



Opsware[®] SAS 6.5 Content Migration 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

Opware SAS Version 6.5

Copyright © 2000-2007 Opware Inc. All Rights Reserved.

Opware Inc. Unpublished Confidential Information. NOT for Redistribution. All Rights Reserved.

Opware is protected by U.S. Patent Nos. 6,658,426, 6,751,702, 6,816,897, 6,763,361 and patents pending.

Opware, SAS Web Client, Model Repository, Data Access Engine, Web Services Data Access Engine, Software Repository, Command Engine, Opware Agent, Model Repository Multimaster Component, and Code Deployment & Rollback are trademarks and service marks of Opware Inc. All other marks mentioned in this document are the property of their respective owners.

Additional proprietary information about third party and open source materials can be found at <http://www.opware.com/support/sas65tpos.pdf>.

Table of Contents

| | |
|---|-----------|
| Preface | 17 |
| <hr/> | |
| Contents of this Guide | 17 |
| Conventions in this Guide | 18 |
| Icons in this Guide | 19 |
| Guides in the Documentation Set and Associated Users | 20 |
| Opware, Inc. Contact Information | 21 |
| Chapter 1: Migrating SAS Content to 6.5 | 23 |
| <hr/> | |
| About the Content Migration Guide | 23 |
| Content Migration Overview | 24 |
| Software Migration Tool Overview | 25 |
| Content Migration Process | 25 |
| Full Migration | 26 |
| Complete Migration | 30 |
| Migration of Software Installation Dependencies | 30 |
| Configuration Tracking in Opware SAS 6.5 | 31 |
| Prerequisite for the Software Migration Tool | 31 |
| How to Run the Software Migration Tool | 32 |
| Content Migration Best Practices | 33 |
| Command Syntax for the Script | 33 |
| Content Migration Example | 35 |
| Special Cases for Software Tree Migration | 40 |

| | |
|---|-----------|
| Service Level Migration | 41 |
| Template Migration | 41 |
| Inheritance and Negative Overrides | 41 |
| Migration of Software Tree Permissions | 43 |
| Synchronizing Permissions After a Content Migration | 44 |
| RPM Metadata Migration | 45 |
| Running the RPM Metadata Migration Script | 45 |
| ISM Tool Runtime Migration | 47 |
| Running the ISM Tool Runtime Migration Script | 48 |
| Migration of RPMs Created by ISMTool | 49 |
| Migrating RPMs Created by ISMTool | 49 |
| Compliance Data Migration Tool | 50 |
| Running the Compliance Data Migration Tool | 50 |
| Command Syntax For the Script | 51 |
| Chapter 2: Server Tracking in the SAS Web Client | 53 |
| Server Tracking | 53 |
| Ways to Locate, List, and Display Servers | 53 |
| Tracked Server Properties | 54 |
| Supported Operating Systems for Managed Servers | 55 |
| Server Lists | 57 |
| Types of Server Lists | 57 |
| Server Pool | 58 |
| Manage Servers List | 59 |
| Filters on the Manage Servers List | 60 |
| My Servers | 62 |

| | |
|---|-----------|
| Overview of My Servers | 62 |
| Adding Servers to My Servers | 62 |
| Removing Servers from My Servers | 63 |
| Server Search | 64 |
| Searching for a Server By Using the Search Box..... | 64 |
| Ways to Use Search..... | 65 |
| Using the Search Feature | 65 |
| Rules for Server Search and for Creating Dynamic Groups | 70 |
| Line Break Workaround for Server Search | 79 |
| Conditions for Searching with Multiple Rules..... | 80 |
| Server Searching by IP Address | 81 |
| Example: Server Search | 82 |
| Searching for a group of servers..... | 83 |
| Server Identification..... | 84 |
| Overview of Server Identification | 84 |
| Ways that Servers are Identified by Opsware SAS | 85 |
| Customer Accounts in Opsware SAS..... | 86 |
| Associating Servers with Customers | 88 |
| Server Histories and Reports | 90 |
| Overview of Server History..... | 90 |
| Viewing Server History | 92 |
| Time Stamp for Server Operations..... | 92 |
| DCI Reporting | 93 |
| Hardware Information for Managed Servers | 94 |

| | |
|--|------------|
| Viewing Managed Server Hardware Information | 94 |
| Chapter 3: Server Management in the SAS Web Client | 97 |
| Overview of Server Management | 97 |
| Communication Between Managed Servers and Opsware SAS | 99 |
| Network Address Translation (NAT) for Managed Servers..... | 99 |
| Key Terms of Opsware SAS Managed Server Communication..... | 99 |
| Locating the Management IP Address of a Managed Server | 100 |
| Code Deployment and Static NAT | 102 |
| Setting the Primary IP Address of a Server..... | 102 |
| NAT Table Mapping and Managed Servers | 103 |
| Server Groups | 104 |
| Overview of Server Groups | 104 |
| Types of Server Groups | 107 |
| Public Group Modeling | 108 |
| Using Server Groups Tabs | 111 |
| Ways to Create Server Groups | 113 |
| Adding a Server to Static Groups, Both Public and Private | 127 |
| Adding a Server to Dynamic Groups, Both Public and Private | 127 |
| Removing Servers from Static Groups, Both Public and Private..... | 128 |
| Removing Servers from Dynamic Groups, Both Public and Private ... | 129 |
| Moving Servers from One Static Group to Another | 130 |
| Duplicating Server Groups | 130 |
| Rules for Deleting Server Groups..... | 132 |
| Server Life Cycle | 133 |

| | |
|--|------------|
| OS Provisioning and the Server Life Cycle | 133 |
| Server Properties | 139 |
| Server Management Tasks Related to the Server Life Cycle | 141 |
| Changing the Use and Stage Values for Servers..... | 143 |
| Editing the Properties of a Server..... | 143 |
| Tasks Associated with Deactivating a Server | 145 |
| Deactivating a Server..... | 145 |
| Deleting a Server from Opsware SAS | 146 |
| Overview of Cloning a Server..... | 146 |
| Server Management Jobs | 148 |
| My Jobs | 148 |
| My Jobs Display Information | 149 |
| My Jobs in the SAS Web Client..... | 150 |
| My Jobs in SAS Client | 150 |
| Viewing Job Details in the SAS Web Client..... | 151 |
| Viewing My Job Details from inside the SAS Client..... | 153 |
| Server Management Scheduling and Notification | 154 |
| Time-Outs for Server Management Jobs..... | 157 |
| Custom Fields for Servers | 158 |
| Overview of Custom Fields for Servers | 158 |
| Changing the Value for a Custom Field | 159 |
| Custom Attributes for Servers | 160 |

| | |
|---|------------|
| Overview of Custom Attributes for Servers..... | 161 |
| Managing Custom Attributes | 161 |
| Adding Server Custom Attributes | 162 |
| Editing Server Custom Attributes | 163 |
| Deleting Server Custom Attributes..... | 163 |
| Service Levels..... | 164 |
| Overview of Service Levels..... | 164 |
| Adding a Service Level to the SAS Web Client | 165 |
| Editing a Service Level | 166 |
| Ways to View the Service Level for Servers..... | 166 |
| Assigning a Server to a Service Level | 168 |
| Removing a Server from a Service Level..... | 169 |
| Network Configuration..... | 170 |
| Overview of the Server Network Configuration | 170 |
| Configuring Networking for an Opsware Managed Server..... | 171 |
| Chapter 4: Software Provisioning Setup | 175 |
| Software Tree | 176 |
| Overview of Software Tree | 176 |
| Example: Software Tree | 179 |
| Guidelines for Setting Up the Software Tree | 180 |
| Software Tree Hierarchy and Navigation..... | 181 |
| Software Reconciled onto Servers | 182 |
| The Time to Reconcile..... | 183 |
| Node Management in the Software Tree | 183 |

| | |
|--|------------|
| Software Provisioning Setup Tasks..... | 183 |
| Node Adding in the Software Tree | 184 |
| Node Editing in the Software Tree | 186 |
| Node Deleting in the Software Tree | 191 |
| Node Copying in the Software Tree | 193 |
| Node Moving in the Software Tree | 195 |
| Node Configuration Tracking Policy Management..... | 197 |
| Node History | 197 |
| Viewing Node History | 198 |
| Software Attached to Nodes | 199 |
| Overview of Software Attached to Nodes..... | 200 |
| Modeling Software in Nodes | 200 |
| Software Configuration Settings | 203 |
| Viewing Software Attached to a Node | 203 |
| Adding a Package to a Node | 204 |
| Removing a Package from a Node | 207 |
| Changing the Installation Order of Software..... | 209 |
| Software Inheritance from Other Nodes | 210 |
| Changing the Override Values of Inherited Software | 211 |
| Dependencies Between Nodes for Software Installation | 212 |
| Viewing Software Installation Dependencies | 213 |
| Adding Software Installation Dependencies | 214 |
| Removing Software Installation Dependencies | 215 |
| Custom Attributes Set for the Environment..... | 215 |

| | |
|--|------------|
| Overview of Custom Attributes and Folders | 215 |
| Managing Custom Attributes | 216 |
| Adding Custom Attributes for a Node | 216 |
| Editing Custom Attributes for a Node | 217 |
| Deleting Custom Attributes for a Node | 218 |
| Working with Templates | 219 |
| Overview of Working with Templates | 219 |
| Templates, Folders, and Inheritance | 220 |
| Template Inheritance | 221 |
| Attachments: Local, Inherited, and Blocked | 221 |
| Blocking and Reattaching Inherited Attachments | 224 |
| Folders and Templates | 226 |

| | |
|--|-----|
| Overview of Folders and Templates | 226 |
| Templates and Folders Application..... | 227 |
| Creating Templates | 227 |
| Creating Folders | 230 |
| Operating System Management for Templates or Folders | 232 |
| Adding an Operating System to a Template..... | 232 |
| Removing an Operating System from a Template..... | 234 |
| Adding an Operating System to a Folder | 236 |
| Removing an Operating System from a Folder | 238 |
| Adding Applications to Templates or Folders..... | 240 |
| Editing or Removing Applications for Templates or Folders | 242 |
| Adding Patches to Templates or Folders..... | 243 |
| Editing or Removing Patches for Templates or Folders | 244 |
| Adding Service Levels in Templates or Folders | 245 |
| Editing or Removing Service Levels in Templates or Folders | 247 |
| Copying Templates and Folders | 247 |
| Deleting Templates or Folders..... | 248 |
| Folder and Template Inheritance..... | 249 |
| Blocking vs. Removing Attachments..... | 249 |
| Blocking an Attachment from Being Inherited | 249 |

Chapter 5: Package Management **251**

| | |
|---|------------|
| Overview of Package Management | 251 |
| Container Packages and Installable Packages..... | 252 |
| Supported Operating Systems and Package Types | 252 |
| AIX Packages | 255 |

| | |
|--|------------|
| Overview of AIX Packages | 255 |
| LPP Metadata | 255 |
| HP-UX Packages..... | 256 |
| Overview of HP-UX Package | 257 |
| Depot Metadata | 259 |
| Preparing for HP-UX Package Management..... | 259 |
| Example: Commands - Converting a Depot | 260 |
| Example: File - Script to Split a Depot by Product | 260 |
| Example: File - Script to Split a Depot by Bundle | 260 |
| Linux Packages..... | 260 |
| RPM Metadata..... | 261 |
| Solaris Packages | 261 |
| Windows Packages | 263 |
| Microsoft Installer Packages | 263 |
| Microsoft Hotfixes, Security Patches, and Service Packs..... | 265 |
| ZIP Packages | 265 |
| ZIP Package Support..... | 266 |
| ZIP Packaging | 266 |
| ZIP Packages Creation | 266 |
| ZIP Packages Upload..... | 267 |
| Package Installation and Remove Scripts Definition..... | 267 |
| Editing Properties for ZIP Packages..... | 267 |
| Info-Zip Compatible ZIP Packages | 268 |
| Windows Performance for Uploading Packages | 269 |
| Package Management Tasks | 270 |

| | |
|---|-----|
| Displaying Packages | 270 |
| Filters for Displaying Packages | 272 |
| Searching for Packages | 272 |
| Viewing Packages Assigned to Nodes | 274 |
| Overview of Uploading a Package | 275 |
| Uploading a Package..... | 276 |
| Character Encoding for Package Metadata and Scripts | 278 |
| Overview of Overwriting a Package | 280 |
| Overwriting a Package..... | 281 |
| Editing Package Properties | 281 |
| Overview of Deleting a Package | 284 |
| Deleting a Package | 285 |
| Overview of Deprecating a Package | 286 |
| Deprecating a Package..... | 287 |
| Downloading a Package | 288 |

Chapter 6: Software Provisioning **291**

| | |
|---|------------|
| Installing and Uninstalling Software | 291 |
| Overview of Installing and Uninstalling Software | 291 |
| Software Package Management in the Model Repository | 292 |
| Ways to Install Software | 292 |
| Opsware SAS Categories for Software | 293 |
| Software Provisioning Functionality | 293 |
| Platform-Specific Reconcile | 294 |
| Software Installation and Uninstallation Options | 294 |

| | |
|--|------------|
| Script Error Conditions | 294 |
| Inheritance and the Install and Uninstall Software Wizards | 295 |
| Installing Software with the Install Software Wizard | 295 |
| Package Installation on Servers with Low Disk Space | 300 |
| Uninstalling Software with the Uninstall Software Wizard | 302 |
| Template Installation | 305 |
| Overview of Installing Templates | 306 |
| Installing Templates with the Install Templates Wizard | 306 |
| Chapter 7: Configuration Tracking Using the Software Tree | |
| 309 | |
| Overview of Configuration Tracking | 309 |
| File Types Supported | 310 |
| Change Detection | 310 |
| Actions Performed | 311 |
| Backups Performed | 311 |
| Email Automated Configuration Tracking and Logging Actions | 312 |
| Email Notification List | 313 |
| Configuration Tracking Policies | 314 |
| Ways to Create Tracking Policies | 316 |
| Configuration Tracking Policy Targets and Wildcards | 316 |
| Special Considerations for Directory and Wildcard Targets | 317 |
| Configuration Tracking Policy Limits | 318 |
| Node-Based Tracking Policies | 320 |

| | |
|--|------------|
| Creating Node-Based Policy Entries..... | 321 |
| Tracking Policy Deployment..... | 324 |
| Reconciling a Node's Configuration Tracking Policy..... | 324 |
| Viewing a Node's Tracking Policy..... | 325 |
| Editing a Node's Configuration Tracking Policy..... | 326 |
| Editing a Node's Configuration Tracking Policy Entry..... | 327 |
| Disabling a Node's Configuration Tracking Policy Entries..... | 328 |
| Deleting an Entry in a Node's Configuration Tracking Policy..... | 328 |
| Re-enabling a Tracking Policy Entry..... | 329 |
| Chapter 8: Operating System Provisioning | 331 |
| Supported Operating Systems for OS Provisioning..... | 331 |
| OS Installation with the SAS Web Client..... | 333 |
| Servers and the Server Pool..... | 334 |
| Ways to Install Operating Systems on Servers..... | 335 |
| Installing an OS by Using a Template..... | 336 |
| Installing an OS by Using a Custom Installation..... | 341 |
| Recovering when an OS Installation Fails..... | 343 |
| Network Configuration for Servers after OS Provisioning..... | 346 |
| Requirements for Reprovisioning Solaris and Linux Servers..... | 346 |
| Reprovisioning a Solaris or Linux Server..... | 347 |
| Appendix A: Reconcile | 349 |
| Overview of Reconcile..... | 349 |
| Ways to Perform Reconcile..... | 350 |
| Reconcile Operations and Guidelines..... | 351 |

| | |
|---|------------|
| Reconcile Operations | 351 |
| Reconcile and Package Metadata..... | 352 |
| Installation and Uninstallation Order..... | 353 |
| Software Installation Order for Adopted Software..... | 353 |
| Patches and Reconcile | 354 |
| Preview Reconcile | 354 |
| Types of Reconcile..... | 355 |
| Reconcile on Supported Operating Systems..... | 355 |
| AIX Reconcile | 357 |
| HP-UX Reconcile | 357 |
| Solaris Reconcile | 358 |
| Linux Reconcile..... | 358 |
| Specifying Scripts During Reconcile..... | 358 |
| Reconcile Output | 358 |
| Assigning to and Removing Servers from Nodes | 359 |
| Assigning Servers to Nodes | 359 |
| Removing Servers from Nodes | 359 |
| Reconcile Software Wizard | 360 |
| Directly Reconciling Servers | 360 |
| Index | 363 |

Preface

Welcome to the Opsware Server Automation System (SAS) – an enterprise-class software solution that enables customers to get all the benefits of the Opsware data center automation platform and support services. Opsware SAS provides a core foundation for automating formerly manual tasks associated with the deployment, support, and growth of server and server application infrastructure.

