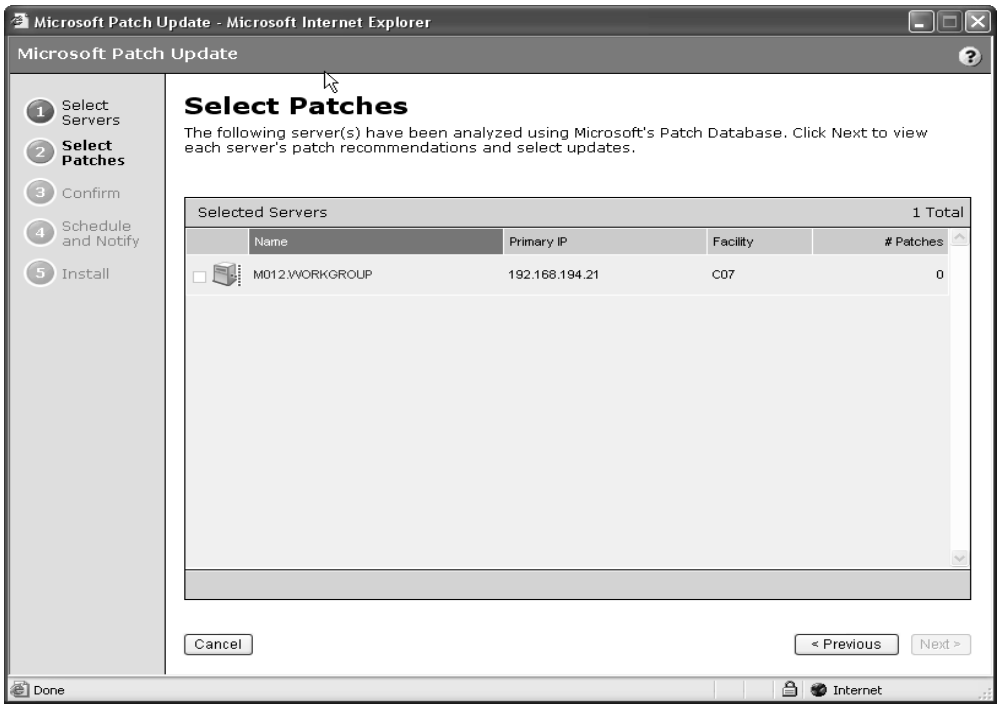

7. Select the patches that you want to install and then click Next. The Confirm page appears.

8. Review your selections.

9 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

If you decide to install the patches immediately by clicking install, a progress bar appears.

10 Click View Details for information about the installation.

Chapter 9: Reconcile

IN THIS CHAPTER

This chapter provides a technical overview of reconcile, which is the process that the Opsware System uses to install and remove software and to enforce its model-based approach. This chapter contains the following sections:

- Overview of Reconcile
- Ways to Perform Reconcile
- How Reconcile Works
- Reconcile on Supported Operating Systems
- About Reconcile and Scripts
- Reconcile Output
- Assigning to and Removing Servers from Nodes
- Reconcile Software Wizard

This chapter explains how reconcile and test reconcile work, and how different software types are treated during reconcile. It also provides information about the relationship of the installation and uninstallation wizards to the reconcile process, and explains the output from reconcile. The wizards are implemented using reconcile, and it is important to understand the reconcile process to understand what happens when a user invokes an installation/uninstallation wizard.

This chapter also explains how to reconcile selected servers through a specific reconcile wizard rather than through the wizards that are used to install and uninstall patches, software, or install templates.

Overview of Reconcile

The Opsware System uses a model-based approach, which provides a high degree of change management control and detailed records of all changes made. These features in turn make it possible to apply and enforce policies even within large and heterogeneous environments.

One central mechanism that the Opsware System uses to implement its model-based approach is called *server reconcile*. Server reconcile orchestrates the installation and uninstallation of all software, including applications and patches, on servers that the Opsware System manages. See "How the Opsware Model Affects Server Management" on page 33 in Chapter 2 for more information.

In the Opsware System, the installation and uninstallation of software begins by first making a change to the model of the server contained in the Model Repository. When you use any of the install or uninstall software wizards, for example, you cause the Opsware System to update the model of a server by adding or removing software nodes from the model. See "Application Provisioning Setup" on page 307 in Chapter 6 for more information.

During the process of using a wizard, if you decide to install or uninstall software immediately, the Opsware System begins a reconcile session. You can also schedule the installation or uninstallation for a later time, and the Opsware System starts the reconcile session at the time that you request. The reconcile session determines what needs to be done to install or uninstall the requested software, tests the results of those actions, and then initiates the tasks necessary to install or uninstall the software, such as downloading packages and initiating system utilities on the servers.

This *indirect* method of installing and uninstalling software – first changing the model of the software, and then changing the server to match the model – allows you to install software more safely and consistently. See "Application Provisioning" on page 385 in Chapter 7 for information about how to use the install and uninstall software wizards.

Ways to Perform Reconcile

Most of the time, users do not actually encounter the term *reconcile* when they use the standard Opsware System installation and uninstallation wizards. The wizards are designed to automate and simplify the reconcile process. Reconcile occurs as a result of using any of the following wizards:

- Install OS (See “Operating System Provisioning” on page 237 in Chapter 4 for more information.)
- Install/Uninstall Patch (See “Patch Management Subsystem” on page 403 in Chapter 8 for more information.)
- Install/Uninstall Software (See “Application Provisioning” on page 385 in Chapter 7 for more information.)
- Install Template (See “Application Provisioning” on page 385 in Chapter 7 for more information.)

In addition, you can invoke reconcile directly in order to install or uninstall software. You can do so by attaching or removing servers from nodes and then directly invoking a reconcile session on the servers by running the Reconcile Software Wizard.

Most often, users perform this *direct* reconcile when changes have been made to nodes, and the users want the servers that are already attached to these nodes to be reconciled to match the modified node.

How Reconcile Works

This section provides information on how reconcile works within the Opsware System and contains the following topics:

- Reconcile Overview
- About Reconcile and Package Metadata
- About Installation and Uninstallation Order
- Software Installation Order for Adopted Software
- About Patches and Reconcile
- About Preview Reconcile
- Types of Reconcile

Reconcile Overview

Reconcile works by comparing what is actually installed on a server to the software that should be installed on the server according to the server’s model. The Opsware System then determines what operations are required to make the server conform to its model.

To make this determination, the Opsware System queries the Opsware Agent on the server and examines the server's model to assemble the following data:

- A list of all the software (including patches) that is installed on the server. Although the Opsware System records this information in the Model Repository, the Opsware System uses the Opsware Agent on the server to compile a list of all installed software. Querying the Opsware Agent is necessary in case any software was installed manually, without using the Opsware System.

Opsware Inc. recommends that users install all software through the Opsware System. Manually installed software, however, can be adopted by the Opsware System. See "Software Installation Order for Adopted Software" on page 446 in this chapter for more information.

- A list of all the software that should be installed on the server, and the proper installation for the software. The Opsware System obtains this information from the Model Repository.
- A list of software on the server that was installed through the Opsware System. The Opsware System requires this information because it cannot uninstall software that was not installed through the Opsware System. This list is obtained from the Model Repository.

After this information is obtained, reconcile determines what actions are required to make the software match its model. Occasionally, the installation and uninstallation of software involves consequences that the user might not have anticipated. The installation of some software, for example, might require that other incompatible or outdated software be removed. In other cases, the installation of software might require the installation of other software components that were not explicitly requested, but are required by the packages that the user has selected.

The Opsware System then calls the server's native utilities to carry out the installation or uninstallation of software. (For example, on a Solaris server, the Opsware System uses the Solaris utility `pkgadd` to install software, and on AIX the Opsware System uses `installp`. See Table 9-1. See "Utilities Used During Reconcile Sessions" on page 448 in this chapter for information about a list of these utilities.) If the reconcile operation requires both uninstallation and installation of software, the uninstallation occurs before the installation.

About Reconcile and Package Metadata

Reconcile relies on metadata about software that is going to be installed and removed, and this metadata is obtained when the software is first uploaded to the Package Repository (either though the command line interface, the Upload Patch Wizard, or though package management). Metadata includes information such as *name* and *version*, and this information varies depending on the package type. (The filename alone is not sufficient to identify the package in the Opsware System.)

In most cases, this metadata is obtained automatically from the system utilities that run on the Software Repository server. The Software Repository server, however, is installed on a Unix-based machine. In the case of Windows packages and patches, the metadata must be entered manually, because the Windows utilities that obtain that type of information do not run on Unix.

About Installation and Uninstallation Order

One of the most important determinations that the reconcile process makes is the order in which software should be installed or uninstalled.

During reconcile, the Opsware System determines the correct installation or uninstallation order based on the following factors:

- If a node contains more than one package, the order in which the packages appear in the node affects the order in which they are installed
- The installation order dependencies between roles, if any

  When users create nodes and associate the nodes with software, they can express installation order dependencies with other nodes (for example, the software in one node should be installed before the software contained in another node, according to how users define the installation order dependencies.)
- The order in which software should be uninstalled

  When the Opsware System installs software on a server, it records the order in which the software was installed (as determined by reconcile) in the Model Repository. The software is uninstalled in the opposite order in which it was installed. This record of software installation order is maintained as long as the Opsware System manages the server. The record contains data about the order from all reconcile operations, not just individual reconcile operations.

Software Installation Order for Adopted Software

The Opsware System does not uninstall software that was not installed through the Opsware System. Software that was manually installed, however, can be *adopted* by the Opsware System in a node. If a user attaches a package that was manually installed on a server, assigns the server to the node, and performs a reconcile, the software becomes adopted by the Opsware System. (The software is not actually reinstalled during this operation.)

The Opsware System does *not* adopt Solaris patches. For example, if you uninstall a Solaris patch that was adopted into the Opsware System, the patch will not be uninstalled.

When software is adopted by the Opsware System, it is uninstalled according to when it was adopted, not when it was originally installed. The following example illustrates the order in which adopted software is uninstalled:

1. A server has a package that is installed on it before the server comes under management by the Opsware System.
2. A user installs three packages on this server using the Opsware System.
3. After these installations, a user decides to adopt the package by creating a node for the package, assigning that node to the server, and reconciling the server through the Reconcile Software Wizard. (The software is not reinstalled when the server is reconciled; instead, the package is simply adopted.)
4. If these packages are later uninstalled, the package that was adopted is uninstalled in the reverse order according to when it was adopted. The adopted package will, therefore, be uninstalled after the three packages, even though the package was in fact on the server before the three packages were installed.

About Patches and Reconcile

All software in the Opsware System is associated with nodes. Ordinarily, these nodes are created explicitly by users during application provisioning setup. Patches, however, are treated differently, in order to expedite the process of patch management. When a patch is first uploaded it is not immediately associated with a node. The first time that a patch is applied to a server, however, the node is created behind the scenes. This node is not part of the ordinary Software Tree and does not display in the tree.

If a user wants to add an installation order dependency to a patch, however, the user does in fact create a node for the patch. The node is used to express the installation order dependency. See “Application Provisioning Setup” on page 307 in Chapter 6 for more information.

About Preview Reconcile

Before any changes are committed to a server, the Opsware System first performs a preview reconcile. The preview reconcile allows users to see exactly what happens to the server as a result of the software that they requested to be installed or uninstalled. (This information displays individually for each server that is selected for reconcile.)

Preview reconcile shows what packages will be installed and what packages will be removed. If a package is removed or installed as a result of another package being installed, the user is informed of the reason that the package must be removed or installed.

In some cases, installation and uninstallation require reboots. This information also displays during the preview reconcile.

When you use the install and the uninstall patch wizards, the Opsware System does not perform a preview reconcile.

Types of Reconcile

The two types of reconcile are partial and full.

During a partial reconcile, the Opsware System only reconciles servers based on the nodes that the user has currently selected. For example, if a user has assigned a server to two nodes through the Install Software Wizard and then proceeds with the software installation, the server is reconciled only with those two nodes. If any other nodes have been assigned or removed through other means, such as nodes assigned through the managed servers list, these nodes are not reconciled.

During a full reconcile, a server is reconciled with all of the nodes that it has been assigned to. (If any nodes have been detached from the server, reconcile also uninstalls the software associated with those nodes.)

If any nodes were changed since they were attached to the server – for example, if patches were added to the nodes or if software was removed from the node, these changes are committed to the server during a full reconcile.

Reconcile on Supported Operating Systems

This section provides information on reconcile on supported operating systems and contains the following topics:

- Reconcile on Supported Operating Systems Overview
- AIX Reconcile
- HP-UX Reconcile
- Solaris Reconcile
- Linux Reconcile

Reconcile on Supported Operating Systems Overview

After the Opsware System determines what packages need to be installed or removed to complete the reconcile operation, reconcile uses a set of standard system utilities to complete the operation. Table 9-1 shows the utilities used during the reconcile session.

Table 9-1: Utilities Used During Reconcile Sessions

SOLARIS	LINUX	AIX	HP-UX	WINDOWS
patchadd (installs patches)	RPM (installs and removes software)	installp (installs software)	swinstall (installs software)	msiexec.exe (installs and uninstalls MSI packages)
patchrm (removes patches)		installp -u (removes software)	swremove (removes software)	unzip.exe (extracts info-zip compatible zip archives)
pkgadd (installs software)		inutoc (generates a table of contents of packages to be installed)	swlist (copies individual packages into one large depot)	mbsacli.exe (obtains the Microsoft patch inventory of a system)

Table 9-1: Utilities Used During Reconcile Sessions

SOLARIS	LINUX	AIX	HP-UX	WINDOWS
pkgrm (removes software)			swmodify (used to convert older format packages to a newer package format)	qchain.exe (used to install many Microsoft patches in succession without the need for intermediate system reboots)
RPM (installs and removes software)		RPM (installs and removes software)		

See “Package Management” on page 267 in Chapter 5 for information about the package types that the Opsware System supports.

AIX Reconcile

AIX software is delivered in LPPs, which are collections of filesets. When a server is reconciled and the Opsware System determines that the reconcile requires filesets to be installed, the Opsware System downloads the entire LPP that contains the filesets from the Software Repository to the server. If the filesets that the reconcile requires are contained in more than one LPP, the additional LPPs are also downloaded.

When you uninstall AIX filesets, the reconcile operation also uninstalls dependent filesets. The list of dependent filesets that are uninstalled appears in the reconcile status messages. The list of dependent filesets to be uninstalled does *not* appear in the Preview reconcile.

HP-UX Reconcile

HP-UX software is delivered in depots, which are collections of filesets. When a server is reconciled and the Opsware System determines that the reconcile requires filesets to be installed, the Opsware System downloads the entire depot that contains the filesets from the Software Repository.

If the filesets that the reconcile requires are contained in more than one depot, the additional depots are also downloaded. The depots are then combined into one large depot, from which the filesets will be installed.

HP-UX filesets often have dependencies on other filesets that the user has not specifically requested (or that have been included in a software node). These filesets can, however, be included in the HP-UX depot. By downloading the entire depot (instead of just the requested filesets), the Opsware System is able to install any additional filesets that are required by the filesets associated with the nodes that are being reconciled. Combining the individual depots into one large depot allows the underlying installation utilities to locate all the filesets that require installation.

HP-UX 10.20 does not support the option `-x show_superseded_patches`; therefore, if you install a superseding patch, it removes the superseded patches. For example, if patch B supersedes patch A, the reconcile operation reports that patch A was removed when patch B is installed.

Solaris Reconcile

Solaris patches do not contain metadata that identify what cluster they belong to after the patch clusters are installed. During the reconcile process, however, the Opsware System records the fact that the patches installed belong to a given patch cluster. This allows the Opsware System to identify patch clusters that are installed on the servers. The Opsware System can use this information to uninstall patch clusters.

Solaris patch clusters cannot be adopted by the Opsware System; if the patch clusters were not installed through the Opsware System, it is not possible to determine if a patch on a server originated from a patch cluster. The Opsware System can, however, adopt individual patches.

Linux Reconcile

RPM is the only package type that the Opsware System uses on the Linux operating system. When software is installed, the -i option is always used; when software is removed, the -e option is always used.

About Reconcile and Scripts

When users upload software, they have the option of specifying scripts that should be run when software is installed or uninstalled. Reconcile executes these scripts on the servers local shell. For these scripts, users can elect to have reconcile react to a non-zero return

code from the script by aborting reconcile operation when the non-zero return code is received. If a non-zero return code is encountered in a post-install script, the Opsware System does not "roll back" or uninstall any software that has already been installed; again, the reconcile process simply halts when the non-zero return code is encountered (if this is the option that was selected for the script.)

These options are set when the software is uploaded to the Opsware System, and can be edited through the patch management and package management interfaces.

Reconcile Output

Reconcile provides detailed feedback about what occurs during the reconcile process, and what changes have been made on the servers selected for reconcile. The Opsware System provides individual output for each server that has been selected for the reconcile operation. The output is the same for all wizards that use reconcile.

The output from a reconcile operation consists of the following types of data:

- A list of all software installed and uninstalled. If software is installed or uninstalled that was not specifically requested but is required for the reconcile operation, the output specifies why the software was added or removed.
- If any pre- or post-install scripts are executed, the first 1000 bytes of the scripts' `stdout` and `stderr` are displayed, as well as the return code for the script.
- Any reboots required by the reconcile operation.
- The output from the utilities that the reconcile operation used to install and uninstall the software. In some cases, these utilities might report errors. For example, a user might request AIX filesets to be installed that are dependent on other filesets that are not available to the Opsware System. This error is reported as part of the reconcile output.

Assigning to and Removing Servers from Nodes

Servers are most often assigned and removed from software nodes by using the installation or uninstallation wizards for patches and software, or the installation wizard templates. You can, however, manually assign or remove a server from a node.

If you manually assign or remove a server from a node, the software for that node is not installed or uninstalled until you use the Reconcile Wizard on that server.

Assigning Servers to Nodes

Perform the following steps to assign servers to nodes:

1. From the navigation panel in the Opsware Command Center, click Servers ➤ Server Search.
2. Use Server Search to find the server or servers that you want to assign to software nodes. The servers must all be running the same version of the same operating system.
3. From your search results list, select the server or servers that you want to assign to a software node.
4. From the Server drop-down menu, select AssignNode. The Assign Node Wizard appears.
5. Navigate to and select the node to which you want to assign the server.
6. Click the Assign button.

Removing Servers from Nodes

Perform the following steps to remove servers from nodes:

1. From the navigation panel in the Opsware Command Center, click Servers ➤ Server Search.
2. Use Server Search to find the server or servers that you want to remove from software nodes. If you are selecting multiple servers, the servers must have at least one node in common.
3. From your search results list, select the server or servers that you want to remove from a software node.
4. From the Server drop-down menu, select Remove Node.
5. Select the check boxes from the nodes that you want to remove and click the Remove button.

Reconcile Software Wizard

Servers are usually reconciled as a result of running any of the OS provisioning, patch management, or application provisioning wizards. Reconcile can, however, be directly invoked on a selected server or group of servers through the Reconcile Software Wizard. This Wizard enables some of the “power user” flexibility that the other wizards hide.

Most often, users perform this *direct* reconcile when changes have been made to nodes, and the users want the servers that are already attached to these nodes to be reconciled to match the modified node.

Additionally, you can use the Reconcile Software Wizard to ensure that the server conforms exactly to its model. For example, if a node is deleted from the Software Tree and a server is attached to that node, the software for the node can only be uninstalled from the server by using the Reconcile Wizard.

You can select one or multiple servers to reconcile. If you select multiple servers, you can only reconcile software that is common to all of the servers you have selected (for example, you can only choose nodes that all servers have in common). You can, however, perform a full reconcile on a group of servers even if they do not have any software in common.

Directly Reconciling Servers

Perform the following steps to directly reconcile servers:

1. From the navigation panel in the Opsware Command Center, click Servers ➤ Server Search.
2. Use Server Search to find the server or servers that you want to reconcile.
3. Review the servers that your search returned and select the servers that you want to reconcile.
4. From the Server drop-down menu, select Reconcile. The Reconcile Wizard appears.
5. If you have selected multiple servers, you can either select All Software to perform a full reconcile on all selected servers, or you can select Common Software to reconcile only the software common to all selected servers. If you selected a single server, you can either select to reconcile “Some Software” to choose the software (nodes) that you want to reconcile, or you can select “All Software” to perform a full reconcile.
6. Click Next to continue.

7. If you selected a partial reconcile (All Software or Common Software), you must now select the check boxes for the software that you want to reconcile your server with. If you selected a full reconcile, a list of all software to be reconciled displays and you can confirm your selection. Confirm your selections and click Preview to continue. A preview reconcile occurs.
8. Review the results of the preview reconcile. Click the Next button.
9. On the Schedule and Notify page, you have the following options:
   - Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
   - Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.
   - Click the Reconcile button to complete the process now.

Chapter 10: Script Execution Subsystem

IN THIS CHAPTER

The following topics are covered in this chapter:

- Script Execution Subsystem
- Initiating Subsystem Operations
- Permissions Required for Subsystem Tasks
- Script Management: Tasks, Tips, and Procedures
- Script Execution: Tasks, Tips, and Procedures
- Script Execution Results: Tasks and Procedures
- Subsystem Error Resolution
- Opsware Custom Extensions

Script Execution Subsystem

This section provides information about the Script Execution Subsystem within the Opsware System and contains the following topics:

- Script Execution Subsystem Overview
- Scripts Types – My Scripts, Shared Scripts, and Ad-Hoc Scripts
- About Subsystem Functionality
- Run Distributed Script Link

Script Execution Subsystem Overview

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 on individual servers, one server after another. With the Script Execution Subsystem, a user performs all script tasks from 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 the script again at a later date.

Script management and execution features are guided by a permission-based system that controls user access to particular types of scripts and tasks.

Scripts Types – My Scripts, Shared Scripts, and Ad-Hoc Scripts

The Script Execution Subsystem supports the three major types of scripts for Unix and Windows operating systems: Unix/Linux shell, Windows batch (.BAT), and Windows Visual Basic (VBScript).

After you create or upload Unix or Windows scripts in the Opsware System, the scripts are stored in the Opsware System in one of two ways:

- As private scripts, accessible only to the user who created them. In the Opsware System, private scripts are called *My Scripts.*

  My Scripts can only be edited, deleted, or executed by the user who created the script. My Scripts are intended for personal use.

- As public scripts, accessible to all Script Execution Subsystem users. In the Opsware System, public scripts are called *Shared Scripts*.

A third type of Opsware System script is created (or uploaded) and then immediately executed by a user. The script is intended for one-time use and is not stored in the Opsware System. In the Opsware System, this type of script is referred to as an *Ad-Hoc Script*. During the Ad-Hoc Script creation and execution process, only one user has access to the script.

After you create a script and store it as a specific type of script in the Opsware System, you cannot convert the script to the other type of script. My Scripts cannot be converted to Shared Scripts (and vice versa).

About Subsystem Functionality

The Script Execution Subsystem provides three basic functions:

- Script management
- Script execution
- Viewing script execution results

The Opsware administrator determines the amount of Opsware Command Center functionality a user has access to and provides permissions appropriate to the user's job.

Script Management Functionality

Script management tasks include:

- Viewing contents of stored My Scripts or Shared Scripts
- Creating (or uploading) My Scripts or Shared Scripts for storage in the Opsware System
- Editing or deleting stored My Scripts or Shared Scripts
- Viewing version history of stored My Scripts or Shared Scripts

Script management functionality is handled by two Opsware System components – the Command Engine and the Opsware Command Center. The Command Engine handles the entry of scripts into the Opsware Model Repository (the script storage location in the Opsware System) and the versioning of stored scripts. The Opsware Command Center provides the user interface for script management activities. It provides tools that allow users to create or upload scripts for storage. It also allows users to create Ad-Hoc (one-time-use, not stored) scripts for immediate execution.

Script Execution Functionality

Script execution tasks include:

- Executing a My Script or Shared Script, stored in the Opsware System, on one or more servers
- Creating (or uploading) an Ad-Hoc Script and then immediately executing it on one or more servers

Script execution functionality is handled by three Opsware System components – the Opsware Command Center, the Command Engine, and the Opsware Agent. The Opsware Command Center provides the user interface for script execution activities. Script execution tasks are automated by a wizard that leads users through the following script execution steps:

1. Select scripts.
2. Select servers.
3. Specify execution options.
4. Confirm settings.
5. Execute scripts across one or more servers.

During script execution on the servers, the Command Engine runs a script that issues an execution command to the Opsware Agent on each server. Each Opsware Agent handles script execution and sends execution results to the Command Engine.

Execution of a script on the managed servers cannot be rolled back.

Script Execution Results Functionality

Execution results display immediately after a script runs and can be viewed any time after a script is executed. The functionality that displays execution results is handled by two Opsware System components—the Command Engine script and the Opsware Command Center. The Command Engine Script enters the execution results data into the Model Repository. The Opsware Command Center retrieves and displays execution results data from the Model Repository and provides tools for the user to download execution results data (output and error files) as a zip file.

Initiating Subsystem Operations

In the Opsware Command Center, you can initiate a Script Execution Subsystem operation in the following three ways:

- *Scripts* link
- *Run Distributed Script* link
- By selecting Server ➤ Run Script from the Managed Servers page

The following topics provide details about the initiating scripts from the Scripts link and from the Run Distributed Script Wizard on the Home page. See "Server Management Tasks Related to the Server Life Cycle" on page 72 in Chapter 2 for information about how to initiate scripts from the Managed Servers page.

The Scripts link displays tools for managing scripts. Select this link to create or upload scripts that you want to save or run, and to view, edit, or delete scripts that have already been stored in the Opsware System. The Scripts link is located on the navigation panel under Software. See Figure 10-1.

Figure 10-1: The Scripts Link Is Used to Create, Run, Upload, Edit, Delete, or View Scripts

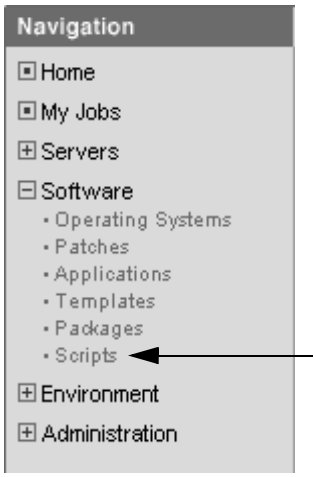

Run Distributed Script Link

The quickest way to execute a script that is stored in the Opsware System, or to create and immediately execute an Ad-Hoc Script, is to select the Run Distributed Script link in the Tasks panel of the Opsware Command Center. See Figure 10-2.

Figure 10-2: Click the Run Distributed Scripts Link to Launch the Script Execution Wizard

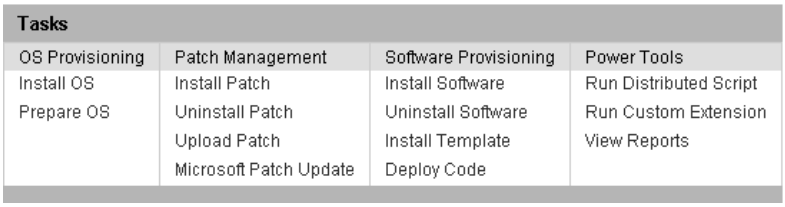

You can view the status of scripts in the My Jobs area of the Opsware Command Center Home page, which displays information about jobs that have run, are currently running, or are scheduled to run, including script execution jobs. Through My Jobs, you can display script execution results. The name of the job is also a hyperlink to a pop-up window that allows you to change scheduling information for that job. In addition to the name, start time, and status, the number of servers affected by the job also appears.

The hyperlink to script execution jobs is called *Run Script*. A particular execution event can be identified by the start time of the execution event. Click any *Run Script* name to view the results of that particular script execution if it has completed. Otherwise, information about when the job is scheduled to run appears instead. See Figure 10-3.

Figure 10-3: The My Jobs Panel That Shows a List of Executed Scripts

You can display a complete list of My Jobs by clicking the See All link in the My Jobs area of the home page, or by clicking the My Jobs link from the navigation panel.

Permissions Required for Subsystem Tasks

The amount of access a user has to Script Execution Subsystem features depends on the permissions the Opsware administrator provides. The following three main types of permissions control Script Execution Subsystem functionality:

- Access to script management features

  Opsware System permissions that provide script management access are *Scripts* and *Edit Shared Scripts*.

- Access to script execution features

  Opsware System permissions that provide script execution access are *Wizard: Run Scripts* and *Run My Script as Root*.

- Access to customer and facility-specific Opsware-managed servers

  Specific customer and facility *read/write* permissions are required to access the managed servers used to execute scripts.

Before you start working with the Script Execution Subsystem, make sure that you have the permissions that you need to perform your tasks. Only features associated with the permissions you have are available on the Opsware Command Center navigation panel and the home page. To obtain permissions for unavailable subsystem features, contact your Opsware administrator. Additional permissions might be needed if your responsibilities change or if you need to perform troubleshooting tasks.

Table 10-1 provides an overview of the most common script management, script execution, and script execution results tasks, and displays the permissions required to perform a task.

In the table, tasks are categorized according to functionality (script management, execution, and execution results). Because each task is performed on a specific type of script (My Scripts, Shared Script, or Ad-Hoc Script), the table also shows the tasks according to the type of script involved.

For example, if a user's job is to execute a Shared Script, the table shows that the user would need *Wizard: Run Scripts* permission. This permission allows Shared Scripts to be executed on the servers as root or local system. If a user's job is to execute a My Script as root or local system, the user would need *Wizard: Run Scripts*, *Scripts,* and *Run My Script As Root* permissions. In both of these examples, the user would also need *read/write* permissions for specific customer and facility servers.

Table 10-1: Permissions Required for Script Tasks

SCRIPT TASK	APPLICABLE SCRIPT TYPE	REQUIRED PERMISSIONS FOR SCRIPT TASK AND TYPE	COMMENTS
SCRIPT MANAGEMENT TASKS			
View list of available scripts.	My Script, Shared Script	• Scripts	
Create or upload and store a script.	My Script	• Scripts	
Edit, delete, or view a stored script.	My Script	• Scripts	

Table 10-1: Permissions Required for Script Tasks

SCRIPT TASK	APPLICABLE SCRIPT TYPE	REQUIRED PERMISSIONS FOR SCRIPT TASK AND TYPE	COMMENTS
View version history of a stored script.	My Script	• Scripts	
Create or upload and store a script.	Shared Script	• Scripts • Edit Shared Scripts	
Edit, delete, or view a stored script.	Shared Script	• Scripts • Edit Shared Scripts	
View version history of a stored script.	Shared Script	• Scripts • Edit Shared Scripts	
SCRIPT EXECUTION TASKS			
Execute a script (as root or local system).	Shared Script	• Wizard: Run Scripts • Read/Write permissions for specific customer and facility servers	A Shared Script always executes on a server as root or local system.
Execute a script (requires a password).	My Script	• Wizard: Run Scripts • Scripts • Read/Write permissions for specific customer and facility servers	With *Wizard: Run Scripts* and *Scripts* permissions, execution of a My Script requires the use of a password.

Table 10-1: Permissions Required for Script Tasks

SCRIPT TASK	APPLICABLE SCRIPT TYPE	REQUIRED PERMISSIONS FOR SCRIPT TASK AND TYPE	COMMENTS
Execute a script (as root or local system).	My Script	• Wizard: Run Scripts • Scripts • Run My Script As Root • Read/Write permissions for specific customer and facility servers	With *Run My Script As Root* permission, no password is required. Without this permission, the user can still execute the script, but only with a password.
Create (or upload) and then execute an *Ad-Hoc Script* (requires a password).	Ad-Hoc Script	• Wizard: Run Scripts • Scripts • Read/Write permissions for specific customer and facility servers	With *Wizard: Run Scripts* and *Scripts* permissions, a password is required to execute an Ad-Hoc Script.
Create (or upload) and then execute an *Ad-Hoc Script* (as root/local system).	Ad-Hoc Script	• Wizard: Run Scripts • Scripts • Run My Script As Root • Read/Write permissions for specific customer and facility servers	With *Wizard: Run Scripts*, *Scripts*, and *Run My Script As Root* permissions, an Ad-Hoc Script executes on the servers as root/local system (without a password).

Table 10-1: Permissions Required for Script Tasks

SCRIPT TASK	APPLICABLE SCRIPT TYPE	REQUIRED PERMISSIONS FOR SCRIPT TASK AND TYPE	COMMENTS
SCRIPT EXECUTION RESULTS TASKS			
View execution results data.	My Script, Shared Script, Ad-Hoc Script	• Wizard: Run Scripts	The *Wizard: Run Scripts* permission allows a user to view results information for any executed script.

Script Management: Tasks, Tips, and Procedures

Script management tasks that the subsystem supports include:

- Creating a Script
- Viewing the Scripts List
- Editing and Deleting a Script
- Viewing Script Version History

The following topics provide information about, and procedures for, performing these script management tasks.

Creating a Script

You can create scripts in the Opsware Command Center or by uploading a script into the Opsware System.

Script Creation Tips

The following information provides important script creation details:

- The Opsware System supports the three major types of scripts for Unix and Windows operating systems: Unix/Linux shell, Windows batch (.BAT), and Windows Visual Basic (VBScript).
- 4 MB is the maximum size allowable for a script.

- When you create a Unix shell script with a language other than the Bourne (sh) shell, use the sh-bang (#!) format at the top of the script to specify the correct command interpreter. The command interpreter needs to be present on the Opsware-managed server.

  For example, if you are using Perl, the beginning of the script would contain the following line:

  ```
  #!/usr/bin/perl
  ```

  The following example shows a short Perl script (it displays "hello world"):

  ```
  #!/usr/bin/perl
  print "hello world\n"
  ```

  VBScripts are executed by the VBScript interpreter on the Windows server.

- To access command line parameters with Unix shell commands, use the following convention: `$1 $2...`
- To access command line parameters with Windows .BAT, use: `%1 %2...`
- Script lines do not need to be terminated in a specific way. But with Windows scripts, the Opsware System converts all `\n` to `\r\n`. With Unix scripts, all `\r\n` are converted to `\n`.
- Scripts should be written to send error output to standard error (file descriptor #2).
- Scripts should use the standard convention of returning a zero code to indicate success. For other return codes, there is no standard code system to follow. Create unique non-zero return codes to handle each type of error.

How to Create or Upload a Script

Perform the following steps to create or upload a My Script or a Shared Script:

1. From the navigation panel, select Software ➤ Scripts.
2. The Scripts page has two tabs: My Scripts and Shared Scripts.
   - To create or upload a My Scripts, click the My Scripts tab ➤ New Script. The New Script page appears.

- To create or upload a Shared Script, click the Shared Scripts tab ➤ New Script. The New Script page appears, as Figure 10-4 shows.

Figure 10-4: Scripts: New Script Page

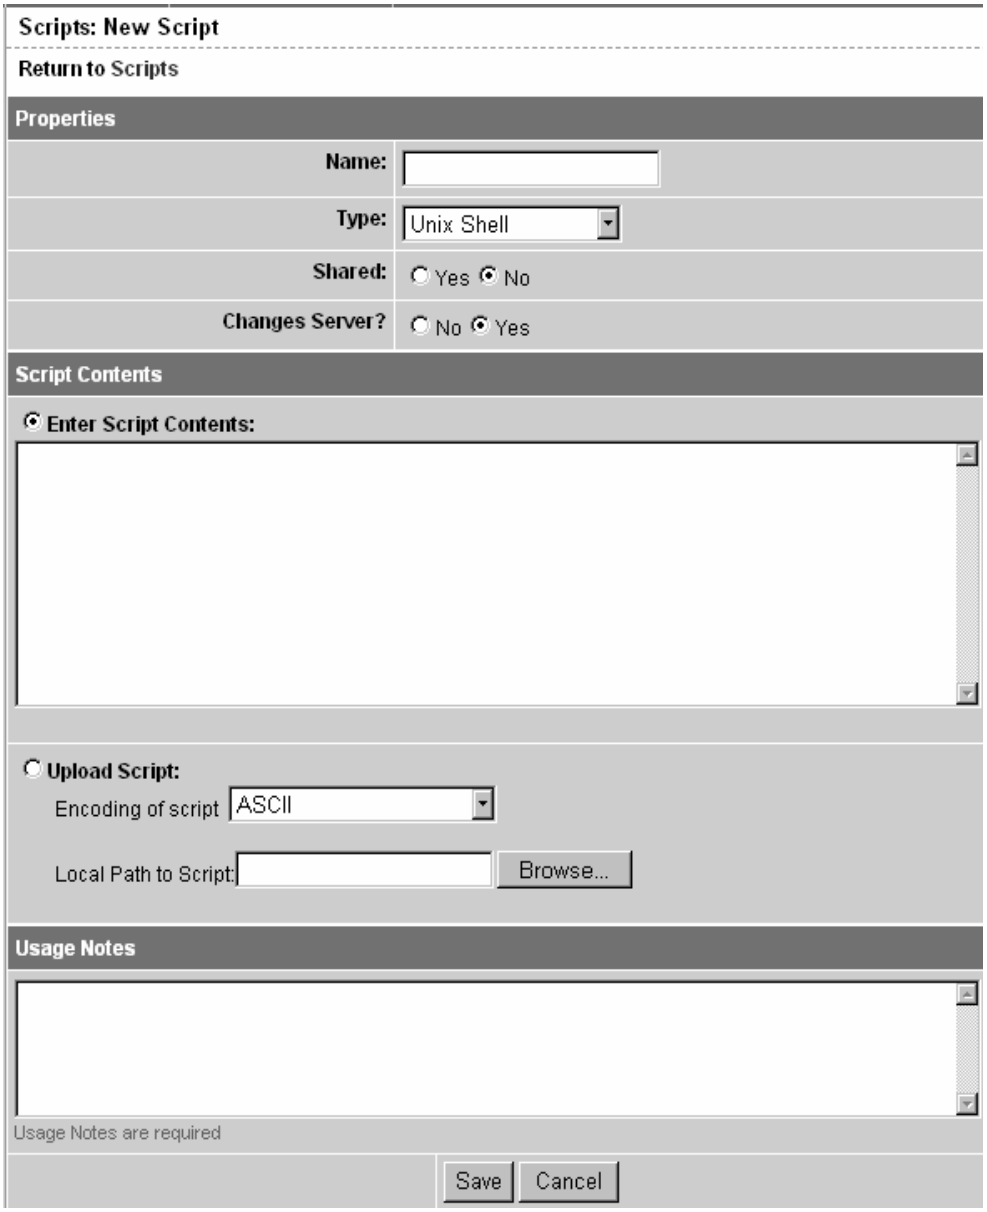

3 On the Scripts: New Script page, enter the following data under Properties:

- Enter the name of the script. The name must be a unique shared-use or personal-use name.
- Select the script type: Unix shell, Windows .BAT, or Windows VBScript.
- If you are creating a Shared Script, select Yes next to Shared. If you are creating a My Script, select No.
- In the Changes Server field, select No if the script does not modify the server, and Yes if it does. If you select Yes, locked servers cannot be selected to run that script.

4 Under Script Contents, enter or upload a script by performing one of the following tasks:

- To upload a script, click Upload Script. In the Local Path to Script box, either manually enter the path to the script, or click Browse to locate the script.

When you upload a script, you must select an encoding scheme for the script from the list so that the Opsware System can convert the bytes inside the script into UTF-8 format by using the encoding scheme with which the script was created.

- To create a script, click Enter Script Contents and manually enter the script in the text box.

The script editor does not recognize tabs and its functionality is browser-dependent.

5 In the Usage Notes section of the page, enter script details or other descriptive information.

6 Click Save to store the script. The script is saved in the Model Repository.

The Scripts page appears and confirms that the script is now stored. The script is included in the list of available Shared Scripts or My Scripts.

Viewing the Scripts List

After you save a script in the Opsware System, you can view it on the list of stored My Scripts or Shared Scripts. My Scripts display only to the user who created them, while Shared Scripts display to all users.

Perform the following steps to view the list of stored My Scripts or Shared Scripts:

1. From the navigation panel, select Software ➤ Scripts.
2. The Scripts page has two tabs: *My Scripts* and *Shared Scripts*.
   - To view the list of My Scripts, click the *My Scripts* tab.
   - To view the list of Shared Scripts, click the *Shared Scripts* tab.

   A list of My Scripts or Shared Scripts displays. Each script name is also a link.
3. To view a script, click its name.

Editing and Deleting a Script

After you save a script in the Opsware System, you can edit or delete the script. Before you edit or delete a script, you might want to view the script properties, contents, usage notes, and change log.

My Scripts are accessible only to the user who created them and can only be edited or deleted by this user.

How to Edit a Script

Perform the following steps to edit a stored script:

1. From the navigation panel, select Software ➤ Scripts.
2. The Scripts page has two tabs: *My Scripts* and *Shared Scripts*.
   - To locate a My Script, click the *My Scripts* tab.
   - To locate a Shared Script, click the *Shared Scripts* tab.

   A list of scripts displays. Each script name is also a link.

3 To view the script, click the name. The View Script page appears, displaying script properties, contents, and usage notes, as Figure 10-5 shows.

Figure 10-5: The View Script Page

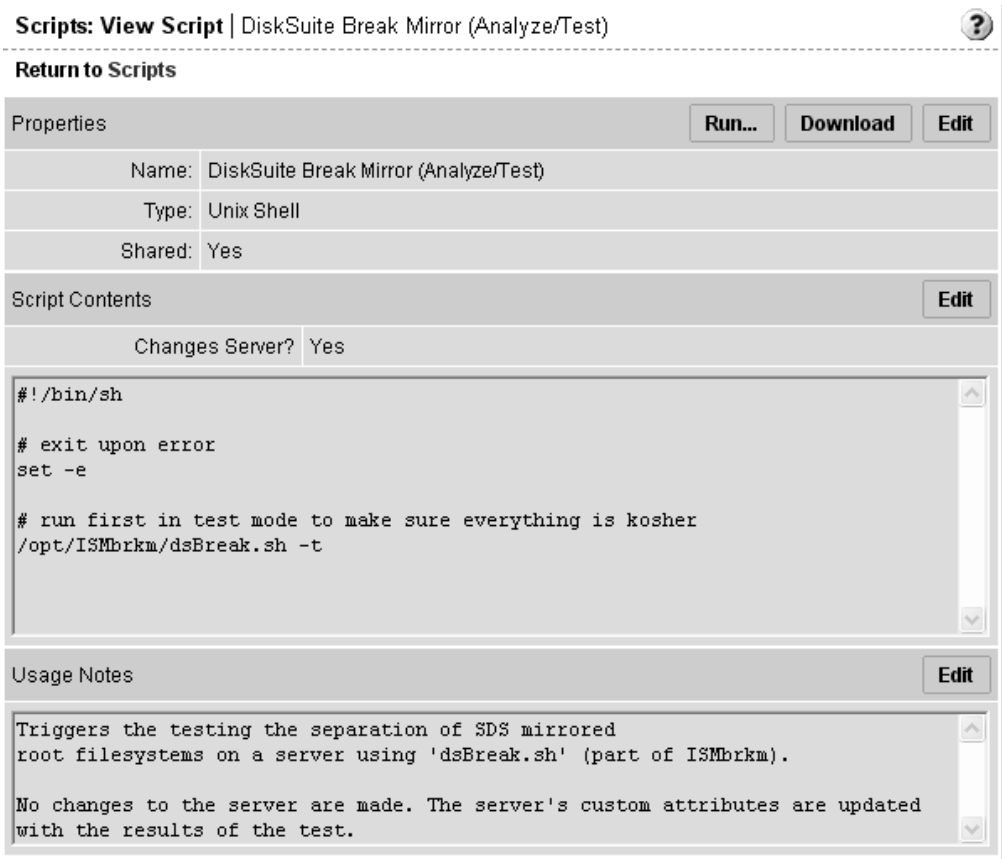

4 Click the Edit button in the Properties panel for details about script properties.

5 On the Scripts: Edit Script page, Change Log information is now available. Click the *Change Log* tab to view the current script's version history.

6 Use the tabs on the Scripts: View Script page to edit script contents or the script name, and enter usage notes. See Figure 10-6.

Figure 10-6: View Scripts Page That Shows Properties, Contents, Usage Notes, and Change Log Tabs

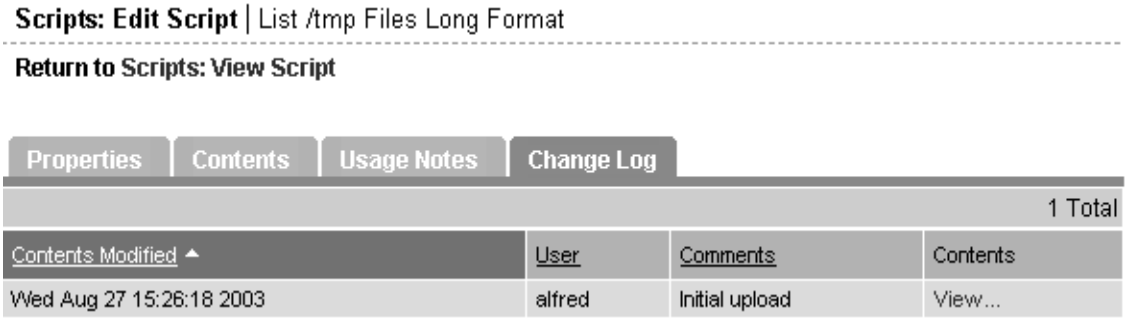

- To edit script contents, click the *Contents* tab. In the Edit Contents panel, click the Edit Script Contents radio button and edit the contents of the script.

  Instead of manually editing script contents, you can also upload new script contents. The uploaded script overwrites current script contents. To upload, click the Upload & Overwrite Script radio button and enter the location of the script you want to upload.

  After you edit script contents or upload new content, enter change log comments in the text box below the Edit Contents panel. Change log comments are required when you edit a script. When you are finished, click the Save button.

- To edit the name of the script, click the *Properties* tab. Edit the name that currently displays in the Name box and click the Save button.

- To enter usage notes, click the *Usage Notes* tab and enter information on the Edit Usage Notes panel. When finished, click the Save button.

How to Delete a Script

Perform the following steps to delete a stored script:

1 From the navigation panel, select Software ➤ Scripts.

2 The Scripts page has two tabs: *My Scripts* and *Shared Scripts*.

- To locate a My Script, click the *My Scripts* tab.
- To locate a Shared Script, click the *Shared Scripts* tab.

A list of scripts displays. Each script name is also a link.

3. To review a script before you delete it, click the name. To return to the list of scripts, click the Scripts link in Return to Scripts (located at the top of the Scripts: View Script page).
4. On the Scripts page, select the scripts that you want to delete by clicking the box located to the left of the script name. You can delete more than one script at a time.
5. Click the Delete button.
6. The Delete Scripts confirmation window appears.
   - To review the list of scripts you selected for deletion, click the View Details link.
   - To cancel the entire operation, click the Cancel button.
   - To delete the selected scripts, click the Delete button.

After you delete a script, you can still view the results of executions performed with that script.

Viewing Script Version History

Version history for a script is maintained in a change log, which is stored with other My Script or Shared Script management information. Each time a script is modified, new script version information is created and stored.

How to View Version History for a Script

Perform the following steps to view script version history:

1. From the navigation panel, select Software ➤ Scripts.
2. The Scripts page has two tabs: *My Scripts* and *Shared Scripts*.
   - To locate a My Script, click the *My Scripts* tab.
   - To locate a Shared Script, click the *Shared Scripts* tab.

   A list of scripts displays. Each script name is also a link.
3. To view the script, click the name. The Scripts: View Script page appears and displays script properties, content, and usage notes.
4. Click the Edit button in the Properties panel.
5. On the Scripts: Edit Script page, click the *Change Log* tab to access the change log and view version history for the current script.

The change log provides the following script version information:

- Date and time the script is modified
- Users who modified the script
- Comments associated with each script modification
- Script contents for that modification

Script Execution: Tasks, Tips, and Procedures

This section provides information about script execution tasks, tips, and procedures and contains the following topics:

- Script Execution Wizard
- Script Execution Tips
- How to Execute a My Script or a Shared Script
- How to Create and Execute an Ad-Hoc Script

Script Execution Wizard

A Script Execution Wizard automates script set up and execution processes and steps the user through the following execution stages (in the order shown):

1. Select Script: You can select only one My Script or Shared Script for each execution run. You can select scripts in the Wizard or from the Scripts page by selecting Software ➤ Scripts in the navigation panel.
2. Select Servers: YOu can select one or more servers from the displayed list of available servers. Only servers running an operating system applicable to the selected script are shown (for example, only servers running Unix are shown for a Unix shell script).
3. Specify Options: Runtime data, information, and execution options are entered at this stage.
4. Confirm Settings: This stage allows the user to review settings prior to execution. Click the Next button.
5. On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

6 Run Script: While the script executes, progress information displays about the total run and each server. When execution ends, summary information displays. After execution, script contents, output data, and error data can be immediately reviewed and downloaded.

The default amount of output data that the Opsware System stores is 10 KB. This amount can be modified by the Opsware Administrator.

Script Execution Tips

The following information provides important script execution details:

- No specific assumptions should be made about the execution environment. No particular environment variables are set.
- When executing processes on Unix servers, make sure long running processes (Web servers, databases, and so forth) are started as daemons. Also, make sure server processes properly daemonize themselves.
- For Windows servers, do not start anything that causes a window to open and wait for input.

How to Execute a My Script or a Shared Script

Perform the following steps to execute a My Script or a Shared Script:

1 Click the Run Distributed Script link in the Tasks area of the Opsware Command Center home page to launch the Run Distributed Script Wizard.

2 On the Overview page, make sure that you select Select Saved Script.

3 Click the Start button.

Alternatively, you can initiate the script execution process by clicking Software ➤ Scripts on the navigation panel, and then selecting a script and clicking the Run button.

4. To list Shared Scripts, click the *Shared Script* tab. Or, to list My Scripts, click the *My Scripts* tab.
5. From the script list, locate the script and select the script's radio button located to the left of the script name.
6. Click Next to continue (or click Previous to return to the previous step).
7. At the Select Servers page, browse or search for a server or servers to use.
   - To browse, click the *Browse* tab to obtain a list of servers available to you. At the top of the page, use the Customers and Facilities filters to narrow your selection.
   - To search, click the *Search* tab. On the Search page, indicate search criteria by selecting the appropriate check box. Default search criteria might be available to help narrow the search for specific customers and facilities.

Only servers are listed that use the operating system appropriate for the type of script that you want to run. For information about a server, click the name of the server. See Figure 10-7.

Figure 10-7: List of Servers That Are Available to Run the Script

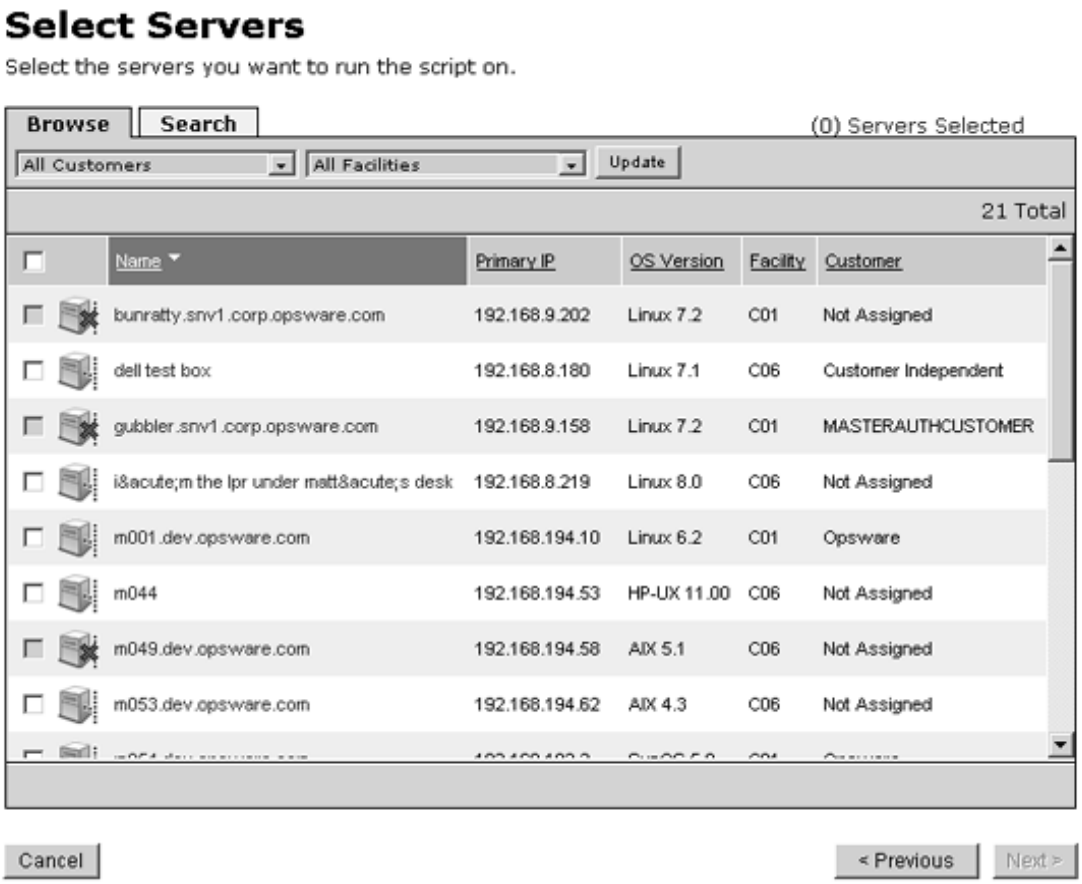

8 To sort the server list by server name, IP address, OS version, facility, or customer, click the heading of the column that you want to sort. For example, to sort by customer, click the heading titled Customer.

9 From the server list, specify the servers that you want to use. Select one or more servers for script execution.

10 Click Next to continue (or click Previous to return to the previous step).

11 On the Specify Options page, enter the runtime information, as Figure 10-8 shows.

Figure 10-8: The Specify Options Page

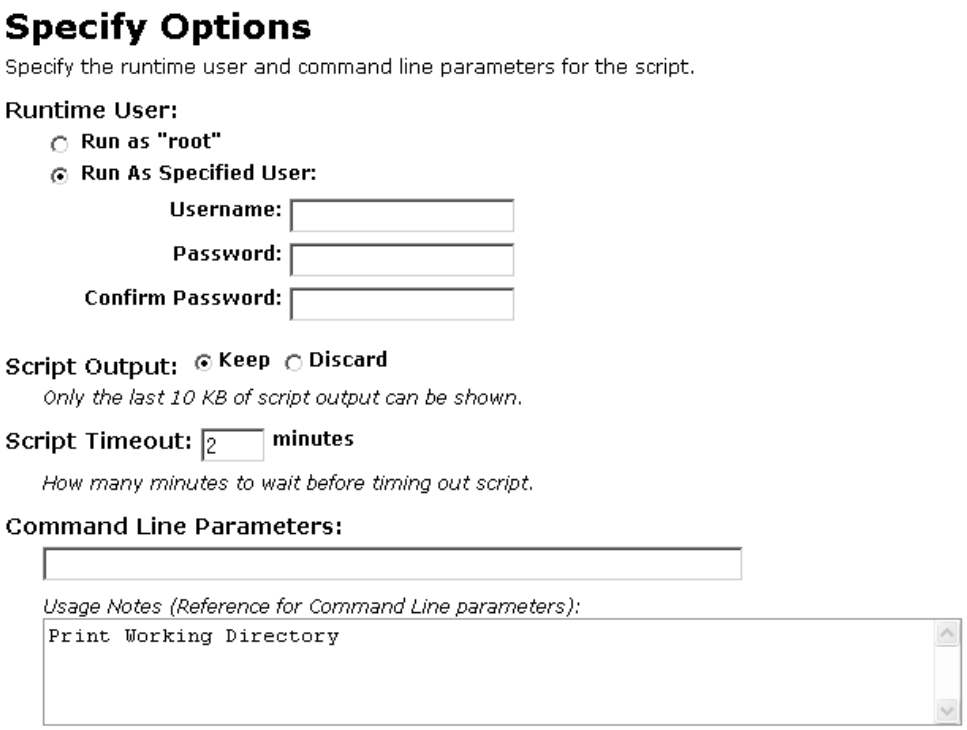

- Specify the runtime user. If you have the appropriate permission, you are able to execute the script as root or local system without entering a password. Otherwise, enter a username and password for the servers you intend to execute on. The username and password you use must be the same across all servers.
- Specify if you want to keep or discard script output. Only a portion of script output is saved (the amount of script output that is saved is configurable by the Opsware administrator, but the default is 10 KB).
- Enter a script timeout value in minutes.

  This value is the amount of time a script has to complete execution activities on a server. If the script is not finished when the timeout value is reached, the script is stopped by the Opsware System and a script error occurs.

  Select a timeout value that is greater than the time required for execution to complete.

- Enter command line parameters.

  Use the same syntax as when entering a script on a command line. For Unix scripts, use Bourne (sh) shell syntax. For Windows scripts, use `cmd.exe` parameter syntax.

- Add usage notes.

12 Click Next to continue (or click Previous to return to the previous step).

13 On the Confirm Settings page, review your script execution settings before you proceed. The Script to be run panel provides information about your script and the Servers panel displays the servers selected for script execution. The Confirm Settings page displays execution settings, as Figure 10-9 shows.

Figure 10-9: The Confirm Settings Page

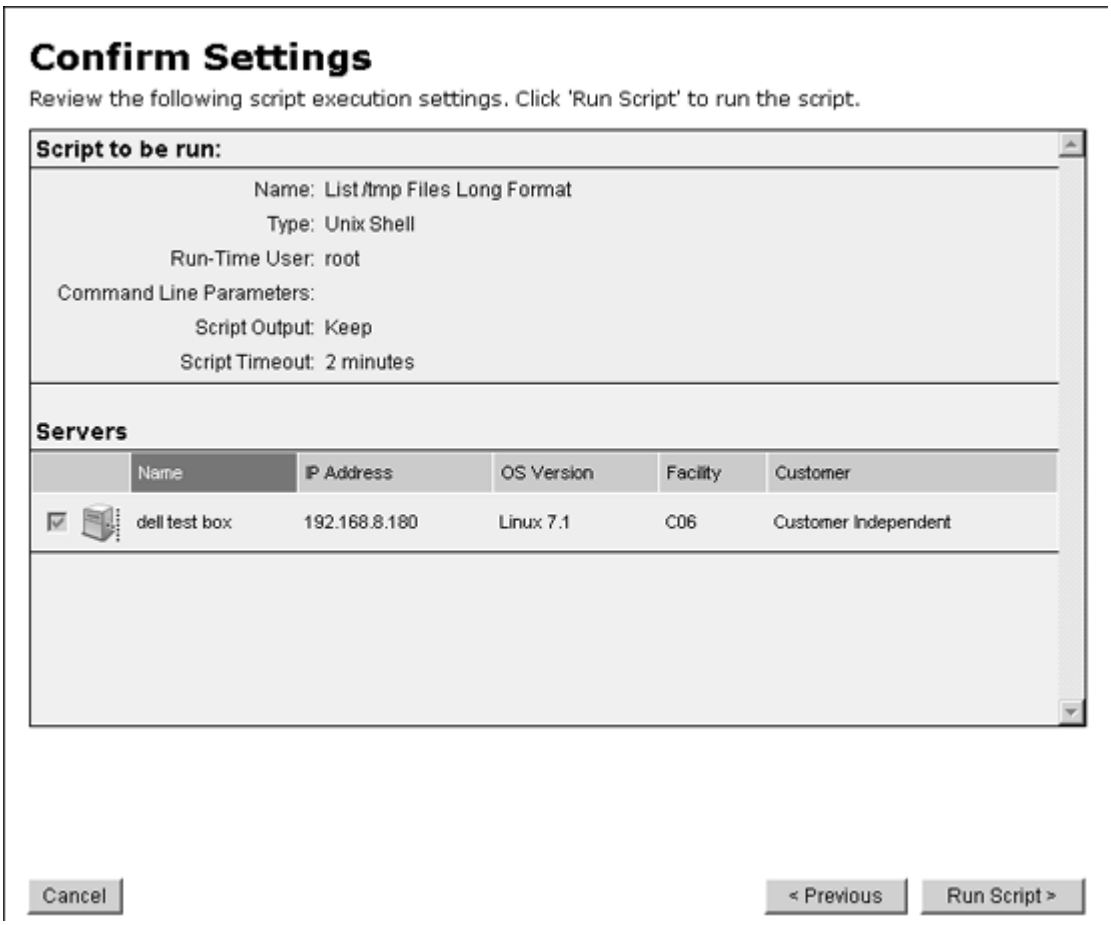

If you discover that script or server changes are needed, click the Previous button as many times as you need to return to the Script or Select Servers page.

14 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.
- Run Script: While the script executes, progress information displays about the total run and each server. When execution ends, summary information displays. After execution, script contents, output data, and error data can be immediately reviewed and downloaded.

If you select to run the job at that time, a progress bar appears that shows the progress of the script execution.

15 After script execution, summary details on each server and the total job display. An ID number at the top of the page identifies the job.

- To see script output, summary details and information about errors that might have occurred, click View Details.
- For each server, you can click Download to obtain a zip file (`results.zip`) that contains script execution output and error data.

  The data is in two files called `stdout` (output data) and `stderr` (error data), in a directory named <*servername*>. For example, for a server named m0043, the output and error files would be located as follows: `m0043/stdout` (output file for server named m0043) and `m0043/stderr` (error file for server named m0043).

16 Click Close to exit the wizard.

How to Create and Execute an Ad-Hoc Script

An Ad-Hoc script is written (or uploaded) and then executed without being stored in the Opsware System.

Ad-Hoc script setup and execution are handled by the same wizard that steps the user through the processes used for stored scripts. However, the initial steps differ because Ad-Hoc script creation is integrated with execution activities.

Perform the following steps to create or upload an Ad-Hoc script, and then execute the script:

1. Click the Run Distributed Script link in the Tasks area of the Opsware Command Center home page to launch the Run Distributed Script Wizard.
2. On the Overview page, click Define Ad-Hoc Script ➤ Start. The Define Script page appears, as Figure 10-10 shows.

Figure 10-10: The Define Script Page

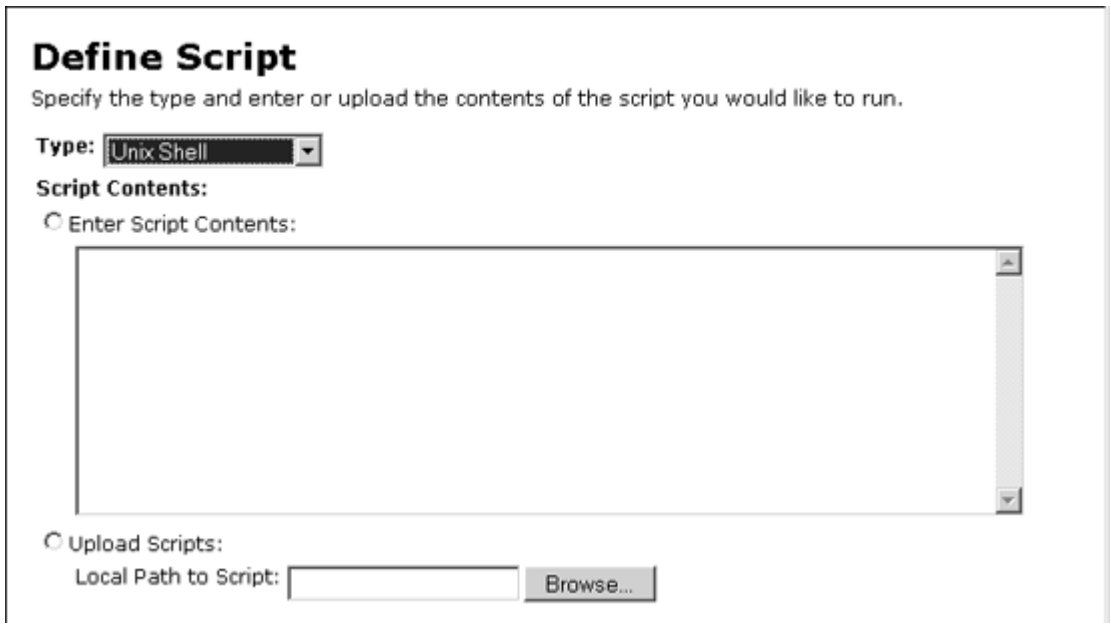

3. Select the type of script that you are creating (Unix shell, Windows VBScript, or Windows .BAT).
4. For Script Contents, indicate if you are entering script contents or uploading the script:
   - To upload a script, click Upload Script. In the Local Path to Script box, either manually enter the path to the script or click Browse to locate the path.
   - To create a script, click Enter Script Contents and enter the script in the text box.

The Scripts editor does not recognize tabs and its functionality is browser-dependent.

5. Click Next to continue (or click Previous to return to the previous step).
6. On the Select Servers page, browse or search for the server or servers to use.
   - To browse, select the *Browse* tab to obtain a listing of servers available to you. At the top of the page, use the Customers and Facilities filters to narrow your selection.
   - To search, select the *Search* tab. At the Search page, indicate search criteria by selecting the appropriate check box. Default search criteria might be available to help narrow the search for specific customers and facilities.

   Only servers are listed that use the operating system appropriate for the type of script that you want to run. For information about a server, click the name of the server.
7. From the list of servers, select one or more servers for script execution.
8. Click Next to continue (or click Previous to return to the previous step).
9. On the Specify Options page, enter the following runtime information:
   - Specify the runtime user. If you have the appropriate permission, you can execute the script as root or local system without entering a password. Otherwise, enter an appropriate username and password.
   - Specify if you want to keep or discard script output. Only a portion of script output is saved (the amount of saved script is configured by the Opsware administrator).
   - Enter a script timeout value in minutes.

     This is the amount of time a script has to complete execution activities on a server. If the script is not finished when the timeout value is reached, the script is stopped by the Opsware System and a script error occurs.

     Select a timeout value that is greater than the time required for execution to complete.
   - Enter command line parameters.

     Use the same syntax as when you enter a script on a command line. For Unix scripts, use Bourne (sh) shell syntax. For Windows scripts, use `cmd.exe` parameter syntax.

10 Click Next to continue (or click Previous to return to the previous step).

11 On the Confirm Settings page, review your script execution settings before you proceed to run the script. The Script to be run panel provides information about your script and the Servers panel displays the servers selected for script execution.

If you discover that script or server changes are needed, click the Previous button as many times as you need to return to the Script or Select Servers page.

12 On the Schedule and Notify page, you have the following options:

- Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
- Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.
- Run Script: While the script executes, progress information displays about the total run and each server. When execution ends, summary information displays. After execution, script contents, output data, and error data can be immediately reviewed and downloaded.

If you select to run the job at that time, a progress bar appears that shows the progress of the script execution.

13 After script execution, summary details display for the total job and each server. An ID number at the top of the page identifies the job.

- To see script output, summary details and information about errors that might have occurred, click View Details.
- For each server, you can click Download to obtain a zip file (`results.zip`) that contains script execution output and error data.

  The data is in two files called `stdout` (output data) and `stderr` (error data), in a directory named <*servername*>. For example, for a server named m0043, the output and error files would be located as follows: `m0043/stdout` (output file for server named m0043) and `m0043/stderr` (error file for server named m0043).

14 Click Close to exit the wizard.

Script Execution Results: Tasks and Procedures

This section provides information about script execution results and contains the following topics:

- Viewing Execution Results Immediately After Script Execution
- Viewing Execution Results Stored in the Opsware System

During execution, the current status of script execution events at each server displays on the Run Script page. When script execution activities finish, the page immediately displays a summary of execution results for each server and for the entire run. See Figure 10-11.

Figure 10-11: Script Execution Progress Page

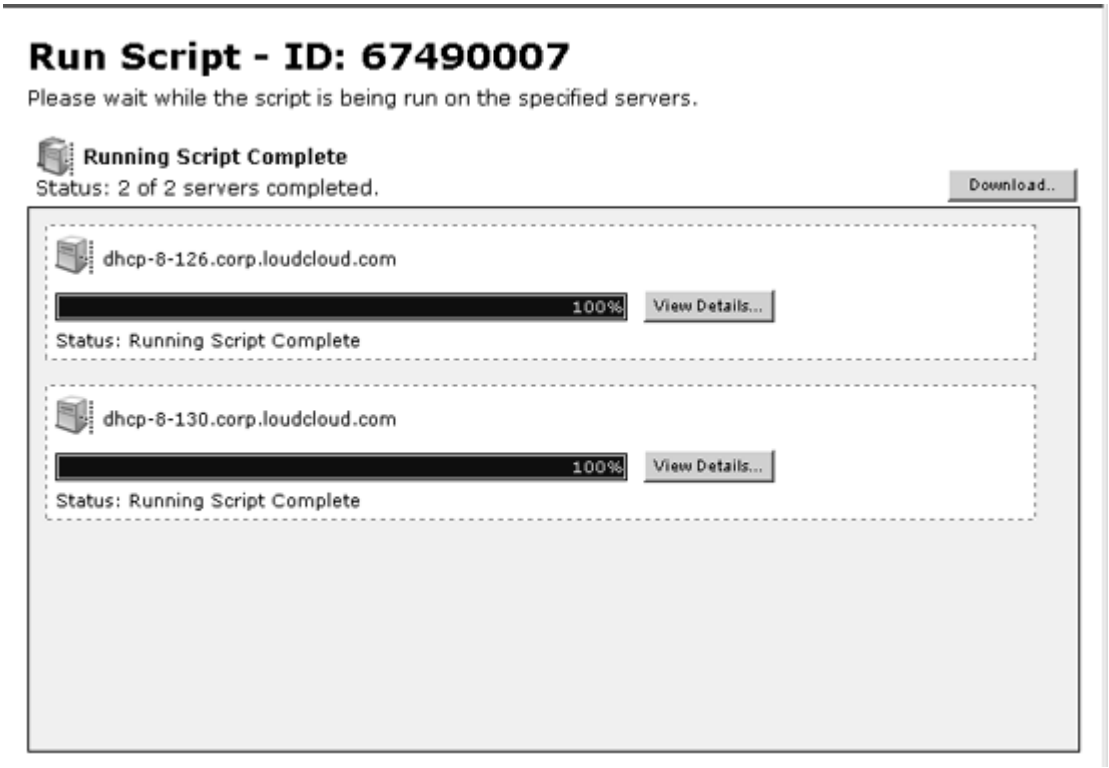

Script execution data and information are stored in the Model Repository and are later accessible through the My Jobs feature for the user who performed the execution.

Viewing Execution Results Immediately After Script Execution

After a script runs, execution data and information for a specific server are available by clicking the View Details button for that server. Information that is displayed includes script execution output, errors (if there are any) and summary data (such as the script name and contents, and the date of the run). See Figure 10-12.

Figure 10-12: Execution Results After Script Execution

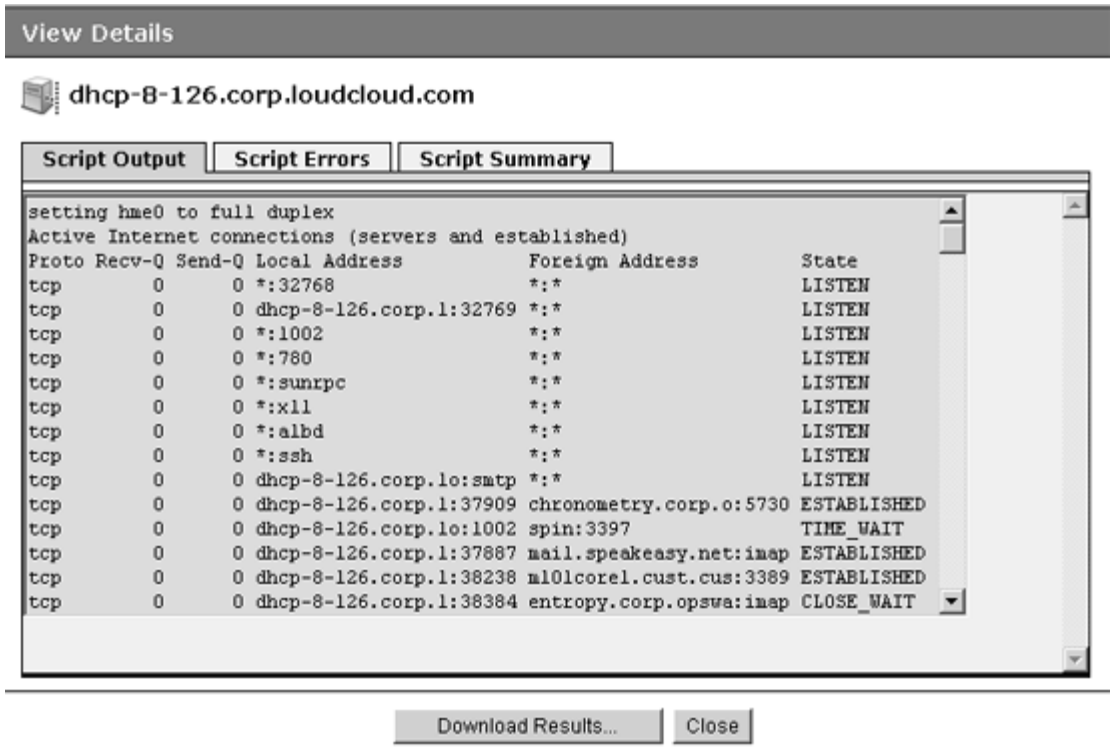

A zip file (`results.zip`) is available that contains standard script execution output and error data in two files (`stdout` and `stderr`), which are located inside a directory name <*`servername`*>. You can download the zip file by clicking the Download button.

Viewing Execution Results Stored in the Opsware System

After script execution, script data results are available through the My Jobs feature.

The My Jobs list is available either at the My Jobs panel (Opsware Command Center home page) or the My Jobs page (opened by clicking the My Jobs link at Opsware Command Center's navigation panel). In the list, executed scripts are identified from other types of Opsware System jobs by the name Run Script. See Figure 10-13.

Figure 10-13: My Jobs Panel on the Opsware Command Center Home Page

You can view Information about a specific executed script (started on a particular day and time) by clicking the Run Script link for the specified day and time. See Figure 10-14.

Figure 10-14: Execution Results for a Particular Script Execution Event

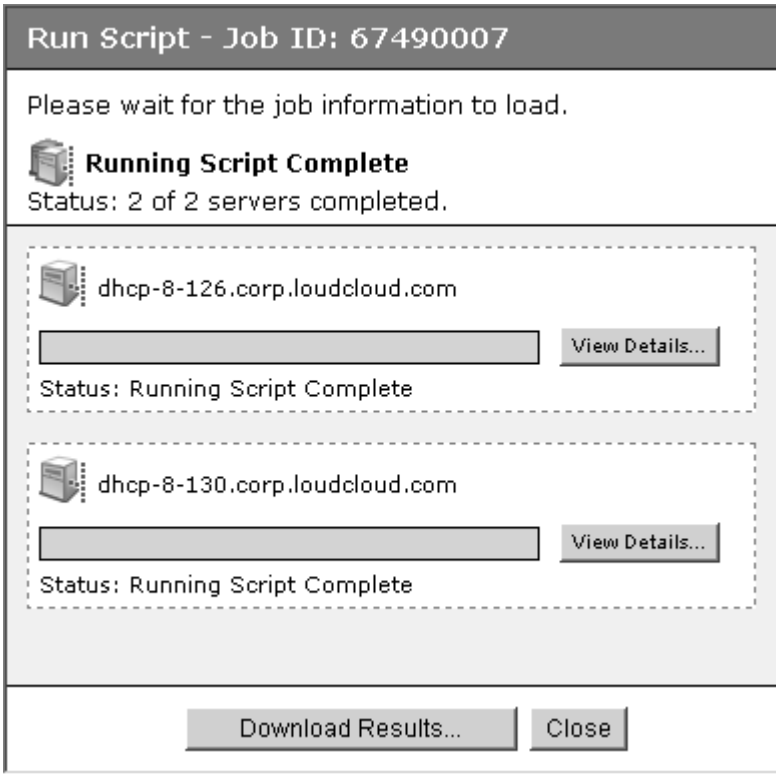

Script results data are held in the Model Repository for a limited amount of time (30 days, for example). The Opsware administrator can configure this time period.

Subsystem Error Resolution

The following section describes the events that might occur during Script Execution Subsystem activities and provide suggestions for resolving the problems. This section contains the following topics:

- Resolving Script Uploading Errors
- Resolving Script Timeout Events
- Investigating Script Non-Zero Return Codes
- Investigating Server Authentication Errors

- Investigating Partial Executions

Resolving Script Uploading Errors

- If errors occur when a user uploads a script, the Opsware Command Center error pages display information about the event.
- The maximum size for a script is 4 MB. Verify that the script being uploaded does not exceed this size.
- If communication with the Opsware Command Center is lost during the script upload process, the script should be uploaded again.

Resolving Script Timeout Events

When a script runs on an Opsware-managed server, if the script hangs or is not finished before the script timeout value is reached, an Opsware System error occurs and the script is stopped on the server.

The script timeout value is entered on the Specify Options page of the Run Distributed Script Wizard. The timeout value is the amount of time before a script times out.

Check that the length of time it takes for script execution does not exceed the current timeout value. Creation of debug output can also be added into the script for troubleshooting purposes. If the timeout event is due to the script hanging, then further examination of the script and server should occur.

Investigating Script Non-Zero Return Codes

If execution of a script on an Opsware-managed server returns a non-zero code, an error is reported. (A zero return code indicates successful – for example, normal – operation.) Information about the error is available to the user immediately following the error event and script execution. Information is also available after execution through the My Jobs feature.

Depending on how a script is written, a non-zero return code might indicate a fatal error.

Investigating Server Authentication Errors

If an authentication error occurs at an Opsware-managed server, verify that the user correctly enters the password. Also, verify that the correct username and password are being used.

If authentication fails during execution of a script, an Opsware System error is raised for the server running the script. Information about the error is displayed in the Opsware Command Center for that server during and immediately following script execution. Script execution information is also available through the My Jobs feature.

Investigating Partial Executions

If a script runs on only some of the selected servers, this might be due to an Opsware Command Engine failure or because someone else has locked the managed servers by, for example, running another Opsware System task that uses the servers. Or, the servers might be unreachable.

Opsware Custom Extensions

Opsware Professional Services can extend Opsware functionality for customers by creating custom extensions to the Opsware System. Opsware Custom Extensions (which are custom Command Engine scripts) extend Opsware System functionality to cover specific customer needs.

In the Opsware System, the Command Engine is a system for running distributed programs across many servers (utilizing the Opsware Agents running on the servers). Opsware System features, such as the Code Deployment Subsystem, use Command Engine scripts to implement part of their functionality

The Custom Extension feature is accessed through the Run Custom Extension Wizard. This Wizard allows a user to choose a custom extension to run, specify necessary input data for the extension, validate the data required to run the extension, run the extension, and view or download the results from the job. When a user runs a custom extension, the job shows up in My Jobs.

When a custom extension is added to one facility, it is automatically propagated to the other facilities in the multimaster mesh of Opsware facilities.

To access the Run Custom Extension Wizard, users must be assigned to a user group that has the permission Wizard: Opsware Extension. When users have this permission, they can run any custom extensions on the servers they have access to in the Opsware Command Center.

Running a Custom Extension

Perform the following steps to run a custom extension:

1. From the Opsware Command Center home page, click the Run Custom Extension link in the Tasks panel.

   Or

   From the navigation panel, click Servers ➤ Managed Servers. The Managed Servers list appears. Select the servers on which you want to run a custom extension and choose Run Extension from the Server menu.

   The Run Custom Extension Wizard appears.

2. Click the Start button to begin.

3. Select the custom extension that you want to run and click the Next button.

   If you have already selected servers from the Managed Servers list, you must ensure that the custom extension that you select can run on the operating systems of the selected servers. Otherwise, an error message appears in this page and you cannot proceed.

   Some custom extensions do not require that you select servers from the Server list. For example, the extension might prompt you in the Specify Settings step to enter server hostnames in a text box. If you already selected servers from the Servers list, an error message appears in the page.

   If you have not already selected servers from the Managed Servers list, the Select Servers page appears.

4. If prompted, select the servers on which you want to run the custom extension and click the Next button. You can find the servers that you want to run a custom extension on by browsing the list or by searching.

   See "Searching with Advanced Search" on page 48 in Chapter 2 for information about how to use the advanced search features in the Opsware System.

The Specify Settings page appears. See Figure 10-15.

Figure 10-15: The Specify Settings Page of the Run Custom Extension Wizard

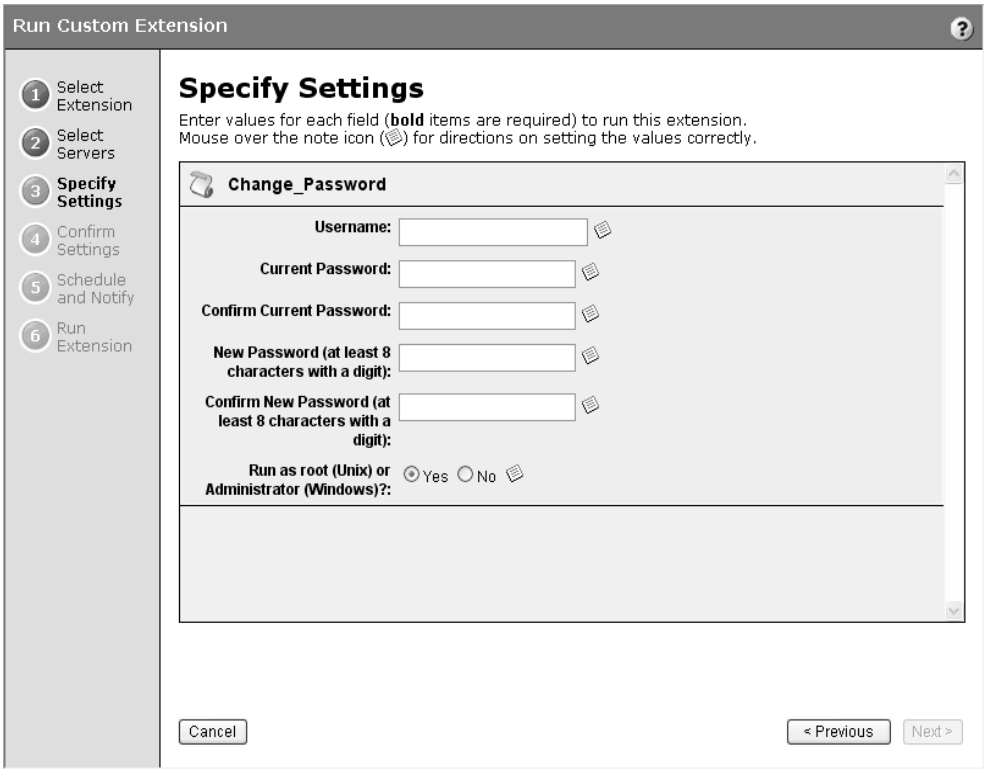

5 Specify the settings for the custom extension and click the Next button.

The settings that appear in the page are unique to the custom extension that is being run. For information about what data to enter in a field, move your mouse pointer over a Note icon.

The Confirm Settings page appears. See Figure 10-16.

Figure 10-16: The Confirm Settings Page of the Run Custom Extension Wizard

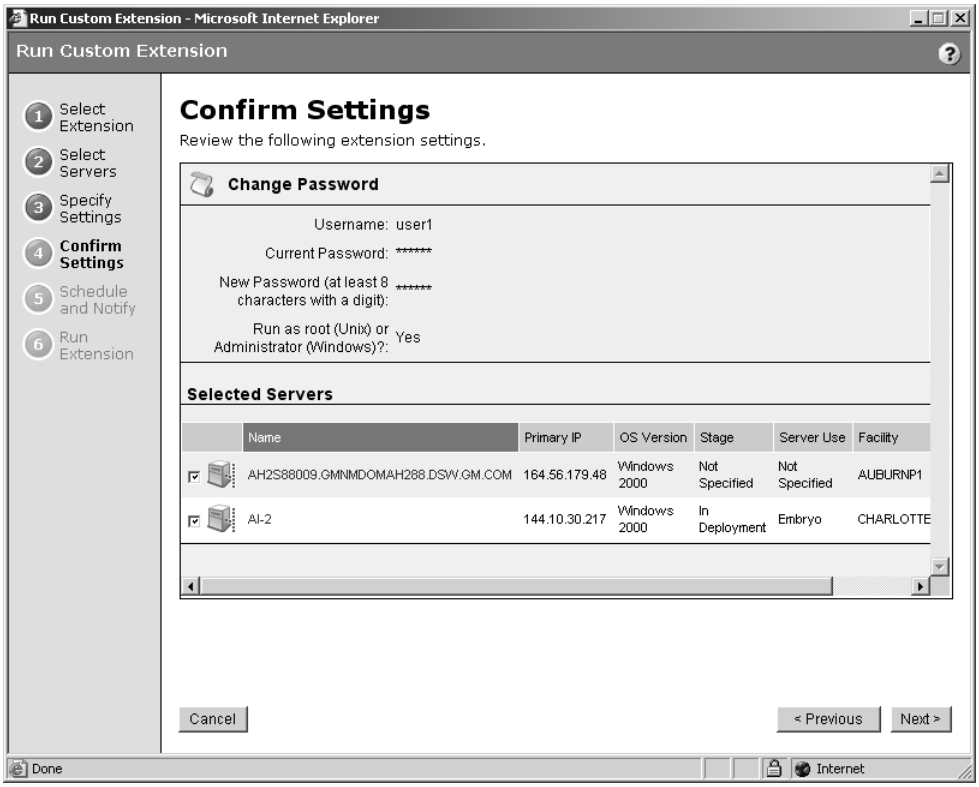

6. Review the values that you entered and the servers that you selected on which to run the custom extension. (You can remove servers from the list by deselecting their check boxes. The list displays the first nine servers on which the custom extension will run. Click the "Show remaining servers" link to display the complete list of select servers.)

7. Click Next.

8. On the Schedule and Notify page, you have the following options:

   - Schedule: Choose either Run Now to execute the operation immediately, or choose Specify Time to schedule the operation for a later time.
   - Notify: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus symbol next to the Recipients field.

- Run Extension: While the custom extension executes, progress information displays about the total run and each server. When execution ends, summary information displays. After execution, script contents, output data, and error data can be immediately reviewed and downloaded.

9 Click Run to start the Custom Extension Wizard.

If you selected to run the job at that time, a progress bar appears for the servers on which the extension is running that shows the progress of the job. Depending upon how the custom extension was written, you will see the progress displayed in one of two ways:

- If the custom extension was written to show progress and status for individual servers, you will see individual progress bars for each server. When the custom extension has finished running, you will see the View Details button. You can click the View Details button to display detailed error information.
- If the custom extension was written to show progress for all servers, then you will see a single progress bars for all servers. When the custom extension has finished running for all servers, you will see the View Details button. You can click the View Details button to display detailed error information.

If an error occurs while the custom extension is running on servers, the progress bar moves to 100% and an error message appears below the bar.

For any servers where the custom extension has failed to run, you will have to either re-launch “Run Custom Extension” from the home page & pick the servers they want to try again -OR- choose those servers from a Managed Server/My Servers/Server Search list & choose the menu item “Run Custom Extension”.

10 (Optional) When the custom extension finishes running, you can click the View Details button to see the results.

The Custom Extension Results window appears. The tabs in the window can vary depending on how the custom extension was implemented. To download the results to a file, click the download link. When you are done viewing the results, click the Close button to close the window.

11 Click the Close button to end the wizard.

Closing the wizard does not stop the custom extension if it is still running. After you close the wizard, you can view the progress of the running custom extension by viewing My Jobs (accessible from the Home page or the navigation panel). Each custom extension job is identified with the name Run Custom Extension. Click the name link to identify which extension was run. See Figure 10-17.

Figure 10-17: Run Custom Extension Jobs

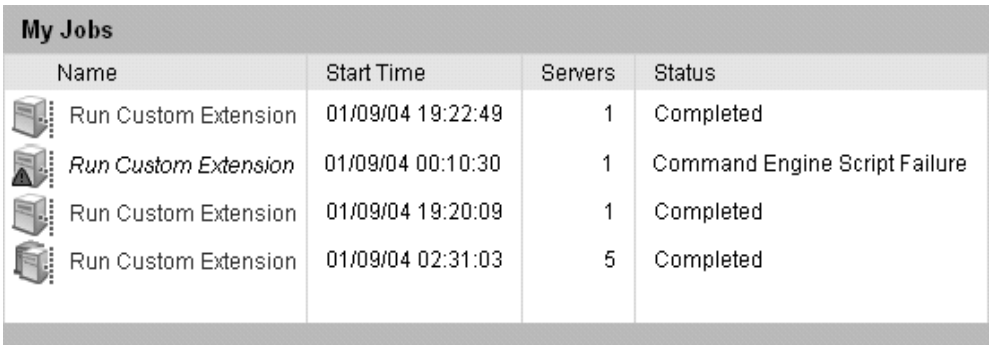

See "Scheduling Server Management Tasks" in Chapter 2 for more information about the My Jobs feature.

Chapter 11: Automated Configuration Tracking

IN THIS CHAPTER

This chapter explains how to use the Opsware System to monitor configuration files and configuration databases and contains the following topics:

Automated Configuration Tracking Overview

The Automated Configuration Tracking feature of the Opsware System allows you to monitor critical configuration files and configuration databases. When the Opsware System detects a change in a tracked configuration file or configuration database, the system can perform a number of actions, including backing up the configuration file or sending an email to a designated individual or group. You use configuration tracking policies to identify the files to be tracked and actions to be taken when change is detected.

The Automated Configuration Tracking feature is designed for flexibility. For example, you can set configuration tracking defaults for a node that contains a particular software application, and all servers attached to that node automatically get those defaults. You can also quickly deploy the common configuration tracking defaults to a large number of servers in your Opsware managed environment or create a specific policy for a single server.

Automated Configuration Tracking allows you to recover from many problems caused by changes to configuration files. Using Automated Configuration Tracking, you can identify which tracked configuration files have changed, thus helping you identify the potential source of a problem. If you back up your configuration files with the Automated Configuration Tracking feature, you can quickly restore the changed configuration files to a previous version.

You can also view a detailed history of all backup activity. This history includes a list of all tracked files that have been backed up and what types of backups occurred. If the backed up configuration files are text-based, you can download the files from the backup history and compare them to determine what specific changes have been made.

The Automated Configuration Tracking feature is not a general-purpose backup solution. Automated Configuration Tracking is designed to monitor text-based configuration files and specific types of configuration databases. The number and size of files that can be monitored on any managed server is limited. See "Configuration Tracking Policy Limits" on page 502 in this chapter for information about for more information.

Configuration Tracking Policies

You use configuration tracking policies to specify which files and configuration databases to monitor for change and what actions to take when change is detected.

A configuration tracking policy consists of one or more configuration tracking policy entries. You create one entry for each target that you want to track. The target specifies the configuration file, directory of configuration files, or configuration database that you want to track. In the case of directories and files, the target is the fully qualified path of the directory or file (unless you are using wildcards). You then specify the action to be taken when change to the target is detected.

Figure 11-1 shows the interface that you use to create a configuration policy entry.

Figure 11-1: Add Entry Interface

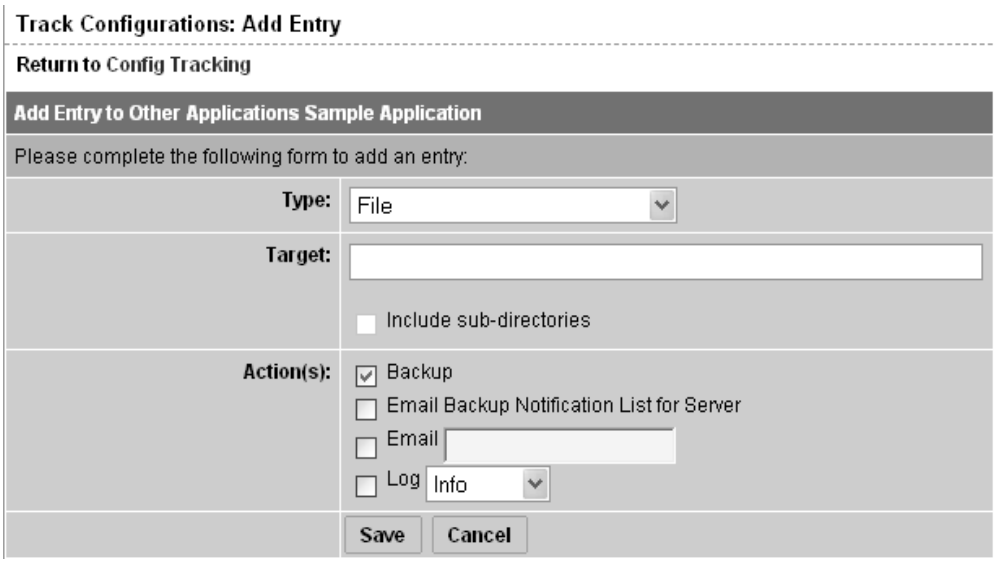

The combination of all the individual tracking policy entries for a server is referred to as a server's configuration tracking policy, and the combination of all the tracking policy entries for a node is referred to as a node's configuration tracking policy (for example, the tracking policy is the combination of tracking policy entries).

Supported Types of Configuration Files and Databases

This section provides information on supported types of configuration files and databases and contains the following topics:

- File Types Supported
- Types of Actions Performed
- Types of Backups Performed
- Configuration Tracking Policy Targets and Wildcards
- Special Considerations for Directory and Wildcard Targets
- Email Automated Configuration Tracking and Logging Actions
- Creating the Email Notification List
- Configuration Tracking Policy Limits
- Methods for Creating Tracking Policies
- Deploying Tracking Policies

File Types Supported

You can use the Automated Configuration Tracking feature with the following types of files:

- Text-based configuration files
- The COM + Registration Database (Windows 2000)
- The IIS Metabase
- Windows Registry keys

Types of Actions Performed

The Opsware System can perform the following actions when change is detected in a tracked configuration file or configuration database:

- Back up
- Send email to addresses specified in the policy entry
- Send email to a designated notification group specified by a custom attribute
- Create an entry in the server's standard system log (The syslog on Unix servers and the event log on Windows servers)

Types of Backups Performed

If you selected backup as the action for a tracked configuration file, the two general types of backups that can occur are incremental backups and full backups.

Incremental Backups

During an incremental backup, only targets that have changed since the last backup (and that have been selected to be backed up) are backed up.

An incremental backup occurs automatically when the Opsware Agent detects change in a tracked file that is selected to be backed up. (The Opsware Agent checks for change every four hours.)

Incremental backups also occur before and after you restore a previous version of a backed up configuration file to a server. These backups allow you to roll back the restored files.