This guide describes how to use Opsware SAS, 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, reconciling servers, and configuration tracking. 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: Migrating SAS to Version 6.5: Provides an overview of the Content Migration process, explains how to run the Software Migration Tool, and describes the impact that the Software Migration Tool has on the SAS Web Client and SAS Client.

Chapter 2: Server Tracking in the SAS Web Client: Provides information about server asset tracking, server lists, server search, server histories and reports.

Chapter 3: Server Management in the SAS Web Client: Provides information about all aspects of server management including server groups, server life cycle, server locking, and service levels.

Chapter 4: 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 5: Software 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 6: Software Provisioning: Provides information about installing and uninstalling software and using templates for installation.

Chapter 7: Configuration Tracking Using the Software Tree: Provides information about configuration tracking policies, the supported types of configuration files and databases that can be backed up, how changes are detected, node-based tracking policies, and reconciling node-based tracking policies.

Chapter 8: 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.

Appendix A: Reconcile: Reviews the server reconciliation process and how it works, with 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 the Reconcile Software Wizard.

Conventions in this Guide





This guide uses the following typographical and formatting conventions.

| NOTATION | DESCRIPTION |
|-------------|---|
| Bold | Identifies field menu names, menu items, button names, and inline terms that begin with a bullet. |

| NOTATION | DESCRIPTION |
|----------------|--|
| Courier | Identifies text that is entered or displayed at the command-line prompt, such as Unix commands, Opsware SAS commands, file names, paths, directories, environment variable names, contents of text files that are viewed or edited with a text editor, source code in a programming language, and SQL (database) commands. |
| <i>Italics</i> | Identifies document titles, DVD titles, web site addresses. Used to introduce new terms when they are first defined in a document and for emphasis. |

Icons in this Guide

This guide uses the following icons.

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

Guides in the Documentation Set and Associated Users

- The *Opsware® SAS User's Guide: Server Automation* is intended for system administrators responsible for all aspects of managing servers in an operational environment. It describes how to use Opsware SAS, introducing the system and the user interface. It provides information about managing servers, remediating servers, script execution, configuration tracking, deploying and rolling back code, and agent deployment. It also explains how to use the Opsware Global Shell and open a Remote Terminal on managed servers.
- *Opsware® SAS User's Guide: Application Automation* is intended for system administrators responsible for performing the day-to-day functions of managing servers. It reviews auditing and compliance, software packaging, visual application management, application configuration, and software and operating system installation on managed servers.
- The *Opsware® SAS Administration Guide* is intended for administrators responsible for monitoring and diagnosing the health of the Opsware SAS core components. It also documents how to set up Opsware user groups and permissions.
- The *Opsware® SAS Deployment and Installation Guide* is intended for advanced system administrators responsible for planning all facets of an Opsware SAS installation. It documents all the main features of Opsware SAS, scopes out the planning tasks necessary to successfully install Opsware SAS, explains how to run the Opsware Installer, and details how to configure each of the components. It also includes information on system sizing and checklists for installation.
- The *Opsware® SAS Policy Setter's Guide* is intended for system administrators responsible for setting up OS provisioning, configuration tracking, code deployment, and software management.
- The *Opsware® SAS Content Utilities Guide* is intended for advanced system administrators responsible for importing content such as software packages into Opsware SAS. It documents the following command-line utilities: OCLI 1.0, IDK, and DET (CBT).
- The *Opsware® Platform Developer's Guide* is intended for software developers responsible for customizing, extending, and integrating Opsware SAS. It documents how to create Web Services, Java RMI, Python, and CLI clients that invoke methods on the Opsware API.

Opsware, Inc. Contact Information

For more information, see the Opsware, Inc. main web site and phone number:

- <http://www.opsware.com/index.htm>
- +1 (408) 744-7300

For links to the latest product documentation and software downloads, see the Opsware Customer Support site:

- <https://download.opsware.com/opsw/main.htm>

For troubleshooting information, see the Opsware Knowledge Base:

- <https://download.opsware.com/kb/kbindex.jspa>

To contact Opsware Customer Support, see the following email address and phone number:

- support@opsware.com
- +1 (877) 677-9273

Chapter 1: Migrating SAS Content to 6.5

IN THIS CHAPTER

This section discusses the following topics:

- About the Content Migration Guide
- Content Migration Overview
- Software Migration Tool Overview
- Prerequisite for the Software Migration Tool
- How to Run the Software Migration Tool
- Command Syntax for the Script
- Content Migration Example
- Special Cases for Software Tree Migration
- Migration of Software Tree Permissions
- RPM Metadata Migration
- ISM Tool Runtime Migration
- Migration of RPMs Created by ISMTool
- Compliance Data Migration Tool

About the Content Migration Guide

After upgrading an Opsware SAS core to 6.5 from a release earlier than 6.0, your SAS Web Client will contain legacy features that are deprecated as of Opsware SAS 6.0:

- Software Tree for software provisioning
- The Reconcile operation for servers and software provisioning
- Templates for software provisioning and OS provisioning
- The Package Management feature in the SAS Web Client

- Configuration Tracking using the Software Tree
- OS provisioning by using the SAS Web Client

These features will continue to be available in your version of the SAS Web Client until you complete the content migration process as described in this chapter of the *Opsware® SAS Content Migration Guide*.

Use the content in this guide for these legacy features until you migrate fully to using the new features in the 6.0 version of the SAS Client.

Content Migration Overview

After upgrading your core to Opsware SAS 6.5, perform the following content migration tasks:

- After upgrading an Opsware SAS core from 6.5, run the software migration tool to migrate the software tree nodes and package information to the Library in the SAS Client. See “Software Migration Tool Overview” on page 25 in this chapter for information about running the Software Migration Script.
- After upgrading an Opsware SAS core to 6.5, run the RPM metadata migration script to import the metadata of any pre-existing RPMs to the Model Repository. See “RPM Metadata Migration” on page 45 in this chapter for information about running the RPM metadata script.
- After upgrading an Opsware SAS core to 6.5, run the ISM Tool runtime migration script to migrate the runtime packages of 3.2 and earlier ISMs to 3.3. See “ISM Tool Runtime Migration” on page 47 in this chapter for more information.
- After upgrading an Opsware SAS core to 6.5, run the `ismpolicyrpm-migrate.sh` script to update the RPMs in the Opsware core that were created by ISMTool shipped with Opsware SAS 6.1.1 or earlier. See “Migration of RPMs Created by ISMTool” on page 49 in this chapter for more information.
- After upgrading an Opsware SAS core to 6.5, you are required to run the Compliance Data Migration Tool to migrate compliance data to the new database schema in the Opsware core. See “Compliance Data Migration Tool” on page 50 in this chapter for more information.

Software Migration Tool Overview

In Opware SAS 6.0, enhanced software management functionality is available as part of the SAS Client. The Software Management feature in the SAS Client replaces the node-based Software Tree accessible from the SAS Web Client.

See the *Opware® SAS User's Guide: Application Automation* for information about the Software Management feature in the SAS Client.

After upgrading an Opware SAS core to 6.5, you are required to run the Software Migration Tool to migrate Software Tree nodes and package information to the Library in the SAS Client. You are not required to run the Software Migration Tool if you upgrading an Opware SAS core from 6.0, 6.0.1, 6.0.2, 6.1, or 6.1.x to 6.5

Running the Software Migration Tool on your upgraded SAS core allows you to use the new software management features with your existing data. After the migration, your Software Tree data and package information appear in the SAS Client under the Migrated folder in the Library.

See the *Opware® SAS Policy Setter's Guide* for information about the Library in the SAS Client.

See *"Full Migration"* on page 26 for information about the Migrated folder in the SAS Client.

Because the new Software Management functionality and the old functionality are both available after upgrading to Opware SAS 6.5 from a release earlier than 6.0, customers can choose the best time for them to migrate completely their Software Tree data and package information. If you are upgrading the Opware core from 6.0, 6.0.1, 6.0.2, 6.1, or 6.1.x to 6.5, then you are not required to run the Software Migration Tool.

Content Migration Process

The content migration process consists of two primary phases:

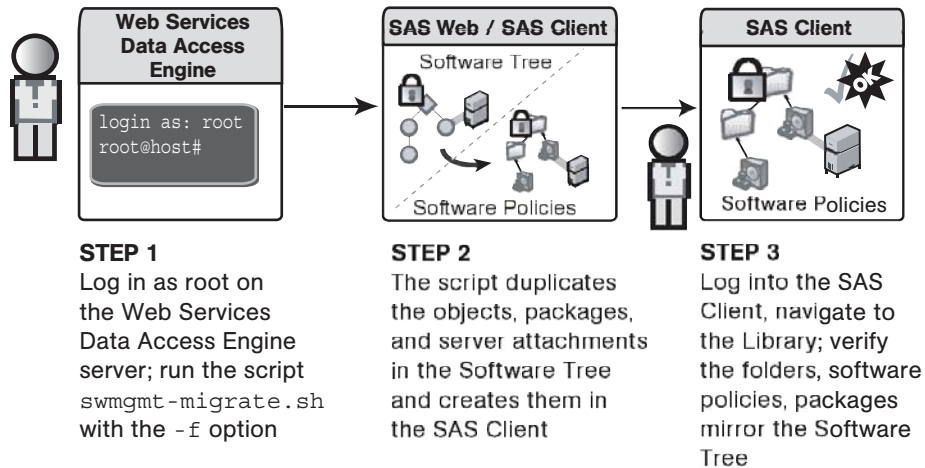
Part 1: Run the Software Migration Tool to duplicate the data and packages in the Software Tree in the SAS Client Library and software policies.

Part 2: Run the Software Migration Tool again on the Opware SAS core to remove the Software Tree data from the SAS Web Client and unlock the packages in the SAS Client Library and software policies.

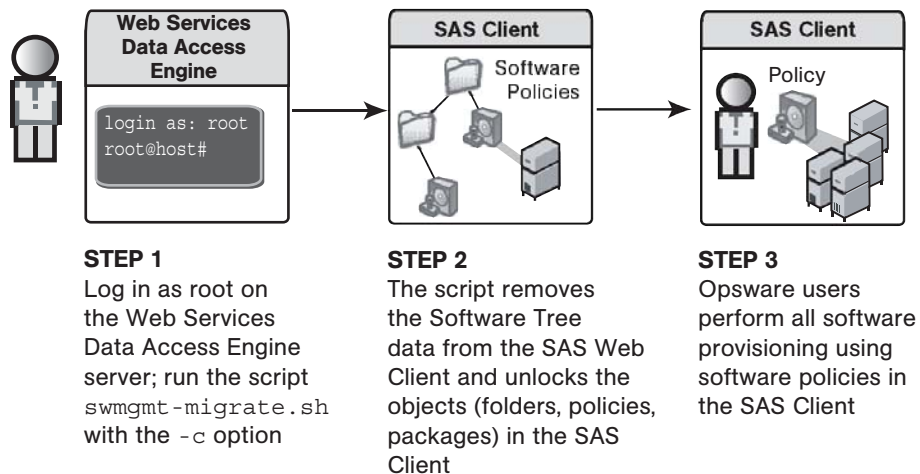
Figure 1-1: Content Migration Process

CONTENT MIGRATION PROCESS

Part 1: Run Software Migration Tool with “full migration”



Part 2: Run Software Migration Tool to **Complete** the Migration



Full Migration

Full migration occurs when you run the Software Migration Tool with the `-f` option.

Impact on the Software Tree

The Software Migration Tool copies all data for the following objects in the Software Tree:

- All nodes in the Software Tree (Software ► Applications in the navigation panel)
- All Service Levels except for the Service Levels under Opware (Service Levels ► Opware)
- All templates except for the templates under Opware (Templates ► Opware)
- All packages in the Package Management feature

After you run the Software Migration Tool, the container and leaf nodes in the Software Tree are locked. Users cannot delete nodes or modify them (which includes changing their properties, package lists, custom attributes, or installation order settings). Users with the manage lock permission for the nodes cannot change the lock state.

With a locked node, users can still:

- Assign, reassign, install, uninstall, and reconcile the nodes against servers and groups
- Utilize configuration tracking for nodes in the Software Tree
- Add child nodes to the Software Tree



Opware Inc. recommends that users do not add child nodes to the Software Tree after the migration script is run and that complete migration should be performed on the SAS core as soon as possible. See “Content Migration Best Practices” on page 33 in this chapter for more information about the preferred way to perform the migration.

Packages included in the nodes are locked and cannot be deleted, replaced, or have their properties modified. With a locked package, users can still:

- Assign or reassign the package to unlocked nodes.
- Install, uninstall, and reconcile nodes with locked packages against servers and groups.

Opware-specific Nodes

When you perform a full migration, the following Opware-specific nodes are not migrated because their software policy equivalents already exist in the SAS Client:

- Templates under /Opware
- Templates under /Opware Tools

- Service levels under /Opware
- Applications under /System Utilities/Opware Tools

Impact on the SAS Client

The Software Migration Tool creates the equivalent objects in the SAS Client Library under the Migrated folder.

In the Migrated folder, the folder hierarchy mirrors the object hierarchy in the Software Tree. Each Software Tree node has a corresponding folder in the SAS Client Library; the path of the folder mirrors the full path of the Software Tree node.

The folders have the same name as the display names of the nodes and the same customer associations as the nodes.

In each folder, the Software Migration Tool creates a software policy. Each software policy has a .p suffix and has properties that mirror the properties of the node it was duplicated from; for example, the OS assignment of the node becomes an OS attribute of the software policy.

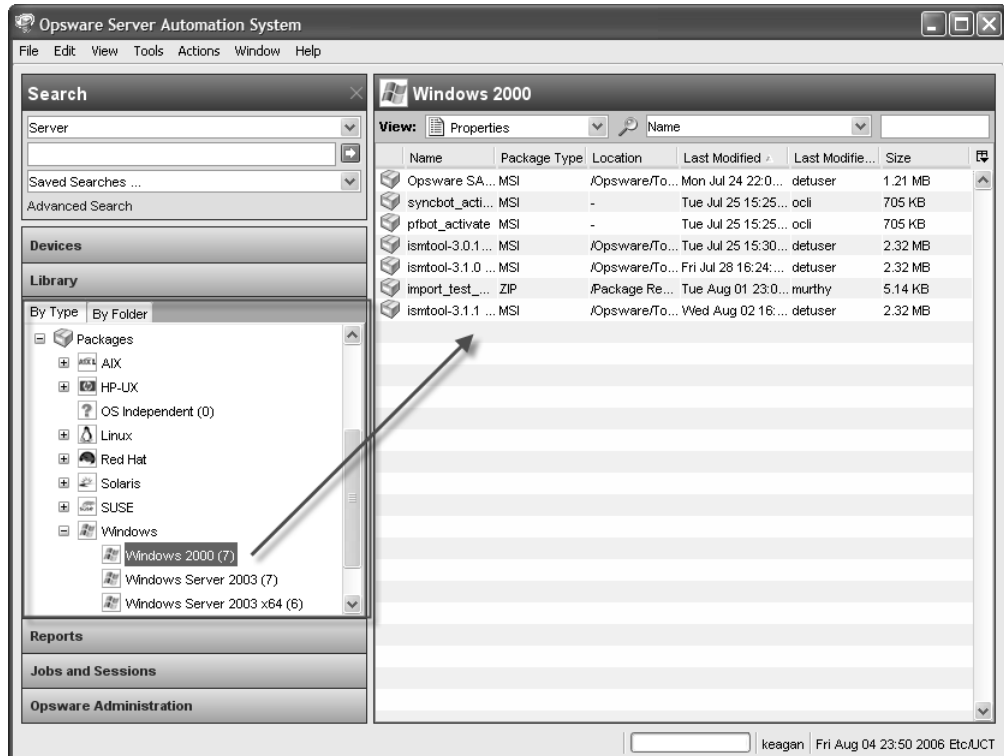
Once the folders and software policies are created in the SAS Client Library, they are marked as in the "in migration" state and they cannot be modified while in this state.