Full Backups

During a full backup, all tracked configuration files that were selected to be backed up are backed up, not just the files that have changed.

Once a week, the Opsware Agent on a server checks to see if any files have changed since the last full backup. If any files have changed, the Opsware System performs a new full backup. If no files have changed, the full backup does not take place.

You can also force the Opsware System to perform a full backup on a server by selecting the *Perform Manual Backup* option. (See "Performing Manual Backups" on page 523 in this chapter for more information.)

See "The Backup History" on page 524 in this chapter for information about Backup types.

Backups are stored in the Software Repository until you delete them. You should delete old backups periodically, especially if you are backing up a large number of files that change frequently. See "Deleting Backups" on page 530 in this chapter for information about the procedure for deleting backups.

Configuration Tracking Policy Targets and Wildcards

If you selected a file or directory as the target type, the target can include wildcards (* to match 0 or more characters, and ? to match a single character). Wildcards cannot be used with other types of targets.

Wildcards used in file targets produce different results than wildcards used in directory targets.

- If you use wildcards in a file target, the Opsware System searches for files that match the pattern that you specify.
- If you use wildcards in a directory target, the Opsware System searches for directories that match the pattern that you specify. The contents of all the matching directories are tracked.

if you selected the "include subdirectories option," the contents of all the subdirectories of all matching directories are also tracked.

Wildcards can be used in any part of the path that you specify in the target (except for the drive letter.) Wildcards used in the target field behave exactly as they do on the selected operating system.

Special Considerations for Directory and Wildcard Targets

In most cases, the Automated Configuration Tracking feature is used to track specific files. Specific files are tracked when you select the File target type and supply the fully qualified path of the filename without using wildcards. (Tracking configuration databases, such as the Windows COM+ Registration Database, also counts as tracking specific files.)

Because the tracking feature is designed to monitor files that you are specifically interested in (for example, configuration files), tracking specific files is the recommended way to use the feature. When you monitor a specific file, the Opsware System is able to keep a more complete record of the file; the Opsware System can, for example, note that the file does not exist on the server.

If you do not know the fully qualified path of a configuration file, you can choose to monitor directories or use wildcard targets. This option is useful in some circumstances, such as when an application dynamically creates configuration files and it is not possible to know the file names in advance. The next sections explain special considerations that apply if you monitor directories or use wildcard targets.

Tracking the Contents of a Directory

The following conditions apply when you select the Directory target type. These conditions apply whether or not you use any wildcards in your target.

- You cannot use the Automated Configuration Tracking feature to delete files that were created in monitored directories. Because the Opsware System was not tracking the file specifically, the Opsware System did not note the file's absence and cannot restore the file to a state before it existed.
- When a specific tracked file is deleted, your selected action is triggered for the file. If a file is deleted inside a monitored directory, the action is triggered only for the directory object and not for the file itself. For example, you have selected the email action, you would receive email that the directory has changed, but you would not receive email about the specific file. (You do receive email about files created inside monitored directories.)
- If you are monitoring a directory tree (by selecting the include subdirectories option), actions are triggered for the file, the file's directory, and all of the parent directories up to the directory specified in the target.

Tracking Files through Wildcard Targets

The following conditions apply only when you select the File target type *and* use wildcards in the target:

- When a file is created that matches a wildcard target, the creation of the file is noted and your selected action is triggered. The fact that the file did not previously exist, however, was not noted, and no entry was created in the backup history about the file's absence. This file cannot be deleted by using the Automated Configuration Tracking feature.

- When a file tracked through a wildcard target is deleted, no action is taken.

Directories Tracked by Wildcard Targets

The conditions discussed in this section apply to directory objects when you select the Directory target type and use wildcards in the target.

The contents of directories tracked by wildcard targets are subject to the same conditions as the contents of directories that are fully specified (for example, not found as a result of a wildcard search). The directory objects tracked through wildcard targets, however, are not tracked in the same way as directory objects with fully specified targets. The differences are the following:

- If you track a specific directory and the directory does not exist on a server, the absence of the directory triggers the action you selected. An entry is made in the backup history about the absence of the directory and this entry can be used to delete the directory. No such entries are made for directories tracked through wildcard targets.
- The deletion of a directory found by a wildcard target does not trigger an action. The creation of a directory that is found by a wildcard target, however, does trigger an action.
- If you are monitoring a directory tree through a wildcard target and any subdirectory is deleted, your selected action is triggered for all the directory's parent directories, up to the directory specified in the target. When subdirectories are created, however, your selected action is triggered both for the top-level directory and for the specific subdirectories.

Email Automated Configuration Tracking and Logging Actions

You can choose to have email sent when a monitored target changes. The following example shows the text of an email generated when a tracked file changed.

```
From: <configurationtracking@yourcompany.com>
Date: Thu Jan 16, 2003  5:40:11 PM US/Pacific
To: <joe@yourcompany.com>
Subject: athena.cust.com: Configuration Tracking CHANGE
notification
Configuration Tracking has detected a CHANGE event
Host: athena.cust.com
Object: /db/file1l1
```


The email specifies the name of the server and the name of the object that changed. The object can be a file, a directory, or a configuration database.

As discussed in the previous section, if you are monitoring a directory target, you receive email about the directory itself and about changes to the files in the directory (except when a file is deleted.) For example, if three new files are created in a directory, you would receive four emails, one for the directory and three for the new files.

If you selected the logging action, an entry is made to the server's standard system log when a change is detected. You select the type of log entry that you want to have written. The Opsware System uses three standard entry types:

- Info
- Warning
- Error

How the entry types are identified is system-dependent. For example, on most systems Warning entries are identified by the word warning. In some systems, however, a number is used to identify the log entry type.

The following example shows a warning log entry written on a Solaris Server:

```
Jan  8 00:05:25 athena.cust.com Configuration Tracking:
[ID702911
local0.warning] Configuration Tracking: /other/otherfile1 :
Event CHANGE occurred
Jan  8 00:05:25 athena.cust.com Configuration Tracking: [ID
702911
local0.warning] Configuration Tracking: /other/otherfile1 :
Event CHANGE occurred
```

Creating the Email Notification List

Sending email to a server's backup notification list is one of the actions that you can select in a tracking policy entry. The email notification list is a list of email addresses that you define for the following custom attribute:

```
backup_notification_email
```

This attribute can be set on the server itself or on the customer to which the server is attached. Setting the attribute at the customer node level allows you to use the same email notification list for all servers that belong to the same customer (assuming that these servers have all been attached to the same customer and do not have the backup_notification_list set on the servers themselves).

Search Order for Email Notification List Attribute

On a server that has a policy that includes the Email Notification List for Server action selected, the server searches for the backup_notification_email attribute in the following order:

- Server
- Customer

After the custom attribute is found, its value is used (for example, the email address of the notification list) and the search is concluded. If, for example, the backup_notification_email attribute is set on a server, the server's email notification list is used, even if the server is assigned to a customer that has a different backup_notification_email attribute.

Format of Email Notification List

The notification list can contain multiple email addresses. The email address must be formatted as a comma-separated list.

Configuration Tracking Policy Limits

Because Automated Configuration Tracking is not a general-purpose backup solution, the Opsware System enforces limits on the number of objects that can be monitored. The limits help keep backup volume from growing too large; excessive backup volume can degrade system performance.

You should also periodically delete backups; See "Deleting Backups" on page 530 in this chapter for more information.

The limits apply to the number and size of objects that can be monitored. In this context, *objects* are files or configuration databases and the directories that make up the path to a file. For example, if you set up a policy entry to track the contents of the /etc/init directory and the directory contains 10 files, the policy entry is tracking 12 objects (the 10 files plus the /etc and /init directories). Similarly, if you set up a policy entry to track the /etc/system file, the policy entry is tracking 2 objects (the /etc directory and the system file).

The following limits apply to the number of objects that can be tracked and to the size of files that can be tracked:

- No more than 800 objects can be tracked on a single server (for example, the server's aggregate tracking policy, which is the combination of all its tracking policy entries, cannot cause more than 800 objects to be tracked.)

- The total number of objects tracked by a single tracking policy entry cannot be larger than 250. (A tracking policy entry is a single item in a tracking policy.)

  For example, if you create a tracking policy entry to monitor the contents of the /etc/init directory, and the directory contains 250 files, you have exceeded the limit for a single tracking policy entry. Because the /etc and /init directories count as objects, the total number of objects monitored by this policy entry exceeds the limit of 250.

- Files greater than 2 megabytes cannot be monitored.

The target you specify in a tracking policy entry can contain wildcards. Exercise caution in using wildcards, because wildcards can cause a large number of files to be monitored.

If you exceed any of these limits, you receive an error message when you attempt to deploy the server's tracking policy, and the deployment fails. See "Deploying Tracking Policies" on page 504 in this chapter for more information.

If a policy exceeds any of these limits after the policy is deployed on a server, no further actions are triggered on that server. Backups, for example, stop taking place. When any limit is exceeded after a policy is deployed, a warning email is sent to the following addresses:

- An administrator's email address that is specified during installation. The email address is referred to as the error email address.
- The server's backup notification email list.

If the server does not have a backup notification list assigned to it, the warning message is sent only to the error email address.

Methods for Creating Tracking Policies

You can create tracking policies in two ways:

- By using Opsware System nodes (See "Node-Based Tracking Policies" on page 505 in this chapter for more information.)
- By creating custom tracking policies for a selected server or set of servers (See "Customizing Configuration Tracking Policies" on page 514 in this chapter for more information.)

The preferred way to create a tracking policy is by using Opsware System nodes. Among other functions, an Opsware System node specifies what software should be installed on the servers that are attached to that node. For example, if a server is assigned to a Sun ONE Web Server node, the software packages associated with that node are installed on the server when the server is reconciled.

Because the node specifies which software packages to install, the tracking policies for the software packages' configuration files are usually created as part of the node.

Deploying Tracking Policies

Tracking policies are not deployed to your servers until you perform a configuration tracking reconcile. (This is not the same as software reconcile). See "Reconciling a Node's Configuration Tracking Policy" on page 513 in this chapter for information about the procedures for performing configuration tracking reconciles. See "Reconciling Customized Tracking Policies" on page 522 in this chapter for information about the procedures for performing configuration tracking reconciles.

How Change Is Detected

All servers that the Opsware System manages have an Opsware Agent installed on them. On servers that use Automated Configuration Tracking, every four hours the Opsware Agent inspects the configuration files and databases that you select to track.

The Opsware Agent computes an MD5 checksum to determine if the contents of a tracked file have changed. (Any change to the contents of the file results in a change to the MD5 checksum.) If the contents of the file have changed, the action that you specify in the tracking policy is performed. For example, if you create an entry in your tracking policy for the /etc/passwd file and select backup as the action to be taken, the file is backed up when the Opsware Agent discovers a change in the /etc/passwd file.

The creation or deletion of a tracked file (or files inside a tracked directory) also counts as change and triggers a policy's action. (There are some exceptions; See "Special Considerations for Directory and Wildcard Targets" on page 498 in this chapter for more information.)

Changes to the properties of a tracked file or directory (such as changes to permissions or timestamps) do not count as a change. When a file or directory is backed up, however, its properties are backed up as well.

The first time that a tracking policy is deployed, all targets are considered changed. The Opsware Agent is encountering the files for the first time, and all of the policy's actions are triggered.

Node-Based Tracking Policies

This section provides information on node-based tracking policies within the Opsware System and contains the following topics:

- Node-Based Tracking Policies Overview
- Creating Node-Based Policy Entries
- Viewing a Node's Tracking Policy
- Editing a Node's Configuration Tracking Policy
- Editing a Node's Configuration Tracking Policy Entry
- Disabling a Node's Configuration Tracking Policy Entries
- Deleting an Entry in a Node's Configuration Tracking Policy
- Re-enabling a Tracking Policy Entry

Node-Based Tracking Policies Overview

You create a node's tracking policy by creating one or more tracking policy entries. After you create a node's policy, you can edit the individual policy entries, you can disable individual policy entries, and you can add new policy entries.

You create nodes by creating child nodes from parent nodes. The child nodes inherit the software policies of their parents. The same principle applies to tracking policies. If the parent node has a tracking policy defined for it, all the node's children inherit the tracking policy. If you make changes to a node's tracking policy, the node's children also inherit the changes.

You can, however, disable all or part of any inherited tracking policy in a child node.

Creating Node-Based Policy Entries

Perform the following steps to create the tracking policy for an existing Opsware System node:

1. From the navigation panel, click Software and then select the relevant type of software node (for example, patches, applications, and so forth), as Figure 11-2 shows.

Figure 11-2: Selecting Node Types

2. Navigate to the node for which you want to create a tracking policy.
3. Click the Config Tracking tab.

4 Click Add Entry. The Track Configurations: Add Entry page appears, as Figure 11-3 shows.

Figure 11-3: Track Configurations: Add Entry

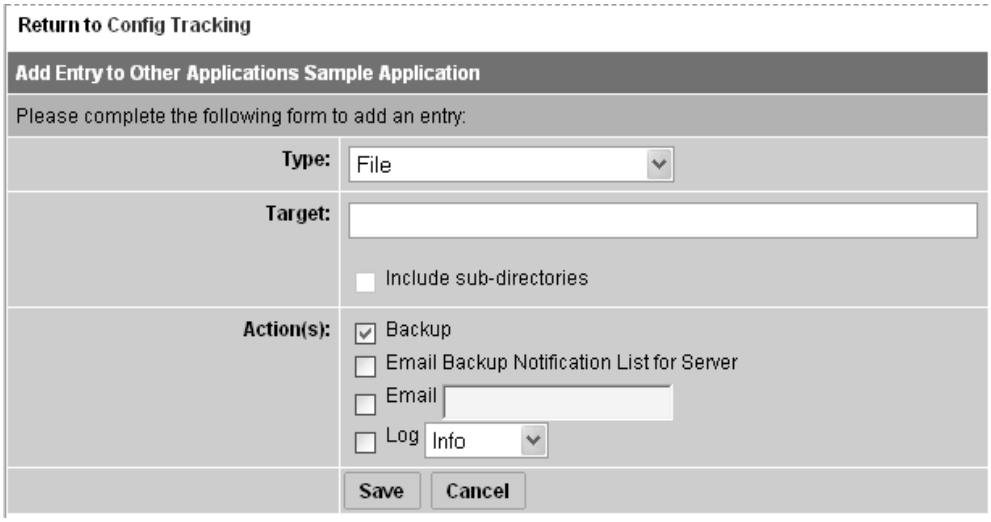

5 Select the type, define the target (when appropriate), and select the actions that you want to take place when a change in the configuration file or database is discovered. Table 11-1 describes each of the selections that you must make.

Table 11-1: Add Entry: Configuration Tracking

FIELD	DESCRIPTION
Type	**File**: monitor the file specified in the target field **Directory**: monitor all the files in the directory specified in the target field. The following types are available only for Windows servers. **Windows Registry**: specify key in target field **IIS Metabase**: entire Metabase is monitored; do not specify target **COM + Registration Database**: entire Registry is monitored; do not specify target.

Table 11-1: Add Entry: Configuration Tracking

FIELD	DESCRIPTION
Target	If you selected the file type, specify the full path (including the drive letter on Windows servers) of the file that you want to monitor. If you selected the directory type, specify the full path of the directory (including the drive letter on Windows servers) that you want to monitor. You also have the option of monitoring subdirectories. (Select the "Include Subdirectories" check box.) If you select the file or directory type, you can use wildcards in the target. (See "Configuration Tracking Policy Targets and Wildcards" on page 498 in this chapter for more information.) If you selected the Windows Registry type, specify the Windows registry key. This key and all its subkeys are backed up. Use standard syntax for Windows Registry keys (For example, `HKEY_LOCAL_MACHINE\SOFTWARE`) If you selected the IIS Metabase or COM + Registration Database type, do not specify the target.
Actions (you can select multiple options)	Backup: back up the specified file, directory, COM + Registration Database, Windows Registry keys, or IIS Metabase. Email Backup Notification List for Server: send an email to the backup notification list for the selected server. See "Creating the Email Notification List" on page 501 in this chapter for more information. Email: send an email to the address or addresses specified in this field when a change is detected. Use a comma-separated list for multiple email addresses. (Not available for Windows Registry.) Log: add an entry to the server's system log when a change is detected. You can choose to write the following types of log entries to the server's system log: Info Warning Error

6 Click Save to add the entry to the tracking policy.

7 If you want to continue to add entries to the tracking policy, click Add Entry and repeat this procedure.

You must perform a Configuration Tracking reconcile to deploy the tracking policy to the appropriate servers.

Viewing a Node's Tracking Policy

Perform the following steps to view a node's tracking policy:

1 From the navigation panel in the Opsware Command Center, click Software then select the type of software to which the node belongs.

2 Navigate to the node whose tracking policy you want to examine.

3 Click the Config Tracking tab. The entries that make up the tracking policy for the node display. You can verify that an entry in a tracking policy was inherited by looking at the Inherited? field, as Figure 11-4 shows.

Figure 11-4: Viewing a Node's Configuration Policy

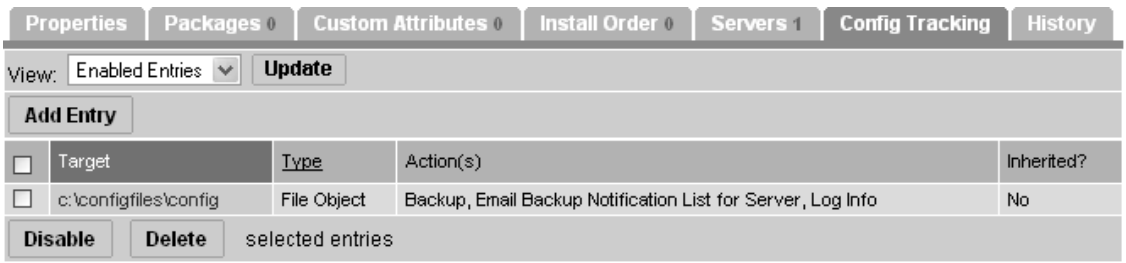

Tracking policy entries are identified by target. If a node inherits two or more policy entries from its parents that have the same target, these entries display as a single entry when you view the child's tracking policy. The single entry, however, combines all of the actions that you selected for the target.

If the child also has a policy entry for the same target, the child's policy displays as a separate entry. For example, if a child node inherits three policy entries from its parents for the /etc/passwd target and the child also has its own policy entry for /etc/passwd, two entries display when you view the child's policy.

Editing a Node's Configuration Tracking Policy

You can edit tracking policy entries in the following ways:

- You can make changes to tracking policy entries that are not inherited from another node. For example, if the policy entry specifies that a configuration file should be backed up, you can add an action, such as sending an email when a change is noticed.
- You can disable a tracking policy entry that has been inherited from another node. (You cannot delete an inherited tracking policy.)
- You can delete a tracking policy entry that is not inherited. After an entry is deleted, it cannot be undeleted. (You cannot disable a tracking policy that is not inherited.)
- You can re-enable a previously disabled tracking policy entry.

If the node has any child nodes, the child nodes inherit all the changes that you make to the parent node.

When you make a change to a node's tracking policy, that change is not immediately made on the servers attached to the node. You must perform a tracking policy reconcile for the changes to take effect.

Editing a Node's Configuration Tracking Policy Entry

You can make changes only to tracking policy entries that are not inherited from other nodes.

Perform the following steps to edit a node's tracking policy:

1. From the navigation panel, click Software and then select the relevant type of software node (for example, patches, applications, and so forth)
2. Navigate to the node that has the tracking policy entries that you want to edit.
3. Make sure that enabled entries display. If disabled entries display, choose enabled entries from the View drop-down menu and then click Update.
4. Click the link in the target field for a non-inherited policy entry that you want to edit. The Edit Entry page appears, as Figure 11-5 shows.

Figure 11-5: Track Configurations: Edit Entry Page

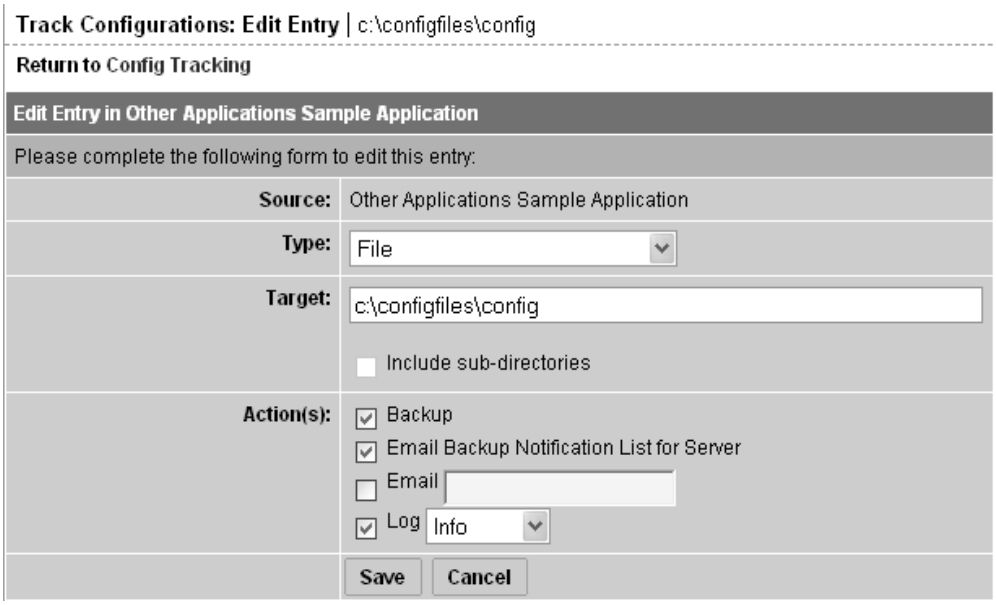

5 Make the desired changes to the tracking policy entry. You can change the type, change the target, clear existing actions, and select new actions.

6 Click Save to commit the changes.

You are returned to the tracking policy page for the selected Node.

You must perform a configuration tracking reconcile to deploy the changes to the appropriate servers.

Disabling a Node's Configuration Tracking Policy Entries

You can only disable inherited tracking policy entries.

Perform the following steps to disable policy entries:

1 From the navigation panel, click Software and then select the relevant type of software node (for example, patches, applications, and so forth).

2 Navigate to the node that has the tracking policy entries that you want to disable.

3. Make sure that Enabled Entries display. If disabled entries display, select enabled entries from the View drop-down menu, and then click Update.
4. Select the inherited tracking policy entry or entries that you want to disable. You can select all displayed entries by clicking the first check box in the list.
5. Click the Disable button.
6. Click the Disabled Entries button to commit the changes.

You must perform a configuration tracking reconcile to deploy the changes to the appropriate servers.

Deleting an Entry in a Node's Configuration Tracking Policy

You can only delete tracking policy entries that are not inherited.

Perform the following steps to delete policy entries:

1. From the navigation panel, click Software and then select the relevant type of software node (for example, patches, applications, and so forth).
2. Navigate to the node that has the tracking policy entries that you want to delete.
3. Make sure that Enabled Entries display. If disabled entries display, select enabled entries from the View drop-down menu, and then click Update.
4. Select the (non-inherited) tracking policy entry or entries that you want to delete. You can select all displayed entries by clicking the first check box in the list.
5. Click the Delete button.
6. Click Delete Entries button to commit the changes.

You must perform a configuration tracking reconcile to deploy the changes to the appropriate servers.

Re-enabling a Tracking Policy Entry

Perform the following steps to re-enable a disabled tracking policy entry:

1. From the navigation panel, click Software and then select the relevant type of software node (for example, patches, applications, and so forth).

2. Navigate to the node that has the tracking policy entries that you want to enable.
3. Select disabled entries from the drop-down box and then click Update to display the disabled entries.
4. Select the disabled tracking policy entry or entries that you want to enable. You can select all displayed entries by clicking the first check box in the list.
5. Click the Enable button.
6. Click the Enabled Entries button to commit the changes.

You must perform a configuration tracking reconcile to deploy the changes to the appropriate servers.

Reconciling a Node's Configuration Tracking Policy

You must perform a configuration tracking reconcile to apply the node's tracking policy to the servers attached to the node. You must do this when you create a policy, when you make any changes to a policy, or when you attach or detach a server from a node.

Performing a configuration tracking reconcile on a server also automatically enables the Configuration Tracking feature on that server. See "If you want Configuration Tracking to remain disabled on a server, be careful not to perform a configuration tracking reconcile on the server. (You can, however, still perform a regular reconcile.)" on page 536 in this chapter for more information.

Perform the following steps to perform a reconcile for a node-based configuration tracking policy:

1. From the navigation panel, click Software and then select the relevant type of software node (for example, patches, applications, and so forth).
2. Navigate to the Node whose tracking policy you want to reconcile.
3. Click the Servers tab. A list of the servers attached to the Node displays.
4. Select the servers whose tracking policy you want to reconcile. (You can select the checkbox at the top of the list to select all the servers in the list.)

5 Under the Configuration Tracking drop-down menu, select "Reconcile Tracking Policies," as Figure 11-6 shows.

Figure 11-6: Reconcile Tracking Polices

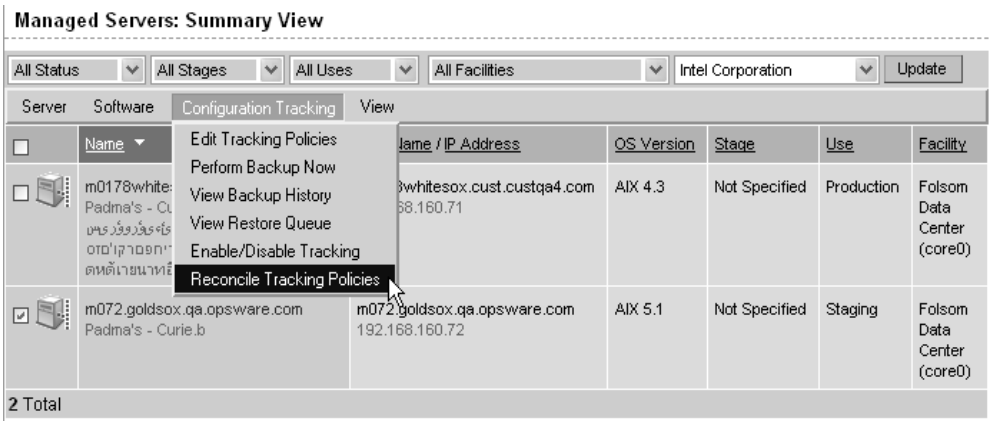

At this point, the Opsware System performs a test reconcile to ensure that the reconcile can be performed without errors. The progress of the test reconcile process displays. Click the View Details button if you want to see information about what changes will be made when the reconcile is performed.

6 If the test reconcile is completed without errors, click Reconcile.

The progress of the reconcile process displays. Click View Details if you want to see information about what changes were made during the Reconcile process.

Customizing Configuration Tracking Policies

This section provides information on customizing configuration tracking policies within the Opsware System and contains the following topics:

- Customizing Configuration Tracking Policies Overview
- Node-Based Entries and Server-Based Entries
- Customizing Multiple Servers
- Adding Or Editing Customized Tracking Policy Entries
- Disabling Customized Tracking Policy Entries
- Enabling Customized Tracking Policy Entries

Customizing Configuration Tracking Policies Overview

You can customize the tracking policy for a server or selected group of servers. Ordinarily, a server gets its tracking policy from the nodes to which it is attached. In some cases, however, you might need to customize the tracking policy for a particular server or set of servers. If, for example, an application on one server is using an optional configuration file, you can customize the tracking policy for that server so that the optional configuration file is monitored.

Node-Based Entries and Server-Based Entries

A tracking policy entry created for a specific server or selected group of servers is called a server-based tracking policy entry. A tracking policy entry that a server obtains from a node that it is attached to is called a node-based tracking policy entry.

Table 11-2 shows the differences between the actions that you can perform on node-based tracking policy entries and server-based tracking policy entries.

Table 11-2: Entry Types

ENTRY TYPE	EDIT	DELETE	ENABLE/DISABLE
Server-based	Yes	Yes	No
Node-based	No	No	Yes

The restrictions for node-based policy entries hold true only when you are customizing the tracking policy for a server or set of servers. You cannot edit a policy entry on a server that is obtained from a node. You can, however, edit the policy of the node itself (See “Editing a Node’s Configuration Tracking Policy” on page 510 in this chapter for information about selecting the node and editing the node’s tracking policy.).

To determine if a tracking policy entry is node-based or server-based, check the Source column in the Track Configurations: Customize Tracking Policies page. Figure 11-7 shows the policy entries for a single selected server.

Figure 11-7: Viewing Configuration Tracking Policy

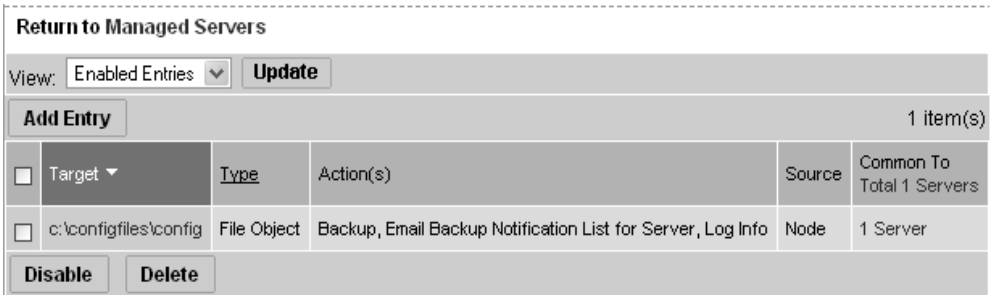

You can edit only tracking policy entries that have been created on the server level. You cannot edit tracking policy entries that have been obtained from a node.

Customizing Multiple Servers

When you customize a tracking policy, you can select one or multiple servers. Here are some considerations when you customize the policies of multiple servers.

- When you select a group of servers that contain both Unix-based (for example, Solaris, HP/UX, Linux, and so forth) and Windows-based servers, you cannot add tracking policy entries. (You can always add entries when you select a single server.)
- When you select a group of servers made up entirely of either Unix-based servers or entirely of Windows-based servers, you have the option of making new tracking policy entries. The policy entries are added to all servers that you select.
- When you select a group of servers made up entirely of either Unix-based servers or of Windows-based servers, you have the option of editing a tracking policy entry. Any policy entry that can be edited can also be added to all selected servers (if it is not already common to all of them).

It can take 20 seconds or longer to customize polices when you select six or more servers.

Adding Or Editing Customized Tracking Policy Entries

Use the following procedure to add or edit tracking policy entries for a server or set of servers running the same general type of operating system. If you select a set of servers running both Windows and Unix operating systems, you cannot add tracking policy entries.

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click Server Search to search for the server whose policy you want to edit. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)
3. Select the server or set of servers that you want to add server-based tracking policy entries to, as Figure 11-8 shows.

Figure 11-8: Selecting Servers

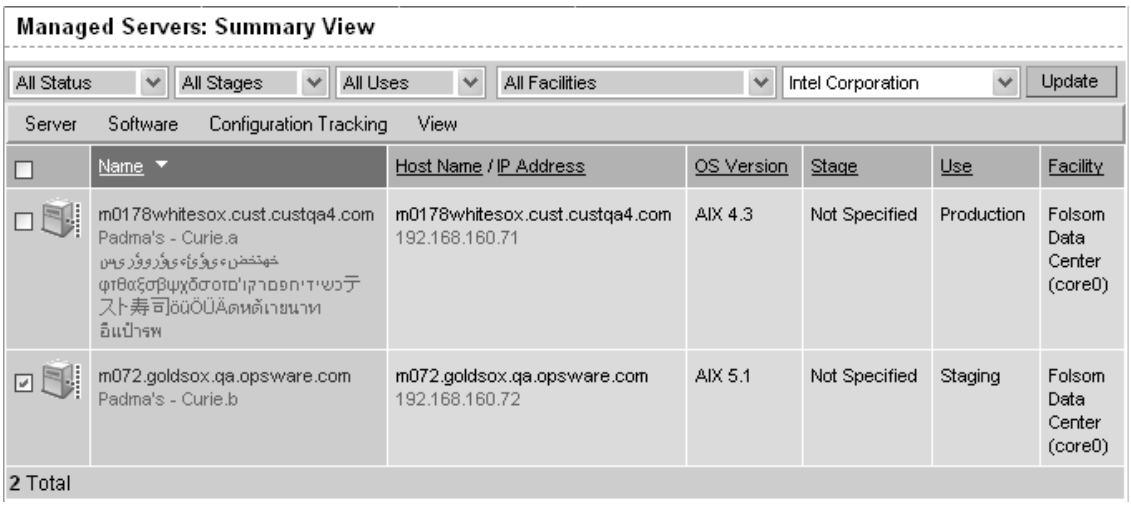

4 From the Configuration Tracking drop-down menu, select Edit Tracking Policies. The Configuration Tracking: Edit Tracking Policies page appears, as Figure 11-9 shows.

Figure 11-9: Configuration Tracking: Edit Tracking Policies Page

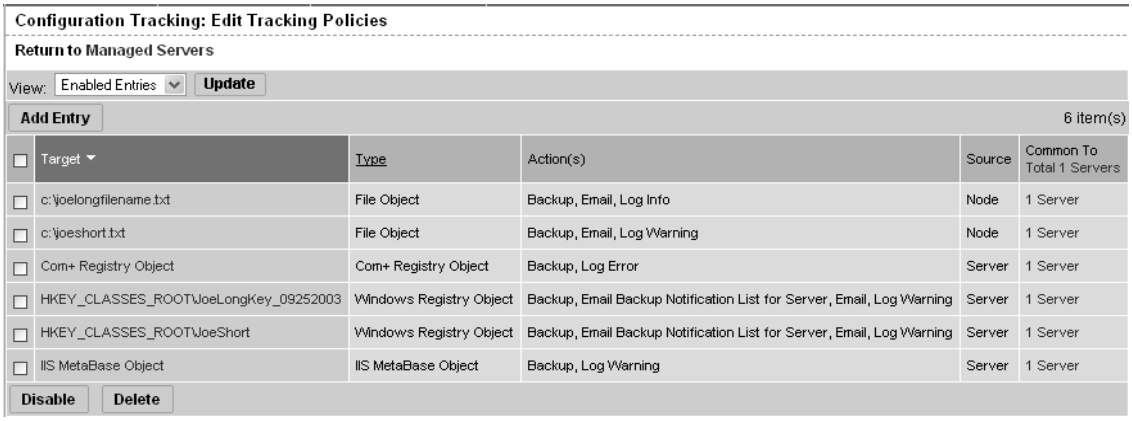

5 Click the Add Entry button to create a new policy entry. To edit an existing entry, click the link for the entry in the Target field.

6 Define the target, select the type, and select the actions to be performed on the target.

Table 11-3 describes each of the selections that you must make.

Table 11-3: Editing Configuration Tracking Policies

FIELD	DESCRIPTION
Type	**File**: monitor the file specified in the target field **Directory**: monitor all the files in the directory specified in the target field. The following types are available only for Windows servers: **Windows Registry**: specify key in target field **IIS Metabase**: entire Metabase is monitored; do not specify target **COM + Registration Database**: entire Registry is monitored; do not specify target.

Table 11-3: Editing Configuration Tracking Policies

FIELD	DESCRIPTION
Target	If you selected the file type, specify the full path (including the drive letter on Windows servers) of the file that you want to monitor. If you selected the directory type, specify the full path of the directory (including the drive letter on Windows machines) that you want to monitor. You also have the option of monitoring subdirectories. (Select the include subdirectories check box.) If you select the file or directory type, you can use wildcards in the target. (See "Configuration Tracking Policy Targets and Wildcards" on page 498 in this chapter for more information.) If you selected the Windows Registry type, specify the Windows registry key. This key and all its subkeys are backed up. Use standard syntax for Windows Registry keys (For example, `HKEY_LOCAL_MACHINE\SOFTWARE`) If you selected the IIS Metabase or COM + Registration Database type, you do not specify the target.

Table 11-3: Editing Configuration Tracking Policies

FIELD	DESCRIPTION
Actions (you can select multiple options)	Backup: back up the specified file, directory, COM + Registration Database, Windows Registry keys, or IIS Metabase Email Backup Notification List for Server: send an email to the backup notification list for the selected server. (Not available for Windows Registry.) See "Creating the Email Notification List" on page 501 in this chapter for more information. Email: send an email to the address or addresses specified in this field when a change is detected. Use a comma-separated list for multiple email addresses. (Not available for Windows Registry.) Log: add an entry to the server's system log when a change is detected. (Not available for Windows Registry.) You can choose to write the following types of log entries to the server's system log: Info Warning Error

7 Click Save to add the entry to the tracking policy.

8 If you want continue to add entries to the tracking policy, click Add Entry and repeat this procedure.

Changes do not take effect until you perform a Configuration Tracking policy reconcile.

Disabling Customized Tracking Policy Entries

You can only disable node-based tracking policy entries.

Perform the following steps to disable customized tracking policy entries:

1 From the navigation panel in the Opsware Command Center, click Servers.

2. Click Server Search to search for the desired servers. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Edit Tracking Policies.

   The Configuration Tracking: Edit Tracking Policies page appears.
4. Select the tracking policy entries that you want to disable.
5. Select the tracking policy entry or entries that you want to disable. (Make sure that all your selected tracking policy entries are node-based.) If all the tracking policy entries displayed are node-based and you want to disable all the displayed tracking policies, you can select the first check box to select all tracking policy entries.
6. Click the Disable button to disable the tracking policy entries that you selected.

Changes do not take effect until you perform a Configuration Tracking policy reconcile.

Enabling Customized Tracking Policy Entries

You can re-enable any previously disabled node-based tracking policy entry.

Perform the following steps to re-enable tracking policies:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click Server Search to search for the desired servers. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Edit Tracking Policies.

   The Configuration Tracking: Edit Tracking Policies page appears.
4. From the View menu, select Disabled Entries and then click Update.
5. Select the tracking policy entry or entries that you want to disable. (Make sure that all your selected tracking policy entries are node-based.) If all the tracking policy entries displayed are node-based and you want to disable all the displayed tracking policies, you can select the first check box to select all tracking policy entries.
6. Click the Enable button to disable the tracking policy entries that you select.

Changes do not take effect until you perform a Configuration Tracking policy reconcile.

Viewing a Server's Tracking Policy

Perform the following steps to view a server's tracking policy:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click Server Search to search for the desired servers. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Edit Tracking Policies. The tracking policy entries display.
4. Click the Tracking Policy link to display the server's tracking policy.

Reconciling Customized Tracking Policies

When you create or edit any customized tracking policy entries, the entries are not deployed to your servers until you perform a configuration tracking reconcile.

A configuration tracking reconcile is not the same as a standard Opsware System reconcile. Performing a standard reconcile does not deploy your tracking policies to your servers.

A configuration tracking reconcile deploys all configuration polices, not just the customized policies. Both node-based polices and customized policies are deployed when you perform a configuration tracking reconcile.

Perform the following steps to reconcile the tracking policy to the servers whose policies you customized:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click Server Search to search for the desired servers. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)

3. From the Configuration Tracking drop-down menu, select Reconcile Tracking Policies.

   The Track Configurations: Preview Reconcile page appears and displays the progress of the test reconcile.

4. After the test reconcile completes successfully, click the Reconcile button to perform the actual reconcile operation.

   The Track Configurations: Reconcile page appears.

5. If you want to see more information about the changes made during the reconcile operation, click the View Details button. Otherwise, click the Done button.

Performing Manual Backups

On a server or set of servers, you can perform a manual backup of all tracked configuration files and databases for which you have selected the backup action. Manual backups can be useful as a precaution before making changes to configuration files. If a problem arises, you can then immediately restore a backed-up configuration file or database to its previous state.

Manual backups are full backups. All tracked configuration files and databases for which the backup action has been selected are backed up, not just the files that have changed.

Perform the following steps to perform a manual backup:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click Server Search to search for the desired servers. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Perform Backup Now.
4. If you do not want to use the default backup name (Manual Backup), type a new name in the backup name field. If you provide a backup name, you can later perform a search for backup names to find this backup point. Backup names do not have to be unique.
5. Click the Start Backup button. The backup progress displays.
6. When the backup is completed, you can click the View Details button to review the list of configuration files or configuration databases that have been backed up.

The Backup History

The Opsware System provides a detailed history of backup activity as well as search capabilities to find backup points by backup name and to find backed up files by filename. This section contains the following topics:

- Viewing the Backup History
- Viewing the List of Backup Events
- Types of Backup Events
- Backup History Search Options
- Backup Info and Backup Manifest

Viewing the Backup History

Perform the following steps to view the backup history:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click Server Search to search for the desired servers. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select View Backup History.

   The Track Configurations: View Group Backup History page displays. This page displays the backup activity for the servers that you select. Figure 11-10 shows a sample backup history.

Figure 11-10: Backup History

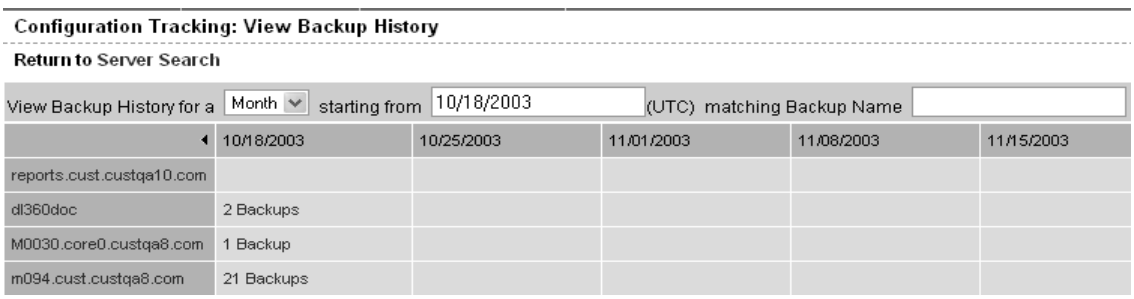

The number of backups that occurred on a particular date displays as a link in the server's date column.

The number of backups refers to the number of backup events (or backup points) when a particular type of backup took place. It does not refer to the number of files backed up. For example, if a server displays 3 backups in a date column, it could refer to the following three backup events:

- A scheduled incremental backup that backed up four files
- A second scheduled incremental backup that backed up 10 files
- A manual backup that backed up 30 files

Viewing the List of Backup Events

To view the list of individual backup events, click the link that displays the number of backups in the desired date column. Figure 11-11 shows a page with a list of backup events.

Figure 11-11: Backup Events

Types of Backup Events

In addition to information such as the date and time the backup occurred, the list of backup events indicates the type of backup. Table 11-4 describes the type of backup events that can occur.

Table 11-4: Types of Backup Events

TYPE	DESCRIPTION
Triggered Full	The automatic weekly backup of all tracked files for which the backup option was selected. This takes places only if any relevant files have changed since the last full backup.
Manual Incremental	A backup of all changed files (for which the backup option was selected) that occurs before and after a restoration or a rollback.
Triggered Incremental	An automatic backup that occurs when the Opsware Agent detects a change to a tracked file for which the backup option was selected. Only changed files are backed up.
Manual Full	A full backup initiated by the user of all tracked files for which the backup option was selected.

Backup History Search Options

By default, you see the backup history for the past week for the servers that you selected.

If you want to display a different date range, you have the following options:

- Display the backup history for different date range by selecting a different option in the "View Backup History for a" box.
- Use the "starting from" field to search for backup history from a past date until the current date.

You can use the matching field to search for a backup name within the selected date range. Wildcards are allowed (the * and ? characters). If you do not type the full name of the backup, you must use a wildcard for any missing parts of the name.

The results of this search show the date and number of backup names on that date that match the man* pattern. When you click the link for the number of backups, a list of the matching backup names displays.

Backup Info and Backup Manifest

When you click a backup name (such as Scheduled Backup) anywhere in the backup history, the Backup Info and Backup Manifest tabs display.

Backup Info Tab

The Backup Info tab displays general information about a backup event. This information includes the name of the backup, the date and time of the backup, and the policies that triggered the backup event.

If the polices are node-based, they are identified by the name of the Node. All customized policies are identified as Server Policy, as Figure 11-12 shows.

Figure 11-12: Track Configurations: View Backup Triggered Backup Page

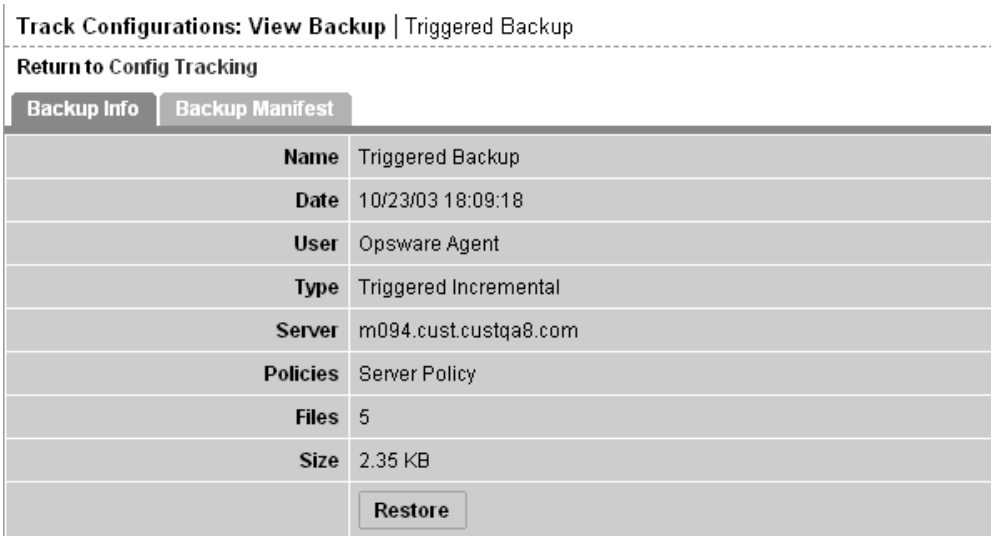

From the Backup Info tab, you can also place the files that were backed up into the Restore Queue. See “How to Restore Backups” on page 534 in this chapter for information about how to use this feature.

Backup Manifest Tab

The Backup Manifest tab displays the list of files that were backed up during the selected backup event, as Figure 11-13 shows.

Figure 11-13: Backup Manifest Tab

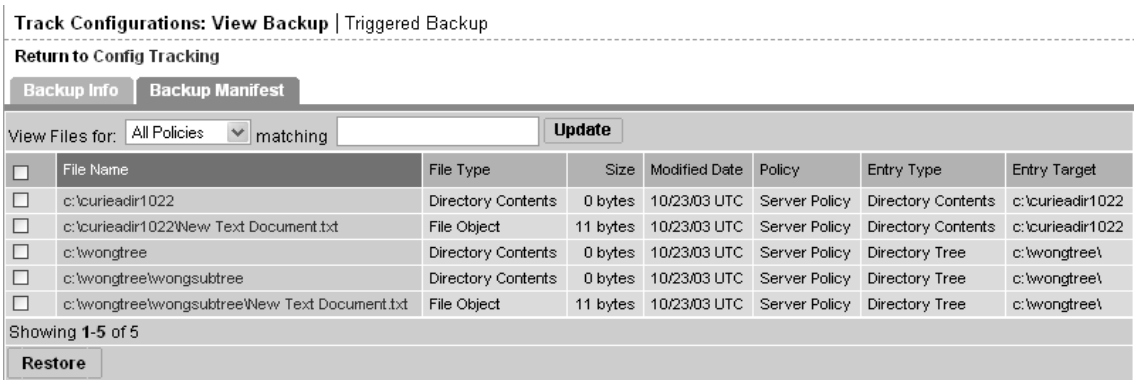

Two types of objects are backed up, as seen in the File Type field.