All history entries for a migrated node are copied to the node's software policy equivalent. Additionally, the Software Migration Tool adds an entry to the history for each node and software policy that a migration event occurred. The history description references the corresponding objects.

After the migration, servers and groups that are attached to Software Tree nodes are also attached to the equivalent software policy objects. For example, if a group of servers has three node attachments, three software policies will be created and attached to the same group of servers.

The Software Migration Tool also duplicates all package information in the Package Management feature in the SAS Web Client (Software ► Packages in the navigation panel). The package information is displayed in the SAS Client Library as shown in Figure 1-2.

Figure 1-2: Packages in the SAS Client Library



Reconcile Behavior After Full Migration

Users can still reconcile servers with nodes that are in the "in migration" state. Likewise, users can remediate servers with the software policy counterparts that are also in the "in migration" state.

However, the two functions are separate: node reconcile in the SAS Web Client operates only on Software Tree node attachments, and software policy remediation in the SAS Client operates only on software policy attachments.

Complete Migration

The Complete Migration step is the second part of the Content Migration process for the Software Tree data. When you run the Software Migration Tool with the `-c` option, the tool performs the following actions:

- Removes all Software Tree nodes along the path that leads to migrated nodes.
- Removes all server attachments to nodes in the SAS Web Client.
- Unlocks the software policy and folder counterparts in the SAS Client.
- Clears the “in migration” state from the software policy objects so that users (with the necessary SAS Client permissions) can modify them.
- Removes packages that are no longer referenced by any node, (whether they are in migration or not) from the SAS Web Client; but they are still available in the SAS Client.

Migration of Software Installation Dependencies

The Software Migration Tool automatically preserves installation order data set for Software Tree nodes in the following ways:

- Within a node, package installation order is based on the order in which the packages were encountered while traversing the Software Tree node hierarchy to get to a leaf node.
- When a Software Tree node contains a software installation dependency, the Software Migration Tool preserves the data about dependent nodes in the software policy equivalent so that when a server is remediated with the software policy, the software is installed in the correct order. The data about dependent nodes is stored internally and is not modifiable in the SAS Client.

See “Dependencies Between Nodes for Software Installation” on page 212 in Chapter 4 for more information about installation order set for Software Tree nodes.

- The the top-level categories in the Software Tree have a default installation order when servers are reconciled. This default installation order based on software category is preserved in the software policy equivalents. The data about the installation order of software categories is stored internally and is not modifiable in the SAS Client.

See “Dependencies Between Nodes for Software Installation” on page 212 in Chapter 4 for more information.

Configuration Tracking in Opware SAS 6.5

As of Opware SAS 6.0, newly installed cores do not provide functionality to set up configuration tracking for Software Tree nodes. Configuration tracking for servers is still available. See the *Opware® SAS User's Guide: Server Automation* for information.

While your upgraded Opware SAS 6.5 core is still in the “in migration” state, you can continue to use the Configuration Tracking feature with Software Tree nodes.

Prerequisite for the Software Migration Tool

Before you run the Software Migration Tool, you must run the `gather_lcrep_stats` procedure, which computes statistics for the LCREP table in the database. For a database with approximately 50M records, this process might take up to 5 hours.

Perform the following steps to collect the statistics for the LCREP table:

- 1** Log on as `root` to the core server running the Opware Command Center (OCC) component.
- 2** Locate the `gather_lcrep_stats.sql` file:
`/opt/opware/twist/migration/gather_lcrep_stats.sql`
- 3** Copy the `gather_lcrep_stats.sql` file from the OCC server to the Model Repository (truth) server. For example, you might copy the file to the following directory on the Model Repository server:
`/u01/app/oracle/admin/truth/scripts`
- 4** Log on as `root` to the Model Repository server.
- 5** Change the permissions mode to `777` on the `gather_lcrep_stats.sql` file, for example:
`chmod 777 \
/u01/app/oracle/admin/truth/scripts/gather_lcrep_stats.sql`
- 6** Become the oracle user:
`su - oracle`
- 7** Run sqlplus:
`sqlplus "/ as sysdba"`
- 8** Enter the following sqlplus commands to execute `gather_lcrep_stats.sql`:
`grant create session to lcrep;
connect lcrep/password_for_lcrep
@/u01/app/oracle/admin/truth/scripts/gather_lcrep_stats.sql
exec gather_lcrep_stats`

- 9 In sqlplus, enter the following commands to revoke the create session privilege from lcrep:

```
connect / as sysdba
revoke create session from lcrep;
```

How to Run the Software Migration Tool

- 1 Log on as root to the server running the Web Services Data Access Engine (twist) component. This component is installed on the same server as the SAS Web Client component.

(Recommended) Run the script from your source core, which is the first core installed in a multimaster mesh. See the *Opware® SAS Planning and Installation Guide* for information about source and target cores.

You only run a migration operation once per multimaster mesh. The multimaster functionality will propagate the change to the rest of the cores in the mesh.

- 2 Locate the script in the following directory:

```
/opt/opware/twist/migration/swmgt-migrate.sh
```

- 3 **(Recommended)** To preview the migration, enter the following command:

```
swmgt-migrate -f -n
```

The script displays a detailed report indicating what will occur when the migration operation is run. The detailed information is important in disclosing what changes will be made and the scope of the changes. Use this information to schedule an appropriate time to perform the actual migration.

The detailed information includes metrics on the amount time the operation might take. Use the metrics as an estimate of how much time in bytes per second (bps) it might take to complete the operation for the core the script is run against. The metrics factor in the requirements of multimaster propagation. These metrics serve as a guideline and should not be interpreted as the exact time the migration operation will run.

- 4 To perform a full migration of the Software Tree data, enter the following command:

```
swmgt-migrate -f
```

Optionally, you can specify the `-s` option to skip the detailed report. The script still displays a basic report.

- 5 At the confirmation prompt, specify that you are ready to proceed with the full migration.

The migration proceeds and displays result information and success/failure details as it runs.

You can access the script logs in the `/var/log/opsware/swmgmt` directory.

- 6 On successful completion of the full migration operation, log into the SAS Client for a core in the mesh and verify the results of the migration.
- 7 Log on as root to the server running the Web Services Data Access Engine (twist) component and rerun the script to complete the migration:

```
swmgmt-migrate -c -f
```



To optimize the execution speed of the migration operation, use the `-f` option.

See “Complete Migration” on page 30 in this chapter for more information.

Content Migration Best Practices

- Make the upgraded release public only after the Software Migration Tool runs and completes successfully.
- After the full migration, existing Software Tree nodes are locked in the SAS Web Client. However, by default, users can still create new nodes in Software Tree. Opware, Inc. recommends that you modify SAS Web Client permissions so that Opware SAS users (excluding Opware System Administrators) do not have write permissions to the Software Tree, Software Provisioning Wizards, Package Management feature, and Templates after the migration.
- To prevent further use of unused packages, the Software Migration Tool moves packages that are not included in any nodes to the Software Repository folder in the SAS Client. Therefore, if a package does not exist in a folder in the SAS Client after the migration, add it to the folders where you need it.

Command Syntax for the Script

```
swmgmt-migrate.sh [options]
```

Table 1-1 provides a description of each option for the Software Migration Tool.

Options `-c`, `-o`, `-p`, and `-r` cannot be combined.

Table 1-1: Command Syntax for the `swmgmt-migrate.sh` Script

| OPTION | DESCRIPTION |
|-----------------|--|
| <code>-c</code> | Completes a migration. See "Complete Migration" on page 30 in this chapter for more information. |
| <code>-f</code> | Migrates all nodes in the Software Tree (Software ► Applications in the navigation panel), Service Levels, and templates, leaving nothing except the top-level categories under Software ► Applications, /Templates/Opware, and /Service Levels/Opware. See "Full Migration" on page 26 in this chapter for more information. Full migration also syncs permissions, sets installation orders, and migrates packages automatically. |
| <code>-m</code> | When used in conjunction with other options, displays all migrated Software Tree nodes. |
| <code>-n</code> | Displays a report on how the migration will be performed; does not perform the actual migration. |
| <code>-o</code> | Sets global ordering by taking a snapshot of existing installation order for nodes and recording the global order numbers by using custom attributes. |
| <code>-p</code> | Synchronizes permissions by granting equivalent software policy permissions for user groups that have existing software tree-related permissions. See "Synchronizing Permissions After a Content Migration" on page 44 in this chapter for more information about when to use the <code>-p</code> option. |
| <code>-r</code> | Rolls back a migration by undoing the "in migration" Software Tree nodes. "Completed" nodes are removed; therefore, it's not possible to roll them back. |
| <code>-s</code> | Performs the migration in silent mode (no preview). |

Table 1-1: Command Syntax for the `swmgmt-migrate.sh` Script

| OPTION | DESCRIPTION |
|---------------------|---|
| <code>--help</code> | Shows usage and examples for the <code>swmgmt-migrate.sh</code> script. |

Command Syntax Examples

To perform a full migration, enter the following command:

```
swmgmt-migrate -f
```

To roll back the last full migration, enter the following command:

```
swmgmt-migrate -r -f
```

To complete a migration, enter the following command:

```
swmgmt-migrate -c -f
```




To optimize the execution speed of the migration operation, use the `-f` option, as shown in the examples above.

Content Migration Example

The following example illustrates the process for migrating Software Tree nodes to Opware SAS 6.5.

In this example, the Software Tree has two container nodes (node A and node B) and one leaf node (node C) in the Application Servers category of the Software Tree. Node C has packages X and Y and node B has package Z.

Figure 1-3: Software Tree Before Running the Migration Script




In this example, these three nodes have the properties shown in Table 1-2:

Table 1-2: Node Properties Before Migration

| NODE | TYPE | CUSTOM ATTRIBUTES | PACKAGES/PATCHES | CUSTOMER |
|------|-----------|-------------------|------------------|-------------|
| A | Container | Yes | None | Independent |
| B | Container | Yes | Z | Independent |
| C | Leaf | Yes | X, Y | Marketing |

An Opware administrator runs the Software Migration script to migrate node C as shown by Figure 1-4:

Figure 1-4: Full Migration of Node C



The script migrates everything that formulates the path to node C, which means it also migrates container nodes A and B.

In the SAS Client, the Software Migration script creates folders that correspond to container nodes A and B. Additionally, because nodes A and B have custom attributes and packages defined, the script creates software policies that correspond to nodes A and B to capture the custom attribute and package information.

Software policy B includes A as a nested policy so that the custom attribute inheritance is captured.

For leaf node C, the script creates software policy C and folder C so that the Marketing customer constraint can be placed on folder C. The net effect is that software policy C is constrained to the customer Marketing.

Software policy C includes software policy B to capture custom attribute and package Z associations due to inheritance in the Software Tree. Lastly, the packages X and Y are associated to software policy C. The existing packages (X and Y) are associated with both Software Tree nodes and software policies.

At this point, all the nodes, folders, software policies, and packages involved are in an "in migration" state. All objects in the Software Tree and in the SAS Client are "locked" and "in migration" mode.

The following table lists the properties of the software policies created as a result:

Table 1-3: Node Properties After Full Migration

| OBJECT | TYPE | CUSTOM ATTRIBUTES | PACKAGES | CUSTOMER CONSTRAINTS |
|--------|-----------------|-------------------|----------|----------------------|
| A | Folder | N/A | N/A | Independent |
| A | Software Policy | Those of node A's | None | N/A |
| B | Folder | N/A | N/A | Independent |
| B | Software Policy | Those of node B's | W | N/A |
| C | Folder | N/A | N/A | Marketing |
| C | Software Policy | Those of node C's | X, Y | N/A |

If the "complete the migration" operation is done with the migration script, then both the Software Tree notes and software policy representation would be "unlocked" so that modifications are allowed for involved objects in both areas. Then the Software Tree branch would be entirely removed, leaving just the software policy representation.

Table 1-4: Software Policy Properties After Complete Migration

| OBJECT | TYPE | CUSTOM ATTRIBUTES | PACKAGES | CUSTOMER CONSTRAINT |
|--------|-----------------|-------------------|----------|---------------------|
| A | Folder | N/A | N/A | Independent |
| A | Software Policy | Those of node A's | None | N/A |
| B | Folder | N/A | N/A | Independent |
| B | Software Policy | Those of node B's | Z | N/A |
| C | Folder | N/A | N/A | Nike |
| C | Software Policy | Those of node C's | X, Y | N/A |

Figure 1-5: Software Policy Hierarchy for Node C After Complete Migration




Figure 1-5 shows the software policy hierarchy with C migrated. The Software Tree source is completely removed. The software policy equivalents are no longer in migration mode. Packages are reused as they were.

Special Cases for Software Tree Migration

When you migrate the Software Tree data to the SAS Client, you should understand how the following Software Tree objects are migrated because their data equivalents in the SAS Client have unique properties:

- Service Levels
- Templates
- Inheritance and Negative Overrides in the Software Tree

Service Level Migration

Unlike application nodes in the Software Tree, service levels cannot have packages as part of their specification and they can have multiple customer assignments (whereas an application nodes can only have one).

The software policy equivalent of a service level is combination of a folder and a software policy. A folder is needed so that the node's customer assignments can be transformed to multiple customer constraints on the folder.

Template Migration

Templates can include application nodes, OS nodes, patches, and service levels.

When an application node or service level is migrated, all templates that include the nodes still function normally.

If a template includes the installation of an operating system, it is used for OS provisioning. Templates that include OS installation are migrated to the SAS Client to Software Policy Templates or to OS sequences. The migrated OS sequences reside in the /Migrated/Templates folder in the SAS Client Library.

Migrated OS sequences are locked in the SAS Client and cannot be modified until after the migration is completed. See “Complete Migration” on page 30 for information.

If the OS Installation Profile in the SAS Web Client is customer independent, the customer field of the Install OS Task can be empty or set to a specific customer when migrated to the SAS Client.

Inheritance and Negative Overrides

In the Software Tree, inheritance and negative overrides can be set for custom attributes and packages that are assigned to nodes. See “Software Inheritance from Other Nodes” on page 210 in Chapter 4 for information about inheritance and negative overrides for packages.

In the SAS Client, software policies do not support inheritance for custom attributes. Therefore, for custom attributes that use inheritance or negative overrides, the Software Migration Tool simply creates the corresponding custom attributes on the software policies directly.

When a node has negative overrides of packages, the software policy counterpart is a flattened out version of the node hierarchy. The software policy includes all packages that the node would inherit, minus the packages that were negatively overridden.

For example, assume you have two nodes – A (parent node) and B (child node of A) – and node A has packages X and Y. If child node B has negative inheritance on package Y, then a software policy B would be created with straight inclusion of package X but not Y, and it would not nest policy A.

Template Inheritance

Applications, patches, and service levels can be attached to templates. They can be inherited and blocked. See “Template Inheritance” on page 221 in Chapter 4 for more information.

When you run the Software Migration Tool, it handles a blocked node attached to a template in the following way:

- Two objects are created in the SAS Client – a software policy equivalent of the template and a software policy equivalent of the node
- The software policy equivalent of the template does not include the software policy equivalent of the node

For example, if G and H are templates, H is a child of G, and G includes application nodes R and V.

Figure 1-6: Example of Template Inheritance in the SAS Web Client




Figure 1-6 shows the template hierarchy with H migrated. R and V are members of G. H is a child of G. H has an override on V so that it only gets R through inheritance.

Continuing with this example, the software policy equivalent to G, say G' would not have the software policy equivalent of V, say V', as a nested policy, but would have R'. Also, H specifies that V is negatively overridden, so that H only inherits R from G as a result.

Figure 1-7: Example of Template Inheritance After Migration to the SAS Client




Figure 1-7 shows all of G, H, R, and V are migrated. H' only nests R' however. It does not nest G' or V'.

Migration of Software Tree Permissions

As of Opware SAS 6.0, software policies and folders have new permissions for the Software Management feature in the SAS Client.

As part of the Content Migration, Opware SAS user groups that had permission to perform software provisioning in the SAS Web Client are granted the corresponding Software Management permissions in the SAS Client.

Table 1-5: Mapping of SAS Web Client Permissions to SAS Client Permissions

| SAS WEB CLIENT PERMISSION: | SAS CLIENT PERMISSION: |
|---|---|
| <ul style="list-style-type: none"> • Model: Applications • Model: Service Levels • Templates | <ul style="list-style-type: none"> • Manage Application Policies: Read/Write • Manage Folders: Read/Write |

Table 1-5: Mapping of SAS Web Client Permissions to SAS Client Permissions

| SAS WEB CLIENT PERMISSION: | SAS CLIENT PERMISSION: |
|--|---|
| <ul style="list-style-type: none"> • Packages | <ul style="list-style-type: none"> • Manage Folders: Read/Write • Manage Packages: Read/Write • Allow Install Packages • Allow Uninstall Packages |
| <ul style="list-style-type: none"> • Wizard: Install Software • Wizard: Uninstall Software | <ul style="list-style-type: none"> • Allow Attach/Detach Application Policy: Yes • Allow Remediate Application Policy: Yes |
| <ul style="list-style-type: none"> • ISM Controls | <ul style="list-style-type: none"> • Allow Run ISM Controls: Yes |
| <ul style="list-style-type: none"> • Other ► Manage Server Permissions ► Model Public Device Groups | <ul style="list-style-type: none"> • Client Features ► Server Group Permissions ► Model Public Device Groups |

Folder-specific permissions are assigned as appropriate when nodes are migrated to the SAS Client. See the *Opsware® SAS Administration Guide* for information about granting folder permissions in the SAS Client.

Synchronizing Permissions After a Content Migration

After upgrading to Opsware SAS 6.5, you can run the Software Migration Tool with the sync permissions (-p) option when you need to grant a user group that has the SAS Web Client software provisioning permissions the corresponding permissions in the SAS Client.

For example, user group A is added and granted access to all categories in the Software Tree (Software ► Applications in the SAS Web Client navigation panel) but the user group does not have permission to access folder and software policy features in the SAS Client.

Later, you grant user group A access permission to certain folders. However, the members of user group A still cannot operate on these folder because their group does not have the feature permission for the SAS Client. You can manually add these permission for the user group or run the Software Migration Tool with sync permission (-p) option.

RPM Metadata Migration

In Opware SAS 6.1, the RPM Deployment feature is available as part of the SAS Client. Opware SAS allows you to deploy RPM packages on Red Hat Linux and SUSE Linux servers without manually specifying all the dependent packages required for installing the RPM packages. When you deploy a RPM package, Opware SAS determines the dependencies and installation order for the RPM package, and identifies if any conflicts exists between the dependencies. After you resolve the conflicts, Opware SAS installs the RPM packages on the managed server. See *Opware® SAS Policy Setter's Guide* for more information on this feature.

When you upload an RPM package using Opware SAS, the metadata for the RPM package is extracted and stored in the Model Repository. After upgrading an Opware SAS core to 6.5, you must run the RPM metadata migration script to import any pre existing RPM metadata to the Model Repository. You are not required to run this script if you are upgrading from Opware SAS 6.1, 6.1.1, or 6.1.2 to 6.5. Running the RPM Metadata Migration script on your upgraded SAS core thus allows you to use the new RPM Deployment feature with your existing data.

Running the RPM Metadata Migration Script

- 1** Log on as root to the server running the Software Repository (word) component.

Run the script on each core in a multimaster mesh. You should run the script on only one core at a time. After you complete running the script in a core, wait for several minutes before running the script in the next core.

- 2** Locate the script in the following directory:

```
/opt/opware/mm_wordbot/util/import_rpm_metadata
```

- 3** Run the script as follows:

```
/opt/opware/mm_wordbot/util/import_rpm_metadata -a
```

See Table 1-6 for the other options for the script.

- 4** You can verify the output of the script to confirm if the script ran successfully. The output of the script informs you of the number of RPMs imported and the number of RPMs that still needed to be imported. After running the script on each core in a multimaster mesh this final number should be 0.

Figure 1-8: Output of RPM Metadata Migration Script

```
# /opt/opware/mm_wordbot/util/import_rpm_metadata -a
```

```
Downloading journal
Retrieving RPM IDs
(8/8): Step 8                      100%
[#####]
RPMs needing metadata import: 0
Retrieving RPM details
Looking for RPMs on local filesystem
RPMs found on local filesystem: 0
Verifying results
(8/8): Step 8                      100%
[#####]
RPMs needing metadata import: 0
Deleteing remote journal
Deleteing local journal
Completed
```

Command Syntax for the Script

Table 1-6: Command Syntax for the RPM Metadata Migration Script

| OPTION | DESCRIPTION |
|--------------|--|
| -a | Re-imports the metadata of the existing RPM packages to the Model Repository automatically. |
| -f | Forces the re-import of metadata of the existing RPM packages. This option is always used in conjunction with -a. |
| -c FILE | Specifies the Software Repository configuration file. <when should a user use this option?> |
| {ID..... -} | Allows you to specify the Opware IDS of the RPM packages. Only the metadata of the RPM packages whose Opware IDS are specified will be imported into the Model Repository. |
| -s | Performs the migration in silent mode (no preview). This option is always used in conjunction with -a |
| -h | Shows usage and examples for the script. |

ISM Tool Runtime Migration

ISMs created with version 3.x of the IDK are compatible only with version 6.x of Opware SAS. You cannot upload ISMs from IDK 3.x or earlier into Opware SAS 6.x . The runtime packages created using ISM tool 3.2 and earlier will fail to run in an Opware SAS 6.x core and in the managed servers. To migrate ISMs from IDK 3.2 or earlier to 3.3, you must run the `ismrt-migrate.sh` script.

If you are upgrading an Opware SAS core to 6.x and you have ISMS created using IDK 3.2 and earlier, you must run the ISM Tool runtime migration script to migrate the runtime packages of all 2.x - 3.2 ISMs to 3.3. You must run this script after you upgrade to an

Opware core to 6.x. Opware recommends that you run this script before you run the software migration tool. See “How to Run the Software Migration Tool” on page 32 in this chapter for information about the software migration tool.

If you are upgrading an Opware SAS core to 6.x and you have ISMS created using IDK 1.x, then you must rebuild the ISMs and reupload them using IDK 3.3. See *Opware® SAS Content Utilities Guide* for information about uploading ISMs.

Running the ISM Tool Runtime Migration Script

- 1** Log on as root to the server running the Opware Command Center (occ) component.

You must run this script only once per multimaster mesh.

- 2** Locate the script in the following directory:

```
/opt/opware/twist/migration/ismrt-migrate.sh
```

- 3** To replace runtime packages of all the 3.2 and earlier ISMs in the applications nodes and software policies to 3.3, run the following command:

```
ismrt-migrate.sh
```

- 4** To replace runtime packages of all the 3.2 and earlier ISMs in the applications nodes only to 3.3, run the following command:

```
ismrt-migrate.sh -a
```

- 5** To replace runtime packages of all the 3.2 and earlier ISMs in the software policies only to 3.3, run the following command:

```
ismrt-migrate.sh -p
```

- 6** At the confirmation prompt, specify that you are ready to proceed with the migration. The migration proceeds and you can access the script logs in the `/var/log/opware/` directory.

Perform the following steps on Windows servers to remove the `ismruntime-msi`.

- 1** To remove the `ismruntime-msi` from Windows servers, run the `uninstallWinIsmRuntime2Dse.bat` script located in the following directory as a distributed script:

```
/opt/opware/twist/migration
```


- 2 Using the SAS Web Client, reconcile the software on these Windows servers. During reconcile select the "Removed existing Software that should not be present" checkbox.

Migration of RPMs Created by ISMTool

This section describes how to migrate the RPMs in your core that were created by ISMTool shipped with Opsware SAS 6.1.1 or earlier. Unless you are certain that you have no such RPMs, follow the instructions in this section. If these RPMs are not migrated, they cannot be installed on managed servers by Opsware SAS. For details on the ISMTool bugs that require this RPM migration, see "What's Fixed in Opsware SAS 6.1.2" in the *Opsware SAS 6.1.2 Release Notes*.

Migrating RPMs Created by ISMTool

The following instructions explain how to migrate the RPMs by running the `ismpolicyrpm-migrate.sh` script. In a mutlimaster mesh, you only need to run the script once for the entire mesh. If you ran the `ismpolicyrpm-migrate.sh` script when applying the SAS 6.1.2 patch, you do not need to run it again.

To migrate the RPMs, perform the following steps:

- 1 On the core server running the Opsware Command Center (occ) component, log on as root.
- 2 Change to the directory containing the `ismpolicyrpm-migrate.sh` script.
`cd /opt/opsware/twist/migration`
- 3 View a list of the RPMs that will be updated by the migration script by entering the following command:

```
./ismpolicyrpm-migrate.sh --preview
```

This command lists the Opsware IDs of the RPMs.

- 4 (Optional) To view the properties of an RPM, you can search for it by Opsware ID in the SAS Client:
 1. In the Navigation pane, select Library ► By Type ► Packages ► *platform*
 2. For the View, select Properties.
 3. In the field next to the magnifying glass, enter the Opsware ID.
- 5 To migrate the RPMs, enter the following command:


```
./ismpolicyrpm-migrate.sh --update
```

This command searches for all software policies created by ISMTool, locates RPMs associated with those policies, and then sets the following install and uninstall parameters on each RPM:

Upgrade: No

Install Flags: --nodeps

Uninstall Fags: --nodeps

Compliance Data Migration Tool

In Opware SAS 6.5, the Audit Exceptions feature is available as part of the Opware SAS Client. After upgrading an Opware SAS core to 6.5, you are required to run the Compliance Data Migration Tool to migrate compliance data to the new database schema in the Opware core.

Running the Compliance Data Migration Tool on your upgraded Opware SAS 6.5 core allows you to use the new Audit Exceptions feature with your existing data. After the migration, the audit rules become available in the SAS Client for constructing audit exceptions using the audit editor.

See the *Opware® SAS User's Guide: Application Automation* for more information about Audit exceptions.

Running the Compliance Data Migration Tool

- 1** Log on as root to the server running the Web Services Data Access Engine (twist) component. This component is installed on the same server as the SAS Web Client component.

(Recommended) Run the script from your source core, which is the first core installed in a multimaster mesh. See the *Opware® SAS Planning and Installation Guide* for information about source and target cores.

You only run a migration operation once per multimaster mesh. The multimaster functionality will propagate the change to the rest of the cores in the mesh.

- 2** Locate the script in the following directory:

```
/opt/opware/twist/compliance-migration/compliance-  
migration.sh
```

- 3** Run the script as follows:

```
compliance-migration.sh --host localhost --port 1026
```

You are required to specify the user name and password to run the script. It is recommended that you specify the Opware DCML Exchange Tool (DET) user name and password to run the script. See the *Opware® SAS Content Utilities Guide* for more information about DET.

Running the script will migrate all the compliance rules data to the new schema. See Table 1-7 for the list of script options.

- 4** To display a list of options of to run the script without specifying any parameters:

```
/opt/opware/twist/compliance-migration/compliance-
migration.sh
```

- 5** To verify if the script has run successfully perform any one of the following steps:

- Run the script with the `_d` (debug) option.
- Verify the log files to check for exceptions.

The log files can be located at the following location:

```
/var/log/opware/compliance/
```

- Run the script as follows to display the output of the script:

```
compliance-migration.sh --user -passwd .&output.txt
```

Command Syntax for the Script

Table 1-7: Command Syntax for the Compliance Data Migration Tool

| OPTION | DESCRIPTION |
|-------------------------|--|
| <code>-s</code> | Performs the migration in silent mode (no preview). |
| <code>-d</code> | Shows debug information for the script. |
| <code>--host</code> | Specifies the host name to run the script. |
| <code>--port</code> | Specifies the connection port of the host machine. If the port number is not specified, the default port 1026 is used. |
| <code>--protocol</code> | Specifies the connection protocol (http or https) required to run the script. |

Table 1-7: Command Syntax for the Compliance Data Migration Tool

| OPTION | DESCRIPTION |
|---------------------|---|
| --user | The user name required to run the script. |
| --passwd | The password required to run the script. |
| --id | Specifies the ID of the compliance data to migrate. |
| --AuditPoliciesOnly | Allows you to migrate only Audit Policies. |
| --AuditsOnly | Allows you to migrate only Audits. |
| --AuditResultsOnly | Allows you to migrate only Audit Results. |
| --SnapshotSpecsOnly | Allows you to migrate only Snapshot Specifications. |
| --SnapshotsOnly | Allows you to migrate only Snapshots. |

Chapter 2: Server Tracking in the SAS Web Client

IN THIS CHAPTER

This section discusses the following topics:

- Server Tracking
- Server Lists
- My Servers
- Server Search
- Server Identification
- Server Histories and Reports
- Hardware Information for Managed Servers

Server Tracking

This section discusses the following topics:

- Ways to Locate, List, and Display Servers
- Tracked Server Properties
- Supported Operating Systems for Managed Servers.

Ways to Locate, List, and Display Servers

You can locate, list, and display servers in the SAS Web Client in the following four ways:

- By searching when you know the name, host name, or IP address of the server you want to provision or manage.
- By viewing the Manage 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 to determine which servers specific software should be installed on or find all servers that match certain rules; 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 “Hardware Information for Managed Servers” on page 94 in this chapter for more information.
- You can also browse managed servers and group of servers using the SAS Client.

Tracked Server Properties

Every server that Opware SAS manages has the following properties:

- IP addresses, host name, 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. Select the Install List tab from the Manage 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.

Opware SAS 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 Opware Model Repository.

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

Opware SAS is able to determine what software is installed on a server because the Opware Agent communicates with the Opware 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.

For a complete list of all the hardware and software information that Opware SAS tracks for managed servers, see the *Opware® SAS User's Guide: Server Automation*.

Supported Operating Systems for Managed Servers

This section lists the supported operating systems for Opware Agents and the SAS Client.

The following table lists the supported operating systems for Opware Agents, which run on the servers managed by Opware SAS.