- File Object in the File Type field indicates that a file has been backed up. You can use the entry to restore the file.
- Directory Contents in the File Type field indicates that a directory object has been backed up. The directory object is the directory itself, and not the contents of the directory. You can use this entry to restore the directory, but you must select the files inside the directory to restore the contents of the directory.

The Entry Type field can have three possible values. The Entry Type identifies the target type in the policy entry that caused the file or directory object to be backed up.

- File Object in the Entry Type field indicates that the file was backed up as the result of a file target type.
- Directory Contents in the Entry Type field indicates that the file or directory object was backed up as result of a directory target type.
- Directory Tree in the Entry Type field indicates that the file or directory object was backed up as a result of a directory target type with the include subdirectories option selected.

File Info and File Versions

When you click a file or directory name anywhere in the backup history (such as in the Backup Info tab), the File Info and File Versions tabs display, as Figure 11-14 shows.

Figure 11-14: File Info

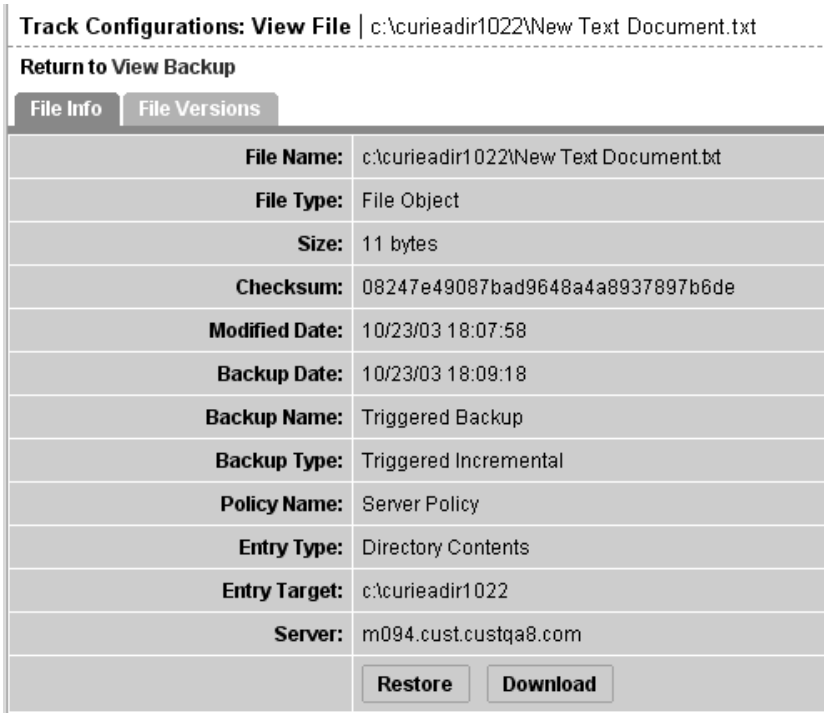

The File Info tab displays information about the specific file that you selected. The Policy Name refers to the policy that caused the file to be backed up. The Policy Name is either the name of the node whose tracking policy triggered the backup, or the Server Policy if the server's customized tracking policy caused the file to be backed up.

You can place the file that you selected in the Restore Queue by clicking the Restore button. See "Restoring Backed Up Files" on page 531 in this chapter for more information.

The File Versions tab displays a list of the backed up versions of the file, as Figure 11-15 shows.

Figure 11-15: File Versions

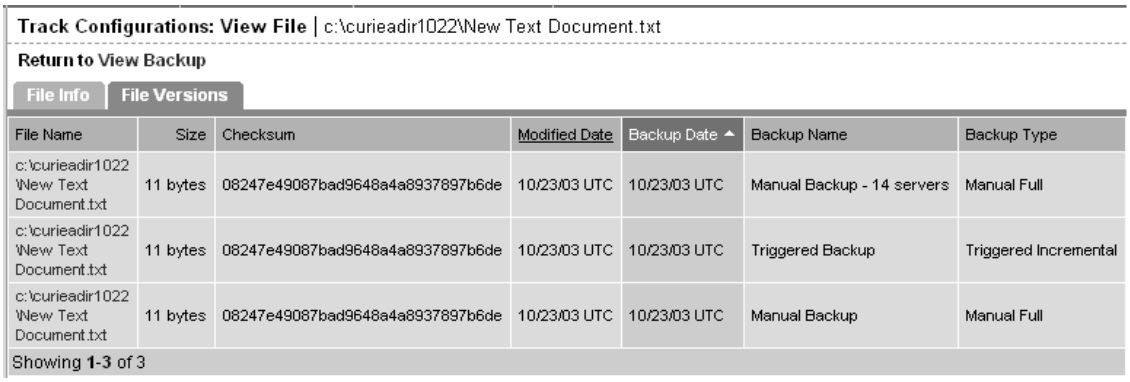

It displays the backup name and backup type of the file. If you click any of the files, the File Info for the file is displayed. You can therefore use the File Versions tab to select a different version of the file, and then restore it or download it from the File History.

Deleting Backups

Backups remain in the backup history until you delete them or the server is deactivated. Backups are stored in the Software Repository. See "Operating System Provisioning" on page 237 in Chapter 4 for information about the Software Repository. You should delete old backup events periodically to reclaim disk storage from the Software Repository.

You can delete entire backup events (identified by a backup name); you cannot delete individual files from the backup history.

Perform the following steps to delete backup points from your backup history:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click Server Search to search for the desired servers. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)
3. Select the server or set of servers that have backup points that you want to delete.
4. Search for a date that contains the backup points that you want to delete. Click the link with the number of backup events that occurred on that date.

5. Select the check boxes for the backup events that you want to delete from the backup history.
6. Click the Delete button.

   The Delete Backup Confirmation page appears
7. Click the Delete Backups button. The backups are now deleted.

Restoring Backed Up Files

You can use the Automated Configuration Tracking feature to restore configuration files or databases that have been backed up. You can restore all files that were backed up during a backup point, or you can select and restore individual files from a backup point. The Opsware System allows you to roll back a restoration (for example, return your server's tracked files to the state immediately before the restoration).

Overview of Restore Procedure

This section presents an overview of the procedures you use to restore backed-up configuration files and databases. See “How to Restore Backups” on page 534 in this chapter for information about the step-by-step procedure.

To restore backed-up configuration files and databases, you view the backup history for a server or set of servers. (You should be familiar with “The Backup History” on page 524 before you perform a restore.)

The backup history displays entries for each date when a backup occurred. As Figure 11-16 shows, the entry for the date displays how many backup points occurred on that date.

Figure 11-16: Backup Points

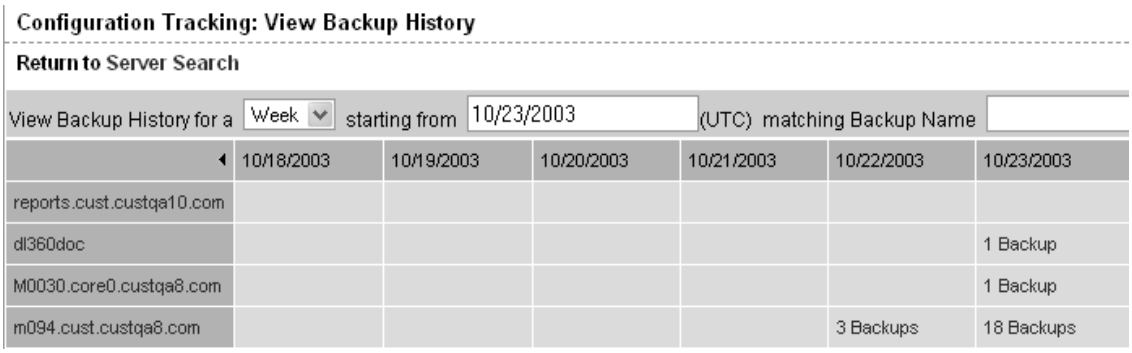

You select the desired backup date, and you can then select one or more backup points that occurred on that date. Before you restore the backups, you can review all files that were backed up at your selected backup points. You can either choose to restore all backed up files from your selected backup points, or you can select the files individually.

To select the backup, click on the link in the date field that displays the number of backups that were performed on that date.

The Restore Queue

This section provides information on the restore queue within the Opsware System and contains the following topics:

- Restore Queue Overview
- Incremental Backups for Restoration and Rollbacks
- Entries for Directories in the Backup History
- Restoring "File Not Found" Entries
- How to Restore Backups
- Rolling Back Restored Files

Restore Queue Overview

When you click the Restore button, the files that you selected are placed in the Restore Queue and the View Restore Queue and Perform Restore page appears. You can then review and select files from all the backup points that you selected.

If you do not want to restore the files at this time, you can perform other actions and the files will remain in the Restore Queue as long as your session is active (even if you leave this page). You can also return to the backup history and select other files to put in the Restore Queue.

When you are ready to restore the backups, return to the Track Configurations: Select a Task page and click the *View Restore Queue and Perform Restore* link.

Incremental Backups for Restoration and Rollbacks

To make rollbacks possible, the Opsware System performs two automatic incremental backups. The first backup occurs immediately before the restoration, and the second occurs immediately after the restoration. Similarly, in order to undo a rollback, the Opsware System performs two automatic incremental backups, one before and one after a rollback.

These backups do not occur if all the files you have selected to restore are identical to the files already on the server. In such a case, the restoration does not in fact change any files on the server, and the rollback option is not available (there are no changes to roll back).

Entries for Directories in the Backup History

As discussed in the "Special Considerations for Directory and Wildcard Targets" on page 498, if you are tracking a directory and the contents of the directory change, the action that you selected is triggered for the directory object itself and for the files that changed. The only exception occurs when a file is deleted from the directory. In that case, the action is triggered for the directory object alone.

If you selected the backup action for the directory target, when the directory contents change, the directory object is backed up as well as the changed files. When a file is deleted, however, only the directory object is backed up.

If you restore files contained in a directory without selecting the entry for the directory object, the directory is re-created on the server if it does not already exist. If you do select the entry for the directory object, however, you can ensure that the directory object is restored with the same properties (such as permissions and timestamp) it had at the selected backup event.

Restoring "File Not Found" Entries

If you are monitoring specific files (as opposed to files monitored in tracked directories and files monitored as the result of wildcard targets), the backup history can contain entries noting "file not found" if the file does not exist on the server, or if the file is later deleted on the server. Similarly, if you are monitoring specific directories (opposed to monitoring directories through wildcard targets), the backup directory can contain "file not found" entries if the directory does not exist on the server, or if the directory is later deleted.

If you select and restore "file not found" entry and the file exists on the server, the file is deleted (it is reverted to the state of the backup event that you selected.)

Exercise caution in restoring entries for deleted directories. If the directory exists on your server and the directory contains files, both the directory and its contents will be deleted.

How to Restore Backups

Perform the following steps to restore backed-up configuration files:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click Server Search to search for the desired servers. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)
3. Under the Configuration Tracking drop-down menu, select View Backup History.
4. Find the date and server that has the backup points that contain the files that you want to restore. See "The Backup History" on page 524 in this chapter for information about finding backup points, viewing backup point details, and viewing lists of backed up files.
5. Select the check box for the backup points that have files that you want to restore.
6. Click the Restore button.

   The files or collection of files from the backup points you selected are placed in the Restore Queue. If you do not want to perform a restore now, you can click the Return to Select a Task link to leave this page. While your session is active, the files remain in the Restore Queue. You can also continue to add more files to the Restore Queue. To return later to the Restore Queue, click the View Restore Queue and Perform Restore from the Track Configurations: Select a Task page.
7. If you are ready to restore files, first review the files in the Restore Queue. If you want to restore all the files in the queue, you can either click the Restore All button or click the check box at the top of the list. If you do not want to restore all of the files in the queue, select the checkboxes for the files that you want to restore.
8. Click the Restore button (or the Restore All button if you want to restore all of the files in the queue).
9. After the restoration is complete, you can click the View Details button for more information about the restored files.

Rolling Back Restored Files

You can roll back any files that you selected to restore. By rolling back the files, you revert the file to the version that was on the server before you performed the restoration (in other words, you undo the restoration.)

The rollback option is not available if you restored any files that are no longer being monitored by Automated Configuration Tracking. You can choose to restore all changed files or select the files individually.

Perform the following steps to roll back restored files:

1. After you restore the backups, click the Rollback button.

   The Track Configurations: Restore Queue page appears, which contains the list of files that you need to revert in order to roll back the tracked configuration files to their previous state.

   This list might not contain all the files that you selected for your original restore. If any of the files that you selected during your original restore are identical to the files already on the server, they do not need to be rolled back. Only the files that changed display.

2. If you want to roll back all the changed files, click Restore All. Otherwise, you can select the individual files that you want to roll back and click the restore button.

   The Track Configurations: Restore Progress page appears.

Exercise caution in restoring entries for deleted directories. If the directory exists on your server and the directory contains files, both the directory and its contents will be deleted.

Enabling and Disabling Configuration Tracking

You can enable or disable Configuration Tracking on any server or set of servers in your Opsware managed environment.

Disabling Configuration Tracking is not the same as disabling an individual tracking policy entry (as discussed in “Disabling Customized Tracking Policy Entries” on page 520). Disabling Configuration Tracking stops all configuration tracking activity on the selected server.

Disabling Configuration Tracking, however, does not change a server's tracking policy in any way. If you later re-enabled Configuration Tracking on the server, the server still has the same tracking policy that it did before you disabled it.

By default, Configuration Tracking is disabled on all managed servers. It is not necessary, however, to manually enable Configuration Tracking in order to turn on the Configuration Tracking feature. Configuration Tracking is automatically enabled on a server when you reconcile its configuration tracking policies, and you must perform a reconcile in order to deploy tracking policies.

If you want Configuration Tracking to remain disabled on a server, be careful not to perform a configuration tracking reconcile on the server. (You can, however, still perform a regular reconcile.)

Perform the following steps to enable or disable configuration tracking:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click Server Search to search for the desired servers. (Alternatively, you can click Managed Servers or Server Groups and then select the server from the server list.)
3. Under the Configuration Tracking drop-down menu, select Enable/Disable. A list that displays the state (Enabled/Disabled) of the servers that you select displays, as Figure 11-17 shows.

Figure 11-17: Enable/Disable Configuration Tracking Page

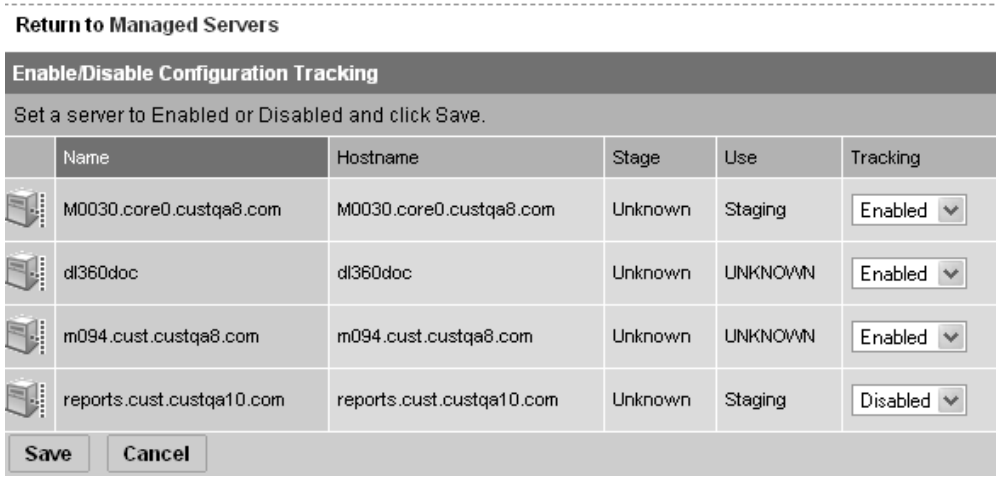

4. Under the Tracking field, select Enabled or Disabled to enable or disable configuration tracking on the selected server.
5. Click Save to commit the changes.

Chapter 12: Code Deployment & Rollback

IN THIS CHAPTER

This chapter discusses how to use the Code Deployment feature in the Opsware Command Center to deploy application code and content to servers managed in the Opsware environment. Topics in this chapter include:

- Opsware Code Deployment Process
- Code Deployment & Rollback Setup
- Performing Services, Synchronizations, and Sequences

You must have specific permissions to deploy code and content by using the Opsware Command Center. Contact your Opsware administrator to obtain the necessary access rights.

Opsware Code Deployment Process

This section provides information on the code deployment process within the Opsware System and contains the following topics:

- Code Deployment Process Overview
- Uploading Code and Content to Staging
- Using Code Deployment & Rollback
- Accessing Code Deployment & Rollback

Code Deployment Process Overview

The Code Deployment & Rollback (CDR) Subsystem in the Opsware Command Center provides tools for deploying new and updated code and content to your operational environment.

Figure 12-1 shows the architecture and process for updating a typical server hosted in an Opsware managed environment.

Figure 12-1: Typical Code and Content Update in the Opsware Managed Environment

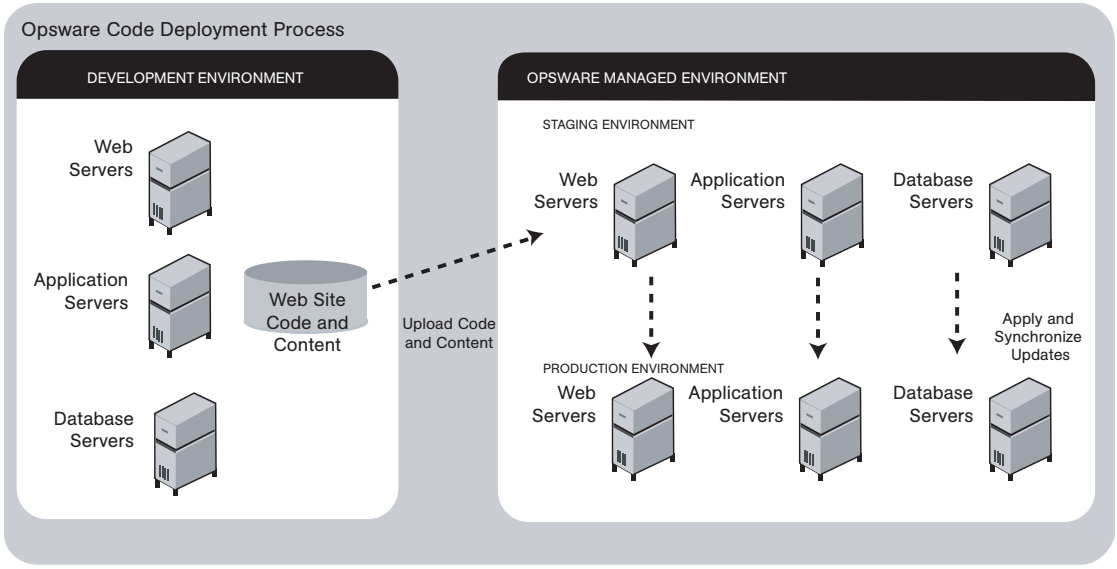

The deployment process involves performing the following high-level tasks:

1. Determining your application code and content deployment requirements and defining the CDR services, synchronizations, and sequences that you need to support them
   - Services are defined for each different type of Web server or application server applications (for example, WebLogic Server) that is installed on the staging and production hosts in your environment.
   - Synchronizations are defined for each service so that you can update files between the source location and one or more destination production hosts that are running the same service.
   - Sequences are optional but can simplify deployment by grouping a collection of service operations and synchronizations that can be performed as a single task.
2. Uploading new or updated code and content to your Opsware staging environment
3. After performing any necessary testing, cutting over to the changed code and content on the staging environment

4. As necessary, performing CDR service operations, such as backing up code and content from your live site
5. Performing CDR operations available to synchronize the updated code and content to your production hosts in the Opsware managed environment
6. To simplify subsequent deployments of new code and content, defining sequences that specify a series of service operations and synchronizations you want to perform as a single action

The code and content deployment process that you follow might be different depending on the architecture of your operational environment and your deployment requirements.

Uploading Code and Content to Staging

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

Figure 12-2 shows an example of a typical development environment and how your uploaded code and content move to the staging environment.

Figure 12-2: How Code and Content Move to the Staging Environment

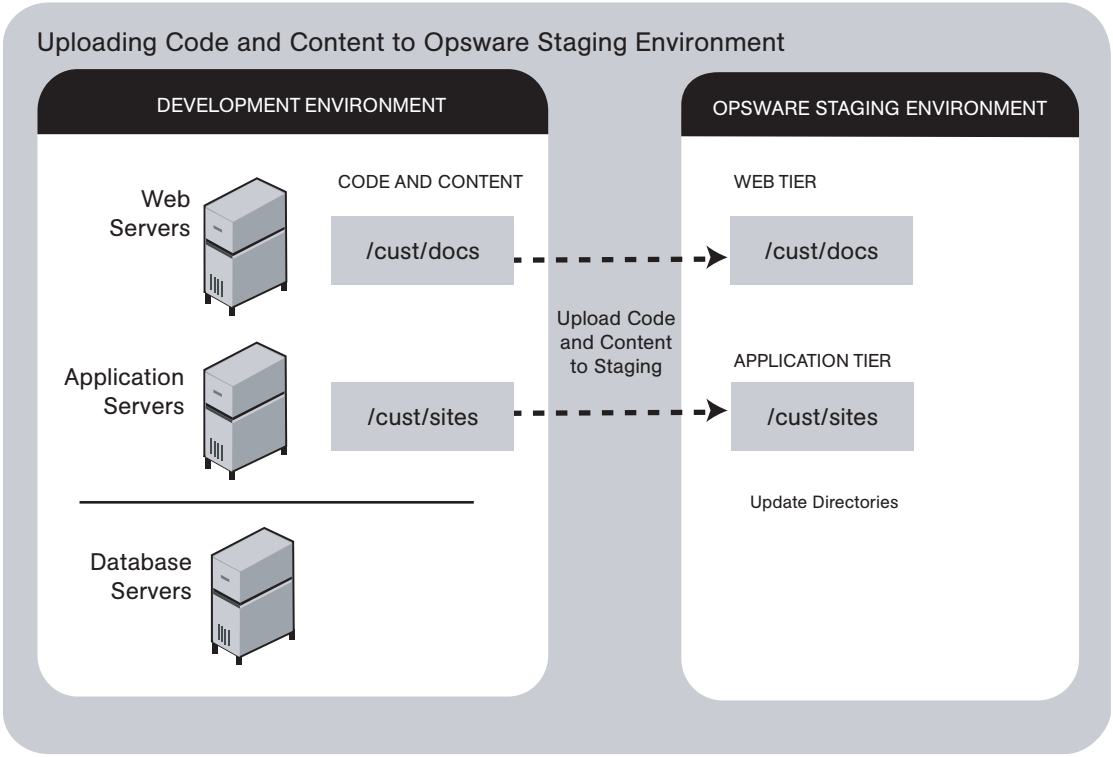

After you upload the files and test your changes, you can synchronize updates to the production hosts running your managed environment. You can run specific synchronizations and perform other service deployment operations by selecting CDR menu options available from the Opsware Command Center navigation panel.

Using Code Deployment & Rollback

After you upload updated code and content to your Opsware-managed staging environment, you can use the CDR operations to cutover to new code and content, perform host synchronizations, and perform other service operations.

CDR uses the following directories to synchronize and cutover code and content for specified hosts:

- **Live directory** – The directory that stores the actual code and content required to run a live site.

- **Update directory** – The directory written to by CDR synchronizations. Stores only the files that changed between the source host Live directory and the Live directories of the destination hosts.
- **Site Previous directory** – This directory holds all the changes necessary to revert the Live directory back to the state it was in before the last cutover. Like the Update directory, the Site Previous directory only stores the files that changed between the current Live directory contents and its previous state.
- **Site Backup directory** – This directory stores a complete backup of the site. The directory is populated when the user issues a Backup service operation.

When you cutover to new code and content, CDR determines the differences between the new code and content in the current Update directory and the Live directory for your site. The files that are different are synchronized to the Live directory. When you synchronize source and destination hosts, CDR moves modified files from the Live directory on a source host to a directory on a destination host.

You cannot use CDR to automate database pushes. However, you can configure CDR so that you can synchronize modified database script files on different hosts.

CDR offers the following features:

- Provides a single tool for deploying code (such as ASP, JSP, and JAR files) and site content (such as HTML, JPEG, GIF, and PDF files). Using a single tool is helpful when the code and content for your site are intermingled.
- Provides direct control over code and content pushes by making it possible to decide what information to update and determine when and how to perform updates.
- Provides flexibility to accommodate frequent updates to staging and production hosts by enabling more frequent pushes in a shorter period of time.
- Allows verification of file changes between staging and production host directories by creating a manifest of updated files. You can verify changes before cutting over to new code and content.
- Provides administrative service operations, including starting and stopping services, and backing up, restoring, and rolling back code and content to return your site to the previous version.

- Lets you push incremental updates to your site so that only files that have changed are pushed to specified locations on staging or production hosts.
- CDR uses the same authentication and navigation that you use in accessing other information and performing other site operations from the Opsware Command Center.

Accessing Code Deployment & Rollback

As with all other features in the Opsware System, the links that you see on the Opsware Command Center Home page and the links that you see in the navigation panel are based on the permissions that you have in combination with the customer you are associated with.

If you do not have permissions for Code Deployment & Rollback, you cannot see the Code Deployment links on the navigation panel, the link called Deploy Code in the Tasks panel of the Opsware Command Center home page appears in italics, and it is not an active link.

If you have Code Deployment & Rollback permissions to no more than one customer, when you expand the Code Deployment section in the navigation panel, you can see a link called Set Customer. Click that link to view the links to the specific Code Deployment functions that you have permissions for in combination with that single customer.

If you have Code Deployment & Rollback permissions to more than one customer, you can see a link called Select Customer. Click that link to display a page that shows the customers you are associated with. Select the customer you want to work with. The CDR Home Page appears, with links to the specific Code Deployment functions that you have permissions for. These links are the same functions that you can find in the navigation panel under Code Deployment.

The navigation instructions and screen captures in this chapter show what a user with permissions to all code deployment functions and access to only one customer can see. Consequently, because your permissions and customers might be different, the available menu selections and features that you see might likewise differ.

Perform the following steps to access Code Deployment & Rollback:

1. If necessary, click the Code Deployment link in the navigation panel to expand the list of CDR options.

2 Click the CDR Home link. The CDR Home Page for [*customer name*] appears, as Figure 12-3 shows.

Figure 12-3: Code Deployment Home Page

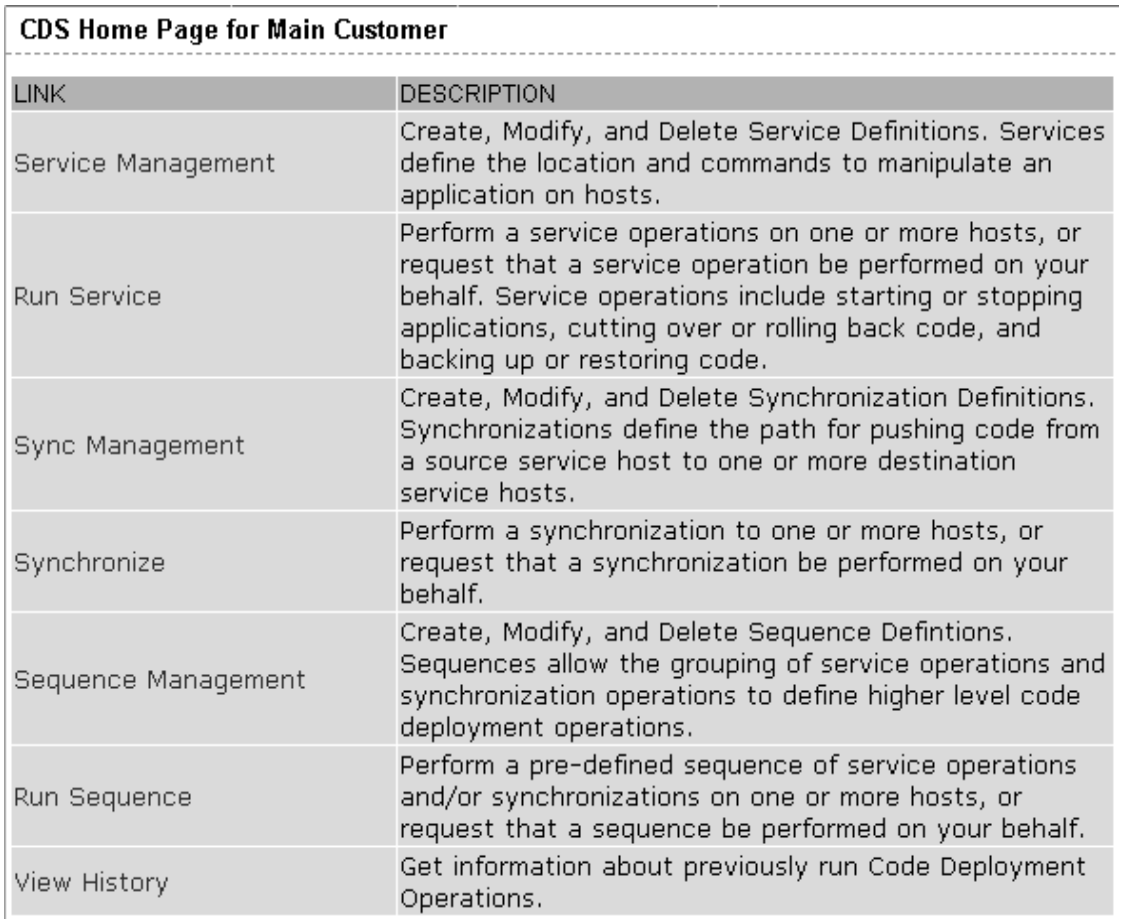

CDS Home Page for Main Customer

LINK	DESCRIPTION
Service Management	Create, Modify, and Delete Service Definitions. Services define the location and commands to manipulate an application on hosts.
Run Service	Perform a service operations on one or more hosts, or request that a service operation be performed on your behalf. Service operations include starting or stopping applications, cutting over or rolling back code, and backing up or restoring code.
Sync Management	Create, Modify, and Delete Synchronization Definitions. Synchronizations define the path for pushing code from a source service host to one or more destination service hosts.
Synchronize	Perform a synchronization to one or more hosts, or request that a synchronization be performed on your behalf.
Sequence Management	Create, Modify, and Delete Sequence Defintions. Sequences allow the grouping of service operations and synchronization operations to define higher level code deployment operations.
Run Sequence	Perform a pre-defined sequence of service operations and/or synchronizations on one or more hosts, or request that a sequence be performed on your behalf.
View History	Get information about previously run Code Deployment Operations.

Depending on your access permissions, the following CDR options appear:

- Service Management – create, modify, or delete service definitions that define the location and commands to manipulate an application on hosts associated with each application instance running in your operational environment
- Run Service – perform a service operation or request that one be performed
- Sync Management – create, modify, or delete synchronization definitions associated with code pushes
- Synchronize – perform a synchronization or request that one be performed

- Sequence Management – create, modify, or delete sequences of operations
- Run Sequences – perform a selected sequence or request that one be performed
- View History – view information stored in an operations log to determine the status of particular deployment operations, and whether they completed successfully.

3 Choose the CDR operations that you want to perform, selecting options from the navigation panel or from the CDR home page.

Code Deployment & Rollback Setup

This section provides information on how to set up and support sites that use CDR for code and content pushes. It contains the following sections:

- Code Deployment & Rollback Overview
- Code Deployment Configuration Checklist
- Deployment and CDR Configuration Procedures
- CDR Configuration Steps
- Determining Your Code and Content Deployment Requirements
- Planning Your CDR Configuration
- Preparing Opsware Host Machines
- Creating or Verifying Directories on Hosts
- Populating Initial Content in Directories
- Setting up Access Control for CDR
- Defining CDR Services, Synchronizations, and Sequences
- Defining and Modifying CDR Services
- Defining a Service
- Running Pre- and Post-Synchronization Scripts
- Modifying a Service
- Deleting a Service
- Creating and Modifying CDR Synchronizations
- Defining a Synchronization

- Modifying a Synchronization
- Deleting a Synchronization
- Creating and Modifying CDR Sequences
- Defining a Sequence
- Modifying a Sequence
- Deleting Sequences
- Verifying and Troubleshooting CDR Configuration

Code Deployment & Rollback Overview

In configuring CDR for a specific site, you first need to install and configure required software on each of the host machines used in your Opsware managed staging and production environment. Then you define the set of services, service operations, and staging and production server synchronizations and sequences to make available.

By selecting Service, Synchronization, and Sequence options from the CDR menus, users can either perform operations or request that other authorized users perform them. (Permissions to perform specific CDR operations depend on the code deployment user groups to which individual users are assigned.)

See the *Opsware System 4.7 Administration Guide* for information about how to create users and assign the Opsware Command Center permissions.

The instructions provided in this chapter are intended to be platform-neutral. However, platform-specific information and examples are provided where necessary.

The preparation of host machines, directory configuration, and testing should all be carried out during scheduled maintenance windows because modifications made to production machines might cause downtime for the live site.

Code Deployment Configuration Checklist

Before you set up your site to use CDR, collect the following information about the site:

- Names of all host machines used for a site and their designation for use as staging, QA, production, and so forth
- All service instances installed for a site (for example, WebLogic, iPlanet Web Server, and so forth)
- All top-level code and content directories that are used by each of the service instances. (Directories are based on the service or service instance and are the same on all host machines where a particular service or service instance is installed.)
- The name of the machine and directory location where site code and content is uploaded. (This is the host and directory location where you upload files from your own development environment, using an Opsware-supported content deployment tool such as OpenDeploy, scp, or rsync over SSH.)

Make sure that you have identified an appropriate process to upload changed code and content from your development environment and check that the appropriate firewall conduits and connections are created to allow uploading changed code and content into the site.

- For a new site deployed in an Opsware managed environment, you should also obtain a copy of your site's current code and content to preload into directories prior to using CDR. Preloading code and content shortens the time required to complete updates the first time that you use CDR to perform synchronizations.

Deployment and CDR Configuration Procedures

The overall process for planning and defining a CDR configuration and using CDR to define services and synchronizations for your site consists of the following tasks:

1. Determine your site's code and content deployment requirements.
2. Define the services and synchronizations that are needed to support your site requirements. Optionally, define any sequences of both services and synchronizations that you would like users to define as sequences so users can perform them in a single step.
3. Upload new or updated code and content to the Opsware staging environment.

4. Cutover to the changed code and content on the staging environment and perform any required testing.

5. As necessary, you can also set up email notification to send requests to select users to perform CDR service operations, such as backing up code and content from their live site and synchronizing the updated code and content to their production hosts.

Your Opsware administrator determines the responsibilities that different users have pertaining to synchronizations and other service operations performed for a specific site.

CDR Configuration Steps

The following summary shows the steps involved in configuring a site to use CDR for code and content updates, code pushes, and other service/synchronization operations.

The sections that follow described each of the steps in detail:

1. Determine your code and content deployment requirements.

   Determine the responsibilities that users who are assigned to perform synchronizations, sequences, and other service operations will have.

2. Plan your CDR configuration.

   Create diagrams of your site's host configuration, specifying synchronization and service descriptions, including any special service operations that you want carried out when a specific synchronization or sequence is performed.

   See "Planning Your CDR Configuration" on page 551 in this chapter for information about how to document your CDR configuration and the services, synchronizations, and sequences you are creating for your site.

3. Set up access control for CDR.

   Have your Opsware administrator create and add users to user groups to create, edit, request, or perform CDR services, synchronizations, and sequences. (User groups that have specific permissions to perform CDR operations are predefined.)

4. Create Services and Synchronizations in CDR.

   Using the service, synchronization, and sequence documentation defined for your site, create each service, synchronization, and sequence in CDR. Assign the user groups required to access each service, synchronization, or sequence when a user logs into the Opsware Command Center.

See "Defining CDR Services, Synchronizations, and Sequences" on page 559 in this chapter for more information.

5 Verify that the following port is accessible between the server you will push code from and the server where you will push code to:

- `telnet <staging_server> 1002`
- `telnet <production_server> 1002`

6 Configure email notification addresses.

Specify the email addresses where notifications are sent when users request that a service operation, synchronization, or sequence be performed on their behalf.

7 Test CDR setup and configuration.

After all services, synchronizations, and sequences are defined, and user accounts and permissions are set up in the Opsware Command Center, test the operations available for each service, synchronization, and sequence defined in CDR. Uploading both code and content changes from your site development environment, verify that CDR can be used to update services on all staging and production hosts for which synchronizations are defined.

Determining Your Code and Content Deployment Requirements

Discover your exact deployment requirements, and determine the responsibilities that users who are assigned to performing synchronizations and other service operations will have.

Depending on the setup of your site, you might want certain users to perform routine content updates to your site and assign responsibility for more critical application code changes to other users who will, for example:

- Perform service operations for your production site
- Synchronize updated code and content to your production site
- Run sequences that perform a sequence of service operations and sequences as a single step

CDR lets you send email requests to specific users, notifying them to perform a synchronization, sequence, or other service operation.

The options that are available to users when they access CDR depend on the user groups and permissions the users have been assigned.

See the *Opsware System 4.7 Administration Guide* for information about how to create users and assign Opsware Command Center permissions.

Planning Your CDR Configuration

Before you can use CDR, define the services and synchronizations you need to update and maintain your site. You define individual services based on each specific Web server or application server application (for example, WebLogic Server) that is installed on the staging and production hosts. You define synchronizations so that you can update files for a given service between the source location and one or more destination production hosts.

To define CDR services and synchronizations, you need to know:

- What the code and content directories are for each host
- Which hosts for your site are staging, production, and QA
- What services (for example, Web server or application server programs) are installed on each server

When you log in to the Opsware Command Center, CDR displays predefined services and synchronization that are available for your site. You see only the services and synchronization that you have authorization to perform because of your user group membership.

The operations that you need to perform are specific to the service (Web server or application server instance) for which you are updating code or content and to the particular host.

Before you use CDR to define services and synchronizations, you should document and diagram your site's configuration. That way, when you start defining services and synchronizations, the process is likely to go more smoothly.

To plan your CDR configuration, complete the following tasks for every instance of a service (for example, WebLogic application server or iPlanet Web server):

1. Create a diagram of your site's host configuration for the service, designating the source and destination hostnames for any synchronizations that you want to define.

Figure 12-4 shows an example of a typical synchronization diagram defined for a specific service, in this case, a WebLogic (jsp) application instance.

Figure 12-4: Service Synchronization Source and Destination Hosts

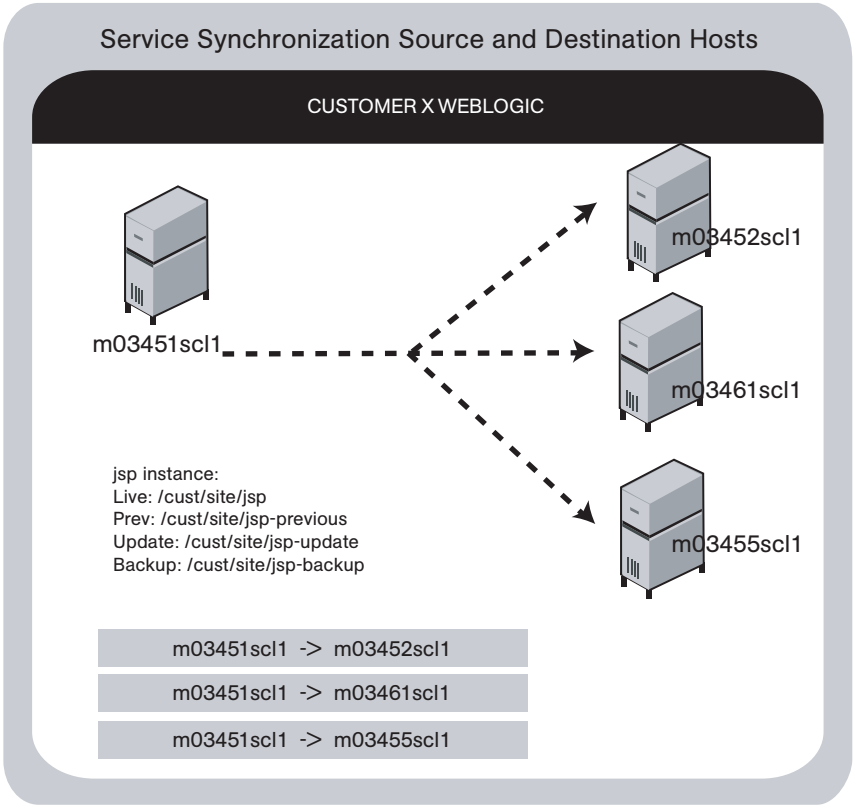

In the diagram, all three arrows are part of the same synchronization and specify update paths from the source to the destination host. The diagram also specifies the site's Live, Previous, Update, and Backup directories used by the instance in performing synchronizations, backup, restore, and rollback operations.

2 List the service directories, scripts, and any special operations or procedures to be performed for any synchronization of code or content for that service.

Table 12-1shows the information that you can specify for a service defined in CDR.

Table 12-1: Table of Information to Specify for Code Deployment Service

SECTION OF PAGE	FIELD NAME
Service	
	Name (required)
	Type (required)
Service Commands	
	Start
	Stop
	Pre-cutover
	Post-cutover
	Pre-rollback
	Post-rollback
	Pre-Sync To Update
	Post-Sync To Update
	Pre-Sync To Live
	Post-Sync To Live
	Pre-backup
	Post-backup
	Pre-restore
	Post-restore
Service Directories	
	Live Directory (required)
	Update Directory (required)
	Backup Directory (required)
	Previous Directory (required)

Table 12-1: Table of Information to Specify for Code Deployment Service

SECTION OF PAGE	FIELD NAME
Service Hosts	
	Hosts
Roles	
	Perform Role Name (required)
	Request Role Name (required)
Service Options	
	CC Operation Requests To

3 Determine the name that you want to give the synchronization when you create it by using CDR, for example, WebLogic Sync (Staging to Production). You should designate the user groups whose members can request or perform the synchronization.

Table 12-2 shows information that you can specify for synchronizations defined in CDR.

Table 12-2: Table of Information to Specify for Synchronizations

SECTION OF PAGE	FIELD NAME
	Name (required)
	Associated Service Name
	Source Host Type (required)
	Source Host (required)
	Destination Host(s) Type (required)
	Destination Host(s) (required)
	Perform Role Name (required)
	Request Role Name (required)
Options	
	CC Operation Requests To
	Strict Synchronization

4 Repeat the process to create additional diagrams for each instance available in your site environment for which you want to be able to push code or content.

Documenting your plans for code and content deployment for your site will simplify the process of defining new services and synchronizations when you access CDR. Distributing this information to other users provides useful documentation of your site's configuration, so that everyone involved in code and content deployment for your site understands what services are defined and what synchronizations are available to push code and content.

Preparing Opsware Host Machines

After you determine the CDR configuration of services and host machines for your site, perform the following tasks:

- Prepare each host machine in that configuration so that CDR can perform the synchronizations and service operations that you define.
- Create or verify the existence of Live, Previous, Update, and Backup directories on all source and destination hosts.

Creating or Verifying Directories on Hosts

Before you perform synchronization, you must create or verify that the Live Directory, Previous Directory, Update Directory, and Backup Directory already exist on all source and destination hosts. Using the list of host machine names that you collected for your site, log in to each of the hosts and check that all the directories already exist or create them.

To determine disk space requirements for a site, you can estimate that CDR requires between two and four times the total size of code and content installed on a particular host, depending on CDR usage factors such as number of changed files between synchronizations and cutovers, use of backup features, and so forth.

- The Live Directory is the directory used for deployed code and content on your site (for example, `/cust/docs` for a Web server, `/cust/site` for an application server like WebLogic).

- The Previous Directory is the directory that records the difference between the current Live Directory's code and content and the version of the site as it existed prior to the last cutover.
- The Update Directory is the directory written to by a CDR synchronization that records the difference between the Live Directories on the source and destination systems.
- The Backup Directory is the directory used to store files when users request a backup copy of the current code and content Live directory.

Ownership of the Live, Previous, Update, and Backup directories is not important because CDR software used to perform service operations and synchronizations runs as root.

Populating Initial Content in Directories

After you create directories on all the hosts designated for a new site deployment (staging, QA, production hosts), you should populate the Live directories on each host with the initial code and content for the site.

Archive a copy of the site files and use a file transfer utility to perform the initial site upload. Using CDR synchronization to initially populate directories has significant overhead and might take longer to perform than directly copying the initial site code and content.

This step (populating directories) only applies to setting up new sites, because current code and content for your site is already available on the Opsware staging and production hosts.

In a Unix environment, you can tar the files and use scp to perform the initial site upload into each host Live directory. In a Windows environment, use the Windows file transfer utility to do the initial site upload when you configure host machines for a new deployment.

Setting up Access Control for CDR

CDR uses the Opsware Command Center authentication to control users' access and ability to perform service operations and synchronizations. Specific permissions to perform code deployment operations are based on a user's membership in predefined

CDR user groups, which an Opsware administrator defines in the Administration section of the navigation panel. Table 12-3, Table 12-4, Table 12-5, and Table 12-6 provide descriptions of permissions that are associated with predefined CDR user groups.

Table 12-3: Special Code Deployment User Groups

CDR USER GROUP	DESCRIPTION
Super-User	Users in this user group can define, request, or perform any code deployment operation on hosts for any customer.
History Viewer	Users in this user group can view a log of operations (service operations, synchronizations, and sequences) that were executed from the Code Deployment subsystem. Viewing this information can help you determine the completion status of particular deployment operations.

Table 12-4: Service User Groups

CDR USER GROUP	DESCRIPTION
Service Editor	Users can define services and modify or delete service definitions.
Service Performer (Production)	These users directly perform or request performance of service operations on hosts designated for use in production.
Service Performer (Staging)	These users directly perform or request performance of service operations on hosts designated for use in staging.
Service Requester (Production)	These users directly request performance of service operations on hosts designated for use in production.
Service Requester (Staging)	These users request performance of service operations on hosts designated for use in staging.

Table 12-5: Synchronization User Groups

CDR USER GROUP	DESCRIPTION
Synchronization Editor	Users can define a synchronization, modify, or delete the synchronization definition.
Synchronization Performer	These users directly perform or request performance of synchronization actions.
Synchronization Requester	These users request performance of synchronization actions.

Table 12-6: Sequence User Groups

CDR USER GROUP	DESCRIPTION
Sequence Editor	Users can define sequences, and modify or delete sequence definitions.
Sequence Performer (Production)	These users directly perform or request performance of sequences of actions on hosts designated for use in production.
Sequence Performer (Staging)	These users directly perform or request performance of sequences of actions on hosts designated for use in staging.
Sequence Requester (Production)	These users request performance of sequences of actions on hosts designated for use in production.
Sequence Requester (Staging)	These users request performance of sequences of actions on hosts designated for use in staging.

When a user submits CDR requests asking that a service operation or synchronization be performed on the user's behalf, an email notification is sent to the individuals assigned to perform the requested service operation or synchronization. See the *Opsware System 4.7 Administration Guide* for information about how to assign users to predefined CDR user groups.

Each user group is created without any users initially added. Your Opsware administrator can add individual users to each CDR user group to control their permissions to request or perform service operations, synchronizations, and sequences.