Table 2-1: Opware Agent Supported Operating Systems

| SUPPORTED OPERATING SYSTEMS FOR OPSWARE AGENT | VERSIONS | ARCHITECTURE |
|---|--|---|
| AIX | AIX 4.3 AIX 5.1 AIX 5.2 AIX 5.3 | POWER POWER POWER POWER |
| HP-UX | HP-UX 10.20 HP-UX 11.00 HP-UX 11.11 HP-UX 11.23 (11i v2) | PA-RISC PA-RISC PA-RISC PA-RISC and Itanium |
| Sun Solaris | Solaris 6 Solaris 7 Solaris 8 Solaris 9 Solaris 10 (Update 1, Update 2, Update 3) | Sun SPARC Sun SPARC Sun SPARC Sun SPARC Sun SPARC, 64 bit x86, 32 bit x86 and Niagara |
| Fujitsu Solaris | Solaris 8 Solaris 9 Solaris 10 | Fujitsu SPARC Fujitsu SPARC Fujitsu SPARC |
| Windows | Windows NT 4.0 Windows 2000 Server Family Windows Server 2003 Windows XP Professional | 32 bit x86 32 bit x86 32 bit x86 and 64 bit x86 32 bit x86 |

Table 2-1: Opware Agent Supported Operating Systems (continued)

| SUPPORTED OPERATING SYSTEMS FOR OPSWARE AGENT | VERSIONS | ARCHITECTURE |
|---|------------------------------------|---------------------------------------|
| Red Hat Linux | Red Hat Linux 7.3 | 32 bit x86 |
| | Red Hat Linux 8.0 | 32 bit x86 |
| | Red Hat Enterprise Linux 2.1 AS | 32 bit x86 |
| | Red Hat Enterprise Linux 2.1 ES | 32 bit x86 |
| | Red Hat Enterprise Linux 2.1 WS | 32 bit x86 |
| | Red Hat Enterprise Linux 3 AS | 32 bit x86 and 64 bit x86 and Itanium |
| | Red Hat Enterprise Linux 3 ES | 32 bit x86 and 64 bit x86 and Itanium |
| | Red Hat Enterprise Linux 3 WS | 32 bit x86 and 64 bit x86 and Itanium |
| | Red Hat Enterprise Linux 4 AS | 32 bit x86 and 64 bit x86 |
| | Red Hat Enterprise Linux 4 ES | 32 bit x86 and 64 bit x86 |
| | Red Hat Enterprise Linux 4WS | 32 bit x86 and 64 bit x86 |
| | Red Hat Enterprise Linux Server 5 | 32 bit x86 and 64 bit x86 |
| | Red Hat Enterprise Linux Desktop 5 | 32 bit x86 and 64 bit x86 |
| SUSE Linux | SUSE Linux Enterprise Server 8 | 32 bit x86 |
| | SUSE Linux Standard Server 8 | 32 bit x86 |
| | SUSE Linux Enterprise Server 9 | 32 bit x86 and 64 bit x86 |
| | SUSE Linux Enterprise Server 10 | 32 bit x86 and 64 bit x86 |
| VMware | ESX Server 3 | 32 bit x86 and 64 bit x86 |



On Red Hat Enterprise Linux 5, Opware does not support SELinux (Security Enhanced Linux). By default, SELinux is enabled on Red Hat Enterprise Linux 5. You must disable the SELinux feature on Red Hat Enterprise Linux 5 for the Opware Agent to function correctly.

The following table lists the operating systems supported for the SAS Client.

Table 2-2: SAS Client Supported Operating Systems

| SUPPORTED OPERATING SYSTEMS FOR SAS CLIENT | VERSIONS | ARCHITECTURE |
|--|---------------|-----------------------|
| Windows | Windows Vista | 32 bit x86 and 64 bit |
| | Windows XP | x86 |
| | Windows 2003 | 32 bit x86 |
| | Windows 2000 | 32 bit x86 |
| | | 32 bit x86 |

Server Lists

This section discusses the following topics:

- Types of Server Lists
- Server Pool
- Manage Servers List
- Filters on the Manage Servers List

Types of Server Lists

The SAS Web Client displays lists for two types of servers, as Figure 2-1 shows.

Figure 2-1: Servers Section in the Navigation Panel



Server Pool: Servers in the Server Pool have registered their presence with Opware SAS but do not have the target OS installed. An OS Build Agent is running on each server so that they can communicate with Opware SAS.

See “Operating System Provisioning” on page 331 in Chapter 8 for information about how to use the Server Pool when you install the target OS on a server.

Manage Servers: The Manage Servers list contains servers on which Opware SAS can perform management tasks, because Opware Agents are installed on them. However, Opware SAS might not have provisioned all 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 **Install OS**.

Server Pool

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

- The host name set by booting the server for the first time over the network or by using an Opware 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 Opware Agent on the server communicated with Opware SAS (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 the *Opware® SAS User’s Guide: Server Automation* for 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.)

Manage Servers List

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

Figure 2-2: The Manage Servers List in the SAS Web Client

Manage Servers (Summary View) 

All Statuses All Stages All Uses
 All Facilities All Manufacturers 
 All Models

Resource Edit View Tasks Configuration Tracking

| <input type="checkbox"/> | Name  | Host Name / IP Address | OS Version | Stage | Use | Facility | Customer |
|--|--|--|------------------------------------|------------------|------------------|----------|-----------------|
|  Public Groups | | | | | | | |
| <input type="checkbox"/> |  m022.dev.opsware.com | m022.dev.opsware.com 192.168.197.91 | Red Hat Enterprise Linux AS 2.1 | Not Specified | Not Specified | C03 | Not Assigned |

By default, servers in the Manage 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 host name or IP address.

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

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

- The name of the server
 - By default, the server's host name appears in this field. However, you can edit it so that it is more descriptive or useful.
- The host name 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 SAS Web Client
- 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.)

Filters on the Manage Servers List

The Manage Servers list displays the following filters that you can specify to qualify the servers that the SAS Web Client displays, as Figure 2-3 shows.

Figure 2-3: Filters in the Manage Servers List

Manage Servers (Summary View)



- **Status:** Specifies the ability of Opware SAS to manage servers. Opware SAS automatically detects the status of servers; a server's status is OK or Not Reachable.
- **OS:** Specifies the operating system on the server, which is obtained by the Opware Agent that is running on the server.
- **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 Opware administrator.
- **Uses:** Specifies how an organization is utilizing servers; for example, a server is a staging server. Users set this property for servers. The Opware administrator can customize the values in this list.
- **Facilities:** Specifies the location of servers. From an SAS Web Client, users can manage servers located in any facility. For example, a user could log in to the SAS Web Client running in facility A and manage the server located in facility B.
- **Customers:** Specifies the customer associated with each server. Your Opware administrator defines the options for customer selections by using the Administration pages.
- **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.

- **Life cycles:** Specifies the various Opsware server life cycle values which include Managed, Available, Build Failed, Installing OS, and Deactivated.



You can change the filters displayed in the Manage Servers page. To change the filters you want to be displayed on the Manage Servers page, click on the icon as shown in Figure 2-4 and specify the filters from the **Edit Filters** Menu.

Figure 2-4: Edit Icon



Figure 2-5 shows the filters that are in the Server Pool list.

Figure 2-5: 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 SAS Web Client in any facility.

My Servers

This section contains the following topics:

- Overview of My Servers
- Adding Servers to My Servers
- Removing Servers from My Servers

Overview of My Servers

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 login to the SAS Web Client running in the same facility, servers stay in My Servers for one year or until you explicitly remove them. Each time that you login to the SAS Web Client, 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 Opware administrator to see the servers in the My Servers area of other Opware users.

Adding Servers to My Servers

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

- 1** From the navigation panel, click Servers ► Manage Servers. The Manage 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 “Using the Search Feature” on page 65 in this chapter for more information. See “Server Searching by IP Address” on page 81 in this chapter for more information.

- 2** Select the servers that you want to add to My Servers.
- 3** Choose **Resource ► Add to My Servers** from the menu above the Manage Servers list. The Add To My Servers window appears, which indicates that you added the chosen number of servers to My Servers.

- 4 Click **Close** 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 2-6 shows.

Figure 2-6: Servers in My Servers

My Servers (Summary View) ?

| Resource | | Edit | View | Tasks | Configuration Tracking | | | |
|--------------------------|--|--|------------|---------------|------------------------|----------|----------|--|
| <input type="checkbox"/> | Name ▲ | Host Name / IP Address | OS Version | Stage | Use | Facility | Customer | |
| <input type="checkbox"/> |  m098.dev.opsware.com | m098.dev.opsware.com 192.168.192.67 | SunOS 5.8 | Not Specified | Production | C03 | Opware | |

You can perform the same server management tasks on servers in My Servers as on the servers in the Manage 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 SAS Web Client. 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 **Edit** ► **Remove from My Servers** from the menu above the Server list.

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

Server Search

This section provides information about Server Search and contains the following topics:


- Searching for a Server By Using the Search Box
- Ways to Use Search
- Using the Search Feature
- Rules for Server Search and for Creating Dynamic Groups
- Line Break Workaround for Server Search
- Conditions for Searching with Multiple Rules
- Server Searching by IP Address
- Example: Server Search
- Searching for a group of 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 top navigation panel to open the Search box, as Figure 2-7 shows.

Figure 2-7: Search Text Box on the SAS Web Client Home Page



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

The search text that you enter can include an asterisk (*) wildcard character. However, the search feature automatically prepends and appends an asterisk to the text.

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


```
192.168.68.6  
host02.coredev-val.sample.com  
192.168.*.19  
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 SAS Web Client displays a message that indicates that no servers were found that matched your query.

See “Using the Search Feature” on page 65 in this chapter for information about how to formulate complex, multiple rule search queries.

Ways to Use Search

By using the SAS Web Client, you can perform 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 or group of 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 select 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 ► Search.

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

- While adding or modifying rules for a dynamic group (By clicking **Search** on the Rules tab for a group).

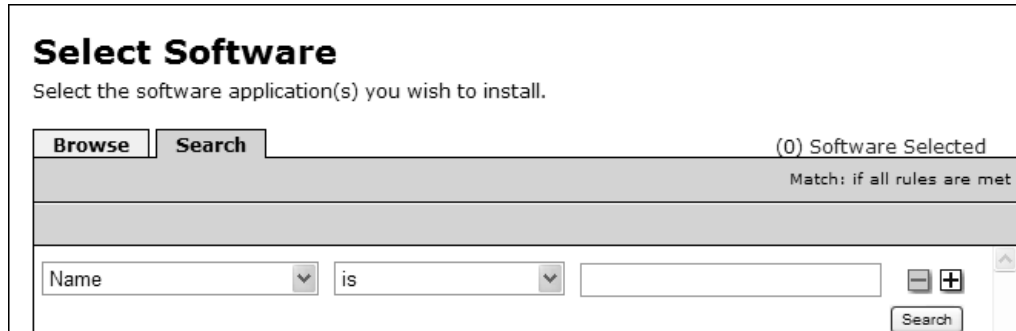
Using the Search Feature

In the Opsware wizards, you can browse for servers, operating systems, patches, applications, templates, and service levels, or use the Search feature to search for these items.

Perform the following steps to search by using the Search feature:

- 1 In an Opware wizard, select the Search tab. The following Search page appears. See Figure 2-8.

Figure 2-8: Search Tab in the Select Servers Step of an Opware Wizard



You can also 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 ► Search. The Search page appears, as Figure 2-9 shows.


Figure 2-9: Search Page



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

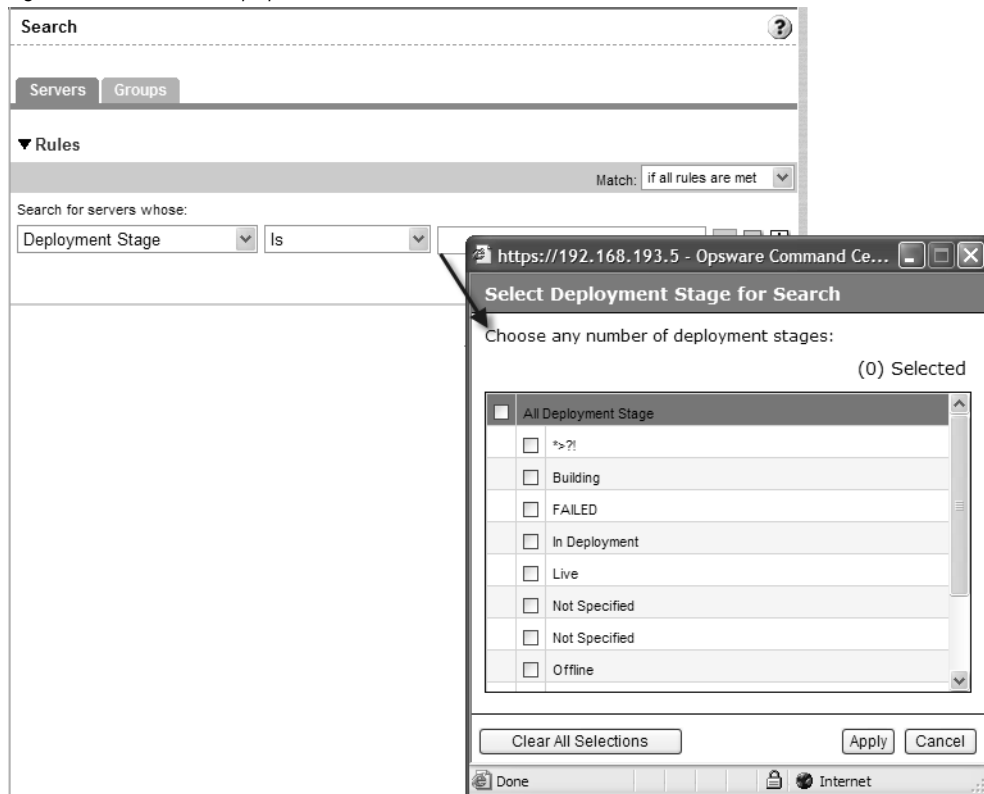
- 2 Specify the rule that you want to search for by selecting it from the first list, as Figure 2-10 shows.

Figure 2-10: Search Rules List in the Search Page



Depending on the rule that you select, a popup window might appear in the page. For example, if you selected Deployment Stage, a popup window showing the stages appears in the page as Figure 2-11 shows:

Figure 2-11: Search Popup Window with Values



You cannot search in Notes that contain line breaks. See “Line Break Workaround for Server Search” on page 79 in this chapter for more information and a workaround.

If you are searching while using an Opsware wizard, the first search rule list might not have all the options. 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 Opsware SAS to search by selecting a value. The operator selected defines how the search text is treated. Negative operators might not be available in all cases.

See “Rules for Server Search and for Creating Dynamic Groups” on page 70 in this chapter for more information about the operator for each search rule.

- 4 Enter the text that you want to search for in the text box or choose a value from the list or popup window. The search text that you enter can include an asterisk (*) wildcard character. The search text is case insensitive. You can also use the SHIFT or CTRL key to select multiple values from the list or popup window.
- 5 (Optional) To add additional rules, click the plus (+) button as Figure 2-12 shows and repeat Steps 2 through 4.

Figure 2-12: Multiple Rules in a Search

The screenshot shows a search interface with the following elements:

- Search** header with a help icon.
- Tabs for **Servers** and **Groups**.
- Rules** section with a dropdown arrow.
- Match:** dropdown menu set to "if all rules are met".
- Search criteria: "Search for servers whose:"
- Rule 1: Deployment Stage (dropdown) **Is** (operator dropdown) In Deployment (text box). Includes a plus (+) button to add more rules.
- Rule 2: Opware Display Name (dropdown) **Contains** (operator dropdown) marketing (text box). Includes a plus (+) button to add more rules.
- Buttons: **Search**, **Save...**, and **Reset**.

- 6 If you specified multiple rules for the search, select whether you want search results only if all rules are met or if any of the rules are met, as Figure 2-13 shows.

Figure 2-13: Operator Controlling Search Results

The close-up shows the **Match:** dropdown menu with the following options:

- if all rules are met
- if any rules are met

By default, search results appear for servers that match all the search rules. If you are searching from an Opware wizard, this field is set to the value if all rules are met; you can change the value when searching for servers, but you cannot change it when searching for patches, software, operating systems, and so on.

- 7** Click **Search**. The list of servers that match your search rules appears in the page, as Figure 2-14 shows.

Figure 2-14: Displayed Search Results

| Results | | | | | | |
|--|---|--------------------------------|---------------|-------------------|----------|----------|
| Resource Edit Tasks Configuration Tracking | | | | | | |
| <input type="checkbox"/> | Name ▲ | Host Name | IP Address | OS Version | Facility | Customer |
| <input type="checkbox"/> |  casela.snv1.corp.opsware.com Web Services test box | casela.snv1.corp.opsware.com | 192.168.9.191 | SunOS 5.9 | C12 | Opsware |
| <input type="checkbox"/> |  chamarel.snv1.corp.opsware.com Web Services test box | chamarel.snv1.corp.opsware.com | 192.168.9.146 | Red Hat Linux 7.3 | C12 | Opsware |

The search results include columns for Name, IP Address, OS Version, Facility, and Customer.

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

Rules for Server Search and for Creating Dynamic Groups

The following table describes the rules that you can use to search for servers or to create dynamic server groups.

Note that anywhere you can enter text, you can enter a wildcard (*) character to broaden your results.

Table 2-3: Rules for Server Search and for Creating Dynamic Groups

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|---|---|
| PROPERTIES | | |
| Agent Discovery Date: The date that the Opsware Agent was installed. | <ul style="list-style-type: none"> Is after Is before | Drop-down lists with the day, month, and year |
| | <ul style="list-style-type: none"> Is within the last | User-entered text |
| | <ul style="list-style-type: none"> Is today | N/A |

Table 2-3: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|--|---|
| Agent Reporting: Whether the Opware Agent is reporting to Opware SAS. | <ul style="list-style-type: none"> • Is • Is not | <ul style="list-style-type: none"> • Has not reported • OK • Registration in progress • Reporting error |
| Agent Status: Whether the Opware Agent is reachable by Opware SAS. | <ul style="list-style-type: none"> • Is • Is not | <ul style="list-style-type: none"> • Not reachable • OK |
| Agent Version: The version of the Opware Agent – such as 14.2.3b. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| Custom Attribute (any): The name of a custom attribute that is associated with the server through attachment or inheritance. | <ul style="list-style-type: none"> • Contains • Is | User-entered text |
| Custom Attribute (local): The name of a custom attribute that is locally attached to the server, | <ul style="list-style-type: none"> • Contains • Is | User-entered text |
| Customer: The customer or account that the server is associated with. | <ul style="list-style-type: none"> • Is • Is not | Popup window of customers |

Table 2-3: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|--|---|
| <p>Deployment Stage: The stage the server performs within a lifecycle environment.</p> | <ul style="list-style-type: none"> • Is • Is not | <ul style="list-style-type: none"> • In Deployment • Live • Not Specified • Offline <hr/> <p>The values that appear in this list are customizable; in addition to the values above, values specific to your environment might appear.</p> <hr/> |
| <p>Facility: The collection of servers managed by an Opware SAS installation.</p> | <ul style="list-style-type: none"> • Is • Is not | <p>Popup window of facilities</p> <hr/> <p>When Opware SAS is running multimaster mode, the list can contain many facilities.</p> <hr/> |
| <p>Group Membership: Whether the server belongs to a group.</p> | <ul style="list-style-type: none"> • Is | <p>Popup window of groups</p> |
| <p>Host name: The host name of the server – such as m004.company.com.</p> | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | <p>User-entered text</p> |
| <p>Name (any): This enables searching for any name or IP address associated with a server.</p> | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | <p>User-entered text</p> |

Table 2-3: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|--|--|
| Notes: The contents of the Notes field from the Properties tab for a server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| Opware Display Name: The user-configurable name for the server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text <hr/> By default, Opware SAS uses the configured host name of the server until a user edits it. <hr/> |
| Server Use: How the server is being used – such as Development, Staging, Production. | <ul style="list-style-type: none"> • Is • Is not | <ul style="list-style-type: none"> • Development • Not Specified • Production • Staging <hr/> You can customize values that appear in this list. In addition to the values above, values specific to your environment might appear. <hr/> |

Table 2-3: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|--|---|--|
| <p>Service Level: The user-defined category that can be used as an organizational tool.</p> | <ul style="list-style-type: none"> • Is • Is attached here or below | <p>Popup window of service levels</p> <hr/> <p>Servers can be associated with multiple service levels.</p> <hr/> <p>See the <i>Opsware® SAS User's Guide</i> for more information about working with service levels.</p> |
| OPSWARE PROPERTIES | | |
| <p>Application Configuration: Whether the server uses the Application Configuration feature.</p> | <ul style="list-style-type: none"> • Is not used • Is used | N/A |
| <p>Code Deployment: Whether the server uses the Code Deployment feature.</p> | <ul style="list-style-type: none"> • Is not used • Is used | N/A |
| <p>Configuration Tracking: Whether the Configuration Tracking feature is monitoring or backing up specific files or configurations on a server.</p> | <ul style="list-style-type: none"> • Is off • Is on | N/A |
| <p>Lifecycle: The server states that are part of bringing a server into Opsware SAS.</p> | <ul style="list-style-type: none"> • Is • Is not | <ul style="list-style-type: none"> • Available • Build Failed • Deactivated • Installing OS • Managed |

Table 2-3: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|--|--|
| Server ID: The internal ID Opware SAS uses to identify the server. | <ul style="list-style-type: none"> • Is • Is not | User-entered text <hr/> In most cases, the Server ID is the same as the MID. <hr/> |
| SOFTWARE | | |
| Attached Software: The software that is assigned or modeled through the Opware reconcile operation – the installation process. | <ul style="list-style-type: none"> • Is • Is attached here or below | Popup window of software |
| Installed Patches: Whether a patch is installed on the server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| Installed Software: The package reported installed on the server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text <hr/> A package does not have to be installed by Opware SAS to be reported as installed on a server. <hr/> |
| OS Version: The OS version defined by OS definitions in the OS Provisioning feature. | <ul style="list-style-type: none"> • Is • Is not | Popup window of operating systems |
| Reported OS: For Windows – the version reported by the OS, for Unix – the version returned by the uname command. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |

Table 2-3: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|--|-------------------|
| Windows Service: The names of the Windows services that are reported to Opsware SAS. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| NETWORK | | |
| DNS Search Domains: The domains configured to be searched in the server's network settings. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| DNS Servers: The IP addresses of the DNS servers configured in the server's network settings. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| Default Gateway: The IP address of the default router. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| IP Address: Any Internet Protocol address for the server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| MAC Address: Any Media Access Control address, which is the network interface card's unique hardware number. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |

Table 2-3: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|---|---|
| WINS Servers: The Windows Internet Naming Servers configured in the server's network settings. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| HARDWARE | | |
| CPU Make and Model: The vendor name and CPU model for the server – such as GENUINEINTEL Intel(R) Pentium(R) 4 CPU 2.60GHz. | <ul style="list-style-type: none"> • Is • Is not | Popup window of CPU makes and models |
| CPU Speed: The Central Processing Unit speed in gigahertz [GHz]. | <ul style="list-style-type: none"> • Does not equal • Equals • Is greater than • Is less than | User-entered text <hr/> A 600 Mhz machine should be entered as 0 . 6. <hr/> |
| Make and Model: The vendor name and server model for the server – such as Compaq - DL360. | <ul style="list-style-type: none"> • Is • Is not | Popup window of server makes and models |
| Number of CPUs: The number of CPUs on the server. | <ul style="list-style-type: none"> • Does not equal • Equals • Is greater than • Is less than | User-entered text |
| RAM: The amount of RAM on the server in megabytes [MB]. | <ul style="list-style-type: none"> • Does not equal • Equals • Is greater than • Is less than | User-entered text <hr/> To enter 1 Gigabyte, type 1024. <hr/> |

Table 2-3: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|--|---|--|
| Serial Number: The serial number of the server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| Storage Make and Model: The vendor name and storage model for the server – such as WDC - WD800BB-75DKA0. | <ul style="list-style-type: none"> • Is • Is not | Popup window of storage makes and models |
| CUSTOM FIELDS | | |
| A Numeric field | <ul style="list-style-type: none"> • Does not equal • Equals • Is greater than • Is less than | User-entered text |
| A String field | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| A URI field | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| A File field | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |

Table 2-3: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---------------------------|----------------------|---|
| A Date field | • Is after | Drop-down lists with the day, month, and year |
| | • Is before | |
| | • Is within the last | User-entered text |
| | • Is today | N/A |

Line Break Workaround for Server Search

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 2-15 shows. For example:

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

Figure 2-15: Line Break Workaround in the Server Search Feature

The screenshot shows a search interface with the following elements:

- Search** header with a help icon.
- Navigation tabs for **Servers** and **Groups**.
- Rules** section with a dropdown arrow.
- Match:** if all rules are met (dropdown menu).
- Search for servers whose:**
 - Field 1: Notes (dropdown menu)
 - Field 2: Is (dropdown menu)
 - Field 3: line1*line2 (text input)
 - Buttons: ... (three dots), - (minus), + (plus)
- Buttons: **Search**, **Save...**, **Reset**

Conditions for Searching with Multiple Rules

When you perform a search with multiple rules, the following conditions apply to the way Opsware SAS provides search results:

- When it evaluates rules, the Search feature considers each rule individually, finding servers (or operating systems, patches, and so forth) that match the individual rule, and then the results of each rule are combined.
- You must select at least one rule, and it must have a value. Default filter rules count as rules with a value. For example, in the MS Updates Wizard, when searching for a server, the Search feature automatically enters a default rule to search only for servers running a Windows OS. Therefore, users are not required to specify any more rules.
- Empty rules 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 Search feature provides default values based on previous selections in the wizard.

For example, in the Select Servers step of the Install Patch Wizard, the search query automatically includes the value for OS version that you specified in earlier steps, as Figure 2-16 shows.

Figure 2-16: Default Values Entered in Searches

The screenshot shows the 'Install Patch' wizard in the 'Select Servers' step. On the left, a navigation pane lists five steps: 1. Select Patches, 2. Select Servers (current), 3. Confirm Selections, 4. Schedule and Notify, and 5. Install. The main area is titled 'Select Servers' and contains a search interface. At the top, there are 'Browse' and 'Search' tabs, and a 'Match: if all rules are met' dropdown. Below this, a section labeled 'Search for servers whose:' contains three search criteria: 'Agent Status' (dropdown), 'Is' (dropdown), and 'OS Version' (dropdown). The 'OS Version' dropdown is open, showing a list of options: 'Not Reachable', 'OK', 'SunOS 5.6' (which is circled in red), 'Build Failed', 'Deactivated', 'Installing OS', and 'Managed'. To the right of each option are minus and plus icons. At the bottom of the search area is a 'Search' button. Below the search area are 'Cancel', '< Previous', and 'Next >' buttons.



The wizards are flexible; for example, when you select multiple servers and applications, Opware SAS will install the correct applications on the servers, even if you selected different OS versions for the servers and applications in the respective steps. Opware SAS also matches the customer association for applications, operating systems, and templates with the customer association of servers.

If Opware SAS cannot find a match, an error message appears at the end of the wizard; therefore, use caution when modifying these default values.

Server Searching by IP Address

Users can search for servers by entering a specific IP address in the Search box in the top navigation or in the Search feature.

Opsware SAS 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 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 Opsware SAS" on page 99 in Chapter 3 for information about how Opsware SAS handles servers that are affected by static NAT.

Example: Server Search

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

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

Searching for a group of servers

- 1 From the navigation panel, click Servers ► Search. The Search page appears. By default, the Servers tab is selected.

Select the Groups tab. The rules for group of servers search appear as Figure 2-17 shows. By default, one search rule is added to the search.

Figure 2-17: Rules for Group Search

The screenshot shows the 'Search' interface with the 'Groups' tab active. Under the 'Rules' section, there is a 'Match:' dropdown set to 'if all rules are met'. Below this, the text 'Search for groups whose:' is followed by a dropdown menu showing 'Group Name' selected. To the right of this dropdown is another dropdown menu showing 'Contains'. To the right of that is an empty text input field. To the right of the text input field are three buttons: an ellipsis (...), a minus sign (-), and a plus sign (+). Below the 'Group Name' dropdown, a mouse cursor is pointing at the 'Group Name' option in the dropdown list. At the bottom right of the rule configuration area is a 'Search' button.

- 2 In the second list, specify how you want Opware SAS to search by selecting either “Contains” or “Is”. The operator selected defines how the search text is treated.
- 3 Enter the text that you want to search for in the text box. The search text that you enter can include an asterisk (*) wildcard character. (The search text is case sensitive. **Search** is not enabled until you enter text in the text box.)



When searching for a group of servers, you can only specify one rule for Group Name. The plus (+) button is disabled.

- 4 Click **Search**. The search results appear as Figure 2-18 shows.

Figure 2-18: Group Search Results

The screenshot displays the 'Search' interface for groups. At the top, there are tabs for 'Servers' and 'Groups'. Below the tabs, there is a 'Rules' section with a 'Match:' dropdown set to 'if all rules are met'. The search criteria are defined as 'Group Name' containing 'test'. A 'Search' button is located to the right of the search criteria. The 'Results' section shows a table with the following data:

| | Name | Location | Members | | Total Servers | Type | Model Attachments | Custom Attributes | Last Use |
|--------------------------|--------------|---------------|---------|-----------|---------------|---------|-------------------|-------------------|-------------------|
| | | | Servers | Subgroups | | | | | |
| <input type="checkbox"/> | a test group | Public Groups | 0 | 0 | 0 | Dynamic | 3 | 0 | 05-22-05 16:00:00 |

Server Identification

This section provides information on server identification within Opware SAS and contains the following topics:

- Overview of Server Identification
- Ways that Servers are Identified by Opware SAS
- Customer Accounts in Opware SAS
- Associating Servers with Customers

Overview of Server Identification

Opware SAS uses the following IDs to track managed servers:

- **MID**: Machine ID. The unique identifier that Opware SAS 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 Opware Agent.

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

See “Example: Opware SAS Swaps a Server’s Hard Disk” on page 86 in this chapter for information about swapping hard disks.

- **Server ID:** The primary key in the Opware Model Repository (database) that represents a given server. The Server ID is used internally in Opware SAS. Generally, users do not need this value for servers to manage them in Opware SAS.
- **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 Opware Agent discovers, typically derived from some property of the server’s chassis. As a common source for this ID, Opware SAS uses an interface’s MAC address or the host ID on Solaris servers, or the serial number for one of the interfaces.

Ways that Servers are Identified by Opware SAS

Servers in the Manage Servers list are identified in the following ways when they register their hardware and software with Opware SAS:

- Opware SAS 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 Opware SAS 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, Opware SAS uses the hardware address by which the server contacted the Opware Build Manager (a component of the OS Provisioning feature).

Example: Opsware SAS Swaps a Server's Hard Disk

The following steps show how Opsware SAS handles swapping a hard disk for a server:

- 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 Opsware SAS. The MID for Server A equals 1230001 and the chassis ID equals XY:96.
- 3** Opsware SAS locates Server A by using the MID.
- 4** Opsware SAS 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 Opsware SAS. The MID for Server B equals 98730001 and the chassis ID equals AB:08.
- 6** Opsware SAS locates Server B by using the MID.
- 7** Opsware SAS updates the data it has for Server B in the Model Repository. It sets the chassis ID equal to AB:08.

Customer Accounts in Opsware SAS

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).

Opsware SAS accommodates these requirements with customer accounts created by your Opsware administrator.

When your Opware administrator creates a customer in Opware SAS, a value for that customer is automatically added to the customer filter in the Manage Servers list, as Figure 2-19 shows.

Figure 2-19: Customer Filter in the Manage Servers List

Manage Servers (Summary View) ?

All Statuses ▾
All Operating Systems ▾
All Stages ▾

Update

All Uses ▾
All Customers ▾
All Manufacturers ▾

All Models ▾
All Customers ▾
Activated ▾

Resource Edit
Configuration Tracking

| | Name | Host Name / IP Address | OS Version | Stage | Use | Customer |
|--------------------------|----------------------|--------------------------------------|-------------------------------------|------------------|------------------|-----------------|
| Public Groups | | | | | | |
| <input type="checkbox"/> | core.tr3.opsware.com | core.tr3.opsware.com 172.16.36.18 | Red Hat Enterprise Linux AS 3 | Not Specified | Not Specified | Opware |
| <input type="checkbox"/> | OPSWARE-STROCOZ | OPSWARE-STROCOZ 172.16.36.43 | Windows 2000 | Not Specified | Not Specified | Not Assigned |

By using customer accounts in the SAS Web Client, 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.

By default, Opware SAS is shipped with the following two customers:

- **Customer Independent:** A global customer in Opware SAS. 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 install an Opware Agent in a server, the server is associated with the Not Assigned customer if IP ranges were not created to automatically associate Opware SAS managed servers with customers. See Figure 2-20.



Opware, Inc. recommends that you associate servers with customers, if necessary, by using the Server Properties pages.

Figure 2-20: Customers List Under Environment in the SAS Web Client

| Customers | | | |
|-----------|------------|--|----------------------|
| | Name | | Name |
| | 12204 | | Corp Test |
| | Big Corp | | Big Corp2 |
| | Test Cust | | Customer Independent |
| | E-Commerce | | Not Assigned |

Associating Servers with Customers


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

- Install Opware Agents on the servers.
See the *Opware® SAS User's Guide: Server Automation* for information.
- Use the OS Provisioning feature to install operating systems on bare-metal servers.
See "Operating System Provisioning" on page 331 in Chapter 8 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 SAS Web Client, 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 the *Opware® SAS Configuration Guide* 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 143 in Chapter 3 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 Opware SAS" on page 99 in Chapter 3 for information about how Opware SAS 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. Opware SAS enforces the relationship between server and facility at hardware registration. See Figure 2-21.

Figure 2-21: Primary IP Addresses in Opware SAS



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 Opware SAS and contains the following topics:

- Overview of Server History
- Viewing Server History
- Time Stamp for Server Operations
- DCI Reporting

Overview of Server History

By using the SAS Web Client, you can view the history of changes made to a server. Each action performed on a managed server is logged in the history with the associated user who performed the action and the time of the action. You can view the history at any time, but you cannot change it. History is read-only. See Figure 2-22 for examples of recorded information.

Figure 2-22: Server History for a Server in the Manage Servers List

Manage Servers: History | chamarel.snv1.corp.opsware.com 

[Return to Manage Servers](#)

| Properties | Network | Membership | Attached Nodes | Installed Packages | Custom Attributes | Config Tracking | History |
|--|-------------|------------|----------------|--------------------------|-------------------|-----------------|---------|
| HISTORY FOR chamarel.snv1.corp.opsware.com BELONGING TO OPSWARE | | | | | | | |
| Show Last: Week 2 Weeks Month Quarter | | | | | | | |
| Description | Modified By | | | Date Modified | | | |
| Added server to server group copy of grY | gshort | | | Thu May 12 01:38:28 2005 | | | |
| Added server to server group Test Dynamic DCNqzse | buildmgr | | | Wed May 11 18:57:30 2005 | | | |
| Added server to server group Test GroupsByNodes DCNuqnj | buildmgr | | | Wed May 11 18:43:43 2005 | | | |

Each History entry contains three pieces of information, as Table 2-4 shows.

Table 2-4: '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. |
| 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) |

Actions Logged in History

Opsware SAS also logs the following actions in the history for each managed server:

- Addition of the server to a node
- Removal of the server from a node
- Reassignment of the server from one node to another
- Addition of a cloned server to a node (when cloning a server, node assignments occur in Opsware SAS and are logged)
- Installation of a template on the server
- Preview reconcile failure or success
- Reconcile failure or success

Time Stamps used in History

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

- The SAS Web Client 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 SAS 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, Opware, Inc. recommends regular backups to enable offline storage of Opware SAS data.

Opware SAS deletes old data from the Opware Model Repository, and does not copy the data before it removes it, but you can retain information for longer periods of time, by using Oracle commands to manipulate scheduled jobs. Contact your Opware, Inc. support representative for assistance in changing these retention periods.

Viewing Server History

Perform the following steps to view the server history:

1 From the navigation panel, click Servers ► Manage Servers. The Manage 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 “Using the Search Feature” on page 65 in this chapter for more information. See “Server Searching by IP Address” on page 81 in this chapter for more information.

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

3 Select the History tab.

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

Time Stamp for Server Operations

Opware SAS maintains a comprehensive audit trail of the software that the Opware users install, configure, and remove from a server. By using the SAS Web Client, you can view the history of the changes made to a server. Entries are generated when actions are performed for managed servers in the SAS Web Client. The history is read-only. See “Server Histories and Reports” on page 90 in this chapter for more information.

The time stamps for system events are determined based on the system clocks of the servers running the Opware core components. To obtain accurate time stamps in server histories (displayed in the SAS Web Client) and Opware component logs, you must:

- Synchronize the system clocks on all servers running Opware components so that all the servers are running with a common time.

- Set the time zone for all servers running Opsware components to Coordinated Universal Time (UTC).

See the *Opsware[®] SAS Deployment and Installation Guide* for more information on facility time requirements.

Additionally, Opsware, Inc. recommends that after installing an Opsware Agent on the server (so that it becomes managed by Opsware SAS), you should synchronize the system clock on the server with the system clocks of the servers running the Opsware core components.

The time stamps appearing in the SAS Web Client server History tab and the Opsware Agent logs are obtained from the Opsware core; however, you might need to review the server's logs (such as, stdout). Having a consistent time stamp in the server's logs and Opsware SAS is essential for effective troubleshooting.

DCI Reporting

The Model Repository maintains precise information about the state and configuration of every server under management. You can access this information through Opsware Data Center Intelligence (DCI) Reporting. DCI Reporting provides dynamic and detailed information about the operational environment, and it includes the following features:

- 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 DCI 1.8 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. Note that the DCI Reporting feature must be installed and running in the facility to view the online documentation.



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

Hardware Information for Managed Servers

The Hardware link of the SAS Web Client 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

See the *Opsware® SAS User's Guide: Server Automation* for the list of server information that the Opsware Agent tracks.

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 information as MAC ID, CPUs, memory, and so on.

To view hardware of managed servers:

- 1** From the navigation panel, click Environment ► Hardware. You see the top level of the Hardware category in the managed environment, as shown in Figure 2-23.

Figure 2-23: Top level hardware node

| | Name | OS Version | Modified | Customer |
|--|--------|----------------|------------|----------------------|
| | SERVER | OS Independent | 03-29-2005 | Customer Independent |

- 2** To view the servers in your managed environment, click the Servers link.

- 3** Drill down to the type of server you want to look at. For example, you might want to look at all Dell POWEREDGE 650s, as shown in Figure 2-24.

Figure 2-24: Hardware home page for Dell POWEREDGE 650s

Hardware > SERVER > DELL COMPUTER CORPORATION > POWEREDGE 650

  **POWEREDGE 650**

No Sub-Nodes

| Properties | Members 107 |
|--|----------------------|
| Cannot edit or delete this Node. This Node is auto-generated. This Node is special and cannot be modified. | |
| Name: | POWEREDGE 650 |
| Customer: | Customer Independent |
| Operating System: | OS Independent |
| Locked: | No |
| Allow Servers: | Yes |
| ID: | 1960003 |

- 4** Next, select the Members tab. You will see a list of all the Dell POWEREDGE 650s in your managed environment.
- 5** 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 Opsware SAS.

Chapter 3: Server Management in the SAS Web Client

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Server Management
- Communication Between Managed Servers and Opware SAS
- Server Groups
- Server Life Cycle
- Server Management Jobs
- Custom Fields for Servers
- Custom Attributes for Servers
- Service Levels
- Network Configuration

This section 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 sections.

Overview of Server Management

Opware SAS manages servers in an operational environment in the following ways:

- Installing Opware Agents on 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, Opware SAS integrates with AIX NIM and HP-UX Ignite installation

technologies to provide a uniform method for OS provisioning across a heterogeneous environment.

See the *Opware® SAS User's Guide: Application Automation* for information about how Opware SAS provisions Microsoft Windows, Red Hat Linux, and Sun Solaris operating systems on managed servers.

See the *Opware® SAS Configuration Guide* for information about how Opware SAS integrates with AIX NIM and HP-UX Ignite installation technologies.

- Automating patch management on Opware SAS-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 the *Opware® SAS User's Guide: Application Automation* for 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 “Software Provisioning” on page 291 in Chapter 6 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 the *Opware® SAS User's Guide: Server Automation* for information about tracking configurations for servers. See Chapter 7, “Configuration Tracking Using the Software Tree” on page 309 of this guide for more information setting configuration tracking policies with the Software Tree.



To manage servers with Opware SAS, you do not need root access on Unix or administrator access on Windows. However, you will need permissions to use specific Opware SAS features, as well as permissions for customers and facilities associated with servers. To obtain these permissions, contact your Opware administrator. For more information, see the Permissions Reference appendix in the *Opware® SAS Configuration Guide*.

Communication Between Managed Servers and Opware SAS

This section provides information about communication between managed servers and Opware SAS and contains the following topics:

- Network Address Translation (NAT) for Managed Servers
- Key Terms of Opware SAS Managed Server Communication
- Locating the Management IP Address of a Managed Server
- Code Deployment and Static NAT
- Setting the Primary IP Address of a Server
- NAT Table Mapping and Managed Servers

Network Address Translation (NAT) for Managed Servers

To manage a server, Opware SAS requires that the server have a unique IP address that is routable from Opware SAS. However, in a large operational environment, all servers might not have unique IP addresses. In this case, Opware SAS supports static Network Address Translation (NAT) for managed servers.

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

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

Key Terms of Opware SAS Managed Server Communication

To understand how Opware SAS communicates with managed servers, you must understand these three terms:

- **Management IP:** The IP address that Opware SAS uses to communicate with the Opware Agent on the server.

During hardware registration for a server, the Opware Agent opens a TCP/IP connection to Opware SAS. The connection contains the source IP (called peer IP) address of the server. By default, Opware SAS 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 & Rollback feature uses the primary IP address to communicate with the server. See “Code Deployment and Static NAT” on page 102 in this chapter for more information.

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



Opware SAS 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.

Locating the Management IP Address of a Managed Server

In the SAS Web Client, 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 SAS Web Client 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 Manage 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 3-1 shows.

Figure 3-1: Management IP Address Information in the Network Tab

| Properties | Network | Membership | Attached Nodes | Installed Packages | Custom Attributes | Config Tracking | History |
|---|---------|--|----------------|--------------------|-------------------|-----------------|---------|
| NETWORK INFORMATION (as of Wed May 18 07:05:57 2005) | | | | | | | |
| Hostname: | | <input type="text" value="core.tr3.opsware.com"/> | | | | | |
| Management IP: | | <input type="text" value="172.16.36.18"/> | | | | | |
| Facility: | | <input type="text" value="DATACENTER1"/> | | | | | |
| Management Interface: | | <input type="text" value="eth0"/> | | | | | |
| Gateway: | | <input type="text" value="172.16.36.17"/> | | | | | |
| DNS Servers: | | <input type="text" value="66.54.32.78"/> <input type="text" value="66.54.0.78"/> | | | | | |
| Search Domains: | | <input type="text" value="tr3.opsware.com"/> <input type="text" value="opsware.com"/> | | | | | |
| CONFIGURATION FOR: eth0 | | | | | | | |
| Use DHCP Settings: | | <input type="text" value="Static"/> | | | | | |
| IP Address: | | <input type="text" value="172.16.36.18"/> | | | | | |
| MAC Address: | | <input type="text" value="00:11:43:D7:2F:5F"/> | | | | | |
| Interface Type: | | <input type="text" value="ETHERNET"/> | | | | | |
| Interface Speed: | | <input type="text" value="(not set)"/> | | | | | |
| Subnet Mask: | | <input type="text" value="255.255.255.248"/> | | | | | |
| CONFIGURATION FOR: eth1 | | | | | | | |
| Use DHCP Settings: | | <input type="text" value="Disabled"/> | | | | | |
| IP Address: | | <input type="text"/> | | | | | |
| MAC Address: | | <input type="text" value="00:11:43:D7:2F:60"/> | | | | | |
| Interface Type: | | <input type="text" value="ETHERNET"/> | | | | | |
| Interface Speed: | | <input type="text" value="(not set)"/> | | | | | |
| Subnet Mask: | | <input type="text"/> | | | | | |
| | | <input type="button" value="Update"/> <input type="button" value="Revert"/> | | | | | |

Figure 3-2 shows the Hardware view in the Manage 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 3-2: Hardware Tab in the Manage Servers List

| Resource | Edit | View | Tasks | Configuration Tracking | | | | |
|--------------------------|----------------------|--------------------------------------|--------------|-----------------------------------|---|---|--|--|
| <input type="checkbox"/> | Name | Host Name / IP Address | Registration | Network Info | Hardware Info | Hardware Components | | |
| <input type="checkbox"/> | Public Groups | | | | | | | |
| <input type="checkbox"/> | core.tr3.opsware.com | core.tr3.opsware.com 172.16.36.18 | OK | 172.16.36.18 (eth0) N/A (eth1) | Man: DELL COMPUTER CORPORATION Mod: POWEREDGE 1850 Mac: 00:11:43:D7:2F:5F Ser: 2YLLP61 | CPUs: Intel(R) Xeon(TM) CPU 2.80GHz, 2793 MHz, 1 MB Intel(R) Xeon(TM) CPU 2.80GHz, 2793 MHz, 1 MB Intel(R) Xeon(TM) CPU 2.80GHz, 2793 MHz, 1 MB Intel(R) Xeon(TM) CPU 2.80GHz, 2793 MHz, 1 MB Memory: 3.9 GB RAM 1.95 GB SWAP Stor: hda (CDROM) sda (SCSI DISK:68.24 GB) | | |

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 Opware Agents in the same NAT domain. Synchronizations cannot be performed between Opware Agents in different NAT domains.

Opware SAS 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 Opware Agents must communicate directly. Users can override the IP address that Opware SAS determined and designate a specific network interface as the management interface.

See the *Opware® SAS User's Guide: Server Automation* 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 Opware Agent-to-Opware 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 ► Manage Servers. The Manage 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 Manage Servers: Properties page appears.

3 Select 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 **Update**.

See "Using the Search Feature" on page 65 in Chapter 2 for more information. See "Server Searching by IP Address" on page 81 in Chapter 2 for more information.

NAT Table Mapping and Managed Servers

Static, one-to-one NAT tables map routable IP addresses between Opware SAS 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 Opware SAS.

Opware SAS does not control these NAT tables and errors can occur if they are modified after Opware-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 Opware SAS side of the NAT mapping is modified, the server becomes unmanageable and might be marked Not Reachable on the Manage Servers: Status page. It stays in this state until the Opware 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 Manage