See the *Opsware System 4.7 Administration Guide* for information about how to add users to CDR user groups.

When a user selects a CDR option, the Opsware System determines the user's user group memberships and determines what service and synchronization actions the user can perform. Depending on user group membership, the user can either (1) perform or request performance of a service management operation or synchronization operation, or (2) request that the operation be performed by users specified in an email notification list.

Defining CDR Services, Synchronizations, and Sequences

By using the list of services and synchronizations that you have planned for your site, you can use Code Deployment & Rollback to create the corresponding service and synchronization definitions in the Opsware Command Center.

You should follow this process:

1. Create all the services required for your site. Each service is defined in terms of the commands such as start or stop that are required for the associated service instance.
2. After you define all services, create all synchronizations that you want to make available. Each synchronization references a specific service and specifies the source host from which the service's directories and files are to be synchronized to one or more destination hosts.
3. After you define services and synchronizations, define sequences to specify a sequence of specific service and synchronization operations that you want to perform as a unit.

See "Planning Your CDR Configuration" on page 551 in this chapter for information about how to define the services and synchronizes that you need to create for a particular site.

Defining and Modifying CDR Services

The CDR Service Management option lets you create new services or modify or delete existing services. For example, if you have a single instance of a service, you can use CDR to define a single CDR service. If you have five instances of a service, you can define five individual CDR services.

You should also create different services to provide control over services performed on staging hosts versus production hosts. For example, you could define a service that only names staging hosts and specify a *perform* user group for users who can perform operations for those hosts. You could then define a second service that names all hosts (both staging and production) and limit the perform user group to selected users.

In CDR, every service is defined down to the level of a single command that needs to run during service operations, such as start, stop, pre-cutover, post-cutover, and so forth. Both services and service instances are defined the same way because conceptually there is no difference between them. For example, you can define an Apache service, or several instances of ATG Dynamo or BEA WebLogic in terms of the scripts required to start or stop the service and scripts to perform at pre-cutover, post-cutover, and so forth.

Defining a Service

If you are invoking Python in a CDR service command, you must invoke Python by using a fully qualified path to `python.exe` in the command. If you are migrating to the current version of the Opsware System from a previous version, you must update any currently defined CDR service commands.

Perform the following steps to define a service:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Click the Service Management option.
3. Click the Define a New Service link. The CDR Service Name and Type page appears, as Figure 12-5 shows.

Figure 12-5: CDR Service Name and Type

4. Specify the name of the service by choosing a name that users can identify with the corresponding application instance, for example, WebLogic (EJB Instance).

5 Specify the type of service that you want to create by selecting the service type from the drop-down list, as Figure 12-6 shows. The drop-down list includes the names of all application instances defined in the Model Repository.

Figure 12-6: CDR Service Commands

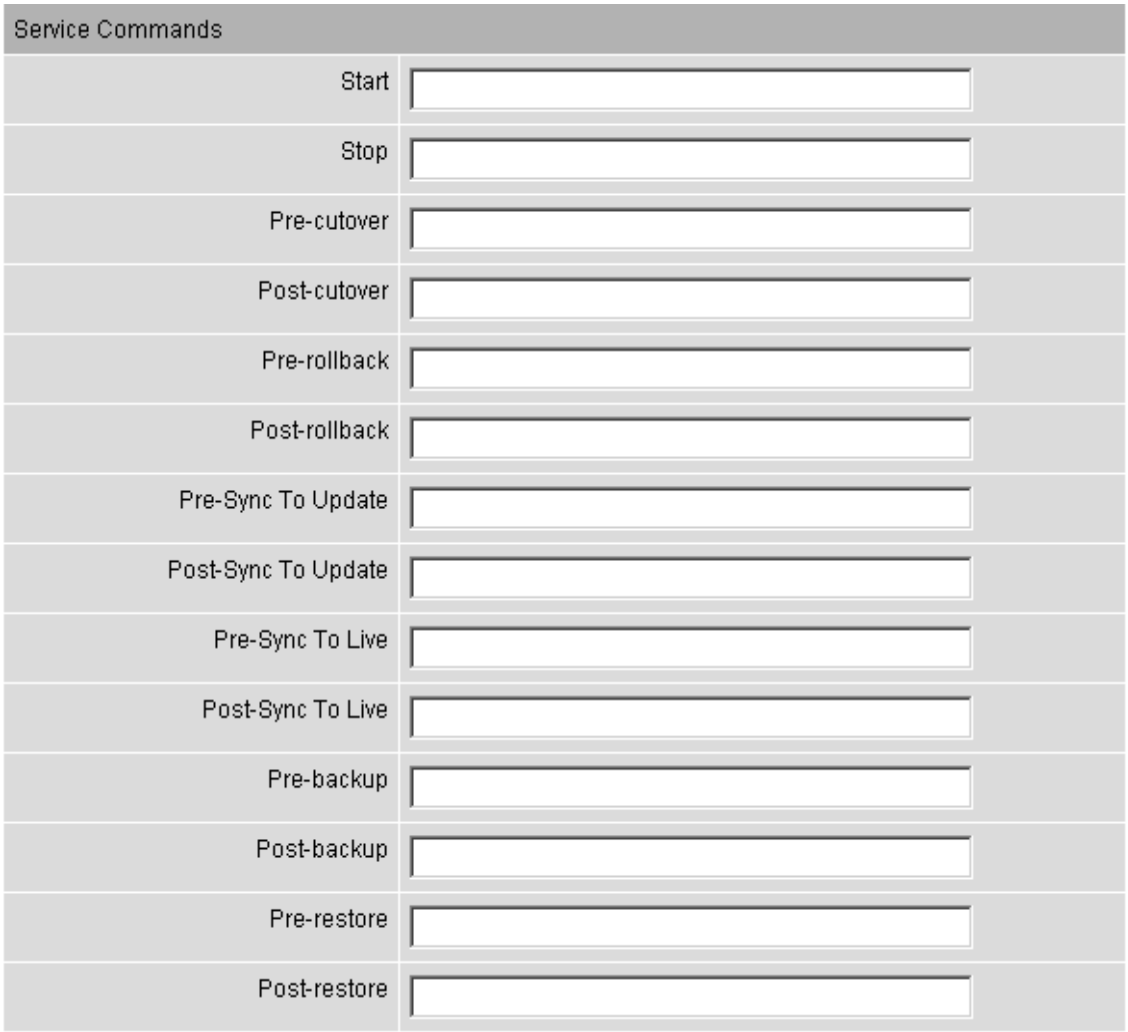

6 In the Service Commands section, enter any commands to perform for the specific service or service instance. In each case, you can enter a single command (specifying a fully qualified path) that is run to effect the operation. The same commands and scripts are applied for all hosts where the service is installed.

- Start and Stop fields – specify single commands or scripts that are executed when users choose the Service Management option to start and stop a specified service.
- Pre-cutover and Post-cutover fields – specify single commands or scripts that are executed before and after a user chooses the Run Service option to cutover code and content changes to Live.
- Pre-Rollback and Post-Rollback fields – specify single commands or scripts that are executed before and after a user chooses the Service Management option to restore code and content in a service's Live directory from the service's Rollback directory on specified hosts.
- Pre-Sync to Update and Post-Sync to Update fields – specify single commands or scripts that are executed before and after a user chooses the Synchronize option to synchronize code and content changes to the Update directory on specified hosts.
- Pre-Sync to Live and Post-Sync to Live fields – specify single commands or scripts that are executed before and after a user chooses the Synchronize option to synchronize code and content changes to the Live directory on specified hosts.
- Pre-Backup and Post-Backup fields – specify single commands or scripts that are executed before and after a user chooses the Service Management option to back up code and content from a service's Live directory to a Backup directory on specified hosts.
- Pre-Restore and Post-Restore fields – specify single commands or scripts that are executed before and after a user chooses the Service Management option to restore code and content from a service's Backup directory to the Live directory on specified hosts.

7 In the Service Directories section (see Figure 12-7), specify the disk locations for Live, Update, Previous, and Backup directories used by the service (common for all hosts where a given service is installed).

- Live Directory – the directory that stores the actual code or content required by a specific service to run a live site.
- Update Directory – the directory written to by CDR synchronizations. Stores files that changed between the source host Live directory and the Live directories on destination hosts where the service is installed.

- Backup Directory – the directory written to by CDR backup operations; used by the Restore option to return a service's Live directories to the code and content of a previous backed up version.
- Previous Directory – the directory written to by CDR cutover operations; used by the Rollback option to return a service's Live directories to the code and content that existed prior to the last performed synchronization.

Figure 12-7: CDR Service Directories

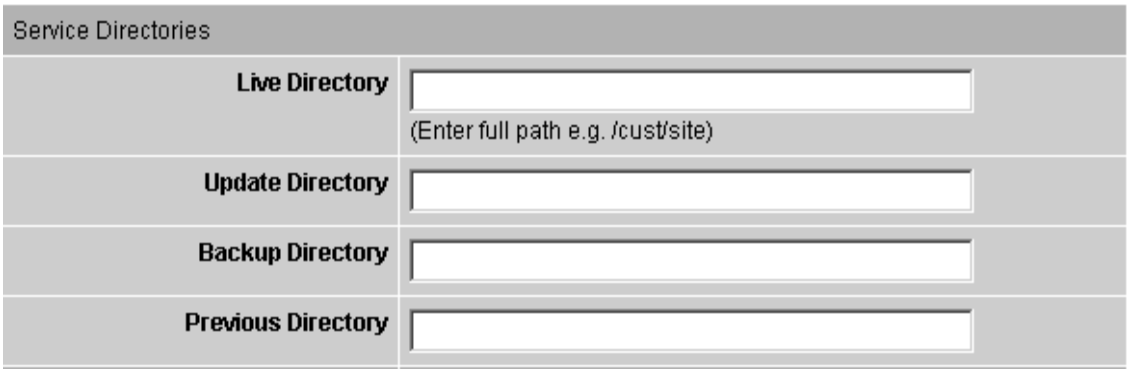

8 In the Service Hosts section, select all hosts on which this service is running. You can use the Shift and Control keys to select multiple hosts. See Figure 12-8.

These servers have a use field that has Code Deployment selected in Server Attributes.

The servers also have a state of OK. If you changed the use of a server by using the Opsware Command Center, click the Refresh button to update the host list.

Figure 12-8: CDR Service Hosts

9 In the Roles section, specify the CDR user groups whose members you want to perform or request operations for the specific service. The Perform Role name determines the user group whose members can perform or request that select staff, or your Operations Center, perform a specific operation associated with the service. The Request Role Name specifies user groups whose members can request only an operation, such as start or stop for a service. See Figure 12-9.

See "Setting up Access Control for CDR" on page 556 in this chapter for information about the description of CDR user groups that you can specify for the Perform Role and Request Role names.

Figure 12-9: CDR Roles for Performers and Requesters

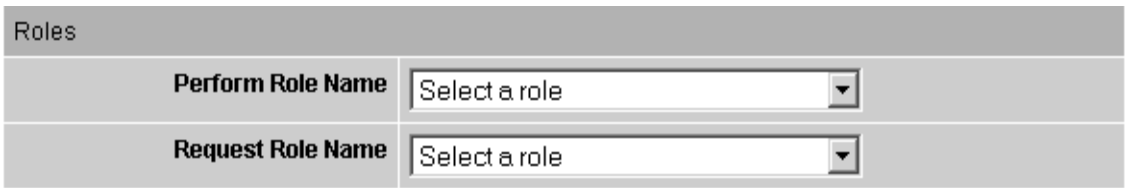

Figure 12-10: Email Addresses to Copy CDR Operation Requests to

10 In the Service Options section (see Figure 12-10), specify any email address contacts that you want to notify for any service operation requests.

Specifying email notifications allows flexibility in assigning requests to select members of your staff or your Operations Center.

11 When you finish making entries to define a new service, click the Save button.

CDR verifies that the service name you specified is unique and then saves the new service definition data in the Model Repository.

To save defined services, you must select at least one hostname and provide entries for the Service Name, Service Type, Start Service, Stop Service, Perform Role, and Request Role fields.

Running Pre- and Post-Synchronization Scripts

Pre- and post-synchronization scripts only run on destination hosts.

On Windows machines, you can use the post-cutover command, for example, to specify a command that performs Windows object registration (among other tasks). In that case, you might define a post-cutover script that (1) lists all files in a directory and (2) passes all files with the .dll extension to `regsvr32.exe` and passes all files with the .msi extension to `msiexec.exe`. Performing these steps registers and un-registers COM objects. A similar script can be developed to de-register and register COM+ objects. This script can be named in CDR and must then be placed on all hosts on which the service could run.

You can specify the instance name as a command line argument in Start and Stop commands or scripts for services that describe instances of the same service running on the same hosts. (You need to create different services for each service instance running on the same hosts because the directories and start and stop script calls used by each instance are different.)

Start and Stop and other service command or script entries: If you need to perform operations that require more than a single command, you should define a sequence of commands in a single script file and then specify that script in the CDR service definition.

Modifying a Service

Occasionally, you need to modify an existing service, for example, to change assigned hosts, make updates to scripts, or make other changes to the attributes of the service.

Perform the following steps to modify a service:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Click the Service Management option.
3. Click the Modify an Existing Service link.
4. Click the name of the service that you want to modify.
5. Update the field entries that you want to modify, and then click OK. A confirmation page appears.

   You can modify all field entries that define a service except for the Service Type field. If you modify the Service Name field to rename a service, CDR confirms that the new name is not already in use.

CDR deletes synchronizations associated with a service when the following modifications are made:

- When a user removes a hostname from the list of hosts defined for a service and that hostname is a source for a synchronization, that synchronization is also removed when the service definition is saved. If that synchronization is used by a sequence, then that sequence is also removed.
- When a user removes a hostname from the list of hosts defined for a service, and that hostname is the last remaining destination for a synchronization, that synchronization is also deleted when the service definition is saved. If that synchronization is used by a sequence, then that sequence is also removed.
- When a user removes a hostname from the list of hosts defined for a service, and that hostname is the last host in a sequence step, then the whole sequence is deleted when the service definition is saved.

Deleting a Service

CDR allows you to delete services and remove their stored definition from the Model Repository.

Perform the following steps to delete a service:

1. If necessary, click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Click the Service Management option.
3. Click the Delete a Service option.
4. Select the check boxes next to the services that you want to delete and click Delete.

   CDR prompts you to confirm the deletion.
5. Click OK. CDR removes the services that you chose to delete.

If you request deleting a service definition, CDR displays a confirmation box that indicates that any associated synchronizations or sequences are also deleted when it deletes the service.

Creating and Modifying CDR Synchronizations

The CDR Sync Management option lets you create, modify, or delete synchronizations so that you can update files for a given service between a source host location and one or more destination hosts. For example, in setting up a synchronization for a WebLogic application server instance, you can create a synchronization to transfer updated files between a staging host and the production host machines used to run your site.

When defining a synchronization, you first select the service, then specify the source and destination hosts you want to synchronize. In addition, you specify the CDR user groups that can perform or request the synchronization. You can also specify options such as addresses for email notification of synchronization requests and how synchronizations are performed (Strict Synchronization transfers updated files and removes deleted files from destination hosts.)

Defining a Synchronization

Perform the following steps to define a synchronization:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Click the Sync Management option.
3. Click the Define a New Synchronization option.
4. Select the service to which you want to add a synchronization.

   CDR displays a page on which you can define a new synchronization, choose source and destination hosts, and specify other synchronization options. See Figure 12-11.

Figure 12-11: Define a New Synchronization Page

5 Specify the name of the synchronization, choosing a name by which users can identify the type of synchronization being performed, for example, WebLogic Sync (Staging to Production).

6 Specify the Source and Destination Host Types, choosing the type from the drop-down lists, which display all values stored in the Model Repository. These values are editable using Server Attributes.

You need to specify a Host Type before any hosts are displayed in the Source or Destination Host lists.

7 Specify the single Source Host for the synchronization from the list of hosts stored in the Model Repository that match the value that the Source Host Type specified.

8 Specify one or more Destination Hosts for the synchronization from the list of hosts stored in the Model Repository that match the value that the Destination Host Type specified.

Use the Shift and Control keys to select multiple destination host machines.

9 In the Perform Role Name and Request Role Name fields, select the CDR user groups that you want to allow to perform or request operations for the synchronization. The Perform Role Name determines the user group whose members can perform, or request that another member perform a particular synchronization. The Request Role Name specifies user groups whose members can request that authorized individuals perform synchronizations.

See “Setting up Access Control for CDR” on page 556 in this chapter for information about a description of CDR user groups that you can specify for the Perform Role and Request Role Names.

10 In the Synchronization Options section, specify any email address contacts that you want notified of any synchronization requests.

11 Click the Strict Synchronization check box to specify that files deleted from the source host are also removed from corresponding directories on destination hosts defined in the synchronization. (Otherwise, if unchecked, the synchronization affects only files that are new or have changed between the source host and destination hosts and files removed from a source host are not removed from destination hosts.)

12 When you finish making entries to define a new synchronization, click the Save button. CDR verifies that the synchronization name that you specified is unique and then saves the new synchronization definition data in the Model Repository.

To save a new synchronization, you must specify a unique synchronization name, the source host and at least one destination host, and user groups that can perform or request synchronizations.

Modifying a Synchronization

Occasionally, you need to modify an existing synchronization, for example, to change source or destination hosts or make other changes to attributes of the synchronization.

Perform the following steps to modify a synchronization:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Click the Sync Management option.
3. Click the Modify an Existing Synchronization option.
4. Click the name of the synchronization you want to modify.
5. Update the field entries that you want to modify, then click OK. A confirmation page appears.

When you modify a synchronization by removing a host from the list of destination hosts, and that host is the last host in a synchronization sequence step, then that sequence is removed.

You can modify all field entries that define a synchronization except for the Source and Destination Host Type fields. If you modify the Synchronization Name field to rename a synchronization, CDR confirms that the new name is not already in use.

Deleting a Synchronization

CDR allows you to delete synchronizations and remove their definition from the Model Repository.

Perform the following steps to delete a synchronization:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Click the Sync Management option.
3. Click the Delete a Synchronization option.

4 Select the check boxes next to the synchronization that you want to delete and click Delete.

CDR prompts you to confirm the deletion.

5 Click OK. CDR removes the synchronizations that you chose to delete.

Deleting synchronizations that are used by a sequence causes that sequence to be deleted.

Creating and Modifying CDR Sequences

The CDR Sequence Management option lets you create, modify, or delete sequences of service operations and synchronizations in a sequence so that you can define meta-operations for CDR.

For example, you can define a sequence to push code from staging to production hosts, stop, cutover, and start a service. A sequence is defined in two parts: the properties of the sequence itself (name, user groups, and so forth) and the steps of the sequence.

Defining a Sequence

Perform the following steps to define a sequence:

1 Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2 Click the Sequence Management option.

3 Click the Create a New Sequence option.

CDR displays a page on which you specify the name of a new sequence, the user groups for the performer and requester for this sequence and email information. See Figure 12-12.

Figure 12-12: Define Sequence Page

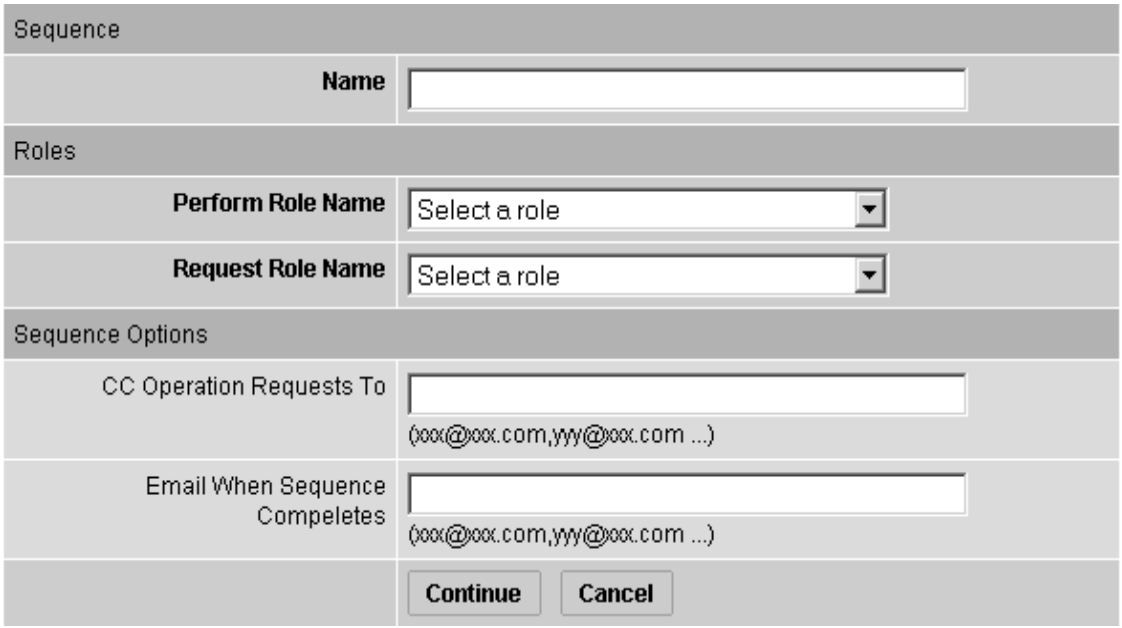

4. Specify the name of the sequence, choosing a name that users can identify with the corresponding operation that this sequence will perform, for example, Push Code to Production.

5. In the Roles section, specify the CDR user groups that you want to allow to perform or request execution of the specific sequence. The Perform Role name determines the user group whose members can perform or request that select members of your staff, or your Operations Center, perform this sequence. The Request Role Name specifies user groups whose members can request a sequence.

   See "Setting up Access Control for CDR" on page 556 in this chapter for information about a description of CDR user groups that you can specify for the Perform Role and Request Role Names.

6. In the Sequence Options section, specify any email addresses to whom you want to send sequence operation requests.

7. You can also specify email addresses to which notifications can be sent when the sequence is performed and completed. The email contains the status of each step of the sequence that was performed and an indication if it ran successfully.

8 Click the Continue button to save the sequence properties.

To save defined sequences, you must provide entries for the Sequence Name, Perform Role, and Request Role fields.

9 A small window pops-up on your screen. Use this window to select the operations to add to the sequence. First select the name of the service that you want to operate on in the Service Name drop-down menu. See Figure 12-13.

If you use a pop-up blocker, this window will not pop up. You can access it by clicking the hyperlinked word popup in the sentence that says, “Add new operations with the popup window.”

Figure 12-13: Sequence Operation Selection Window

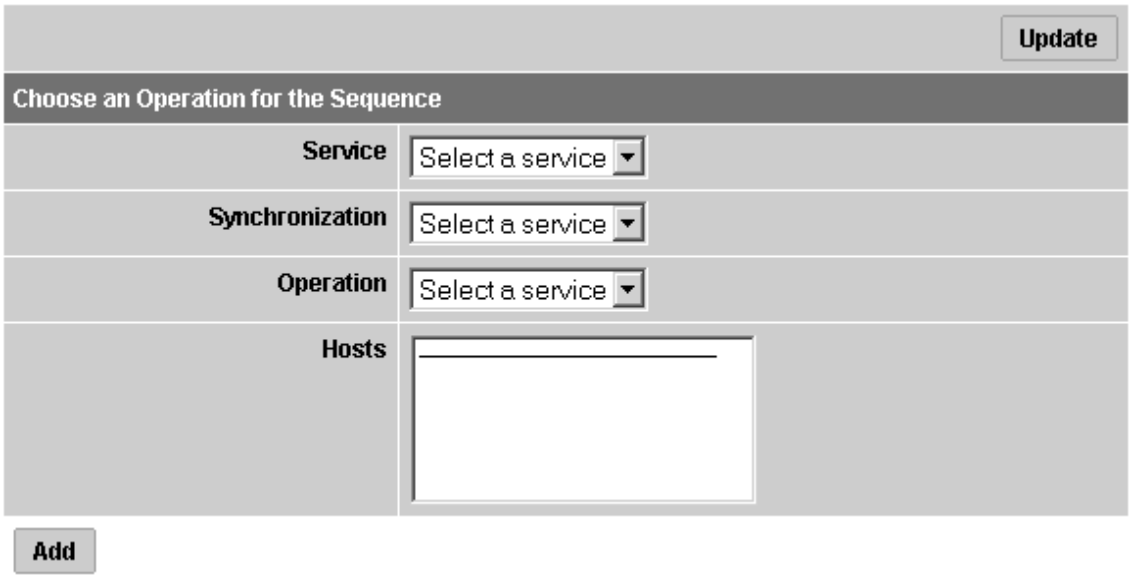

10 To add services or synchronizations to a sequence:

- In both cases, first select a service from the Service drop-down menu, then select a service from the Synchronization drop-down menu.
- To add a synchronization operation, select Synchronize to Update or Synchronize To Live from the Operation drop-down menu, then select one or more destination hosts

for the synchronization from the Hosts select box. Finally, click the Add button. The information about the newly added step appears in the main window.

- To add a service operation, select None from the Synchronization drop-down menu. Then, select the name of the service operation that you want to add from the Operation drop-down menu, and select the hosts that you want to perform the service operation on in the Hosts select box. Finally, click the Add button. The information about the newly added step appears in the main window.

11 Click the Save button to save the sequence.

Modifying a Sequence

Occasionally, you need to modify an existing sequence, for example, to change assigned hosts in a step, add a step, or make other changes to attributes of the sequences.

Perform the following steps to modify a sequence:

1 Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2 Click the Sequence Management option.

3 Click the Modify an Existing Sequence option.

4 Click the hyperlinked name of the sequence that you want to modify.

5 Update the field entries that you want to modify and then click Continue.

6 Edit any of the sequence steps that you want.

7 Click Save to save the changes.

Deleting Sequences

CDR also allows you to delete sequences and remove their stored definition from the Model Repository.

Perform the following steps to delete a sequence:

1 Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2 Click the Sequence Management option.

3 Click the Delete a Sequence option.

4. Select the check boxes next to the sequences that you want to delete and click Delete.
5. CDR prompts you to confirm the deletion.
6. Click OK. CDR removes the sequences that you chose to delete.

Deleting sequences has no impact on defined services or synchronizations.

Verifying and Troubleshooting CDR Configuration

After you set up all the CDR services and synchronizations required for your site, and perform all other setup required on host machines in either your development environment or the Opsware managed environment, verify operation of the complete configuration.

The following list provides the steps to follow to verify your CDR configuration:

1. Log in to the Opsware Command Center with permissions to perform service operations and synchronizations.
2. If necessary, click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
3. Modify files in your staging host's Update (source) directory to enable testing synchronizations.
4. Perform all defined synchronizations. After completing the synchronizations, verify that files, which were modified on your staging source host, were modified correctly in the directories of each destination host.
5. Perform all service operations for each defined service to verify the operations of scripts for starting, stopping, cutting over, backing up, restoring, and rolling back updates. Also verify that all pre- and post-operations were successful.
6. Verify any sequences that you defined, executing each sequence and then checking that the operations complete successfully.

Performing Services, Synchronizations, and Sequences

This section provides information on performing services, synchronizations, and sequences within the Opsware System and contains the following topics:

- Performing Services, Synchronizations, and Sequences Overview
- Synchronization of Site Code and Content
- Performing Synchronizations
- Cutover to Changed Code and Content
- CDR Service Operations
- Starting and Stopping Host Services
- Backing Up Code and Content
- Restoring Code and Content from a Previous Version
- Rolling Back Code and Content to the Previous Version
- Accessing Service Operations in CDR
- Performing Service Operations by Service Name
- Performing Service Operations by Hostname
- Performing Sequences
- Processing Code Deployment Requests from Users
- Performing Synchronizations and Service Operations
- Viewing Status of Previous Operations

Performing Services, Synchronizations, and Sequences Overview

After you upload updated code and content to your Opsware staging environment, you use Code Deployment & Rollback to cutover to new code and content, perform host synchronizations, and perform other service operations.

When you cutover to new code and content, CDR determines the differences between the new code and content in the current Update directory and the Live directory. The files that are different are synchronized to the Live directory. When you synchronize source and destination hosts, CDR moves modified files from the Live directory on a source host to a directory on a destination host.

The code and content footprint that CDR maintains on a host is larger than the storage space for the code and content files of the live site. The amount of storage used beyond the actual size of your site increases with the number of files that you modify and back up.

Synchronization of Site Code and Content

CDR uses the following directories to synchronize and cutover code and content for specified hosts:

- **Live directory** – The directory that stores the actual code and content required to run a live site.
- **Update directory** – The directory written to by CDR synchronizations. Stores only the files that changed between the source host Live directory and the Live directories of the destination hosts.
- **Site Previous directory** – This directory holds all the changes necessary to revert the Live directory back to the state it was in before the last cutover. Like the Update directory, the Previous directory only stores the files that changed between the current Live directory contents and its previous state.
- **Site Backup directory** – This directory stores a complete backup of the site. The directory is populated when the user issues a Backup service operation.

CDR can synchronize updates from the Live directory of a source host to either the Update or Live directories on destination hosts. Decide whether you want to synchronize changed files on the source host to Update or Live directories on the destination hosts:

- Synchronize to Update directories – CDR determines files that have changed by comparing files in the destination host Live directory and the source host Live directory. Updated files are stored in the Update directories of destination hosts.
- Synchronize to Live directories – CDR updates changed files to the Live directory on destination hosts bypassing the Update directory.

If you choose to synchronize directly to Live directories, the Rollback operation does not function properly. Therefore, choose this option only for synchronizations that are not likely to impact site stability.

When you choose to synchronize to Live directories, back up the Live directories first and run the Restore operation to return your site to the previous version.

See "Rolling Back Code and Content to the Previous Version" on page 584 in this chapter for more information.

Figure 12-14 shows an example of synchronization between hosts and how you synchronize updates for each service.

Figure 12-14: Service Synchronization of Hosts Using CDR

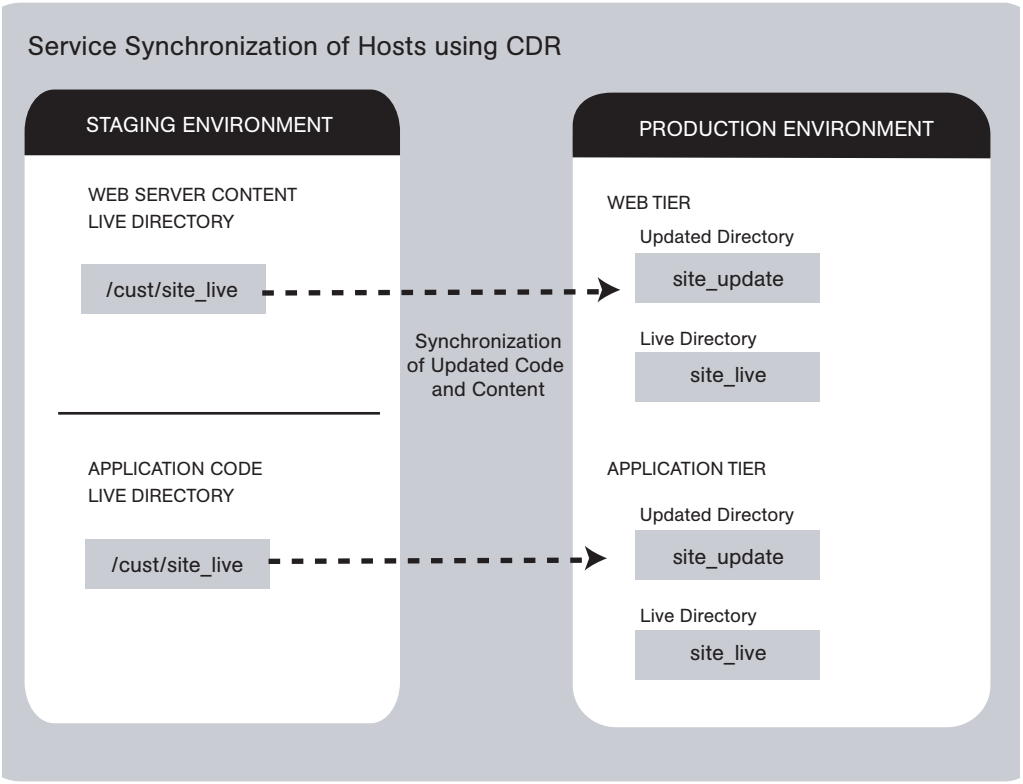

Contact your Opsware administrator for a description of services and synchronizations set up for your site and the specific operations that you need to run when updating code or content for a particular host service.

Performing Synchronizations

Perform the following steps to synchronize updated code and content from one source host to one or more destination hosts:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Click the Synchronize link.

A page appears that displays the synchronizations that you can perform.

Contact your Opsware administrator for a description of synchronizations set up for your site.

3 Select the synchronization that you want to perform by clicking the link. The CDR Synchronize for [*customer name*] page appears, as Figure 12-15 shows.

Figure 12-15: The CDR Synchronize Page

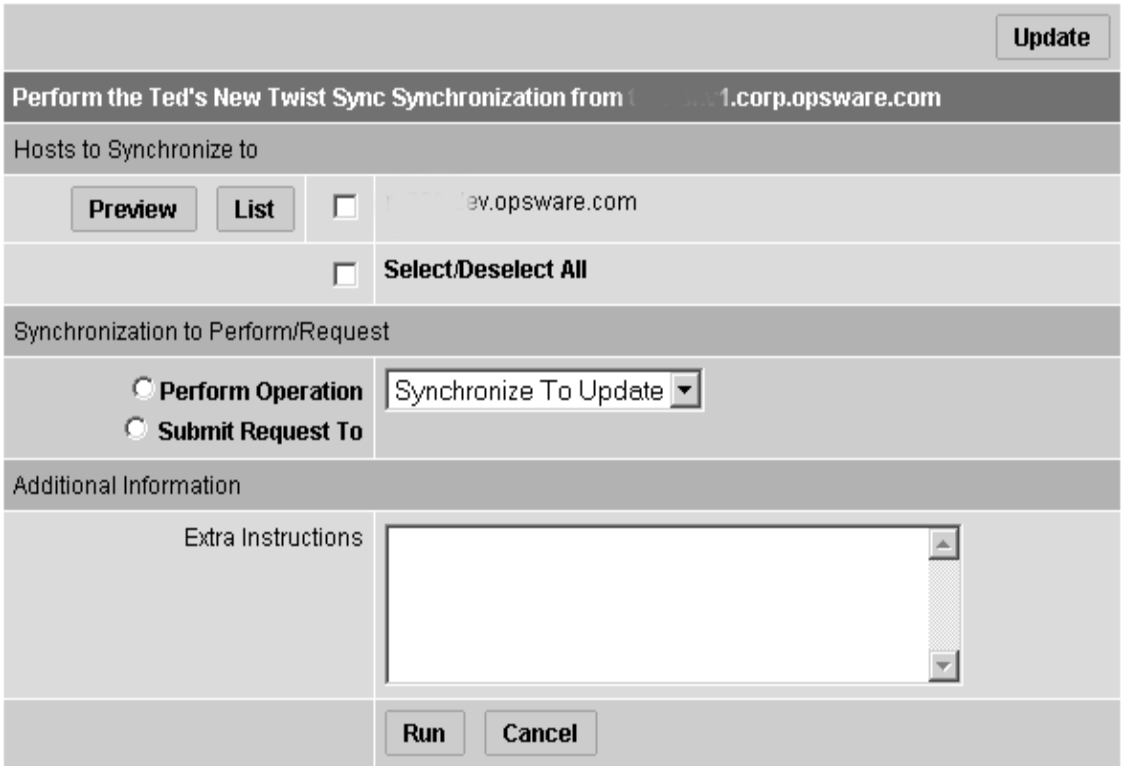

4 Select one or more of the displayed hostnames on which you want to perform synchronization. The hosts that you select are the destination hosts.

Choose the Select/Deselect All option to select all or clear all hostnames.

5 (Optional) To view a list of files that will be created, modified, and deleted on the destination hosts, click the Preview button.

OR

(Optional) To view a list of all the files on the destination hosts, click the List button.

6 Select the Perform Operation option to directly perform a selected synchronization.

Or

Select the Submit Request To option to send a request to users specified to receive email notification for service and synchronization requests. When you submit a request, specify any additional instructions that you want to include for the requested synchronization. For example, these might be instructions such as the time that you want the synchronization performed, verification, or other related services to perform.

The Perform Operation option is only visible when you are a member of an Opsware Command Center user group allowed to directly perform a synchronization.

7 Choose the type of synchronization that you want performed from the drop-down list:

- Synchronize To Update
- Synchronize To Live

See "Synchronization of Site Code and Content" on page 577 in this chapter for information about these options.

8 To initiate the synchronization or send the request, click the Run button.

Cutover to Changed Code and Content

Perform the Cutover operation to make the Update directory and the current live site identical.

When you cutover, CDR performs the following actions:

- Updates the Site Previous directory with files from the Live directory. CDR saves modified files and files-to-be-deleted to the Site Previous directory. The Site Previous directory contains the files necessary to restore the live site to the previous version.

- Determines the differences between the Update directory and the current Live directory. The files that are different are synchronized from the Update directory to the Live directory. See Figure 12-16.

Figure 12-16: Directory and File Updates for Cutover Operations Using CDR

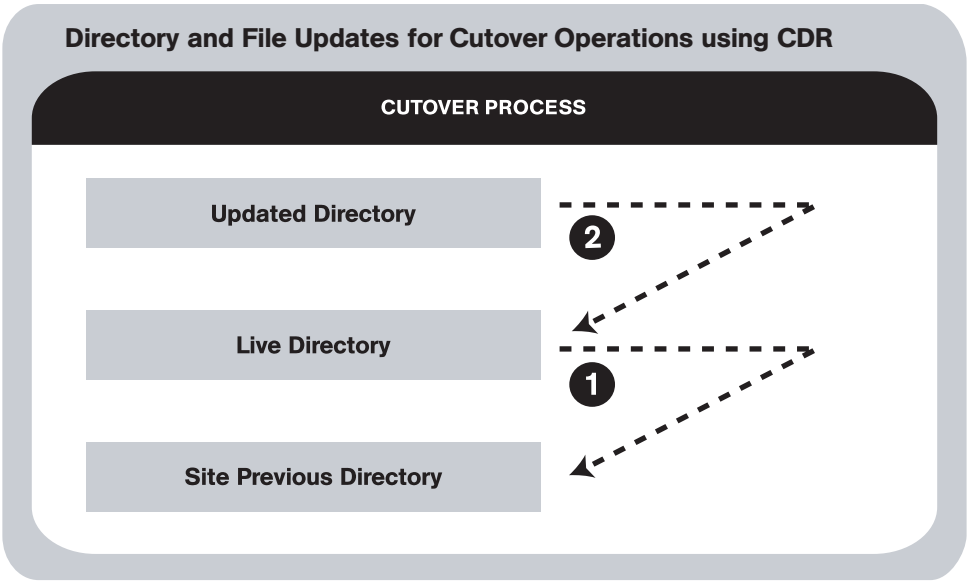

CDR determines file differences between source and destination directories based on file size, modification date and time, ownership, group and permissions attributes.

By using the cutover process, CDR ensures your ability to roll back to the previous version of your code and content if you experience a problem.

Your Opsware administrator can configure a CDR service to run scripts before and after cutting over to updated code and content. For example, before and after cutting over, you might distribute content on geographically disperse servers.

See “Synchronization of Site Code and Content” on page 577 in this chapter for information about a description of these directories. See “CDR Service Operations” on page 582 in this chapter for information about how to roll back code and content to the previous version.

CDR Service Operations

In addition to the Cutover operation, CDR provides a number of service operations:

- Starting and stopping host services
- Backing up code and content
- Restoring code and content from a previous version
- Rolling back code and content to the previous version

Performing these operations might be required depending on the type of code and content changes made or the host services that are affected.

The operations that you need to perform are specific to the service (for example, Web server or application server instance) for which you are updating code or content and the particular host.

Contact your Opsware administrator for a description of services set up for your site and the specific operations that you need to run when updating code or content for a particular host service.

Starting and Stopping Host Services

Start – launch a defined service; for example, starting a Web or application server instance that is running on a specific host.

Stop – shut down a defined service; for example, shutting down a Web or application server instance before cutting over to new or changed code and content on a specific host.

Stopping and starting services might be required depending on the type of code and content changes made or the host services that are affected. Discuss your requirements with your Opsware administrator.

Typically, you only stop and start host services for the hosts in your staging environment. Select members of your staff, or other individuals in your operations center, can stop and start host services for the hosts in your production environment.

Backing Up Code and Content

When you use CDR to back up your site, CDR saves the entire contents of the current Live directory for a specific service in the Backup directory. CDR saves the backup copy to the local disk for the host on which you ran the Backup operation. See Figure 12-17.

You can use CDR to keep only one backup copy at a time for a service.

Figure 12-17: Update Directory to Site Backup Directory for Backup Using CDR

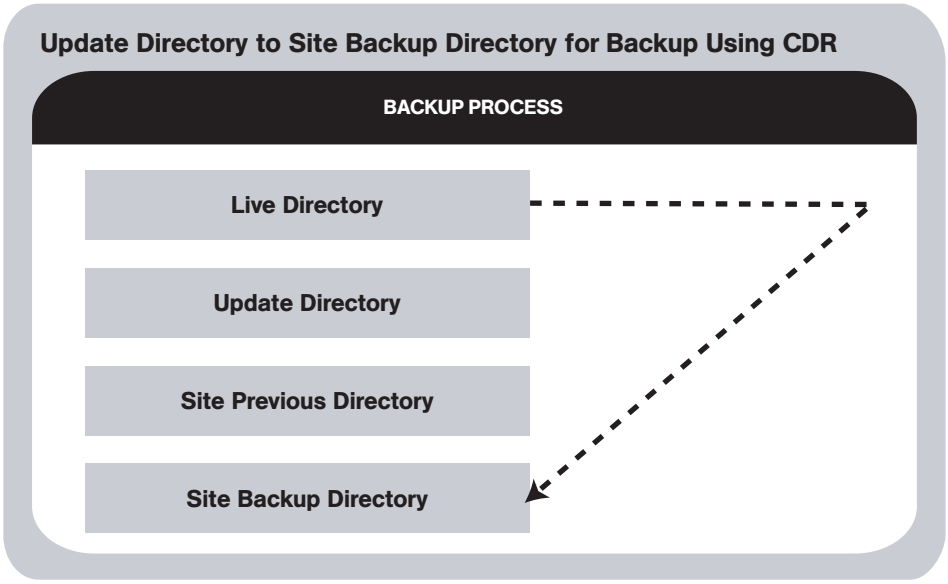

When you run the Restore operation, CDR replaces the Live directory contents with files stored in the Backup directory.

When you reach a high level of site stability, backing up your site is recommended, especially if you plan to make changes to site code and content.

Restoring Code and Content from a Previous Version

The Restore operation restores the previous Live directory by copying the contents of the Backup directory to the Live directory. Restoring code and content from the Backup directory does not change files that are stored in the Update directory. See Figure 12-18.

Figure 12-18: Site Backup Directory to Update Directory for Restoring Using CDR

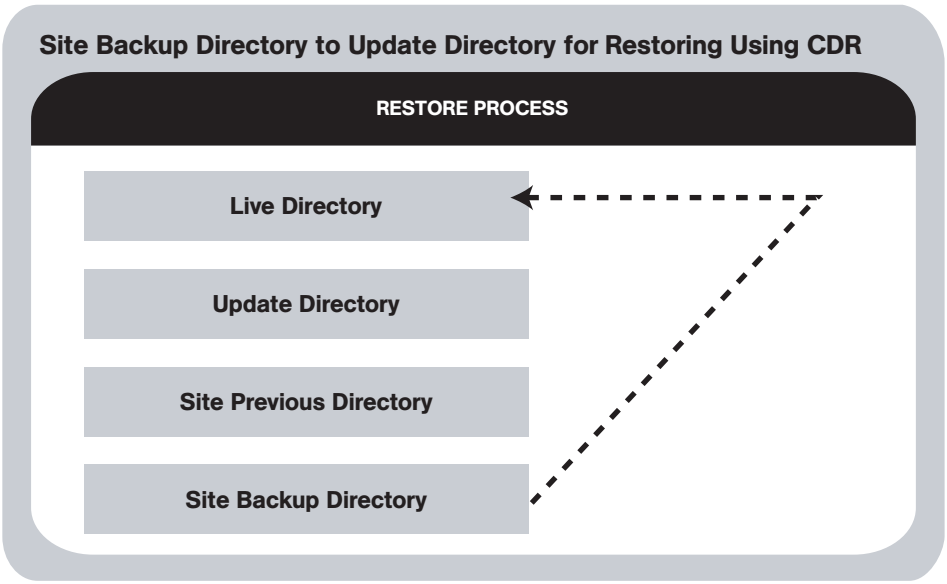

Before you restore code and content, you must have backed up the contents in the Live directory to the Backup directory by performing a Backup operation.

Rolling Back Code and Content to the Previous Version

If you experience a problem after cutting over updated code and content to your production site, you can roll back to the previous version.

Rolling back returns the site to the state it was in prior to the last cutover that you performed. See Figure 12-19.

Figure 12-19: Directory and File Updates for Rollback Operations Using CDR

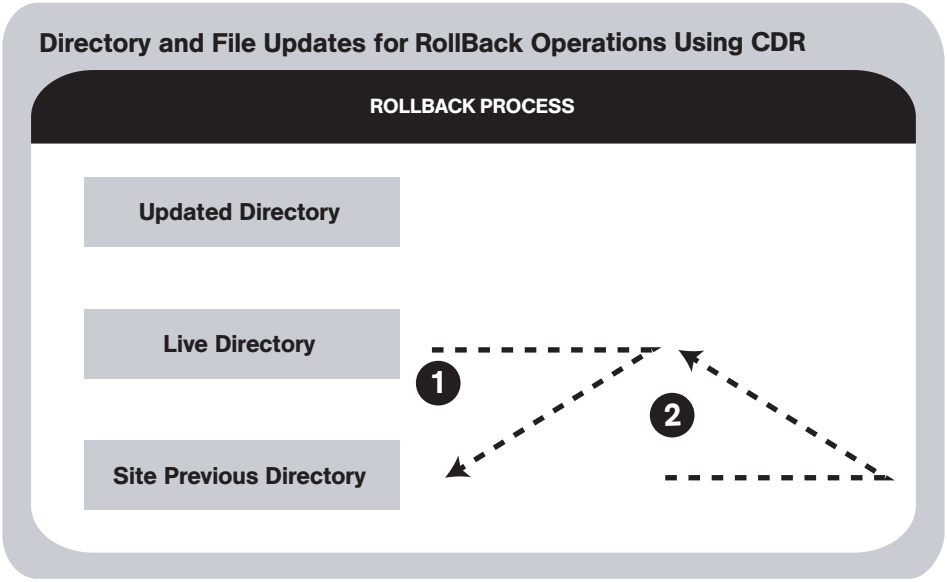

During cutover, CDR updates the Site Previous directory with files from the Live directory. CDR saves modified files and files-to-be-deleted to the Site Previous directory. The Site Previous directory contains the files necessary to restore the live site to the previous version.

During rollback, CDR restores the set of different files (modified files and files that were deleted during cutover) to the Live directory.

If you upload files directly to the Live directory or choose to synchronize directly to Live directories, the rollback operation does not function properly. Under these conditions, back up your Live directory and run the Restore operation to return your site to the previous version.