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 Opware Agent. Both servers eventually become manageable.

Server Groups

This section provides information on group of servers within Opware SAS and discusses the following topics:

- Overview of Server Groups
- Types of Server Groups
- Public Group Modeling
- Ways to Create Server Groups
- Adding a Server to Static Groups, Both Public and Private
- Removing Servers from Static Groups, Both Public and Private
- Duplicating Server Groups
- Rules for Deleting Server Groups

Overview of Server Groups

The Server Groups feature is useful for gathering servers into collections. These groups can be used as a shortcut for performing the same action on all of the servers simultaneously, instead of performing the action on each individual server, one at a time. Server groups can also be used to simply organize groups of servers.

Server groups can be comprised of individual servers as well as other groups of servers.



The My Servers feature can also be used to gather servers and groups of servers, but it has different functionality than the Server Groups feature. You can add individual servers, groups of servers, and nested groups that you access frequently to My Servers. See “My Servers” on page 62 in Chapter 2 for information about using this feature.

Uses for Server Groups

Some recommended uses for server groups include:

- Grouping servers by OS version
- Grouping servers by customer
- Grouping servers by facility
- Grouping servers by deployment stage
- Grouping servers by use (Server Use in the Server Properties page)
- Grouping servers by operational boundaries, for example, grouping together all servers that require identical application configuration
- Grouping servers by access control boundaries, for example, creating server groups that are associated with a specific User Group

Permissions Required for Working with Server Groups

Users must have the permissions shown in Table 3-1 in order to perform specific tasks related to server groups. Only administrators can set permissions.

Table 3-1: Permissions Required for Working with Server Groups

| NAME OF PERMISSION | WHERE SELECTED | ENABLES YOU TO |
|-----------------------------|---|---|
| Manage Servers | | Create, edit, and delete private server groups, both static and dynamic. |
| Manage Public Server Groups | The Manage Servers Permissions section on the Other tab in Users and Groups | Create, edit, and delete public server groups, both static and dynamic. |
| Model Public Server Groups | The Manage Servers Permissions section on the Other tab in Users and Groups | Model public server groups. See “Public Group Modeling” on page 108 in this chapter for more information. |

Table 3-1: Permissions Required for Working with Server Groups (continued)

| NAME OF PERMISSION | WHERE SELECTED | ENABLES YOU TO |
|------------------------|---|--|
| Allow Run Refresh Jobs | The Manage Servers Permissions section on the Other tab in Users and Groups | Add or remove servers from groups before a scheduled job is run. This permission gives you the option to refresh group membership before a job is run so that it is only run on the servers that belong to the group when the job is actually run. |

Characteristics of Server Groups

When using server groups, groups have the following characteristics:

- Individual servers can be included in as many groups as you want, or not included in any groups.
- Adding servers to a group does not remove those servers from the list of all servers that appears when you click Servers ► Manage Servers in the navigation panel.
- Groups can contain servers and subgroups.
- Server groups are hierarchical (they can be nested) with these caveats:
 - Private and public groups cannot be mixed in a hierarchy, but static and dynamic groups can.

See “Types of Server Groups” on page 107 in this chapter for more information about private and public groups.

- The rules for a dynamic group are not inherited from a parent dynamic group to a dynamic subgroup.

See “Dynamic Groups” on page 108 in this chapter for more information about the characteristics of dynamic groups.

- Groups do not inherit modeling data from their parents, including node attachments and custom attributes.
- When you run an operation on a group that contains nested groups, the operation also applies to all the servers in the nested groups below the current group.

- When an Application Configuration operation within the SAS Client is applied to groups, and those groups contain subgroups, the operation does not apply to all the servers in the subgroups. It only applies to the group upon which the operation was directly applied.

Types of Server Groups

There are private groups and public groups, and each can be either static or dynamic.

Private Groups

If you belong to a user group that has access to the Manage Servers list, you can create groups that you alone can see and work with. Only you see your private groups – other Opsware users cannot see them. Private groups behave the same way as public groups, with the exception that modeling is not available for private groups. See “Public Group Modeling” on page 108 in Chapter 3 for information about how the Opsware model affects groups.

When you create your first group, the default type will be Private Static, which can be changed to Private Dynamic, Public Static, or Public Dynamic. When you create a subgroup, the type of group is private if you are in a private group when you create the new group, and the type is public if you are in a public group when you create the new group. Public and private groups cannot be mixed in a hierarchy. In other words, if the parent group is public, the subgroups must be public, and if the parent group is private, the subgroups are also private.

Public Groups

Public groups can be created, edited, or deleted by anybody who has Manage Public Server Groups permissions. Public groups are visible to all users, and can be used by anybody, regardless of who created them, but only users with the Manage Public Server Groups permissions can change the rules that govern dynamic groups.

A link called Public Groups appears at the top of server lists. Clicking that link displays a list of available public groups.

Only public groups can be used for modeling.

Static Groups

Static groups can be either public or private, and no specific permissions are required for static groups. A static group has servers that are added to and removed from the group manually. When using static groups, you first create the group, and then select the servers to populate it.

Dynamic Groups

Dynamic groups contain servers that are added to or removed from the group based on a set of user-defined rules. If the rules are changed or the servers in the environment change, servers will be added to or removed from the group automatically. Rules apply only to the group being created or modified, not to any subgroups.

Once the rules have been created, Opsware SAS will search for servers that match the criteria of that specific group, and add them to the group. When the rules are changed, Opsware SAS will search again, and the resulting group members reflect the changed criteria. Consequently, as servers are added to or removed from management using Opsware SAS, the members of the group will change automatically.

Opsware SAS calculates server group membership each time any of the following actions occur:

- After users add, delete, or change the rules for dynamic server groups.
- When attributes of servers change such that dynamic group membership could change.

Additionally, Opsware SAS automatically recalculates dynamic group membership every hour.

When a user schedules a job to run, dynamic group membership can be determined in either of the following ways:

- Based on the servers in the dynamic group when the job was scheduled.
- Based on the set of servers in the dynamic group when the job actually runs. The membership is recalculated at that time.

Public Group Modeling

With Opsware modeling, the desired state of a server is defined and then applied to servers. In the case of public server groups – static and dynamic – you can define a model consisting of applications, patches, service levels, and custom attributes, which will be applied to all servers in the group. The modeling information is attached to the group, but not to any subgroups.

If the modeling information changes, the servers in the group are not affected until the reconcile operation runs on those servers. If the model has already been reconciled on that server when it is removed from the group, the installed material will be removed at the next reconcile.

Manage Servers Display for Public and Private Groups

As Figure 3-3 shows, server groups appear at the top of the Manage Server list. Public Groups appear as a link at the top of the list. Private groups, if any have been created, appear next, followed by the list of individual servers.

Figure 3-3: Manage Servers Displaying Public and Private Groups – Servers and Group View

| Resource Edit View Tasks Configuration Tracking | | | | | | | |
|---|--|---|----------------------------------|------------------|------------------|----------|----------|
| <input type="checkbox"/> | Name ^ | Host Name / IP Address | OS Version | Stage | Use | Facility | Customer |
| | Public Groups | | | | | | |
| <input type="checkbox"/> | barbar, 0 Servers (Static) | | | | | | |
| <input type="checkbox"/> | barfoo, 0 Servers (Dynamic) | | | | | | |
| <input type="checkbox"/> | dhcp-8-071.corp.opsware.com faisal's desktop do not touch | dhcp-8-071.corp.opsware.com 192.168.8.71 | Red Hat Enterprise Linux AS 3 | Not Specified | Not Specified | C29 | Opsware |

Manage Servers List - Groups Only

When you select **View ► Groups Only** from the menu, information about the number of servers in a group, the number of groups in a group, and whether the group is static or dynamic appears in the list of groups, as Figure 3-4 shows.

Figure 3-4: Manage Servers List – Group Only View

| Manage Servers (Summary View - Groups Only) | | | | | | | | |
|---|---------------|-----------------------|-----------|---------------|--------|-------------------|-------------------|----------|
| OK | | All Operating Systems | | | Update | | | |
| Resource Edit View Tasks Configuration Tracking | | | | | | | | |
| <input type="checkbox"/> | Name ^ | Members | | Total Servers | Type | Model Attachments | Custom Attributes | Last Use |
| | | Servers | Subgroups | | | | | |
| | Public Groups | 0 | 113 | 222 | | 0 | 0 | |
| <input type="checkbox"/> | A Group | 0 | 1 | 0 | Static | 0 | 0 | |

You can modify your views of servers and server groups by selecting Summary, Hardware, Software, and Communications from the **View** menu. You can also elect to further modify your view by selecting Servers and Groups, Servers Only, or Groups Only.

Operations on Server Groups

Any operation that can be done to a server can be done to a server group, because the group acts as a container for a collection of servers, and so provides a shortcut to avoid having to repeat the same operation on each individual server.

When any of the following operations are performed on a group, they are actually performed on the servers within the group, not on the group itself.

- Install
 - Application
 - Patch (Unix patch)
 - OS
 - By Template
- Uninstall
 - Application
 - Patch (Unix Patch)
- Run
 - Script
 - Control
 - Custom Extension
 - Comm Test
- Reconcile
- Customer
- Usage
- SAS Client operations
 - Configure Application
 - Audit Application Configurations
 - Perform Server Audit
 - Create Server Snapshot

Also, some operations allow users to refresh the list of servers in the group (if the user has the Allow Run Refresh Jobs permissions). Users with the correct permissions schedule a job, and before the job is run, Opware SAS will update the members of the server groups upon which the operation is performed. The following actions allow refresh when scheduling a job:

- Install Application
- Install Patch
- Install By Template
- Run Script
- Run Custom Extension
- Reconcile

Using Server Groups Tabs

When you view group properties, the information for the group appears in the page, as Figure 3-5 shows.

You can view group properties in either of the following ways:

- By clicking the This Group link at the top of each server group list.
- By selecting a group in the list and choosing **Resource** ► **Properties** from the menu.

Figure 3-5: Tabs Available for Server Groups

Manage Servers / Public / All Windows Servers / **Properties** ?

Properties Rules Custom Attributes 0 Patches 0 Applications 0 Service Levels 0 History

The following properties describe or restrict use of the current group. Edit

| | |
|---------------------------------------|---------------------|
| Name: | All Windows Servers |
| Description: | |
| Type: | Dynamic |
| Status: | Active |
| Accessibility: | Public |
| Servers (at this level): | 56 |
| Groups (at this level): | 0 |
| Unique servers for all levels: | 56 |
| Date last used (by a job): | 2005-05-23 |
| Server Group ID: | 43380004 |

You can use these tabs to perform the following actions for a server group.

To perform any of these actions, you must have the correct Opware SAS permissions. See “Permissions Required for Working with Server Groups” on page 105 in this chapter for more information.

- **Properties:** This displays information about the group, such as the type of group, the number of servers in the group, and the status of the group. Clicking **Edit** allows you to change the group name and description, and to convert a dynamic group to a static group.



A created static group without servers added yet – an empty static group – can be converted to a dynamic group. Conversely, a static group that has servers cannot be converted to a dynamic group.

- **Rules:** This appears when you are viewing a dynamic group. The rules tab displays the rules used by the dynamic group to determine group membership. The rules apply to the current group only and do not apply to subgroups. Click **Edit** to change the rules for the group.

See “Creating Groups by Using Search” on page 118 in this chapter for more information about specifying rules for a dynamic group.

- **Custom Attributes:** This allows you to set custom attributes for a server group. Click **New** to add an attribute and then click **Edit** to change existing attributes. Custom attributes are not inherited by subgroups within a group hierarchy.

See “Custom Attributes for Servers” on page 160 in this chapter for more information about how custom attributes affect Opware-managed servers.

- **Patches:** This allows you to add Unix patches to a server group. Attaching Unix patches to a server group is like adding a patch to an Opware Template with one exception – groups are not associated with a specific OS or customer. Therefore, you can add a patch for any OS to the group and groups can contain servers running different operating systems. (Patches are always associated with the customer Customer Independent.) Templates and folders are always associated with one OS version and a specific customer.
- **Applications:** Allows you to add applications to a server group. Attaching applications to a server group is similar to adding an application to an Opware Template with one exception – groups are not associated with a specific OS or customer; therefore, you can add an application for any OS or for any customer to the group.



The Reconcile Software wizard is flexible when you run it on a server group; for example, when you reconcile groups of servers, Opware SAS will install the correct applications on the servers in the group, even if the OS versions for the servers and the patches or applications do not all match. Opware SAS also matches the customer association for applications, operating systems, and templates with the customer association of servers. If Opware SAS cannot find a match, an error message (no valid devices found) appears at the end of the wizard.

- **Service Levels:** This allows you to attach service levels to a group. Attaching service levels to a server is like attaching a service level to an Opware Template with one exception – groups are not associated with a specific OS or customer; therefore, you can add a service level for any OS or for any customer to the group.

See the Application Provisioning Setup chapter in the *Opware[®] SAS Configuration Guide* for more information about attaching service levels.

- **History:** This allows you to view the changes to groups. For each action on the group (but not group members), the history displays a description, the date the action occurred, and the user who performed the action (if the group is public).

Ways to Create Server Groups

Server Groups can be created by:

- Using the New Group option from the Resource pull down menu on the Manage Servers page
- Clicking the Create Server Groups icon in the upper right corner of the Copy to Group dialog
- Performing a server search, and saving the resulting list of servers or the rules as a group

Creating Static Groups by Using the New Group Option

Static server groups require servers to be added to them manually. The servers in a static group also must be removed manually.



A created static group without servers added yet – an empty static group – can be converted to a dynamic group. Conversely, a static group that has servers cannot be converted to a dynamic group.

To create static server groups, perform the following steps:

- 1** From the navigation panel, click Servers ► Manage Servers. The Manage Servers page appears.
- 2** From the **Resource** menu, choose **New Group**. The New Private Static Group dialog appears, as shown in Figure 3-6.

Figure 3-6: Default New Private Static Group Dialog

| Name | Type |
|----------------------------|---------|
| Public Groups | Dynamic |
| Content Groups, Dynamic | Dynamic |
| Content Management Servers | Dynamic |
| Dynamic Content Servers | Dynamic |

- 3** To create a Private Static group in the top level of Manage Servers, perform the following steps:
 1. Enter the name of the group in the Save As text box.
 2. Click **Save**.
- 4** To create a Private Static group below another group, perform the following steps:
 1. Navigate to the group below which you want to create a new group.

2. Enter the name of the group in the Save As text box.
3. Click **Save**.

5 To create a Public Static Group, perform the following steps:

1. Click the Public Groups link in the Name and Type window, and navigate to the location in the group hierarchy where you want to create the group.
2. Enter the name of the group in the Save As text box.
3. Click **Save**.



The Save In drop down list is populated according to the location you drill down to in order to create your group. You can use the Save In drop down list to verify that you are in the correct location in the group hierarchy, and you can move to a different location in the hierarchy by selecting it from the Save In list.

Creating Static Groups by Using the Copy to Group Dialog

To create new server groups by copying existing server groups, perform the following actions:

- 1** From the navigation panel, click Servers ► Manage Servers. The Manage Servers page appears.
- 2** Navigate to the servers and groups that you want to copy and click the check boxes next to the servers and groups you want to copy.
- 3** From the **Edit** menu, choose **Copy to Group**. The Copy to Group dialog box appears.
- 4** Navigate to the place in the group hierarchy where you want to copy the servers and group then by clicking the group links displayed in the Name and Type field.
- 5** In the Options field, select either of the following options:
 - **Maintain Hierarchy:** Select this option to copy any server groups exactly as they are.
 - **Expand to a Flat List:** Select this option to copy only the servers within the group to the new group.

Each of those options displays the result of choosing that option by showing the numbers of servers and groups in the new group's hierarchy, or the number of unique servers in the flat list.



If you select the Expand to a Flat List option, and a server is a member of more than one group, that server will only appear in the list once.

6 (Optional) You can also create an entirely new group for the copied servers and groups by clicking the Create New Group icon in the upper right corner of the dialog.

A new dialog appears, prompting you to name your new group.

7 Enter the name of your newly-created group.

8 Click **OK**. The name dialog closes, and the Copy to Group dialog reappears.

9 Click **OK** on the Copy to Group dialog.

Creating Dynamic Groups by Using the New Group Option

Dynamic Server Groups are rule-based, and the servers in dynamic groups will be added or removed automatically based on the rules that you define.

The method for creating dynamic server groups is the same as for creating static server groups. The difference is that when dynamic is selected, you are presented with a page that allows you to define the rules for the group.



Dynamic groups can be converted to static groups, and all servers remain in the group, but all rules will be lost when they are converted.

To create dynamic groups using the New