Accessing Service Operations in CDR

The Service Management option provides a number of service or administrative operations, including starting and stopping host services, backing up, restoring, or rolling back code and content, and cutting over to new site code and content.

To access Code Deployment & Rollback, your Opsware administrator must add you as a member of a user group authorized to use CDR.

You have the option of initiating a service either by selecting a service name or by selecting hostnames. You must select both a service to perform and the hosts on which to perform the service. See Figure 12-20.

Figure 12-20: Run Service Page

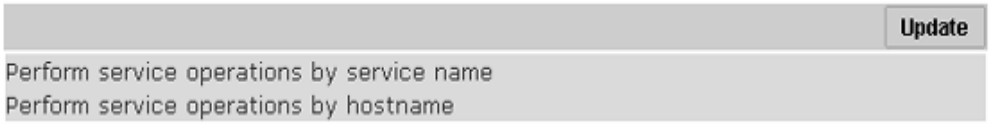

- Initiate a service by selecting the "Perform service operations by service name" option first when you want to perform a service on multiple hosts at the same time. Selecting the service name first shows you all the hosts for which that service is defined.
- Initiate a service by selecting the "Perform service operations by hostname" option first when you want to perform a service on a single host or on a specific host where you know the hostname. Selecting the hostname first shows you all the services that are defined for that host.

Performing Service Operations by Service Name

Perform the following steps to perform service operations by service name:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Select the Service Management option.
3. From the Service Management page, select the "Perform service operations by service name" link.
4. Select a service from a list of services defined for your site. A page that prompts you to select the hosts and the operation that you want to perform appears, as Figure 12-21 shows.

Figure 12-21: Perform Service Operations by Service Name Page

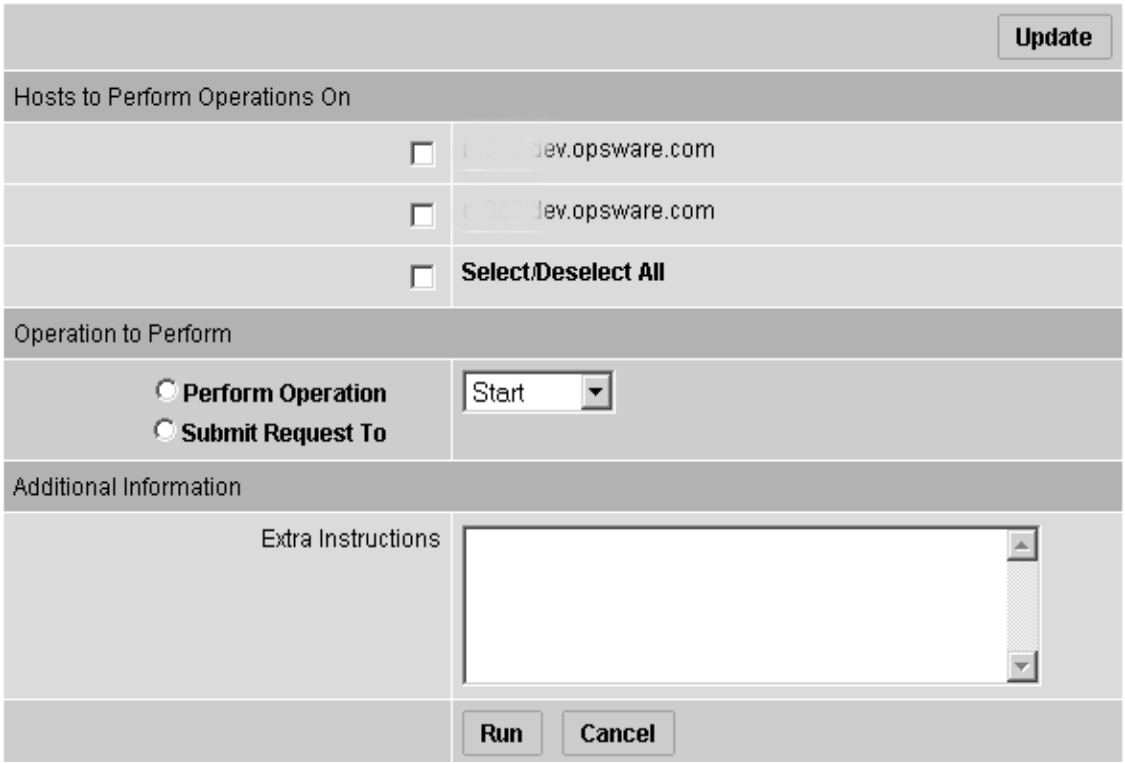

5 Select one or more of the displayed hostnames. You can choose the Select/Deselect All option to select all or clear all hostnames for the operation that you want to perform.

6 Select the Perform Operation option to directly perform a selected operation.

The Perform Operation option is visible only when you are a member of the CDR user group that allows users to directly perform a service operation.

OR

Select the Submit Request To option to send a request for authorized individuals to perform the operation for you. When you submit a request, specify any additional instructions that might be required to perform the requested operation.

7. Choose the type of operation that you want performed from the drop-down list:
   - Start
   - Stop
   - Cutover
   - Rollback
   - Backup
   - Restore

   See "CDR Service Operations" on page 582 in this chapter for information about these operations.

8. To initiate the operation or send the request, click the Run button.

Performing Service Operations by Hostname

Perform the following steps to perform service operations by hostname:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Select the Service Management option.
3. From the Service Management page, select the "Perform service operations by hostname" link.
4. Select a hostname from the list of staging and production hosts available for your site. A page that prompts you to select the service and the operation that you want to perform appears, as Figure 12-22 shows.

Figure 12-22: Perform Service Operations by Hostname Page

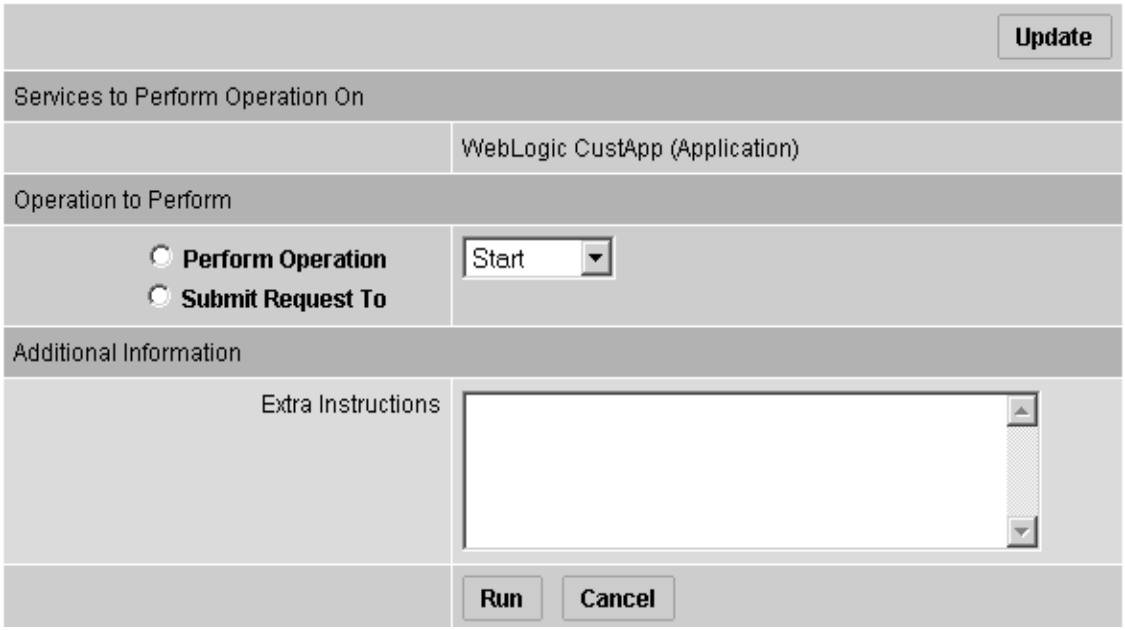

5. Select the service for which you want to perform an operation.
6. Select the Perform Operation option to directly perform a selected operation or select the Submit Request To option to send a request to have authorized individuals perform the operation for you. When you submit a request, specify any additional instructions that might be required to perform the requested operation.

The Perform Operation option is visible only when you are a member of the user group that allows users to directly perform a service operation.

7. Choose the type of operation that you want performed from the drop-down list:
   - Start
   - Stop
   - Cutover
   - Rollback
   - Backup

- Restore

See "CDR Service Operations" on page 582 in this chapter for information about these operations.

8 To initiate the operation or send the request, click the Run button.

Performing Sequences

CDR also allows you to perform service operations and synchronizations that have been set up as a sequence of operations.

Perform the following steps to perform CDR sequences set up for your site:

1 Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2 Click the Run Sequence link.

The CDR Run Sequence for [*customer name*] page appears that shows the sequences that you can run. See Figure 12-23.

Figure 12-23: CDR Run Sequence Page

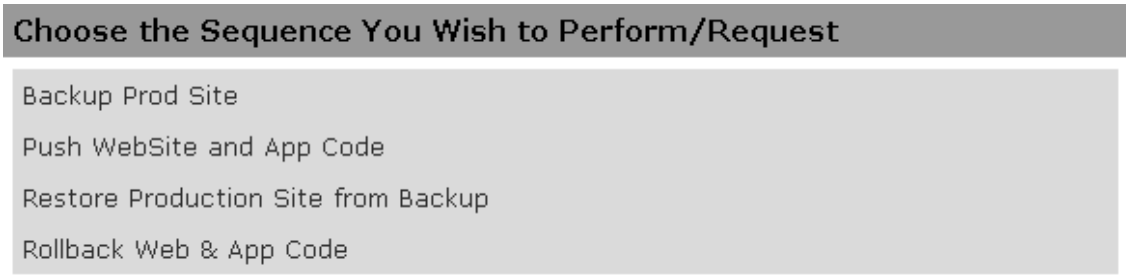

3 Select the sequence that you want to perform by clicking the link. The Run Sequence page appears, as Figure 12-24 shows.

Figure 12-24: Run Sequence Page That Shows Details of a Specific Sequence

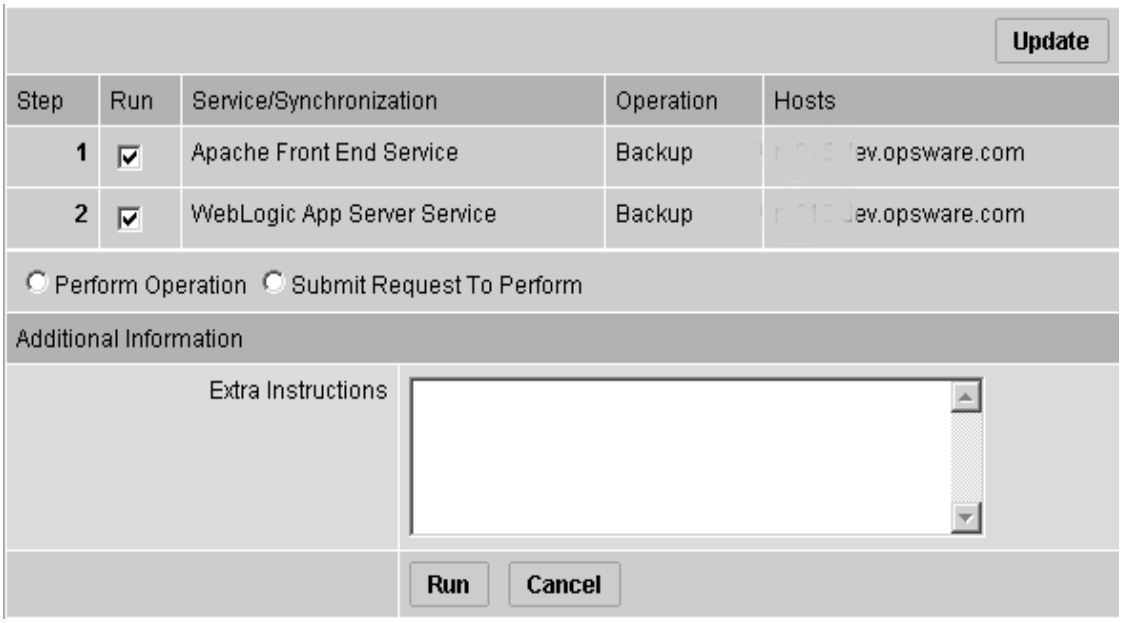

4 Select the Perform Operation radio button to directly perform the selected sequence.

Or

Select the Submit Request To Perform radio button to send a request to the users specified to receive email notification for service, synchronization, and sequence requests. When you submit a request, specify any additional instructions that you want to include for the requested sequence. For example, you might want to include instructions such as the time that you want the sequence to run, verification, or other related services to perform.

The Perform Operation option is visible only when you are a member of a user group allowed to directly perform a sequence.

5 To initiate the sequence or send the request, click the Run button.

Processing Code Deployment Requests from Users

When CDR users request that a service operation or synchronization be performed on their behalf, an email notification is sent to the individuals assigned to perform the requested service operation or synchronization.

Performing Synchronizations and Service Operations

When CDR processes a request to perform a synchronization or service operation, an email notification is sent to the addresses specified to receive those requests.

The following message is a typical example of an email notification request.

```
From:    CDR-tool@opsware.com
To:      opscenter@opsware.com
Date:    Tue, 10 Jul 2001 11:25:13 -0700
Subject: Request To Perform Start operation for opsware.com

Please perform the following request

Requestor: jhancock/opsware.com
Request Time: Jul 10, 2001 11:25:13 AM PDT
Requested Service: Demo Apache
Requested action: Start
Perform on the following hosts:
host1.opsware.com
host2.opsware.com

Extra Instructions:
```

In this case, the email message specifies a request from a user, jhancock, to perform a Start operation on two hosts running the Demo Apache service. The subject line provides a summary of the request. In addition, the email indicates the time that the request was sent.

Perform the following steps to process the request:

1. Identify any special instructions that might need to be carried out for the specific request.
2. Log in to the Opsware Command Center, choose the CDR option, and then select the particular option within Code Deployment & Rollback to perform the requested operation.

3. When you successfully complete the CDR request, you might want to notify the individual user making the request, and all other involved parties, that the requested operation was completed.

If you encounter problems in completing a request and cannot resolve them, contact your Opsware administrator for any specific remedies and follow normal escalation procedures defined for your operational environment.

See "Verifying and Troubleshooting CDR Configuration" on page 575 in this chapter for information about troubleshooting tips.

Viewing Status of Previous Operations

Code Deployment & Rollback maintains a log of operations (service operations, synchronizations and sequences) that were executed. You can view this information to determine the status of particular deployment operations, and whether they have completed successfully. You can also use the My Jobs task area of the Opsware Command Center Home page, or the My Jobs link to view this information.

Accessing the Log

Perform the following steps to access the log:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Select the View History option.

Code Deployment & Rollback displays a page providing a list, most recent to oldest, of operations that are either in progress or those that have been completed. This information displays only for the past 60 days. Each page is limited to a list of 10 operations. A Next link displays at the end of the page if there are more than 10 operations to view. Click the Next link to view subsequent operations. A Previous link is available to return to previous pages.

You need to refresh the page to view the status of any operations initiated after you selected the View History link.

A similar page displays after you select the View History link. See Figure 12-25.

Figure 12-25: CDR View History Page

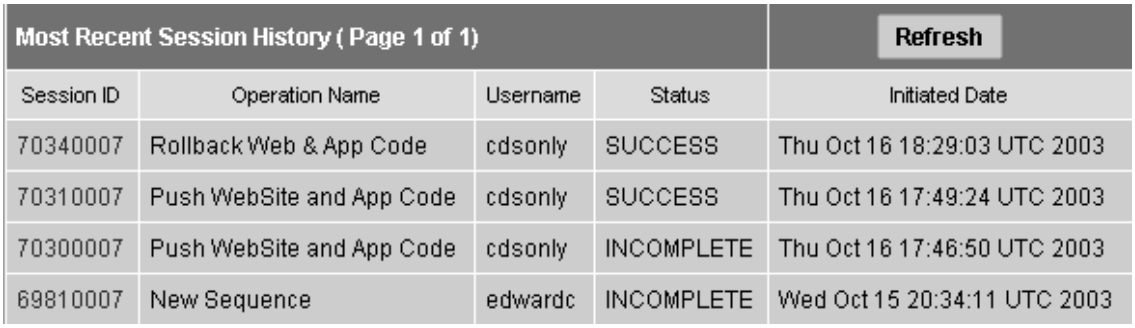

Most Recent Session History (Page 1 of 1)				Refresh
Session ID	Operation Name	Username	Status	Initiated Date
70340007	Rollback Web & App Code	cdsonly	SUCCESS	Thu Oct 16 18:29:03 UTC 2003
70310007	Push WebSite and App Code	cdsonly	SUCCESS	Thu Oct 16 17:49:24 UTC 2003
70300007	Push WebSite and App Code	cdsonly	INCOMPLETE	Thu Oct 16 17:46:50 UTC 2003
69810007	New Sequence	edwardc	INCOMPLETE	Wed Oct 15 20:34:11 UTC 2003

The individual headings of column information included in the table are:

- Session ID – A session is created each time a CDR operation is performed. Click the Session ID to view detailed results of the operation.
- Operation Name – Name of the service, synchronization or sequence as specified when they were first defined.
- Username – User ID of the user that initiated the operation.
- Status – Describes the state of the operation at the end of the sequence. The status message varies depending on the type of operation. Single step operations always result in Success/Failure messages while multiple step operations (Sequence operations) could result in Complete with Error, Incomplete, or Success messages. Table 12-7 describes the possible status messages.

Table 12-7: CDR History Status Messages

STATUS TYPE	DESCRIPTION
Abort	This message displays only if a CDR specific script that is executed in order to complete service, synchronization or sequences fails. Such issues should be escalated to your Opsware Support Representative.
Active	This message displays if an operation is still in progress.
Complete with error	The sequence completed successfully but there were errors reported while the operation was in progress and the user opted to continue rather than cancel the operation.
Failure	The operation (service, sequence or synchronization) did not complete successfully.

Table 12-7: CDR History Status Messages

STATUS TYPE	DESCRIPTION
Incomplete	The sequence resulted in an error and the user opted to cancel the sequence rather than continue.
Success	The operation (service, synchronization or sequence) was completed successfully and no errors were reported. In the case of a sequence, this message means that all the steps were completed successfully.
Initiated Date	The date on which the operation was initiated.

Appendix A: Opsware Command Line Interface

IN THIS APPENDIX

This appendix provides information about how to use the Opsware Command Line Interface:

- Opsware Command Line Interface Installation
- Software Repository OCLI
- Using the OCLI to Access the Software Repository
- Supported Operating Systems and Package Types

Opsware Command Line Interface Installation

To download the Opsware Command Line Interface packages for installation, you must have read permission to the Opsware customer. Contact your Opsware administrator to obtain the necessary access rights.

Before you use the Opsware Command Line Interface (OCLI), you must install the Opsware Agent and OCLI packages on the host from which you want to use the OCLI. See “Installing an Opsware Agent on a Server” on page 146 in Chapter 2 for more information.

Perform the following steps to install an OCLI package:

1. Download the package from the Opsware Command Center:

   Search for the package OCLI. From the Search panel, enter `ocli` in the Search box, select Packages option in the list, and click the Go button. The Packages: Search Packages page appears that displays all packages that match the search criteria. More than one page might result from the search. Use the navigation bar at the bottom of the page to move from page to page.

Each operating system and operating system version has different packages.

2. Click the package name for the OCLI that you want to download. The Packages: Edit Properties [*package name*] page appears.
3. Click the Download button to save the package locally.
4. Copy the OCLI package to each host on which you want to use the OCLI.
5. Perform the following steps, which vary by operating system:

   For Unix

   1. The file downloads as a non-executable file. Change the file mode to executable.
   2. Execute the package as root by entering the following command at the prompt:

      `<package_name> -d <installation_directory>`

Specifying the directory (-d `<installation_directory>`) in which to install the OCLI is optional. If you do not specify the installation directory, the OCLI is installed in the current directory.

   3. Include the file `login.csh` or `login.sh` in your environment, depending upon which shell you use.
      - For the shells csh, tcsh, and other variants, enter the following command at the prompt:

        ```
        source <installation_directory>/ocli/login.csh
        ```

      - For the shells sh, bash, ksh, and other variants, enter the following command at the command line:

        ```
        .<installation_directory>/ocli/login.sh
        ```

   4. Include `/opt/OPSW/bin` in your PATH.
      - For the shells csh, tcsh, and other variants, enter the following command at the prompt:

        ```
        setenv PATH /opt/OPSW/bin:${PATH}
        ```


- For the shells sh, bash, ksh, and other variants, enter the following command at the command line:

  ```
  export PATH=/opt/OPSW/bin:${PATH}
  ```

For Windows

1. Execute the package as Administrator.

   ```
   <package_name> -d <installation_directory>
   ```

2. Launch a command window and enter the following command at the prompt:

   ```
   set PATH=%PATH%;<installation.dir>\ocli\scripts
   ```

3. In the command window, enter the following command at the prompt:

   ```
   set PATH=%PATH%;%SYSTEMDRIVE%\Program
   Files\Loudcloud\lcpython15
   ```

Software Repository OCLI

This section provides information on the Software Repository OCLI and contains the following topics:

- Software Repository OCLI Overview
- File Transfer Commands
- Syntax for the Commands

Software Repository OCLI Overview

You can use the Opsware Command Center to manage packages in the Software Repository. See “Package Management” on page 267 in Chapter 5 for more information.

As a backup to access the Software Repository and for bulk uploads and downloads, you can use the OCLI.

You can only upload or download packages for the customer associated with the server from which you are running the OCLI. Contact your Opsware administrator to obtain the necessary access rights. If you must upload or download a package for a different customer, use the Opsware Command Center to change the customer association for the server.See “Editing the Properties of a Server” on page 75 in Chapter 2 for more information.

Servers *cannot* be associated with Customer Independent; therefore, if you need to upload a package associated with Customer Independent, you must upload it from a server associated with the Opsware customer. Associating a server with the Opsware customer can be a security issue; therefore, you should control the access to this server while it is associated with the Opsware customer.

The interface to each command is a CLI that begins with 'o' and has a prefix denoting the category of operation that it performs.

The commands and their associated interfaces are available on these operating systems that Opsware supports: Solaris, Linux, AIX, HP-UX, Windows NT, Windows 2000, and Windows 2003.

All commands support standard POSIX-style command line options (single dash, single letter, such as -h) and GNU-style command line options (double dash, multiple letters, such as `--help`).

File Transfer Commands

COMMAND	DESCRIPTION
`oupload`	Upload a file to the Software Repository.
`odownload`	Download a file from the Software Repository.

Syntax for the Commands

`oupload [`*`options`*`] `*`filenames`*

The filename can contain a relative or absolute local file path.

`odownload [`*`options`*`] filenames [`*`localpath`*`]`

The *localpath* can contain a relative or absolute local file or directory path.

Using the OCLI to Access the Software Repository

Perform the following steps to use the OCLI to access the software repository:

1. After fully testing a package, upload the package to the Software Repository by entering the following command at the prompt:

   ```
   oupload --pkgtype <package_type> --customer <customer> --os
   <operating_system> <source_path>
   ```

If a value for an option contains spaces, you must enclose the value in quotation marks.

For RPM packages, always remember to upload the source files after uploading a package. Uploading the source files is important from a maintenance perspective because it allows users to modify packages at a later date.

2. After you upload the files, verify that they exist on the Software Repository by using the Search Panel function in the Opsware Command Center. Select package from the drop-down list, and use * as the name of the file.

After you upload a package, define the appropriate node in the Software Tree for the new package and attach the package to the node. See “Application Provisioning Setup” on page 307 in Chapter 6 for more information.

3. (Optional) To download a package from the Software Repository by using the OCLI, enter the following command at the prompt:
   `odownload [options] <filename> <local_path>`

See "Options Common to All Commands" on page 603 in this appendix for information about a description of the options for the `odownload` command.

See "Unique Options for the oupload Command" on page 608 in this appendix for information about a description of the options for the `oupload` command.

Example: Using the OCLI

To upload `iPlanet_Web_Server-4.1sp19-LC~0.sparc64.rpm` for the customer Opsware and the operating system Solaris 5.8, enter the following command at the prompt:

```
oupload --pkgtype RPM --customer Opsware --os "SunOS 5.8"
iPlanet_Web_Server-4.1sp19-LC~0.sparc64.rpm
```

If a value contains spaces, you must enclose the value in quotation marks.

Options Common to All Commands

ARGUMENTS	VALUES	DESCRIPTION
`--customer <value>` (`-C=X`)	String (customer name, wildcards accepted) or integer (customer ID)	Specifies the customer of the file. Specifying this option is required unless you are using `--patchtype` in `opupload`. When you upload an AIX LPP file, or an HP-UX Depot that contains patches, it is associated with "Customer Independent" regardless of the customer you enter by using the `-c` option. When you upload an AIX Maintenance Level set of LPPs, you must associate them with "Customer Independent" so that all base filesets and update filesets contained in it are associated with the same customer.
`--feedback` (`-Q`)	N/A	Displays feedback while the command runs. By default, this option is enabled. Cannot specify this option with `-q`

ARGUMENTS	VALUES	DESCRIPTION
`--fr <value>` (`-f=X`)	• Alphanumeric • Period • Hyphen • Default = theword	Specifies the hostname or IP address of the Software Repository
`--frport <port>` (`-F=X`)	Integer Default = 1003	Specifies the port of the Software Repository
`--fullhelp` (`-H`)	N/A	Displays full help information Cannot specify this option with `-h` or `–V`
`--help` (`-h`)	N/A	Displays abbreviated help information Cannot specify this option with `-H` or `–V`
`--nofeedback` (`-q`)	N/A	Does not display feedback while the command runs Cannot specify this option with `–Q`

ARGUMENTS	VALUES	DESCRIPTION
`--os <type>` `(-o=X)`	String (OS name, wildcards accepted) The following are the allowable values: AIX 4.3 AIX 5.1 AIX 5.2 AIX 5.3 HP-UX 10.20 HP-UX 11.00 HP-UX 11.11 OS Independent Red Hat Enterprise Linux AS 2.1 Red Hat Enterprise Linux AS 3.0 Red Hat Enterprise Linux ES 2.1 Red Hat Enterprise Linux ES 3 Red Hat Enterprise Linux WS 3 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 SUSE Linux Enterprise 8.0 SUSE Linux Standard 8.0 SUSE Linux Enterprise 9.0 SunOS 5.6 SunOS 5.7 SunOS 5.8 SunOS 5.9 Windows 2000 Windows 2003 Windows NT 4.0	Specifies the operating system of the package Specifying this option is required. If a value has a space in the name, enclose the entire name in quotes. For Fujitsu Solaris 2.8, use the value for SunOS 5.8. For Fujitsu Solaris 2.9, use the value for SunOS 5.9. See “Allowable integer values for -os option” in this appendix for more information.

ARGUMENTS	VALUES	DESCRIPTION
`--timeout <value>` (`-z=X`)	Integer Default = 60	Sets the timeout to the server in seconds
`--truthgw <value>` (`-g=X`)	• Alphanumeric • Period • Hyphen • Default = spin	Specifies the hostname or IP address of the Data Access Engine
`--truthgwport <port>` (`-G=X`)	Integer Default = 1004	Specifies the port of the Data Access Engine
`--verbose` (`-v`)	N/A	Displays debug information
`--version` (`-V`)	N/A	Displays version information for the OCLI Cannot specify this option with `-h` or `–H`.

Allowable integer values for --os option

STRING NAME (OS NAME)	INTEGER VALUE (ID)
AIX 4.3	870007
AIX 5.1	10001
AIX 5.2	260007
AIX 5.3	40007
HP-UX 10.20	230007
HP-UX 11.00	1070007
HP-UX 11.11	1080007
OS Independent	1
Red Hat Enterprise Linux AS 2.1	960007
Red Hat Enterprise Linux AS 3.0	430007
Red Hat Enterprise Linux ES 2.1	10730013
Red Hat Enterprise Linux ES 3.0	10720013
Red Hat Enterprise Linux WS 3.0	270022
Red Hat Enterprise Linux 6.2	140000
Red Hat Enterprise Linux 7.1	210022
Red Hat Enterprise Linux 7.2	950007
Red Hat Enterprise Linux 7.3	410007
Red Hat Enterprise Linux 8.0	420007
SUSE Linux Standard Server 8.0	20007
SUSE Linux Enterprise Server 8.0	10030
SUSE Linux Enterprise Server 9.0	20032
SunOS 5.6	130000
SunOS 5.7	90000
SunOS 5.8	150001
SunOS 5.9	920007

STRING NAME (OS NAME)	INTEGER VALUE (ID)
Windows 2000	120000
Windows 2003	10007
Windows NT 4.0	8000

Unique Options for the ouplоad Command

ARGUMENTS	VALUES	DESCRIPTION
`--patchtype <type>` (`-a=X`	• AIX LPP • HP-UX Depot • Windows Hotfix • Windows Service Pack • Solaris Patch • Solaris Patch Cluster	Cannot specify this option with `-c`.

ARGUMENTS	VALUES	DESCRIPTION
`--pkgtype <type>` (`-t=X`)	• AIX LPP • HP-UX Depot • RPM • Windows Hotfix • Windows MSI • Windows Service Pack • Solaris Package • Solaris Patch • Solaris Patch Cluster • Microsoft Patch Database • OS Provisioning Install Hooks • Windows Zip File	Specifies the type of file. Specifying either this option or the `-patchtype` option is required. Wildcards are accepted. The OCLI does *not* support uploading response files for the Solaris Package package type. Use Opsware Command Center to associate a response file with a Solaris Package. See "Editing Package Properties" on page 300 in Chapter 5 for more information. If a value contains spaces, you must enclose the value in quotation marks.

The following table includes the values that must be supplied interactively based on the package type.

PACKAGE TYPE	OPTIONS	DATA TYPE
Windows Hotfix	N/A	N/A
Windows MSI	Product version Product name	Free form text
Windows OS Service Pack	Service pack level	Free form text
AIX LPP	N/A	N/A
HP-UX Depot	N/A	N/A
RPM	N/A	N/A
Solaris Package	N/A	N/A
Solaris Patch	N/A	N/A
Solaris Patch Cluster	N/A	N/A
Unknown	N/A	N/A
Microsoft Patch Database	N/A	N/A (Can only be uploaded for Windows 2000)
OS Provisioning Install Hooks	N/A	N/A
Windows Zip File	N/A	N/A
Windows Utility	N/A	N/A (Can only be uploaded for Windows 2000)

Supported Operating Systems and Package Types

Each operating system that the Opsware System supports has a list of package types. The Opsware System supports these package types on the supported operating systems, as the following table shows.

OPERATING SYSTEM	PACKAGE TYPE	ADDITIONAL METADATA*
AIX	AIX LPP (contains an update fileset or base filesets)	N/A
	RPM	N/A
HP-UX	HP-UX Depot (contains products and filesets)	N/A
Linux	RPM	N/A
Solaris	Solaris Patch	N/A
	Solaris Patch Cluster (contains patches)	N/A
	Solaris Package	None
	RPM	N/A
Windows	Windows Hotfix	Service Pack Level
	Windows MSI	Product version and name
Windows Utility	Windows OS Service Pack	Service Pack Level
OS Independent	Unknown	N/A

* For certain package types, the Opsware Command Center requires that you provide additional metadata for the package.

Specifying an Encoding Scheme in the Opsware Command Line Interface (OCLI)

In Opsware System 4.0.1, Japanese Edition, the Opsware Command Line Interface (OCLI) was updated with new arguments to allow you to specify an encoding scheme when uploading and downloading packages and for the customer display name.

New Arguments for the oupload Command

You only need to enter these arguments when you want to override the default settings in the LANG environment variable in your shell.

The following table shows the new arguments for the `oupload` command.

ARGUMENT	VALUE	DESCRIPTION
--filename-encoding (-e)	String	Specifies the character set of the file name When specifying non-ASCII characters in the value for the --customer argument, include the -e argument on the command line to tell the Opsware System which character set to use when communicating with the Opsware System Model Repository database.
--metainfo-encoding (-E)	String	Specifies the character set of the meta-information in the package

New Arguments for the odownload Command

You only need to enter this argument when you want to override the default settings in the LANG environment variable in your shell.

The following table shows the new arguments for the `odownload` command.

ARGUMENT	VALUE	DESCRIPTION
--filename-encoding (-e)	String	Specifies the character set encoding in which to save the file name

Appendix B: Agent Upgrade Tool

IN THIS APPENDIX

This appendix provides the following information about how to use the Agent Upgrade Tool:

- Opsware Agent Upgrade Tool Overview
- Prerequisites for Using the Opsware Agent Upgrade Tool
- Upgrading the Opsware Agent on Managed Servers
- Commands for the Opsware Agent Upgrade Tool
- Options for the Opsware Agent Upgrade Tool
- Examples of Options for the Opsware Agent Upgrade Tool
- Examples of Options for the Opsware Agent Upgrade Tool

Opsware Agent Upgrade Tool Overview

After you upgrade the Opsware System running in a facility, you should upgrade the Opsware Agents on every managed server to the new version, so that you can utilize the new features in the newly-upgraded core.

The Opsware System features continue to work on a managed server even when it is running an older Opsware Agent. However, new features in the new versions might not be available for that server.

Refer to the Release Notes for the new version for information about the compatibility of new features with older agents.

You can upgrade the Opsware Agents on managed servers in the following ways:

- Use the Opsware Agent Installer (a command line interface) to install a new Opsware Agent on one server at a time.

  See “Overview of Server Assimilation” in Chapter 2 for information about how to use the Opsware Agent Installer.

- Use the Opsware Agent Upgrade Tool to upgrade Opsware Agents on groups of servers. Running the tool upgrades deployed Opsware Agents on managed servers. You can run the script simultaneously on many servers to upgrade large groups of Opsware Agents.

The Opsware Agent Upgrade Tool has the following characteristics:

- It is a command line interface that provides a flexible mechanism for selecting servers to upgrade, and for monitoring and reviewing upgrade operations.
- You can use it to upgrade many Opsware Agents on managed servers simultaneously.
- It runs within your preferred Unix shell, allowing it to leverage the power of standard Unix shells and text processing tools.
- You can use it to upgrade a server in any facility running an Opsware System. You can run it from an Opsware shell attached to any Opsware System in any facility.

The Opsware shell is a program that authenticates users in the Opsware System, starts the user's normal Unix shell (as specified in the standard password database). Using the Opsware shell allows the user to run the Opsware Agent Upgrade Tool in this facility.

Prerequisites for Using the Opsware Agent Upgrade Tool

- For Solaris and Linux servers that can connect to the Opsware System in a facility, and which have the Opsware Agent already installed and running, install the Opsware Shell RPM by downloading the `opsh` package from the Opsware Command Center.

  Installing the opsh RPM installs the Opsware shell and the Opsware Agent Upgrade Tool in the directory /opt/OPSWopsh/bin.

  See "Downloading a Package" in Chapter 5 for information about how to download a package in the Opsware Command Center.

- You need the correct permissions to upgrade Opsware Agents. Run the Opsware shell by specifying the Opsware admin (username admin) and password to ensure that you have the appropriate permission. (Contact your Opsware administrator to obtain the password.)

  When you start an Opsware shell to run the Opsware Agent Upgrade Tool, the user name and password are authenticated by the Opsware System.

Upgrading the Opsware Agent on Managed Servers

Step 1 is required before you can do anything else.

1. After you install the opsh RPM on a Unix server that can connect to the Opsware System, enter the following command as root to start the Opsware shell:

   ```
   opsh [username@]facility-domain
   ```

   See "Commands for the Opsware Agent Upgrade Tool" for a description of this command.

2. (Optional) To obtain information about the current Opsware Agents running on the managed servers before you upgrade them, enter any of the following commands and options:

   ```
   opsh_agent query server-options
   ```

   (Enter the previous command if you want to view a report of the Opsware Agent versions running on the servers before you upgrade them.)

   ```
   opsh_agent verify server-options schedule-options \ agent-version
   ```

   (Enter the previous command if you want to verify the versions of the Opsware Agents running on the managed servers before you upgrade them.)

3. To upgrade agents on specified servers, enter the following Opsware Agent Upgrade Tool commands and options:

   ```
   opsh_agent stage server-options schedule-options \
   [--always] agent-version
   ```

   (Enter this command if you want to download the package for the Opsware Agent to the managed server before you run the upgrade.)

   ```
   opsh_agent upgrade server-options schedule-options \
   [--always] agent-version
   ```

4. (Optional) To review the status of the Opsware Agent upgrade, enter the following command and option:

   ```
   opsh_agent review session-id
   ```

Commands for the Opsware Agent Upgrade Tool

- `opsh [username@]facility-domain`

  This command starts an Opsware shell and authenticates the user name against the Opsware facility running at the specified domain.

If you do not specify a user name, the currently logged in user name is used. The Opsware shell prompts for a password.

A new Unix shell (which is attached to the specified Opsware core-domain) is started. (The password database for the user specifies which Unix shell to use.)

- `opsh_agent query` *`server-options`*

  This command must be run from an Opsware shell started with the `opsh` command.

  This command queries the reported version of Opsware Agents and any staging status for the specified servers by examining data in the Model Repository.

  One line is printed to stdout for each server that shows device ID, IP address, current Opsware Agent version, and any staging status.

  You can specify the servers by using the --device, --customer, --facility, and --os options.

- `opsh_agent stage server-options schedule-options \`

  `[--always] agent-version`

  You must run this command from an Opsware shell started with the `opsh` command.

  This command contacts the Opsware Agent on each specified server and instructs it to download the package for the specified version of the Opsware Agent from the Software Repository.

  If the download is successful, the staging status is written to the Model Repository for the server.

  To download the package to the server even when this command was entered previously (recorded in the Model Repository), specify the --always option.

  One line is printed to stdout for each server that shows the device ID, IP address, and a success or failure indicator.

  You can specify the servers by using the --server, --customer, --facility and --os options.

  A session is started and the session ID displays for later review. After the session ID displays, you can type CTRL-C and review the session later using the opsh_agent review command.

- `opsh_agent upgrade server-options schedule-options \`

  `[--always] agent-version`

You must run this command from an Opsware shell started with the `opsh` command.

This command contacts the Opsware Agent on each specified server and instructs it to upgrade to the specified version. If the necessary package has not been downloaded on the server already (the `opsh_agent` stage command was entered), the package is downloaded from the Software Repository.

If the upgrade is successful, the package is removed from the server and the staging status is deleted from the Model Repository.

To upgrade the Opsware Agent even when the specified version of the Opsware Agent was already installed on the managed servers, enter the `--always` option. (The Model Repository records when Opsware Agents are upgraded on servers.)

One line is printed to stdout for each server that shows the device ID, IP address, and a success or failure indicator.

You can specify the servers by using the `--server, --customer, --facility,` and `--os` options.

A session is started and the session ID displays for later review. After the session ID displays, you can type CTRL-C and review the session later by using the `opsh_agent` review command.

- `opsh_agent verify server-options schedule-options \`

  `agent-version`

  You must run this command from an Opsware shell started with the `opsh` command.

  This command contacts the Opsware Agent on each specified server to verify that it is running the specified version.

  One line is printed to stdout for each server that shows the device ID, IP address, the word OLD, NEW, or CURRENT and the actual Opsware Agent version running on the server.

  You can specify the servers by using the `--server, --customer`, `--facility,` and `--os` options.

  A session is started and the session ID displays for later review. After the session ID displays, you can enter CTRL-C and review the session later by using the `opsh_agent` review command.

- `opsh_agent review session-id`

You must run this command from an Opsware shell started with the `opsh` command; although, not necessarily the same Opsware shell from which the original command was started.

This command attaches to a running opsh_agent stage, opsh_agent upgrade or opsh_agent verify session running on the Command Engine. It prints the same output to stdout that the original command would have printed if the user had not typed CTRL-C and terminated the command. If the session is complete, it shows the same results that were shown when the session completed.

Options for the Opsware Agent Upgrade Tool

Server-options: `--server/-S svr-spec --customer/-C cust-spec --facility/-F fac-spec --os/-O os-spec`

If more than one of the `--customer`, `--facility`, or `--os` options is specified, only servers that match all options are selected. Any servers specified by using the `--server` option are added to (or subtracted from) the list specified by combining the `--customer`, `--facility,` and `--os` options.

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--server`	`-S`	`svr-spec`	Server by device ID, IP address, or system name
`--customer`	`-C`	`cust-spec`	All servers associated with the customer specified by the customer ID or name
`--facility`	`-F`	`fac-spec`	All servers in the facility specified by the facility ID or name
`--os`	`-O`	`os-spec`	All servers running the operating system specified by the OS name

Schedule-options: `--when/-W when-time --until/-U until-time`

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--when`	`-W`	`when-time`	Start time for a stage, upgrade, verify, or test operation in the format: MM/DD/YYY-HH:MM If the `--when` option is used, the operation starts at the specified time, but the command displays a session ID and returns immediately. When you schedule an operation, you use the `review` command to review the results after the operation has run. If the `--when` option is not specified, the operation starts immediately and the command displays the output of the command.
`--until`	`-U`	`until-time`	End time for a stage, upgrade, verify or test operation in the format: MM/DD/YYY-HH:MM If the `--until` option is specified, the operation stops processing servers at the specified time. Any servers that are not complete are left in a consistent state; this might require that the session run past the specified time.

Miscellaneous options: `--ip/-I --always/-A`

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--ip`	`-I`	(N/A)	Display IP addresses instead of hostnames
`--always`	`-A`	(N/A)	Always stage or upgrade servers even if the current version is staged or upgraded
`--parallel`	`-P`	`Concurrency`	The maximum of concurrent commands (Recommended default = 10)
`--theword`	`-T`	`hostname`	The hostname or IP address to use when contacting the Software Repository from a server

Examples of Options for the Opsware Agent Upgrade Tool

The following table provides examples for running the Opsware Agent Upgrade Tool.

EXAMPLE	DESCRIPTION
`--server "1 2"`	Selects servers 1 and 2
`--facility "Y Z"`	Selects all servers (for all customers) in facilities Y and Z
`--customer "-A -B" --facility Z`	Selects all servers in facility Z except those owned by customers A and B
`--server "1 2 -3 -4" --customer "A B" --facility "Y Z"`	Select servers 1 and 2 as well as all servers owned by customers A or B which are in facilities Y or Z except servers 3 and 4
`--server "1 -2" --customer "A B" --facility "-Y -Z" --os "SunOS 5.8"`	Select server 1 and all servers owned by customers A or B, except those in facilities Y or Z and which are SunOS 5.8 machines excluding server 2

Example Commands and Output for Agent Upgrade Tool

```
# cd /opt/OPSWopsh/bin
# ./opsh admin@core2.cust.com
admin@core2.cust.com's password:
#
# ./opsh_agent verify --os "SunOS*" 14a.2.12.18
Session 37802500101L
Device ID Name/IP address Version Result Status Reason
410101L core2-1.core2.cust.com  14a.2.12.18 CURRENT SUCCESS
^C
Interrupted review of running session 37802500101L
Use review 37802500101L command anytime to review session status
#
# ./opsh_agent review 37802500101L
Session 37802500101L
Device ID Name/IP address Version Result Status Reason
410101L core2-1.core2.cust.com 14a.2.12.18 CURRENT SUCCESS
670101L dhcp-174.core2.cust.com 14a.2.12.16 OLDER SUCCESS
1460100L emb218-37.core0.cust.com 14a.2.12.18 CURRENT SUCCESS
20100L core0-1.core0.cust.com 14a.2.12.18 CURRENT SUCCESS
10100L core0-2.core0.cust.com 14a.2.12.21 NEWER SUCCESS
210100L m022.core0.cust.com 14a.2.12.18 CURRENT SUCCESS
Session 37802500101L completed.
```


Appendix C: OS Installation Integration

IN THIS APPENDIX

This appendix discusses how to integrate the Opsware System with operating system installation technologies. Topics in this appendix include an overview of OS installation technologies and how to integrate the Opsware System with the following technologies:

- OS Installation Technologies
- OS Installation Integration
- Integration High-Level Steps
- Integration with Red Hat Kickstart
- Integration with Solaris Jumpstart
- Integration with Windows OS Installation Technologies
- Example: Integration with Windows NT and Symantec Ghost
- Integration with Network Installation Management and AIX
- Integration with Ignite-UX and HP-UX

The Opsware System includes the Opsware OS Provisioning Subsystem, which allows you to install the following versions of Sun Solaris, Linux, and Microsoft Windows operating systems:

- Windows NT 4.0, 2000, and 2003
- Red Hat Linux 7.1, 7.2, and Advanced Server 2.1
- SUSE Linux Standard Server 8.0 and SUSE Linux Enterprise Server 8.0
- Sun Solaris 2.6, 7, 8, and 9

See Chapter 3, “Operating System Provisioning” on page 237 for the instructions to set up the Opsware System to provision servers with Solaris, Linux, and Windows.

Alternatively, you can integrate the Opsware System with an OS installation technology that is already functioning in your operational environment. Using a third-party installation technology enables installing an OS by using vendor utilities and automatically installing the Opsware Agent, which registers servers with the Model Repository.

The Opsware OS Provisioning Subsystem does not provision HP-UX or AIX operating systems. Follow the procedures in this appendix to integrate the Opsware System with Network Installation Management (NIM) to provision AIX and Ignite-UX to provision HP-UX.

OS Installation Technologies

Operating system (OS) vendors provide automation technology for installing their operating systems. OS installation technologies follow this general process:

1. Boot the server from boot media.
2. Partition and format the target disks.
3. Install the OS onto the target disks.
4. Reboot the server from the newly installed OS.

For example, Solaris JumpStart follows this process:

1. Boot the server from the network.
2. Partition and format the server's disks.
3. Install the OS, optionally installing additional patches and packages.
4. Reboot the server from the newly installed OS.

OS installation technologies provide a way to invoke customer-supplied code at the end of the OS installation process. After the operating systems are installed, users can run scripts or programs to customize servers. The Opsware System uses this integration point to automatically install the Opsware Agent and register the servers with the Model Repository.

OS Installation Integration

This section provides information on OS installation integration and contains the following topics:

- OS Installation Integration Overview
- Modeling Operating Systems
- How the Opsware System Assimilates Servers

OS Installation Integration Overview

The Opsware System integrates with OS installation technologies, including third-party and vendor-provided OS bootstrapping technologies. These OS installation technologies give you the ability to set up unattended OS installations. You boot a system and the software installs automatically.

Integrating with these technologies provides a uniform method for OS installation because the Opsware System conforms the installed OS to the data model in the Model Repository of what OS software should be installed.

Servers conform to the model because the Opsware System installs patches, Service Packs, or Hotfixes, installs other software (such as SSH), removes software, turns processes on and off (such as turning off the FTP server), and updates configurations.

Modeling Operating Systems

To manage servers' operating systems by using the Opsware System, you model the operating systems in the Opsware Command Center. An OS model is represented as a template that specifies a set of nodes assigned to servers when the template is applied to the servers. The Node assignments enable the Opsware System to conform the OS on the servers to the model of what should be installed on servers.

When you model an OS:

- Patches, Service Packs, and software that are specified in the model are installed after the server's OS is installed. Modeling an OS ensures that all newly provisioned servers have required patches, security fixes, and utility software installed automatically.

- Software that is not part of the model can optionally be removed.

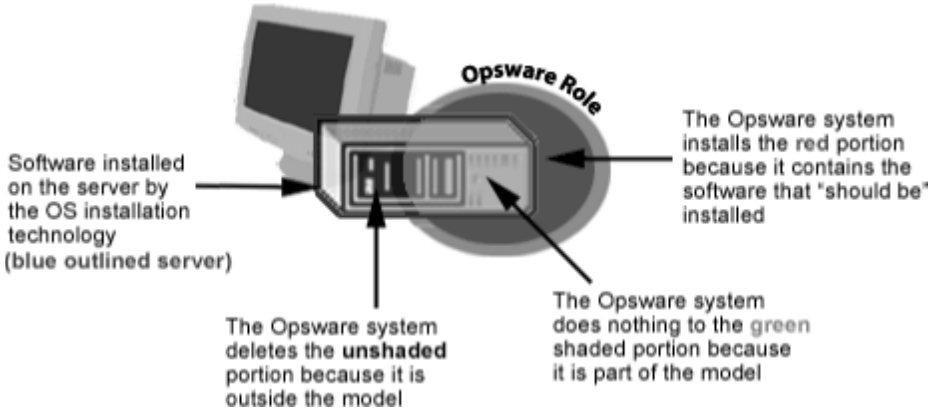

When you do not model an OS in the Opsware System:

- The Opsware Agent can be automatically installed, and the server can be registered with the Model Repository.
- The Opsware System removes only software that was previously installed by the Opsware System and does not disrupt an OS that an OS installation technology installed.
- However, the Opsware System cannot manage the OS software until a model of the OS is created and the server is associated with that model.

How the Opsware System Assimilates Servers

The Opsware System assimilates servers by automatically installing the Opsware Agent and registering the servers with the Model Repository.

After an OS is installed with an installation technology, a user can run a post-installation script to customize the server. In the script, the user can include logic to invoke the Opsware Agent Installer, so that the following actions are performed:

1. The Opsware Agent Installer installs an Opsware Agent on the server.
2. The Opsware Agent gathers information about the server (such as hardware attributes) and registers the server with the Model Repository.
3. The Opsware System attaches the server to the appropriate hardware, facility, and customer nodes.

   The server is associated with the default facility for the local instance of the Opsware System.

If IP address ranges are specified for customers in Opsware Command Center, the server is associated with the customer through its IP address.

If the server's IP address does not fall within a specified IP address range, the server is associated with the default IP range group (the "Default" IP range group). The default group is associated with the "Not Assigned" customer.

4 The Opsware System assigns the server to an OS Template, which designates sets of node assignments. This information enables the Opsware System to conform the OS to a model of what should be.

For example, all Windows 2000 servers can be attached to an OS Template that specifies that a given service pack should be applied to the servers.

If an OS Template is not specified, the Opsware System attaches the server to a generic Template that is empty (it does not contain software).

5 (Optional) The Opsware System reconciles the server, which conforms the server with the model by installing patches (Service Packs and Hotfixes for Windows), installing software (such as, the latest version of SSH), removing software, turning processes on and off (such as turning off the FTP server), and updating configurations.

Integration High-Level Steps

This section provides information on how to integrate the Opsware System into a network environment and contains the following topics:

- Integrating the Opsware System
- Opsware Agent Installer Commands and Options
- Opsware Agent Installer Options
- Examples: Opsware Agent Installer Command and Options

Integrating the Opsware System

Regardless of the OS installation technology used in the managed environment, you integrate the Opsware System in this general way:

1 Make the Opsware Agent Installer binary available to the server, through NFS mount, by copying the installer to a local disk, by using wget, or through another method.

The Opsware Agent Installer binary must be available to the post-installation scripts when they run.

2. Include logic in a post-installation script or file (Windows) for the OS installation technology to assimilate the server by running the Opsware Agent Installer with the appropriate options.
   - Unix: Write out a script that runs on reboot, assimilates the server, and then removes itself.
   - Windows 2000: Run commands in a post-installation script by using the [`GuiRunOnce`] section in unattended installation files or in the file sysprep.inf.
   - Windows NT: Insert an entry in the Registry to call a batch file that runs once at startup and specifies the Opsware Agent Installer options in a post-installation script.

The post-installation script or commands must arrange for the server to be assimilated the first time it reboots after OS installation. The Opsware System cannot assimilate a server until it reboots after OS installation. Before the reboot, the server is not running the final OS. Therefore, it is not appropriate to register the server with the Model Repository.

3. If you plan to manage the server's OS by using the Opsware System (the `--template` and `--reconcile` options are specified with the Opsware Agent Installer), model the OS in Opsware Command Center. To model the OS in the Opsware Command Center, create Nodes and Templates for the OS and upload the OS packages to the Software Repository.

   See Chapter 6, "Application Provisioning Setup" on page 307 for information about how to model an operating system in the Opsware Command Center. See Chapter 5, "Package Management" on page 267 for information about how to upload packages to the Package Repository.

Opsware Agent Installer Commands and Options

In the post-installation script or file, specify the correct Opsware Agent Installer for the operating system and the appropriate options for the installation environment.

Unix Executable

```
opsware-agent-<version>-<system_name>-<system_version>
```


Windows Executable

`opsware-agent-<version>-<system_name>-<system_version>.exe`

Opsware Agent Installer Options

When you use the Opsware Agent Installer CLI, you can include the following options to control the way that the Opsware Agent is installed on a server.

OPTION	DESCRIPTION
`--clean` (`-c`)	Removes any machine-specific identifying material from the server. Specifically, removes the machine ID file (MID), and all machine-specific cryptographic material. Use this option when a server is deactivated and deleted from the Opsware Command Center and needs to be returned to service at a later time.
`-f`	Forces Opsware Agent installation and removes the target installation directory if it exists. REQUIREMENT: When using the `-f` option, you must run the Opsware Agent Installer as root on Unix operating systems and as the administrator on Windows operating systems.
`--logfile`	Specifies the path to the Opsware Agent Installer log file. By default, the current directory is set as the path. By default, the log file has the following filename: `opsware-agent-installer-<date>.log`
`--loglevel <level>`	Sets the log level for log messages. With this option, specify one of the following levels: `error`, `warn`, `info`, `trace`, or `none`. The level `error` logs the least detail. The level `trace` logs all messages. By default, the log level is set to the log level `info`.

OPTION	DESCRIPTION
`-o`	Logs all output to `stdout` instead of a log file. This option is invoked automatically if the default log file or the log file passed with the `--logfile` option cannot be created, for example, when running the Opsware Agent Installer from non-writeable media, such as a CD-ROM.
`--reconcile <type>`	Reconciles the server against any nodes assigned to the server. The <*type*> can be `full` or `addonly`. `full` – All nodes in a category are selected and reconcile removes software that the Opsware System did not install. `addonly` – Software installed outside of the Opsware System is not removed. WARNING: When assimilating a server that is already functioning in the operational environment, use caution when specifying the option `--reconcile`. If you specify this option, you might inadvertently uninstall software from the server.
`--rpmbin <path>`	Specifies the path to the RPM binary to use for RPM operations. Use this option, when RPM is already installed on the server, to point the Opsware Agent at the RPM binary. Use the `--withrpm` option to install RPM if a usable instance of RPM is *not* already installed. NOTE: It is unnecessary to use this option with the `--withrpm` option.
`-s`	Starts the Opsware Agent after installing it. By default, the Opsware Agent Installer does not start the Opsware Agent.

OPTION	DESCRIPTION
`--template <ID>`	Assigns the nodes contained in the template to the server. `<ID>` can be an ID or a full name of a template. If this option is specified with the `--reconcile` option, the Opsware System assigns the nodes in the template to the server before reconciling the server. WARNING: When assimilating a server that is already functioning in the operational environment, use caution when you specify the option `--template`. If you specify this option, you might inadvertently uninstall software from the server.
`--withmsi`	Installs MSI 2.0 along with the Opsware Agent. If MSI 2.0 is already installed, this option has no effect. Works with Windows NT 4.0 Service Pack 6a, Windows 2000, and Windows 2003.
`--withrpm`	Installs the RPM handler with the Opsware Agent. By default, an Opsware Agent is not installed with this option. Opsware Inc. recommends that you always include the `--withrpm` option when you install Opsware Agents on Solaris servers. NOTE: Use the `--withrpm` option only with the Opsware Agent Installers for these operating systems: Solaris 5.6, 5.7, 5.8, and 5.9, and AIX 4.3 and AIX 5.1. On Solaris, RPM 3.0.6 is installed in the directory `/opt/OPSWrpm` and the RPM database is installed in the directory `/var/opt/OPSWrpm/lib/rpm`. On AIX, RPM 3.0.5 is installed in the directory `/opt/freeware` and the RPM database is installed in the directory `/var/opt/freeware/lib/rpm`.
`--workdir <path>`	Specifies the path to the Opsware Agent Installer temporary working directory. Use this option if the default working directory causes problems with installation.

Examples: Opsware Agent Installer Command and Options

Enter the following command and options to install the Opsware Agent for Solaris 5.7 in the default directories and log results of the installation in the log file:

```
% opsware-agent-1.0.0-solaris-5.7 --logfile opsware-agent-
installer-[current_date].log --loglevel info
```

Enter the following command and options to install the Opsware Agent for Windows NT 4.0 in the default directories and log results of the installation in the log file:

```
% opsware-agent-1.0.0-win32-4.0.exe --logfile opsware-agent-
installer-[current_date].log --loglevel info
```

Integration with Red Hat Kickstart

Red Hat Kickstart uses a configuration file that contains a number of distinct sections. The section `%post` contains a set of commands to run after the OS installation is complete.

Perform the following steps:

1. Copy the Opsware Agent Installer to local disk storage, for example `/var/tmp,` on the server that is being assimilated.

Copy the Opsware Agent Installer to a directory that will not be empty when the server reboots. The operating system might empty the contents of the `/tmp` and `/var/tmp` directories when the server reboots.

2. Create an init script that includes the following logic to invoke the Opsware Agent Installer on reboot, and then remove itself:

```
cp /<volume>/<agent_installer> <installer_local_path>
cat > /etc/rc.d/rc3.d/S99zAgentInstaller << EOF
#!/bin/sh
<installer_local_path> <agent_installer_options>
if [ $? -eq 0 ]; then
   rm -f /etc/rc.d/rc3.d/S99zAgentInstaller
   rm -f <installer_local_path>
```



```
fi
EOF
```

The script `S99zAgentInstaller` removes itself and the Opsware Agent Installer binary when the Opsware Agent Installer returns zero. If the installation returns an error, the script `S99zAgentInstaller` will not remove itself from the server and will attempt installation on the next reboot. Alternatively, you can rerun the init script manually (you must be root).

Example File: init Script for Kickstart

The following example assumes that the Opsware Agent Installer binary is available by using an NFS volume mounted at /sw.

```
cp /sw/opsware-agent-5.1.14-linux-7.2 /var/tmp/opsware-agent-
installer
cat > /etc/rc.d/rc3.d/S99zAgentInstaller << EOF
#!/bin/sh
/var/tmp/opsware-agent-installer --template 12340002 --settime
--reconcile full
if [ $? -eq 0 ]; then
   rm -f /etc/rc.d/rc3.d/S99zAgentInstaller
   rm -f /var/tmp/opsware-agent-installer
fi
EOF
```

Integration with Solaris Jumpstart

Solaris Jumpstart uses a profile that is capable of running a post-installation script (also referred to as a finish script). The post-installation contains a set of commands to run after the OS installation is complete.

When you integrate the Opsware System with Solaris Jumpstart, use a post-installation script to perform the following actions:

1. Copy the Opsware Agent Installer to local disk storage, for example `/var/tmp,` on the server being assimilated.

Copy the Opsware Agent Installer to a directory that will not be empty when the server reboots. The operating system might empty the contents of the `/tmp` and `/var/tmp` directories when the server reboots.

2 Create an init script that invokes the Opsware Agent installer on reboot, and then remove itself.

Example File: Jumpstart Finish Script

The following example assumes that the Opsware Agent Installer binary is installed in the Jumpstart configuration directory.

For use in a post-installation script, Jumpstart automatically sets the variable `$SI_CONFIG_DIR` to refer to the Jumpstart configuration directory. See the *Solaris Advanced Installation Guide* for information.

```
#!/bin/sh
#
# finish script which adds agent installer during Jumpstart
#
AGENT=opsware-agent-5.1.35-solaris-5.8
AGENT_START_SCRIPT=/etc/rc3.d/S99zAgentInstaller
TEMPLATE_ID=12345

# copy agent installer to client's /var/tmp
# client's filesystem is mounted as /a during jumpstart
cp $SI_CONFIG_DIR/$AGENT /a/var/tmp/opsware-agent-installer
chmod 0755 /a/var/tmp/opsware-agent-installer

# setup a script to run the installer on reboot
touch /a/$AGENT_START_SCRIPT
chmod 711 /a/$AGENT_START_SCRIPT
chown root:sys /a/$AGENT_START_SCRIPT

cat >> /a/$AGENT_START_SCRIPT << EOF
#!/bin/sh
exec > /var/tmp/`basename $AGENT_START_SCRIPT`.log 2>&1
set –x
/var/tmp/opsware-agent-installer \
        --template $TEMPLATE_ID \
```



```
        --settime \
        --reconcile addonly \
        --decommission \
        --logfile /var/tmp/opsware-agent-installer.log

if [ \$? -eq 0 ]; then
  cp $AGENT_START_SCRIPT /var/tmp/
  rm -f $AGENT_START_SCRIPT
fi
EOF
```

Integration with Windows OS Installation Technologies

This section provides information on integrating with Windows OS installation technologies within the Opsware System and contains the following topics:

- Windows OS Installation Integration Process
- Example: Integration with Windows 2000 and Symantec Ghost
- Running the Opsware Agent Installer by Using Sysprep
- Example File: Preparing a Windows 2000 System for Imaging
- Example Batch File: Running the Agent Installer for Windows 2000

Windows OS Installation Integration Process

Integrating the Opsware System with Windows OS installation technologies follows this general process:

1. Install a Windows operating system (2000 or NT) on a server. Install the OS manually (by using a CD or from a network) or use a vendor OS installation technology, for example:
   - Imaging by using Symantec Ghost

     Symantec Ghost uses imaging technology to install operating systems onto servers. An image is a sector-by-sector copy of the entire contents of a disk. The image is installed in its entirety. You cannot use Symantec Ghost to selectively install parts of an image.
   - Imaging by using PowerQuest Drive Image and PowerQuest DeployCenter 5.0

Deploy or upgrade Windows workstations or servers by remotely deploying an exact image of a hard disk.

- Remote installation by using Microsoft Remote Installation Services (RIS)

  RIS is a program for installing Windows 2000, and applications and Service Packs. RIS is made up of individual services that are combined to enable the remote installation of Windows 2000.

- Microsoft Systems Management Server (SMS)

  SMS deploys applications, software updates, and operating systems over simple or advanced enterprise networks.

2. Set up the Opsware Agent Installer to run the first time a system reboots after the OS installation. The Opsware Agent Installer assimilates a server by installing the Opsware Agent, reconciling the server (optional), and registering the server with the Model Repository.

   - On Windows 2000, use Windows 2000 System Preparation Tool (Sysprep) to modify the Windows registry to run the Opsware Agent Installer. Sysprep is used to prepare Windows 2000 System Images as part of an automated deployment.
   - On Windows NT, modify the Windows registry to run the Opsware Agent Installer once at startup after the Windows installation.

     Use the Windows registry entry to specify a batch file that lists the Opsware Agent Installer and associated options to run. You can specify the Opsware Agent Installer options to apply to a specific Windows server in a variety of ways; for example:

     - Maintain a file that maps server IP addresses or Ethernet MAC addresses to Opsware Agent Installer options and look up the server information in this file from the Opsware Agent Installer batch file.
     - If installing Windows from an image, prompt the user for the Opsware Agent Installer options when the server is running DOS before the Windows image is installed. Save the user input in an options file on the network. The options file will be specific to the Windows server being built. Read the option file from the Opsware Agent Installer batch file.

Example: Integration with Windows 2000 and Symantec Ghost

Integrating the Opsware System with Windows 2000 and Symantec Ghost follows this two-phased process.

Phase 1: Create an image

1. Manually install Windows 2000 on a server and customize the installation as required.
2. Use the Microsoft utility sysprep to remove all machine-specific configuration, such as the Windows security identifier (SID), network configuration, and so forth.
3. Use Symantec Ghost to take an image of the server and save the image to a network share.

   See your Symantec Ghost documentation for information about this process.

Phase 2: Provision a server by using the image

1. Boot the server by using an MS-DOS boot diskette that contains Symantec Ghost (a DOS application).
2. Run Symantec Ghost and install the image created in Phase 1.
3. Reboot the server.

   Windows runs a mini-setup wizard to configure the server.
4. Answer the prompts to the wizard. You can automate this process so that the answers to the wizard are pre-answered.

   After the wizard finishes, the system reboots again and is ready for Opsware assimilation.
5. Assimilate the server by running the Opsware Agent Installer and passing it the appropriate options.

   See “Opsware Agent Installer Commands and Options” on page 630 in this appendix for a description of each option.

   Automatically running the Opsware Agent Installer when using Symantec Ghost requires additional integration tasks. See “Running the Opsware Agent Installer by Using Sysprep” on page 639.

Running the Opsware Agent Installer by Using Sysprep

A standard Windows 2000 unattended installation file has a [GuiRunOnce] section. During the installation process, this section automatically adds the section's entries into the computer's RunOnce registry subkey. When the computer's first user logs on, the computer executes any commands in the RunOnce registry entry, and then removes the commands from the registry.

Using Sysprep allows you to arrange for the Opsware Agent Installer to run automatically.

- When using Sysprep to prepare a system for imaging, use the [GuiRunOnce] section in the file sysprep.inf to specify options for the Opsware Agent Installer when the server reboots after the setup wizard runs.

  For example:

```
[GuiRunOnce]
"net use z: \\yourshare\software"
"z:opsware-agent-5.1.14-win32-5.0.exe --template 56780002
--settime --reconcile full"
```

  The [`GuiRunOnce`] commands install the Opsware Agent each time the image is installed. However, the Opsware Agent Installer always runs with the same options. For example, if the commands include the `--template` and `--reconcile` options, you must create an image for each OS Template to which you plan to attach servers.

  Use the [`GuiRunOnce`] section to run a batch file that specifies the Opsware Agent Installer options. Retrieving the options from a batch file allows you to specify the options when the image is installed (rather than when the image is created). For example, you could specify which `--template` and `--reconcile` options to use in a batch file.

  For example:

```
[GuiRunOnce]
"net use z: \\yourshare\tools"
"z:install-opsware-agent.cmd"
```

Example File: Preparing a Windows 2000 System for Imaging

The following example sysprep.inf file prepares the Windows 2000 system prior to imaging a disk. The file edits the Windows 2000 registry to call a batch file. The batch file is setup to run when the server reboots after image installation.

```
[unattended]
OemSkipEula = Yes

[Guiunattended]
OEMSkipRegional = 1
OEMSkipWelcome = 1
```



```
AdminPassword = PASSWORD
AutoLogon=Yes
AutoLogonCount = 2
TimeZone = 90

[UserData]
Computername = *
orgName = Opsware Inc.
ProductID = XXXXX-XXXXX-XXXXX-XXXXX-XXXXX <--replace with your
Windows 2000 CD product code
FullName = shadow

[LicenseFilePrintData]
AutoMode = PerServer
AutoUsers = 9999

[Networking]
InstallDefaultComponents = Yes

[Identification]
joinworkgroup=Embryo

[GuiRunOnce]
"net use z: \\imagestore.example.com\winimages"
"z:\tools\agent-install"
```

Example Batch File: Running the Agent Installer for Windows 2000

The following example batch file runs the Opsware Agent Installer and specifies the Opsware Agent Installer options. This batch file uses third-party freeware tools to perform DOS command line parsing and determine the server MAC address. The tool uses this value to locate the server-specific Opsware Agent Installer options.

This example batch file uses the following third-party tools:

- NBMAC (available from http://www.kostis.net/en)
- LMOD (available from http://home.mnet-online.de/horst.muc/)

```
@echo off
rem Find the appropriate arguments file
set prov_dir=z:
set params_dir=%prov_dir%\params
set tools_dir=%prov_drive%\tools
```

```
rem construct a unique temporary file name
set mac_file=%tmp%\%random%.bat
%tools_dir%\nbmac > %mac_file%

%tools_dir%\nbmac | lmod set mac_addr=[$1] > %mac_file%
call %mac_file%
del %mac_file%
find "rem mac_addr %mac_addr%" %params_dir%\*.arg

set foundfile=
for %%i in (%params_dir%\*.arg) do call findit.cmd %%i
if "%foundfile%"=="" goto notfound
echo Executing post-install commands from %foundfile%
copy %foundfile% %tmp%\opost.cmd
call %tmp%\opost.cmd

del %foundfile%
del %tmp%\opost.cmd
goto _end

:notfound
echo No post-install commands found for mac address %mac_addr%
goto _end

:_end
```

Example: Integration with Windows NT and Symantec Ghost

This section provides information on integrating the Opsware System with Windows NT and Symantec Ghost and contains the following topics:

- Integrating with Windows NT and Symantec Ghost Process
- Example File: Preparing a Windows NT System for Imaging
- Example Batch File: Running the Agent Installer for Windows NT
- Example File: Configuring Machine-Specific Settings for Windows NT

Integrating with Windows NT and Symantec Ghost Process

Integrating the Opsware System with Windows NT and Symantec Ghost follows this two-phased process:

Phase 1: Create an image

1. Manually install the Windows NT operating system on a server and customize the installation as required.
2. Run a Registry modification file from diskette that prepares the Windows NT system prior to imaging the disk. By editing the Windows NT registry, the file sets up a post-installation batch file to run when the server reboots. The batch file indirectly specifies the Opsware Agent Installer options to run and the SID and hostname changes.

   See “Example File: Preparing a Windows NT System for Imaging” on page 644.
3. Use Symantec Ghost to take an image of the server and save the image to a network share.

   See your Symantec Ghost documentation for information about this process.

Phase 2: Provision a server by using the image

1. Boot the server by using a MS-DOS boot diskette that contains Symantec Ghost (a DOS application).

   When the server is running under DOS before Symantec Ghost runs, prompt the user for the Opsware Agent Installer options to run after the image is installed. Save the user input in an arguments file on the network, so that the file contains:

   - Identifying information for a server
   - The Opsware Agent Installer options to run on that server

   See “Opsware Agent Installer Commands and Options” on page 630 in this appendix for a description of each option.
2. Run Symantec Ghost and install the image created in phase 1.
3. Reboot the server.
4. (AUTOMATED) The post-installation batch file runs and assimilates the server by running the Opsware Agent Installer and passing it the appropriate options.

   The post-installation batch file looks for the correct arguments file to run for that server. Each arguments file contains the MID and other identifying data of the server where it should be run. When the post-installation batch file finds the correct file, it runs the Opsware Agent Installer commands in the file.

   See “Example Batch File: Running the Agent Installer for Windows NT” on page 644.

The post-installation batch file can call a file that configures machine-specific settings for the NT 4 system (for example, assigns a domain-unique SID and hostname).

See "Example File: Configuring Machine-Specific Settings for Windows NT" on page 645.

Example File: Preparing a Windows NT System for Imaging

The following example file prepares the Windows NT system prior to imaging a disk. The file (`nt4-sys-prep.reg`) edits the Windows NT registry to call a batch file, `post-install.cmd`. The batch file is setup to run when the server reboots after image installation.

```
nt4-sys-prep.reg
REGEDIT4
[HKEY_LOCAL_
MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\RunOnce]
"OpswarePrep"="Z:\\tools\\post-install.cmd"

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows
NT\CurrentVersion\Winlogon]
"AutoAdminLogon"="1"
"DefaultUserName"="Administrator"
"DefaultPassword"=""
```

Example Batch File: Running the Agent Installer for Windows NT

The following example shows a post-installation batch file that runs and assimilates a server by running the Opsware Agent Installer and any other commands in the post-installation arguments file.

The example file (`post-install.cmd`) accesses a server-specific file on the network by using a server's MAC address to find the correct file.

In this example file, `%foundfile%` is the arguments file. The arguments file is copied to the file `opost.cmd` and the file is run. The arguments file contains commands such as:

```
z:\agent\agentnt.exe --template 12340002 --settime --reconcile
full
```


The short form filename, which conforms to the 8.3 DOS filename conventions, is identical to the long form filename (such as `Opsware-agent-5.1.14-win32-4.0.exe`).

```
post-install.cmd
@echo off
rem Find the appropriate post-installation arguments file
set prov_drive=z:
set params_dir=%prov_drive%\params
set tools_dir=%prov_drive%\tools

rem construct a unique temporary file name
set mac_file=%tmp%\%random%.bat
%tools_dir%\nbmac > %mac_file%
%tools_dir%\nbmac | %tools_dir%\lmod set mac_addr=[$1] > %mac_
file%
call %mac_file%
del %mac_file%
find "rem mac_addr %mac_addr%" %params_dir%\*.arg

set foundfile=
for %%i in (%params_dir%\*.arg) do call %tools_dir%\findit.cmd
%%i
if "%foundfile%"=="" goto notfound
echo Executing post-install commands from %foundfile%
copy %foundfile% %tmp%\opost.cmd
call %tmp%\opost.cmd
del %foundfile%
del %tmp%\opost.cmd
goto _end

:notfound
echo No post-install commands found for mac address %mac_addr%
goto _end

:_end
```

Example File: Configuring Machine-Specific Settings for Windows NT

The following example shows a batch file that configures machine-specific settings for an NT 4 system after the image is installed. This batch file uses third-party freeware tools to perform DOS command line parsing and reset the SID and Windows hostname.

This example batch file uses the following third-party tools:

- LMOD (available from

  home.mnet-online.de/horst.muc/)

- NEWSID (available from

  www.sysinternals.com)

```
nt4-sh.bat
@echo off
rem Find the tools
set prov_drive=z:
set tools_dir=%prov_drive%\tools
set shtmp=d:\temp

rem construct a unique hostname
echo.|time > %shtmp%\_rnd.dat
echo.|date >> %shtmp%\_rnd.dat
%tools_dir%\crc32 %shtmp%\_rnd.dat | %tools_dir%\lmod set _
hostname=OPSW-[$2] > %shtmp%\_hostname.bat
call %shtmp%\_hostname.bat

rem remove temporary files
del /f %shtmp%\_rnd.dat
del /f %shtmp%\_hostname.bat

rem change sid and hostname
%tools_dir%\newsid /a %_hostname%
```

Integration with Network Installation Management and AIX

This section provides information about how to integrate the Opsware System with network installation management (NIM) and AIX and contains the following topics:

- Integration with NIM and AIX Overview
- Example File: NIM Customization to Install the Opsware Agent
- Example File: NIM Customization to Increase the Partition Size

Integration with NIM and AIX Overview

You can use AIX Network Installation Management (NIM) to manage the installation of the Base Operating System (BOS) and optional software on one or more servers running the AIX operating system.

Using NIM, you can install a group of machines with a common configuration or customize an installation for the specific needs of a given machine. The NIM environment is made up of client and server machines. A server provides resources (for example, files and programs required for installation) to another machine. A machine that is dependent on a server to provide resources is known as a client.

- For information about setting up and managing a NIM environment, see AIX Network Installation Management Guide and Reference and NIM: From A to Z in AIX 4.3 (an IBM Corporation Redbook) documentation from IBM Corporation.
- After you integrate the Opsware System with a NIM environment, you can use Opsware to manage AIX packages and install AIX applications on servers.

See the following topics:

- Chapter 6, "Managing Nodes on the Software Tree" on page 315 of this guide.
- "AIX Package Management" on page 272 in Chapter 5 for information about how the Opsware System manages AIX base and update filesets
- "About AIX Patches" on page 418 in Chapter 8 for information about how the Opsware System manages APARs

To integrate the Opsware System with NIM and AIX, perform these tasks:

1. Copy the Opsware Agent Installer to local disk storage, for example `/var/tmp`, on the server being assimilated.

Copy the Opsware Agent Installer to a directory that will not be empty when the server reboots. The operating system might empty the contents of the /tmp and /var/tmp directories when the server reboots.

2. Create a NIM customization script that performs these actions:
   - Increases the partition size to accommodate the Opsware Agent
   - Invokes the Opsware Agent Installer on reboot, and then removes itself

   You can accomplish this step in any of the following ways:
   - Create a customization SCRIPT to run after installation that creates a script that will run the Opsware Agent Installer after first reboot.

- Create a customization FB script to run at first reboot that runs the Opsware Agent Installer.
- Add the logic to run the Opsware Agent Installer and increase the partition size to an existing NIM customization script.

3. Define a NIM script resource on the NIM Master server for the customization script that you created.
4. During NIM BOS installation on a NIM client, specify the Opsware Agent customization script to run after installation.

Example File: NIM Customization to Install the Opsware Agent

The following example shows a NIM customization script that runs at the end of the NIM installation process. Running this script requires that the Opsware Agent Installer binaries are accessible in an NFS-exported directory; for example, `/export/nim/opsw_bin`.

```
#!/bin/sh
#
# Copyright (c) 2002 by Opsware, Inc
# All rights reserved.
#
# Setup Opsware Agent Installer
#

OS_VER=`uname -v`"."`uname -r`
TEMPLATE_ID="$1"

BASE=/
AGENT_SRC_DIR=/export/nim/opsw_bin
AGENT_DST_DIR=/var/lc/bootstrap
AGENT_START_SCRIPT=/etc/rc.d/rc2.d/S99zAgentInstaller
AGENT_OPTS="-os --settime --decommission --logfile
$AGENT_DST_DIR/install.log"

#
# mount the agent installer directory
#
NFS_HOST=nim.dev.opsware.com
MNT=/mnt.$$
mkdir ${MNT}
mount $NFS_HOST:$AGENT_SRC_DIR ${MNT}
AGENT_SRC_DIR=${MNT}

#
```



```
# use latest agent
#
AGENT=`ls $AGENT_SRC_DIR/opsware-agent-*-aix-$OS_VER | tail -1`
AGENT=`basename $AGENT`

if [ "$TEMPLATE_ID" ]; then
AGENT_OPTS="--template $TEMPLATE_ID --reconcile addonly $AGENT_
OPTS"
fi

echo "Agent installer version: $AGENT"
echo "Agent installer options: $AGENT_OPTS"

#
# copy over the agent installer
#
umask 022
mkdir -p $BASE/$AGENT_DST_DIR
cp $AGENT_SRC_DIR/$AGENT $BASE/$AGENT_DST_DIR/opsware-agent-
installer
chmod 555 $BASE/$AGENT_DST_DIR/opsware-agent-installer

#
# setup a start script to run the agent installer upon reboot
#
touch $BASE/$AGENT_START_SCRIPT
chmod 711 $BASE/$AGENT_START_SCRIPT
chown root:sys $BASE/$AGENT_START_SCRIPT

cat >> $BASE/$AGENT_START_SCRIPT <<EOF
#!/bin/sh

(
$AGENT_DST_DIR/opsware-agent-installer $AGENT_OPTS
rm –f $AGENT_START_SCRIPT
) 2>&1 | tee -a $AGENT_DST_DIR/`basename $AGENT_START_
SCRIPT`.log
EOF

umount ${MNT}
rmdir ${MNT}
```

Example File: NIM Customization to Increase the Partition Size

The following example shows a NIM Customization FB script that increases the partition size to accommodate the Opsware Agent on the NIM client.

```
#!/bin/sh
chfs -a size='2097152' /tmp
chfs -a size='2097152' /var
crfs -v jfs -g'rootvg' -a size='2097152' -m'/opt' -A''`locale
yesstr |
awk -F: ' {print $1}'`'' -p'rw' -t''`locale nostr | awk -F:
'{print
$1}'`'' -a frag='4096' -a nbpi='4096' -a ag='8'
mount /opt
chfs -a size='2097152' /opt
```

The part of the script that starts with `crfs` and ends with `-a ag='8'` is a single line.

Integration with Ignite-UX and HP-UX

This section provides information about how to integrate the Opsware System with Ignite-UX and HP-UX and contains the following topics:

- Integration with Ignite-UX and HP-UX Overview
- Example File: Ignite Configuration File
- Example File: Ignite Script to Invoke Opsware Agent Installer

Integration with Ignite-UX and HP-UX Overview

You can use Ignite-UX to facilitate installing and recovering HP-UX on HP computer systems in your computing environment.

Ignite-UX uses configurations that control the installation of the HP-UX operating system on servers. Using Ignite-UX, you can perform the following tasks:

- Create and reuse standard system configurations.
- Archive a standard system configuration and use that archive to replicate systems.
- Create customized processes to allow interactive and unattended installs.
- Recover OS and applications after crashes and hardware failures.

After running an Ignite-UX install session, you have a working HP-UX client system.

For information about how to set up and manage Ignite-UX, see the *Ignite-UX Administration Guide* from Hewlett-Packard Company.

After you integrate the Opsware System with an Ignite-UX system, you can use Opsware to manage HP-UX packages and install HP-UX applications on servers.

See the following topics:

- Chapter 6, "Managing Nodes on the Software Tree" on page 315 of this guide.
- "HP-UX Package Management" on page 273 in Chapter 5 for information about how the Opsware System manages HP-UX products and filesets
- "About HP-UX Patches" on page 419 in Chapter 8 for information about how the Opsware System manages patch products and patch filesets

To integrate the Opsware System with Ignite-UX and HP-UX, perform these tasks:

1. Copy the Opsware Agent Installer to local disk storage, for example `/var/tmp`, on the server that is being assimilated.

Copy the Opsware Agent Installer to a directory that will not be empty when the server reboots. The operating system might empty the contents of the /tmp and /var/tmp directories when the server reboots.

2. Create a script that invokes the Opsware Agent Installer on reboot, and then removes itself.
3. Save the script on the Ignite server.
4. Create a configuration file that includes the script in a `post_config_script` clause and add the configuration file to a configuration in the Ignite configuration INDEX in the directory `/var/opt/ignite/INDEX`.

   OR

   Add a `post_config_script` clause that includes the script to an existing configuration file referenced in the Ignite INDEX.

Example File: Ignite Configuration File

The following example shows how you might create a configuration file that includes the script in a `post_config_script` clause:

```
post_config_script+=”/var/opt/ignite/scripts/add_agent_
installer”
```

Example File: Ignite Script to Invoke Opsware Agent Installer

The following example shows an Ignite script that runs at first reboot to install the Opsware Agent on a server running HP-UX.

```
#!/bin/sh
#
# Copyright (c) 2002 by Opsware, Inc
# All rights reserved.
#
# Setup Opsware Agent Installer
#

OS_VER=`uname -r | cut -c3-`
TEMPLATE_ID="$1"

BASE=/
AGENT_SRC_DIR=/var/opt/ignite/clients
AGENT_DST_DIR=/var/lc/bootstrap
AGENT_START_SCRIPT=/sbin/rc3.d/S99zAgentInstaller
AGENT_OPTS="-os --decommission --logfile $AGENT_DST_DIR/
install.log"

#
# mount the agent installer directory
#
NFS_HOST=ignite.dev.opsware.com
MNT=/mnt.$$
mkdir ${MNT}
mount -F nfs $NFS_HOST:$AGENT_SRC_DIR ${MNT}
AGENT_SRC_DIR=${MNT}

#
# use latest agent
#
AGENT=`ls $AGENT_SRC_DIR/opsware-agent-*-hpux-$OS_VER | tail -1`
AGENT=`basename $AGENT`

if [ "$TEMPLATE_ID" ]; then
AGENT_OPTS="--template $TEMPLATE_ID --reconcile addonly $AGENT_
OPTS"
fi
```



```
echo "Agent installer version: $AGENT"
echo "Agent installer options: $AGENT_OPTS"

#
# copy over the agent installer
#
umask 022
mkdir -p $BASE/$AGENT_DST_DIR
cp $AGENT_SRC_DIR/$AGENT $BASE/$AGENT_DST_DIR/opsware-agent-
installer
chmod 555 $BASE/$AGENT_DST_DIR/opsware-agent-installer

#

# setup a start script to run the agent installer upon reboot
#
touch $BASE/$AGENT_START_SCRIPT
chmod 711 $BASE/$AGENT_START_SCRIPT
chown root:sys $BASE/$AGENT_START_SCRIPT

cat >> $BASE/$AGENT_START_SCRIPT <<EOF
#!/bin/sh

case "\$1" in
start_msg)
print "Installing Opsware Agent ($AGENT):"
exit $OKAY
;;
esac
(
$AGENT_DST_DIR/opsware-agent-installer $AGENT_OPTS
rm –f $AGENT_START_SCRIPT
) 2>&1 | /usr/bin/tee -a $AGENT_DST_DIR/`basename $AGENT_START_
SCRIPT`.log
EOF

umount ${MNT}
rmdir ${MNT}
```


Appendix D: Communication Test Troubleshooting

This appendix provides troubleshooting information for Communication Test errors. It is designed to help you discover why some of your managed servers might have unreachable agents and to help troubleshoot potential errors.

The Communication Test performs the following diagnostic tests to determine if an agent is reachable:

- **Command Engine to Agent (AGT)** – Determines if the Command Engine can communicate with the agent. The Command Engine is the Opsware system component that enables distributed programs to run across many servers. The Command Engine handles the entry of scripts into the Opsware Model Repository (the script storage location in the Opsware System) and the versioning of stored scripts.
- **Crypto Match (CRP)** – Checks that the SSL cryptographic files that the agent uses are valid.
- **Agent to Command Engine (CE)** – Verifies that the agent can connect to the Command Engine and retrieve a command for execution.
- **Agent to Data Access Engine (DAE)** – Checks whether or not the agent can connect to the Data Access Engine and retrieve its device record. 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.
- **Agent to Software Repository (SWR)** – Determines if the agent can establish an SSL connection to the Software Repository. The Software Repository is the central repository for all software that the Opsware system technology manages. It contains software packages for operating systems, applications, databases, customer code, and software configuration information.
- **Machine ID Match (MID)** – Checks that the machine ID (MID) on the server matches the MID registered in the Model Repository for the agent.

When the test run finishes, it returns results that show either success or failure for each test run on each server. For each failed test, the nature of failure is listed by error type in the error details column of the Communication Test window. In some cases, the failure of one test might prevent other tests from being executed.

For more information on the Opsware Agent, Opsware managed servers, and how to run Communication Test, see Chapter 2, Server Management.

Command Engine to Agent (AGT)

This test checks that the Command Engine can initiate an SSL connection to the agent and execute an XML/RPC request.

The thirteen possible results are:

- Command Engine to Agent (AGT) – OK
- Command Engine to Agent (AGT) – Untested
- Command Engine to Agent (AGT) – Unexpected error
- Command Engine to Agent (AGT) – Connection refused
- Command Engine to Agent (AGT) – Connection timeout
- Command Engine to Agent (AGT) – Request timeout
- Command Engine to Agent (AGT) – Server never registered
- Command Engine to Agent (AGT) – Realm is unreachable
- Command Engine to Agent (AGT) – Tunnel setup error
- Command Engine to Agent (AGT) – Gateway denied access
- Command Engine to Agent (AGT) – Internal gateway error
- Command Engine to Agent (AGT) – Gateway could not connect to server
- Command Engine to Agent (AGT) – Gateway timeout

Command Engine to Agent (AGT) – OK

No troubleshooting necessary.

Command Engine to Agent (AGT) – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the Command Engine cannot contact the agent, then no other tests are possible.

What Can I Do If a Test Is Not Run During an AGT Test?

First resolve all tests that failed, and then run the Communication Test again.

Command Engine to Agent (AGT) – Unexpected error

This result indicates that the test encountered an unexpected error.

What Can I Do If I Get an Unexpected Error?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware Support.

Command Engine to Agent (AGT) – Connection refused

This result indicates that the Command Engine is receiving a TCP reset packet when it attempts to connect to the agent on port 1002. The likely cause is that the agent is not running. A firewall might also be blocking the connection.

What Can I Do If the Connection is Refused During an AGT Test?

1. Log in to the server and confirm that the agent is running. For this information, see "Verifying That an Agent is Running" on page 680.
2. If the agent is not running, restart the agent. For these instructions, see "Restarting an Opsware Agent" on page 681.
3. From the managed server, use netstat to confirm that a socket is in listen mode on port 1002. If not, stop and restart the agent.
4. From the server itself, use Telnet to connect to the IP address of the server where the agent is installed and port (1002) that the agent is listening on. If this does not succeed, stop and restart the agent.
5. Verify that the Management IP address that Opsware is using to reach the server is the correct address. For this information, see "Checking Management IP of a Managed Server" on page 681. If the IP addresses do not match, stop and restart the agent, then rerun the test.
6. If the previous steps are performed and the test still fails, the problem is likely caused by either a software-based firewall on the server itself or an external firewall blocking the connection.

Command Engine to Agent (AGT) – Connection timeout

This result indicates that the Command Engine is not receiving any reply packets when it attempts to initiate a TCP connection to the agent on port 1002. The likely cause is that the server is not running, or that the IP address that Opsware is using to reach the agent is incorrect. (A firewall might also be blocking the connection.) To check the IP address that Opsware is using to reach the agent, see "Checking Management IP of a Managed Server" on page 681.

What Can I Do If the Connection Times Out During an AGT Test?

Follow the same steps used to resolve this issue specified in "What Can I Do If the Connection is Refused During an AGT Test?" on page 657.

Command Engine to Agent (AGT) – Request timeout

This result indicates that the Command Engine is able to successfully complete a TCP connection to the agent on port 1002, but no response is received from the agent in response to the XML-RPC request. The likely cause is that the agent is hung.

What Can I Do If the Request Times Out During an AGT Test?

1. Log in to the server and restart the agent. For these instructions, see "Restarting an Opsware Agent" on page 681.
2. Check to see whether or not some other process is consistently utilizing an excessive amount of the CPU on the server where the agent is installed. Also check to see if the system is performing slowly due to a lack of available memory and/or excessive file IO. In any of these cases, the system might be performing too slowly to permit the agent to respond to the test in a timely manner.

Command Engine to Agent (AGT) – Server never registered

This test indicates that the server being tested has neither been registered with the Command Engine, nor can it communicate with the Command Engine. The cause of this could be any number of reasons similar to those in the Agent to Command Engine (CE) test. It is also possible (but unlikely) that the agent was installed but never started.

What Can I Do If the Server Has Not Been Registered with the Command Engine During an AGT Test?

To troubleshoot this error, use the following procedures:

1. Ensure that the agent is running. For these instructions, see "Verifying That an Agent is Running" on page 680.

2. Ensure that the agent can contact the Command Engine.
3. If the agent is in a satellite data center, ensure that its Gateways are properly configured and that it is properly configured to use those Gateways. For these instructions, see “Checking Network Gateway Configuration” on page 682.
4. If the agent is not in a satellite:
   - Ensure the hostname “way” (no quotes) resolves to its valid IP address. For information on how to do this, see “Resolving Hostname” on page 683.
   - Verify that a connection can be established to port 1018 of way. Use the command “telnet way 1018" (or equivalent).

One (or more) of the above checks will fail. To solve that failure, refer to the corresponding error code for the Agent to Command Engine (CE) test on page 662, or to the realm connectivity and configuration test.

Command Engine to Agent (AGT) – Realm is unreachable

The satellite realm where the managed server is located is unreachable. This means that a path of tunnels between the Gateways in the Opsware core and the realm of the managed server cannot be established.

What Can I Do If the Realm Is Unreachable During an AGT Test?

This error could be due to a network problem, a malfunctioning or failed gateway, or a Gateway misconfiguration. Contact Opsware support for assistance in troubleshooting the Gateway network.

Command Engine to Agent (AGT) – Tunnel setup error

The Command Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do If I Get a Tunnel Setup Error During an AGT Test?

Contact your Opsware administrator.

Command Engine to Agent (AGT) – Gateway denied access

The Gateway is working but refused to proxy the connection on behalf of the agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the Command Engine access to the agent.

What Can I Do If the Gateway is Denied Access During an AGT Test?

Contact your Opsware administrator.

Command Engine to Agent (AGT) – Internal gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do If There is an Internal Gateway Error During an AGT Test?

Contact your Opsware administrator.

Command Engine to Agent (AGT) – Gateway could not connect to server

The Gateway could not establish a connection to the agent. This might be because the agent is not running, or because a firewall might be blocking the connection.

What Can I Do If the Gateway Couldn't Connect to the Server During an AGT Test?

If you suspect the agent is not running, see "Verifying That an Agent is Running" on page 680. To make sure that the Gateway can establish a connection to the IP address of the server where the agent is installed, try to ping the IP address of the server where the agent is installed.

Command Engine to Agent (AGT) – Gateway timeout

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do If the Gateway Times Out During an AGT Test

Ensure that network connectivity is available between the Gateways in the path between the realm of the managed server and the Opsware core.

Crypto Match (CRP)

This test checks that the SSL cryptographic files that the agent uses are valid.

The five possible results are:

- Crypto Match (CRP) – OK
- Crypto Match (CRP) – Untested
- Crypto Match (CRP) – Unexpected error
- Crypto Match (CRP) – Agent certificate mismatch
- Crypto Match (CRP) – SSL negotiation failure

Crypto Match (CRP) – OK

No troubleshooting necessary.

Crypto Match (CRP) – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the agent cannot be reached, then no other tests are possible.

What Can I Do If a Test Is Not Run During a CRP Test?

First resolve all tests that failed, and then run the Communication Test again.

Crypto Match (CRP) – Unexpected error

This result indicates that the test encountered an unexpected error.

What Can I Do If I Get an Unexpected Error During a CRP Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware Support.

Crypto Match (CRP) – Agent certificate mismatch

This result indicates that the SSL certificate that the agent is using (cogbot.srv) does not match the SSL certificate that is registered with Opsware for that agent.

What Can I Do If I Get a Certificate CN Mismatch During a CRP Test?

Use the Recert Agent Custom Extension to issue a new certificate to the agent. For this information, see the appropriate section in Chapter 10 of the *Opsware System 4.7 User's Guide*.

Crypto Match (CRP) – SSL negotiation failure

This result indicates that the agent is not accepting SSL connections for the Opsware core. (The Opsware core is the entire collection of servers and services that provide Opsware services.) The likely cause of this error is that one or more files in the agent crypto directory are missing or are invalid.

What Can I Do If I Get an SSL Negotiation Failure During an CRP Test?

Run the Server Recert custom extension in the "set allow recert flag only" mode on the server, and then Run the Opsware Agent Installer with the "-c" switch.

Reinstalling the agent with the "-c" option ("c" stands for "clean") removes all certs on the server and also removes the MID file, which forces the agent to retrieve a new MID from the Data Access Engine.

- For information on how to run the Recert Agent Custom Extension to issue a new certificate to the agent, see the appropriate section in Chapter 10 in the *Opsware System 4.7 User's Guide*.
- For information on how to install an Opsware Agent using the "-c" switch, see Chapter 2 of the *Opsware System 4.7 User's Guide*, Server Assimilation section.

After you reinstall the agent, run the test again to check if the agent is now reachable.

Agent to Command Engine (CE)

This test checks that the agent can connect to the Command Engine and retrieve a command for execution.

The sixteen possible results are:

- Agent to Command Engine (CE) – OK
- Agent to Command Engine (CE) – Untested
- Agent to Command Engine (CE) – Unexpected error
- Agent to Command Engine (CE) – Connection refused
- Agent to Command Engine (CE) – Connection timeout
- Agent to Command Engine (CE) – DNS does not resolve
- Agent to Command Engine (CE) – Old agent version
- Agent to Command Engine (CE) – Realm is unreachable

- Agent to Command Engine (CE) – No gateway defined
- Agent to Command Engine (CE) – Tunnel setup error
- Agent to Command Engine (CE) – Gateway denied access
- Agent to Command Engine (CE) – Gateway name resolution error
- Agent to Command Engine (CE) – Internal gateway error
- Agent to Command Engine (CE) – Gateway could not connect to server
- Agent to Command Engine (CE) – Gateway timeout
- Agent to Command Engine (CE) – No callback from agent

Agent to Command Engine (CE) – OK

No troubleshooting necessary.

Agent to Command Engine (CE) – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the agent cannot reach the Command Engine, then no other tests are possible.

What Can I Do If a Test Is Not Run During a CE Test?

First resolve all tests that failed, and then run the Communication Test again.

Agent to Command Engine (CE) – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During a CE Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware Support.

Agent to Command Engine (CE) – Connection refused

This result indicates that the agent is receiving a TCP reset packet when attempting to connect to the Command Engine on port 1018. The likely cause is that the agent is connecting to the wrong IP address. In other words, the agent does not know the correct IP address of the Command Engine. It is also possible that a firewall might be blocking the connection.

What Can I Do If the Connection is Refused During a CE Test?

1. Check that the name "way" resolves to its correct IP address. For instructions on how to do this, see "Resolving Hostname" on page 683.
2. Check to make sure there isn't a firewall refusing the connection to this IP address.

Agent to Command Engine (CE) – Connection timeout

This result indicates that the agent is not receiving any reply packets when it attempts to initiate a TCP connection to the Command Engine on port 1018. The likely cause is that the agent is connecting to the "wrong" IP address. In other words, the agent doesn't know the correct IP address of the Command Engine. A firewall might also be blocking the connection.

What Can I Do If the Connection Times Out During a CE Test?

Follow the same steps specified in "What Can I Do If the Connection is Refused During a CE Test?" on page 664.

Agent to Command Engine (CE) – DNS does not resolve

This result indicates that the agent cannot resolve the hostname "way" to a valid IP address. In other words, the agent does not know the correct IP address of the Command Engine.

What Can I Do If the Command Engine Name Does Not Resolve During a CE Test?

Log in to the server and use a command such as Telnet to confirm that the hostname "way" can resolve (for example: "telnet way 1018"). If not, check the DNS configuration of the server to make sure that the hostname "way" is configured to its correct IP address. For this information, see "Resolving Hostname" on page 683.

Agent to Command Engine (CE) – Old agent version

This result indicates that the agent was unable to contact the Command Engine, but the test was unable to determine the exact cause because the agent is out of date.

What Can I Do If the Agent is Out of Date During a CE Test?

If this error occurs, it will likely be for one of two reasons: either the hostname of the Command Engine ("way") did not resolve, or the connection was refused.

- If you believe that the hostname of the Command Engine ("way") did not resolve, then see "Agent to Command Engine (CE) – DNS does not resolve" on page 664.

- If you determine that the connection was refused, see "Agent to Command Engine (CE) – Connection refused" on page 663.

Alternatively, you can upgrade the agent to the latest version (contact Opsware support) and re-run the test. For information on how to install an Opsware Agent, see Chapter 2 of the Opsware System 4.7 User's Guide in the Server Assimilation section.

Agent to Command Engine (CE) – Realm is unreachable

The satellite realm where the managed server is located is unreachable. This error means that a path of tunnels between the Gateways in the Opsware core and the realm of the managed server cannot be established.

What Can I Do if the Realm is Unreachable During a CE Test?

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact your Opsware administrator for assistance in troubleshooting the Gateway network.

Agent to Command Engine (CE) – No gateway defined

The managed server is in a satellite realm, but its agent is not properly configured to use a Gateway. Agents located in satellites must use a Gateway to contact the core.

What Can I Do If No Gateway is Defined During a CE Test?

To troubleshoot this error, try the following:

1. Create or open the opswgw.args file on the managed server. The opswgw.args file is located on the managed server at:
   - **Unix/Linux**: /var/lc/cogbot/etc
   - **Windows**: %SystemDrive%\Program Files\Common Files\Loudcloud\cogbot\etc
2. Make sure that this file contains a single line as shown:

```
opswgw.gw_list: <gw_ip_address>:<gw_port>,<gw_up_
address>:<gw_port>
```

Agent to Command Engine (CE) – Tunnel setup error

The Command Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do If A Tunnel Setup Occurs Error During a CE Test?

Contact your Opsware administrator.

Agent to Command Engine (CE) – Gateway denied access

The Gateway is working but refused to proxy the connection on behalf of the agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the agent to access the Command Engine.

What Can I Do if the Gateway is Denied Access During a CE Test?

Contact your Opsware administrator.

Agent to Command Engine (CE) – Gateway name resolution error

The server running the Gateway in the Opsware core was unable to resolve the hostname "way". It must be able to do this in order to proxy connections on behalf of managed servers in satellite realms.

What Can I Do if a Name Resolution Error Occurs on the Gateway During a CE Test?

Log in to the server where the core gateway is located and use a command such as ping or host to confirm that the hostname "way" can be resolved (for example: "host way").

If you cannot connect, contact your Opsware administrator so that you can check the DNS configuration of the core Gateway server.

Agent to Command Engine (CE) – Internal gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do if an Internal Gateway Error Occurs During a CE Test?

Contact your Opsware administrator.

Agent to Command Engine (CE) – Gateway could not connect to server

The gateway could not establish a connection to the Command Engine. The situation might be because the Command Engine is not running, or because the Gateway is resolving the Command Engine hostname ("way") to the wrong IP address. It is also possible that a firewall might be blocking the connection.

What Can I Do if the Gateway Can't Connect to Server During a CE Test?

Check that the name "way" resolves to the correct IP address and that the Gateway can establish a connection to port 1018 at that IP. For more information, see "Resolving Hostname" on page 683 and "Verifying That a Port is Open on a Managed Server" on page 680.

Agent to Command Engine (CE) – Gateway timeout

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do if the Gateway Times Out During a CE Test?

Ensure that network connectivity is available between the Gateways in the path between the realm of the managed server and the Opsware core.

Agent to Command Engine (CE) – No callback from agent

The Command Engine was able to contact the agent, but the agent did not call back to retrieve its command. However, the agent reports that it can connect to a Command Engine. This most likely means that the managed server's name resolution mechanism (for example, DNS) is misconfigured to point the server to a different Opsware core facility than is currently stored for the server by Opsware.

What Can I Do if There is No Callback from Agent?

It is possible that the MID file is missing. If the MID file is missing, it can be recreated easily by creating a file called 'mid' in the correct location which contains the value of the "Server ID" from the Server's Properties page in the OCC.

If the MID file is not missing, then ensure that the server's name resolution mechanism (DNS, NIS, and so on) is properly configured so that the hostnames "spin" and "way" resolve to the appropriate IP addresses or Opsware core services in the same core, and that those hosts can be reached from the server on ports 1004 and 1018, respectively. If this is the case, it is likely that the server has been redirected to a different core recently, and the Agent has not yet registered, which will cause Opsware to update its records. It this issue remains unresolved for more than 12 hours, contact Opsware support.

Agent to Data Access Engine (DAE)

This test checks that the agent can retrieve its device record from Data Access Engine.

The fifteen possible results are:

- Agent to Data Access Engine (DAE) – OK
- Agent to Data Access Engine (DAE) – Untested
- Agent to Data Access Engine (DAE) – Unexpected error
- Agent to Data Access Engine (DAE) – Connection refused
- Agent to Data Access Engine (DAE) – Connection timeout
- Agent to Data Access Engine (DAE) – DNS does not resolve
- Agent to Data Access Engine (DAE) – Old agent version
- Agent to Data Access Engine (DAE) – Realm is unreachable
- Agent to Data Access Engine (DAE) – No gateway defined
- Agent to Data Access Engine (DAE) – Tunnel setup error
- Agent to Data Access Engine (DAE) – Gateway denied access
- Agent to Data Access Engine (DAE) – Gateway name resolution error
- Agent to Data Access Engine (DAE) – Internal gateway error
- Agent to Data Access Engine (DAE) – Gateway could not connect to server
- Agent to Data Access Engine (DAE) – Gateway timeout

Agent to Data Access Engine (DAE) – OK

No troubleshooting necessary.

Agent to Data Access Engine (DAE) – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the agent cannot reach the Data Access Engine then no other tests are possible.

What Can I Do If a Test Is Not Run During a DAE Test?

First resolve all tests that failed, and then run the Communication Test again.

Agent to Data Access Engine (DAE) – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During a DAE Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware Support.

Agent to Data Access Engine (DAE) – Connection refused

This result indicates that the agent is receiving a TCP reset packet when attempting to connect to the Data Access Engine on port 1004. The likely cause is that the agent is connecting to the wrong IP address. A firewall might also be blocking the connection.

What Can I Do If the Connection is Refused During a DAE Test?

1. Check that the name "spin" resolves to its correct IP address. For this information, see "Resolving Hostname" on page 683.
2. Check to make sure that a firewall is not refusing the connection to this IP address.

Agent to Data Access Engine (DAE) – Connection timeout

This result indicates that the agent is not receiving any reply packets when it attempts to initiate a TCP connection to the Data Access Engine on port 1004. The likely cause is that the agent is connecting to the wrong IP address. In other words, the agent does not know the correct IP address of the Command Engine. A firewall might also be blocking the connection.

What Can I Do If the Connection Times Out During a DAE Test?

Follow the same steps specified in "What Can I Do If the Connection is Refused During a DAE Test?" on page 669.

Agent to Data Access Engine (DAE) – DNS does not resolve

This result indicates that the agent cannot resolve the hostname "spin" to a valid IP address. In other words, the agent does not know the correct IP address of the Data Access Engine.

What Can I Do If the Data Access Engine Name Does Not Resolve During a DAE Test?

Log in to the server and use a command such as Telnet to confirm that the hostname "spin" can be resolved (for example: telnet spin 1004"). If not, check the DNS configuration of the server to make sure that the hostname "spin" is configured to its correct IP address. For information on how to resolve a hostname, see "Resolving Hostname" on page 683.

Agent to Data Access Engine (DAE) – Old agent version

This result indicates that the agent was unable to contact the Data Access Engine, and the test is unable to determine the exact cause because the agent is out of date.

What Can I Do If the Agent is Out of Date During an DAE Test?

If this error occurs, it will likely be for one of two reasons: either the hostname of the Data Access Engine ("spin") did not resolve, or the connection was refused.

- If you believe that the hostname of the Data Access Engine ("way") did not resolve, then see "Agent to Data Access Engine (DAE) – DNS does not resolve" on page 669.
- If you determine that the connection was refused, see "Agent to Data Access Engine (DAE) – Connection refused" on page 669.

Alternatively, you can upgrade the agent to the latest version (contact Opsware support) and re-run the test. For information on how to install an Opsware Agent, refer to Chapter 2 of the Opsware System 4.7 User's Guide, in the Server Assimilation section.

Agent to Data Access Engine (DAE) – Realm is unreachable

The satellite realm where the managed server is located is unreachable. This error means that a path of tunnels between the gateways in the Opsware core and the realm of the managed server cannot be established.

What Can I Do if the Realm is Unreachable During a DAE *Test?*

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact your Opsware administrator for assistance in troubleshooting the Gateway network

Agent to Data Access Engine (DAE) – No gateway defined

The managed server is in a satellite realm, but its agent is not properly configured to use a Gateway. Agents located in satellites must use a Gateway to contact the core.

What Can I Do If No Gateway is Defined During a DAE *Test?*

To troubleshoot this error, try the following:

1. Create or open the opswgw.args file on the managed server. The opswgw.args file is located on the managed server at:
   - **Unix/Linux**: /var/lc/cogbot/etc
   - **Windows**: %SystemDrive%\Program Files\Common Files\Loudcloud\cogbot\etc
2. Make sure this file contains a single line as shown:

```
opswgw.gw_list: <gw_ip_address>:<gw_port>,<gw_up_
address>:<gw_port>
```

Agent to Data Access Engine (DAE) – Tunnel setup error

The Data Access Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do if a Tunnel Setup Error Occurs During a DAE *Test?*

Contact your Opsware administrator.

Agent to Data Access Engine (DAE) – Gateway denied access

The Gateway is working but refused to proxy the connection on behalf of the agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the agent to access the Data Access Engine.

What Can I Do if the Gateway is Denied Access During a DAE *Test?*

Contact your Opsware administrator.

Agent to Data Access Engine (DAE) – Gateway name resolution error

The server running the Gateway in the Opsware core was unable to resolve the hostname “spin”. It must be able to do this in order to proxy connections on behalf of managed servers in satellite realms.

What Can I Do if There is a Name Resolution Error on the Gateway During a DAE *Test?*

Log in to the server where the core gateway is located and use a command such as ping or host to confirm that the hostname “spin” can be resolved (for example: “host spin”).

If you cannot connect, contact your Opsware administrator so that you can check the DNS configuration of the core Gateway server.

Agent to Data Access Engine (DAE) – Internal gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do if an Internal Gateway Error Occurs During a DAE *Test?*

Contact your Opsware administrator.

Agent to Data Access Engine (DAE) – Gateway could not connect to server

The gateway could not establish a connection to the Data Access Engine. This might be because the Data Access Engine is not running, or because the Gateway is resolving the Data Access Engine hostname ("spin") to the wrong IP address. It is also possible that a firewall might be blocking the connection.

What Can I Do if the Gateway Can't Connect to Server During a DAE *Test?*

Check that the name "spin" resolves to the correct IP address and that the Gateway can establish a connection to port 1018 at that IP. For more information, see "Resolving Hostname" on page 683 and "Verifying That a Port is Open on a Managed Server" on page 680.

Agent to Data Access Engine (DAE) – Gateway timeout

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do if the Gateway Times Out During a DAE Test?

Ensure that network connectivity is available between the Gateways in the path between the managed server's realm and the Opsware core.

Agent to Software Repository (SWR)

This test checks that the agent can establish an SSL connection to the Software Repository.

There 16 possible results are:

- Agent to Software Repository (SWR) – OK

- Agent to Software Repository (SWR) – Untested
- Agent to Software Repository (SWR) – Unexpected error
- Agent to Software Repository (SWR) – Connection refused
- Agent to Software Repository (SWR) – Connection timeout
- Agent to Software Repository (SWR) – DNS does not resolve
- Agent to Software Repository (SWR) – Old agent version
- Agent to Software Repository (SWR) - Server identification error
- Agent to Software Repository (SWR) – Realm is unreachable
- Agent to Software Repository (SWR) – No gateway defined
- Agent to Software Repository (SWR) – Tunnel setup error
- Agent to Software Repository (SWR) – Gateway denied access
- Agent to Software Repository (SWR) – Gateway name resolution error
- Agent to Software Repository (SWR) – Internal gateway error
- Agent to Software Repository (SWR) – Gateway Could not connect to server
- Agent to Software Repository (SWR) – Gateway timeout

Agent to Software Repository (SWR) – OK

No troubleshooting necessary.

Agent to Software Repository (SWR) – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the agent cannot reach the Software Repository, then no other tests are possible.

What Can I Do If a Test Is Not Run During a SWR Test?

First resolve all tests that failed, and then run the Communication Test again.

Agent to Software Repository (SWR) – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During a SWR Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware Support.

Agent to Software Repository (SWR) – Connection refused

This result indicates that the agent is receiving a TCP reset packet when attempting to connect to the Software Repository on port 1003. The likely cause is that the agent is trying to connect to the wrong IP address. A firewall might also be blocking the connection.

What Can I Do If the Connection is Refused During an SWR Test?

1. Check that the name "theword" resolves to the correct IP address. For this information, see "Resolving Hostname" on page 683.
2. Check to make sure that a firewall isn't refusing the connection to this IP address.

Agent to Software Repository (SWR) – Connection timeout

This result indicates that the agent is receiving a TCP reset packet when attempting to connect to the Software Repository on port 1003. The likely cause is that the agent is connecting to the wrong IP address. In other words, the agent does not know the correct IP address of the Software Repository. A firewall might also be blocking the connection.

What Can I Do If the Connection Times Out During an SWR Test?

Follow the same steps specified in "What Can I Do If the Connection is Refused During an SWR Test?" on page 674.

Agent to Software Repository (SWR) – DNS does not resolve

This result indicates that the agent cannot resolve the hostname "theword" to a valid IP address. In other words, the agent does not know the correct IP address of the Software Repository.

What Can I Do If the Software Repository Name ("theword") Does Not Resolve During an SWR Test?

Log in to the server and use a command such as Telnet to confirm that the hostname "theword" can be resolved (for example: "telnet theword 1003"). If not, contact your Opsware administrator so that you can check the DNS configuration of the server.

Agent to Software Repository (SWR) – Old agent version

This result indicates that the agent was unable to contact the Software Repository, and the test is unable to determine the exact cause because the agent is out of date.

What Can I Do If the Agent is Out of Date During an SWR Test?

If this error occurs, it will likely be for one of two reasons: either the hostname of the Software Repository ("theword") did not resolve, or the connection was refused.

- If you think that the hostname of the Software Repository ("theword") did not resolve, then see "Agent to Software Repository (SWR) – DNS does not resolve" on page 674.
- If you determine that the connection was refused, see "Agent to Software Repository (SWR) – Connection refused" on page 674.

Alternatively, you can upgrade the agent to the latest version (contact Opsware support) and re-run the test. For information on how to install an Opsware Agent, refer to Chapter 2 of the Opsware System 4.7 User's Guide, in the Server Assimilation section.

Agent to Software Repository (SWR) - Server identification error

Whenever an agent makes a request of the Software Repository, the identity of the server is validated to confirm that the server should be allowed access to the information requested. This error indicates that the Software Repository was unable to identify the server being tested, or incorrectly identified that server.

What Can I Do If I Get a Server Identification Error?

The Software Repository identifies servers based on the incoming IP address of the request. To troubleshoot this error, try the following:

1. Check the Device Properties tab for the server in the Opsware Command Center to see if Network Address Translation (NAT) is in use. If it is, make sure that NAT is statically configured, and that only one server is using the NAT address. If multiple servers are using the same IP address, you will need to reconfigure the NAT device. For more information, please refer to Chapter 2 of the Opsware System 4.7 User's Guide.
2. If the agent is installed on a cluster, check that each node in the cluster has a unique IP address at which it can be reached. You might have to add static routes to the server to ensure that connections made from that server to the Opsware core use the unique IP. If NAT is not in use, you can alternately mark the correct interface as the

"primary" interface through the Network Configuration tab for the server in the Opsware Command Center. For more information, please refer to Chapter 2 of the Opsware System 4.7 User's Guide.

3. The server's IP address might have changed recently. If this is the case, stop and restart the agent. For instructions on how to stop and start an agent, see "Restarting an Opsware Agent" on page 681.

Agent to Software Repository (SWR) – Realm is unreachable

The satellite realm where the managed server is located is unreachable. This error means that a path of tunnels between the gateways in the Opsware core and the realm of the managed server cannot be established.

What Can I Do if the Realm is Unreachable During a SWR Test?

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact your Opsware administrator for assistance in troubleshooting the Gateway network.

Agent to Software Repository (SWR) – No gateway defined

The managed server is in a satellite realm, but its agent is not properly configured to use a Gateway. Agents located in satellites must use a Gateway to contact the core.

What Can I Do If No Gateway is Defined During a SWR Test?

To troubleshoot this error, try the following:

1. Create or open the opswgw.args file on the managed server. The opswgw.args file is located on the managed server at:
   - **Unix/Linux**: /var/lc/cogbot/etc
   - **Windows**: %SystemDrive%\Program Files\Common Files\Loudcloud\cogbot\etc
2. Make sure that this file contains a single line as shown:

```
opswgw.gw_list: <gw_ip_address>:<gw_port>,<gw_up_
address>:<gw_port>
```

Agent to Software Repository (SWR) – Tunnel setup error

The Data Access Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do If a Tunnel Setup Error Occurs During a SWR Test?

Contact your Opsware administrator.

Agent to Software Repository (SWR) – Gateway denied access

The Gateway is working but refused to proxy the connection on behalf of the agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the agent to access the Software Repository.

What Can I Do if the Gateway is Denied Access During a SWR Test?

Contact your Opsware administrator.

Agent to Software Repository (SWR) – Gateway name resolution error

The server running the Gateway in the Opsware core was unable to resolve the hostname "theword". It must be able to do this in order to proxy connections on behalf of managed servers in satellite realms.

What Can I Do if a Name Resolution Error Occurs on the Gateway During a SWR *Test?*

Log in to the server where the core Gateway is located and use a command such as ping or host to confirm that the hostname "theword" can be resolved (for example: "host theword").

If you cannot connect, contact your Opsware administrator so that you can check the DNS configuration of the core Gateway server.

Agent to Software Repository (SWR) – Internal gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do if an Internal Gateway Error Occurs During a SWR *Test?*

Contact your Opsware administrator.

Agent to Software Repository (SWR) – Gateway Could not connect to server

The gateway couldn't establish a connection to the Software Repository. This error might be because the Software Repository is not running, or because the Gateway is resolving the Software Repository hostname ("theword") to the wrong IP address. It is also possible that a firewall might be blocking the connection.

What Can I Do if the Gateway Can't Connect to Server During a SWR *Test?*

Check that the name "theword" resolves to the correct IP address and that the gateway can establish a connection to port 1018 at that IP address. For more information, see "Resolving Hostname" on page 683 and "Verifying That a Port is Open on a Managed Server" on page 680.

Agent to Software Repository (SWR) – Gateway timeout

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do if the Gateway Times Out During a SWR Test?

Ensure that network connectivity is available between the Gateways in the path between the realm of the managed server and the Opsware core.

Machine ID Match (MID)

This test checks whether the MID that the agent reported matches that recorded in the Model Repository (Opsware data repository).

You can receive four possible errors from the Machine ID (MID) Communication Test:

- Machine ID Match (MID) – OK
- Machine ID Match (MID) – Untested
- Machine ID Match (MID) – Unexpected error
- Machine ID Match (MID) – MID mismatch

Machine ID Match (MID) – OK

No troubleshooting necessary.

Machine ID Match (MID) – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the agent cannot reach the Model Repository, then no other tests are possible.

What Can I Do If a Test Is Not Run During an MID Test?

First resolve all tests that failed, and then run the Communication Test again.

Machine ID Match (MID) – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During an MID Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware Support.

Machine ID Match (MID) – MID mismatch

This result indicates that the MID that the agent reported does not match the recorded MID in the Model Repository for that agent. The likely cause is that the Command Engine is running the test against the wrong agent.

What Can I Do If the MID is Mismatched During an MID Test?

To troubleshoot this error, try the following:

1. Check the Device Properties tab for the server in the Opsware Command Center to see if NAT is in use for this server. If it is, make sure that static, 1-to-1 NAT is being used. Opsware requires that all managed servers be reachable on a distinct, consistent IP address, so configurations that assign addresses dynamically or use port-based translation are not supported.
2. If the agent is installed on a cluster, check that each node in the cluster has a unique IP address at which it can be reached. You might have to add static routes to the server to ensure that connections made from that server to the Opsware core use the unique IP. If NAT is not in use, you can alternately mark the correct interface as the "primary" interface via the Network Configuration tab for the server in the Opsware Command Center.
3. The IP address might have changed recently. If this is the case, stop and restart the agent. For these instructions, see "Restarting an Opsware Agent" on page 681.

Common Troubleshooting Tasks

The following list of troubleshooting tasks are common to more than one Communication Test error:

- Verifying That an Agent is Running
- Verifying That a Port is Open on a Managed Server

- Restarting an Opsware Agent
- Checking Management IP of a Managed Server
- Checking Network Gateway Configuration
- Resolving Hostname

Verifying That an Agent is Running

To verify that an agent is running on a server, perform the following steps:

1. On Solaris, HP-UX, or AIX, enter this argument at the command line:

   ```
   /usr/ucb/ps auxwww | grep cog
   ```

   You should get this result if the agent is running:

   ```
   /opt/OPSW/bin/python /opt/OPSW/blackshadow/shadowbot/
   daemonbot.pyc --conf /var/lc/cogbot/etc/cogbot.args
   ```

2. On Linux, enter this argument at the command line:

   ```
   ps auxwww | grep cog
   ```

   You should get this result if the agent is running:

   ```
   /opt/OPSW/bin/python /opt/OPSW/blackshadow/shadowbot/
   daemonbot.pyc --conf /var/lc/cogbot/etc/cogbot.args
   ```

3. On Windows, from the Administrative Tools | Services, check to make sure that the 'shadowbotservice' is running.

Verifying That a Port is Open on a Managed Server

For some errors, you will need to verify that the port is open on the server where the agent is installed. To do this, perform the following steps:

1. Check if the port is open.
2. On Solaris, HP-UX, AIX, or Linux enter:

   ```
   'netstat -an | grep 1002 | grep LISTEN'
   ```

   If the port is open on the box, you should get back the following:

   ```
   *.1002    *.*    0      0 24576      0 LISTEN
   ```

3. On Windows, at the command prompt enter:

   ```
   'netstat -an | find "1002" | find "LISTEN"':\
   ```


If the port is open on the box, you should get back the following result:

```
TCP0.0.0.0:10020.0.0.0:0LISTENING
```

4. Confirm that the port is actually open. To do this, from the computer where the agent is installed, Telnet to port 1002 by using both localhost and the external IP address of the server. Performing the Telnet will help you confirm that a 'connection refused' message is being caused by the lack of an open port on the managed server rather than a problem with networking hardware between the core and the managed server.

Restarting an Opsware Agent

For Solaris, Linux, or AIX, perform the following steps:

1. To stop an Opsware Agent on Solaris, Linux, or AIX, execute the following command:

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

2. To restart the Opsware Agent on Solaris, Linux, or AIX, execute the following command:

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

For HP-UX, perform the following steps:

1. To stop an Opsware Agent on HP-UX, execute the following command:

   ```
   /sbin/init.d/cogbot stop
   ```

2. To restart an Opsware Agent on HP-UX, execute the following command:

   ```
   /sbin/init.d/cogbot start
   ```

For Windows, perform the following steps:

1. Execute the following command to stop and start an Opsware Agent:

   ```
   net stop shadowbot
   ```

   ```
   net start shadowbot
   ```

Checking Management IP of a Managed Server

To check the Management IP of a managed server, perform the following steps:

1. To view the management IP of the managed server, log in to the Opsware Command Center.
2. From the Navigation panel, click Servers ➤ Managed Servers.
3. From the Managed Servers list, click the display name of the server for which you want to check the Management IP.
4. Click the Network tab of the server's properties.
5. Check to make sure that the Management IP address matches the IP address of the managed server.

Checking Network Gateway Configuration

To check the network Gateway configuration, perform the following steps:

1. On Solaris, enter this command to check routing table:

   ```
   netstat -rn
   ```

   Your results should look like this:

   ```
   default              192.168.8.1          UG        1   5904
   ```

   where `192.168.8.1` is the IP of the gateway.

2. On Linux, enter this command to check routing table:

   ```
   route -n
   ```

   Your results should look like this:

   ```
   0.0.0.0       192.168.8.1    0.0.0.0       UG   0     0
   0 eth0
   ```

   where `192.168.8.1` is the IP of the gateway.

3. On Windows, enter this command to check routing table:

   ```
   route print
   ```

   Your results should look something like this:

   ```
   0.0.0.00.0.0.0192.168.8.1192.168.8.12020
   ```

   where `192.168.8.1` is the IP of the gateway.

4. In each case, you should also ping 192.168.8.1 (IP) to confirm that you can actually reach the gateway.

Resolving Hostname

All managed servers (those with agents) must be able to resolve unqualified Opsware service names for these Opsware components:

- spin (Data Access Engine)
- way (Command Engine)
- theword (Software Repository)

If you need to make sure one of these hostnames resolve correctly, contact your Opsware administrator to find out what qualified hostname or IP address these service names should resolve to.

1. Try to ping the host. For example, execute the following command if you wanted to resolve the hostname 'way':

   ```
   ping way
   ```

2. If the hostname cannot resolve, you will get the following errors:

   Linux/Solaris/AIX/HP-UX:

   ```
   ping: unknown host way
   ```

   Windows:

   ```
   Ping request could not find host way. Please check the name
   and try again.
   ```

3. If the hostname can resolve, you might get back various permutations of these types of messages (OS independent):

   ```
   way is alive
   ```

   or

   ```
   pinging way (ip) with 32 bytes of data
   ```


Appendix E: Glossary

IN THIS APPENDIX

This appendix and describes the Opsware terminology and the acronyms used within the Opsware System.

Access & Authentication Directory. Stores user account information. Used to authenticate users when they use the Opsware System and control their access to critical Opsware System resources. Implemented using an LDAP directory.

Ad-Hoc Scripts. A script that is created (or uploaded) and then immediately executed by a user. The script is intended for one-time use and is not stored in the Opsware System.

administrator. *See* Opsware administrator.

agent. *See* Opsware Agent.

Agent Installer. An application that installs the Opsware Agent on a server.

Agent Uninstaller. An application that uninstalls the Opsware Agent on a server.

application provisioning. The process of installing an application from the Software Repository onto a selected set of servers. Application provisioning can also involve the automatic execution of installation and post-installation scripts. Applications can be provisioned by using the Install Software Wizard, the Install Template Wizard, or by attaching a server to a node and then reconciling the server.

assimilation. *See* server assimilation.

Automated Configuration Tracking. An Opsware subsystem that allows users to monitor critical configuration files and configuration databases. When the Opsware System detects a change in a tracked configuration file or configuration database, the system can perform a number of actions, including backing up the configuration file or sending an email to a designated individual or group.

available patch. A patch that the patch administrator has tested and marked as available. Only patches that have been marked as available can be installed by anyone other than a patch administrator. (The patch administrator can install an unavailable patch in order to test it.)

available server. A reserve of new, unconfigured Opsware-enabled servers ready for quick deployment. The provisioned server can be moved into the Live environment to replace existing servers, add capacity, or support new applications. While optional, provides faster recovery options in cases of hardware failure.

backup. A feature in Automated Configuration Tracking that performs a backup of a file or database when it detects a change to a tracked configuration file or database. This action is performed only if the backup action is selected in the configuration tracking policy for the file or database.

backup (CDR). Process of saving the entire contents of the current Live directory for a specific service to the Backup directory. Code Deployment & Rollback (CDR) saves the backup copy to the local disk for the host on which the Backup operation was run. Only one backup copy is maintained at any time for a service.

backup event. An event that causes configuration files or configuration databases to be backed up. Types of backup events include manual, full, and triggered.

blocked attachment. An attachment that is not installed when that template is applied. The attachment also does not appear in child templates or folders.

Boot Server. A part of the OS Provisioning Subsystem that 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. A part of the OS Provisioning Subsystem that facilitates communication between the OS Boot Agent and the Command Engine for OS provisioning.

CDR. *See* Code Deployment & Rollback (CDR).

change log. An audit trail of changes made to a node (read-only). Tracks changes made to a node. Identifies who has recently modified the node to add or remove software packages, add or remove operating systems, add or move servers, and create or remove subordinate nodes.

Code Deployment & Rollback (CDR). An Opsware subsystem used to push updated code and content to staging host servers.

Code Deployment Role. A specific role that authorizes access to capabilities and functions with the Opsware Code Deployment & Rollback Subsystem.

Command Engine. The Opsware System component that enables distributed programs to run across many servers. The Command Engine handles the entry of scripts into the Opsware Model Repository (the script storage location in the Opsware System) and the versioning of stored scripts. Command Engine scripts are written in Python and run on the Command Engine server.

Communication Test. A feature that helps in identifying managed servers with unreachable Opsware Agents. A Communication Test lists all servers with unreachable agents, returns specific errors associated with each unreachable agent, and provides troubleshooting information to resolve the error. The Communication Test runs various tests like Command Engine to Agent Communication, Crypto Match, Agent to Command Engine Communication, Agent to Data Access Engine, Agent to Software Repository Communication, and Machine ID mismatch to determine if an Opsware Agent is reachable.

configuration tracking policy. The configuration tracking policy defines the set of files or configuration databases to be monitored, and the actions to be taken when change is detected to a tracked file.

configuration tracking reconcile. Process by which new configuration tracking polices or changes to existing configuration tracking polices are deployed on servers.

custom attributes. Attributes such as miscellaneous parameters and named data values that users can set for servers in the Opsware Command Center. Used when performing a variety of Opsware functions, including network and server configuration, notifications, and CRON script configurations.

custom extension. Custom Command Engine scripts that extend Opsware System functionality to customers to cover their specific needs.

customer. An account within the Opsware System that has access to designated resources, such as servers and software.

cutover. A feature in CDR, that causes the Update directory and current Live directory to be identical. Performed automatically by determining the differences between the Update directory and current the Live directory. The files that are different are synchronized from the Update directory to the current Live directory.

Data Access Engine. The 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.

data center. Legacy term. *See* facility.

Data Center Intelligence Reporting. An interface for mining the data that is contained in the Model Repository about all managed servers.

deactivated server. Server removed from Opsware management even though its history still exists.

deployment. Within CDR, automatically pushes code and content from a staging server to a live network server.

deprecated. A possible state of a package or patch in the Opsware System. A deprecated package or patch can no longer be installed on a managed server, but might still be installed on a server before the patch or package was deprecated.

device. Legacy term. *See* server.

Distributed Scripts. An Opsware subsystem that allows you to manage scripts in your managed environment.

Dormant Opsware Agent. An Opsware Agent that runs in the dormant mode after its installation when the Opsware System core is not available on the network. The dormant agent periodically attempts to contact the core and when the core is available, it performs the initialization tasks to complete its installation.

email notification list. In the Automated Configuration Tracking Subsystem, an email can be sent to the email addresses in the email notification list whenever a change to a tracked file or configuration database is detected.

Environment Tree. The Environment Tree manages characteristics about a customer's unique data center environment, including hardware, location of servers, network infrastructure, application names, business units, and service levels assigned to servers and applications. The information contained in the Environment Tree, combined with the information contained in the Software Tree, is utilized by the Opsware Automation Platform to model and simulate operational changes before they are executed in the production environment.

facility. Collection of servers that a single Opsware Model Repository manages, and the database that stores information about the managed environment.

full backup. During a full backup, all tracked configuration files that were selected to be backed up are backed up (and not just the files that have changed). Full backup is performed if you select backup as the action for a tracked configuration file.

full reconcile. A reconcile process that reconciles a server with all of the nodes that it has been assigned to.

group. An identified association within the Opsware System used to aggregate servers for a specific purpose or node (for example, configuration and backup).

Import Media tool. A utility script included with the Opsware System that is used to import OS media from the Media Server to the Opsware System.

incremental backup. During an incremental backup, only targets that have changed since the last backup (and that have been selected to be backed up) are backed up. Incremental backup is performed if you select backup as the action for a tracked configuration file.

inherited attachment. An attachment that is inherited from an ancestor folder or a template.

initialization. Legacy term. *See* OS Provisioning.

IP Range Groups. A designated set of servers assigned to a customer account, grouped by either a physical or a logical list.

IP Ranges. A designated grouping of servers.

Live directory. In CDR, the directory that stores the actual code and content required to run a live site.

local attachment. An attachment that is attached directly to a folder or a template.

MAC. *See* Media Access Control Address (MAC).

Machine ID (MID). A unique identifier that the Opsware system uses to identify the server. The Opsware System assigns a unique number to the server when it first registers and stores the Machine ID and uses it to identify each server.

managed server. A Server that has an Opsware Agent installed on it and is under the control of a particular Opsware core.

management IP. The IP address that the Opsware System uses to communicate with the Opsware Agent on the server.

manifest. Within CDR, a list of files that indicate the results or preview of an update to be performed. Each entry in the list specifies the file size, last-modified date and timestamp, and the full directory path to the listed file.

Media Access Control Address (MAC). The network interface card's unique hardware number. The MAC is used as the server's physical address on the network.

Media Resource Locator (MRL). A network path in URL format that is registered with the Opsware System. The path defines the installation media for an OS.

Media Server. Contains the vendor-supplied OS media used during OS provisioning over the network. The OS media on the Media Server is accesses over the network by using NFS for Linux and Solaris OS provisioning, and by using SMB for Windows OS provisioning.

MID. *See* Machine ID (MID).

Model Repository. The Opsware data repository that stores information about managed server configurations within the Opsware System. It contains all information necessary to build, operate, and maintain an Opsware-managed site, including a list of all servers

under management, the hardware associated with these servers, including memory, CPUs, storage capacity, and the configuration of these servers, including IP addresses, DNS configuration, and so on.

Modeling and Change Simulation Engine. Opsware System enables users to first model and simulate proposed operational changes to their environment before propagating these changes to production servers and applications. Utilizing the information contained in the Software and Environment Trees, the Modeling and Change Simulation Engine maintains a model of the various hardware and software configurations and other customer characteristics associated with each of the production servers under Opsware System's control.

MRL. *See* Media Resource Locator (MRL).

multimaster mesh. A set of two or more Opsware cores that are linked by synchronizing the data in the Model Repositories at each of the cores. The Model Repositories in each of the cores are continually updated so that they are exact duplicates of each other. All the Opsware cores in a multimaster mesh can be managed through a single Opsware Command Center.

Multimaster Replication Engine. A component of the Opsware System that allows customers to store and maintain a blueprint of the software and environmental characteristics of each facility in multiple locations, so that the infrastructure can be easily rebuilt in the event of a disaster. It also assists in facility migration activities and knowledge sharing across the enterprise.

My Jobs. A page in the Opsware Command Center that displays a list of jobs from the Model Repository such as software installation or server provisioning.

My Scripts. Private scripts that can only be executed by the user who created the script. My Scripts are created for personal use.

name-value pairs. Legacy term. See custom attributes.

node. A hierarchical set of categories or types that classify hardware, software, configuration, or other components of a site's infrastructure. Simplifies server management (for example, servers within the Opsware System) and the software applications and configurations associated with those servers.

node-based configuration tracking policy. A configuration tracking policy defined for a particular software node for a particular application.

OCLI. *See* Opsware command Line Interface (OCLI).

Opsware administrator. Responsible for overall administration, policy, and practices for individuals accessing the Opsware System. Can add users and define access to specific

Opsware System features that allow users to view site information and deploy new code and content to their site.

Opsware Agent. Intelligent software on Opsware-managed servers that is used to make changes to the servers. Depending on the request, might use global Opsware services. Some functions supported include software installation and removal, software and hardware configuration, server status reporting, and auditing.

Opsware Automation Subsystems. Opsware System is made up of a set of Opsware Automation Subsystems. Opsware Automation Subsystems are the components that automate particular IT processes. The Opsware Automation Subsystems include the following functions: Software Provisioning, Patch Automation, Configuration Tracking, Code Deployment and Rollback, Script Execution, and Data Center Intelligence Reporting.

Opsware Command Center. Web-based user interface for managing the Opsware environment.

Opsware Command Line Interface (OCLI). An alternative interface to the Opsware Command Center. The OCLI allows you to perform some actions not possible though the browser-based interface of the Opsware Command Center, such as uploading multiple packages, patches, AIX filesets, and so forth, in a batch operation.

Opsware core. The set of servers that run the components that make up the Opsware System, which includes the Model Repository, the Software Repository, the Data Access Engine, and the Access & Authentication Directory.

Opsware System. The server management application to preserve the knowledge of system administrators, network engineers, and database administrators in a centralized knowledgebase. Automates previously manual tasks associated with the deployment, support, and growth of a data center infrastructure.

OS Build Agent. A part of the OS Provisioning Subsystem that is responsible for registering bare metal servers in the Opsware System and guiding the installation process.

OS media. Installation software for an OS from the software vendor that is distributed on a CD-ROM, or DVD, or can be obtained by downloading the software from the vendor's FTP site.

OS Provisioning. Process of installing a basic set of software components, including an operating system and an Opsware Agent to add a server into the Opsware managed environment. After provisioning is complete, the server is ready to be managed by the Opsware System.

Package Repository. Legacy term. *See* Software Repository.

packages. The collection of executables, configuration, or script files that are associated with an Opsware-installable application or program. In the Opsware system a package contains software package files registered in the Software Repository. Contains software for operating systems, applications (for example, BEA WebLogic, IBM WebSphere), databases, customer code, and software configuration information.

partial reconcile. A reconcile process that only reconciles servers based on the nodes that the user has currently selected.

patch management administrator. Administrator responsible for testing patches and defining patch options, such as installation and uninstallation scripts. A patch cannot be installed by other personnel until the patch administrator has marked the patch available through the Opsware Command Center.

Patch Management. An Opsware subsystem that allows you to upload, test, and deploy patches in a safer and uniform way.

permissions. A designation within a User Role that allows or disallows access to Opsware System features and resources.

preview reconcile. Before the Opsware System installs software on a server, it performs a preview reconcile, and determines what will happen when the actual reconcile is performed (for example, what packages will be installed or removed, what server reboots are required, and so forth.)

primary IP. A locally-configured IP address of the management interface.

privileges. *See* Permissions. *See* User Role.

reconcile. Process of updating the actual software configuration of a server based on the specified configuration stored in the Model Repository.

reconcile output. After a reconcile operation completes, the Opsware System displays the reconcile output for each server that was reconciled. The reconcile output aggregates output from the various installation, uninstallation, or post-installation scripts, messages from the Opsware System, and messages from the system utilities that reconcile uses to perform the installation and uninstallation of packages, operating systems, and patches.

Reconcile Software Wizard. A Wizard that can enable a user to directly invoke the reconcile process on a selected server or a group of servers.

restore. A function of the Automated Configuration Subsystem that allows the user to return the configuration file or database to a previous state, when the backup action for a tracked file or database is selected.

restore. Within CDR, the process of restoring the previous Live directory from the Backup directory to the Live directory.

restore queue. Queue in which configuration files are placed before they are restored to a server.

Role. Legacy term. *See* node.

rollback. Within CDR, returns a site to the state prior to the last cutover. During rollback, restores the set of modified and deleted files to the Live directory.

Script Execution. *See* Distributed Scripts.

sequence. Process within CDR that simplifies deployment by grouping a collection of service operations and synchronizations that can be performed as a single task.

Sequence Editor. In CDR, a predefined User Role to create, modify, or delete a sequence definition.

Sequence Performer(Production). In CDR, a predefined User Role to directly perform or request performance of a sequence action on production hosts.

Sequence Performer (Staging). In CDR, a predefined User Role to directly perform or request performance of a sequence action on staging hosts.

Sequence Requester(Production). In CDR, a predefined User Role to request performance of a sequence action on production hosts.

Sequence Requester(Staging). In CDR, a predefined User Role to request performance of a sequence action on staging hosts.

servers. Any specific hardware. Specific nodes are attached to servers that determine the specific software, configuration, and other server attributes.

server assimilation. The Opsware System assimilates servers that are already functioning in the operational environment, which allows users to deploy and manage new applications installed on those servers. Assimilating servers installs Opsware Agents on the servers and registers them with the Model Repository.

server baselines. Process of defining and provisioning servers with standard configurations. Opsware templates can be used to automate the building of complete server baselines.

Server ID. The primary key in the Opsware Model Repository that represents a given server. The Server ID is used internally in the Opsware System.

server lifecycle. The various server states assigned to a server by the Opsware System. Server states include Unprovisioned, Available, Installing OS, and Managed.

server locking. A feature that allows users with read/write access to lock a managed server to prevent any server-modifying operations from being performed on that server until it is unlocked. The server locking feature also prevents server-modifying scripts from being run on locked servers if the script has been flagged.

server management. Process by which users can manage and track servers in an Opsware-managed environment. The Opsware System forces changes to the operating environment by first changing the centralized configuration information in the Model Repository and then changing the actual configuration of physical servers.

Server Pool. Servers that have registered their presence with the Opsware System but do not have a full operating system installed.

server provisioning. Process of installing a basic set of software components that include the operating system, an Opsware Agent, and other system utilities and debugging tools to manage the server. Configuration is defined in the Model Repository.

server reconcile. A process that compares a designated server image from the Model Repository with a specific server, checking for configuration, content, versions, and so forth, to determine if the live server is current and up-to-date. Includes OS, applications, upgrades, and patches.

Server Search. Feature that allows you to search for servers based on a variety of criteria, including OS version, installed package, customer, and installed patch.

Server Status. Feature that defines server availability. The three major status conditions are USE, STAGE, and STATE.

server-based configuration tracking policy. A configuration tracking policy that is defined for a particular server or group of servers, rather than for a particular software node (application).

service. A host application (for example, BEA WebLogic, Allaire ColdFusion, Microsoft IIS, Apache Web Server, or iPlanet Application Server).

Service Editor. In CDR, a predefined User Role to define and modify or delete service definitions.

Service Performer (Production). In CDR, a predefined User Role to directly perform or request performance of service operations on production hosts (servers).

Service Performer (Staging). In CDR, a predefined User Role to directly perform or request performance of service operations on staging hosts.

Service Requester (Production). In CDR, a predefined User Role to request performance of service operations on production hosts.

Service Requester (Staging). In CDR, a predefined User Role to request performance of service operations on staging hosts.

service-instance. Multiple independent instances of a service running on a host (for example, BEA WebLogic, which can run single or multiple instances).

Service Levels. User-defined categories that are used to group servers in any arbitrary way. For example, a user can group servers by functionality, tier, application, or ontogeny.

Shared Scripts. Public scripts that every Opsware System user can access.

Site Backup directory. In CDR, the directory that stores a complete backup of the Live directory when the user issues a Backup service operation.

Site Previous directory. In CDR, the directory that stores the files that have changed between the current Live directory and its previous state prior to the last cutover. It holds all the changes necessary to revert the Live directory back to the state that it was in before the last cutover.

Software Provisioning. An Opsware subsystem that allows system administrators to install, configure, and remove packaged software in a systematic way across servers that are distributed over many different facilities.

Software Repository. The central repository for all software managed by the Opsware System. It contains software packages for operating systems, applications, databases, customer code, and software configuration information.

Software Tree. The Software Tree records a variety of information for software applications and operating systems, including data about how changes to a given software application might impact other existing applications.

synchronization. Process within CDR to move modified files from a directory on a source host to a directory on a destination host.

Synchronization Editor. In CDR, a predefined User Role to create, modify, or delete a synchronization definition.

Synchronization Performer. In CDR, a predefined User Role to directly perform or request performance of a synchronization action.

Synchronization Requester. In (CDR), a predefined User Role to request performance of a synchronization action.

template. Used to install a set of (usually related) applications through a single invocation of a wizard.

template inheritance. Process by which templates and folders inherit all attachments of the folder they reside in. Inheritance is propagated from parent (folder) to child (template or folder) and to all children of children.

Update directory. The directory that CDR writes to when synchronizing modified files in source and destination hosts. After synchronization, the Update directory is different from the current Live directories. After cutover, the Update directory and current Live directory are identical.

user. An individual with access to the Opsware environment. The level of access is provided by the assignment of user nodes to individual users by the Opsware administrator.

User Directory. Legacy term. *See* Access & Authentication Directory.

User Role. Defined permissions and access to specific Opsware resources. Roles are granted to assigned users. (Opsware administrator access only.)

Web Service APIs. A Web services interface to facilitate the integration of operations and business support systems with Opsware System. The Opsware Web Services APIs allow other IT systems, such as customers' existing monitoring, trouble ticketing, billing, and virtualization technology, to exchange information with the Opsware System.

Web Services Data Access Engine. A Web services interface to the Model Repository that provides increased performance to other Opsware System components.

Wizard. Graphical user interface that groups a series of data collection operations, actions, and jobs into a logical, easy-to-understand workflow presentation.

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

Z

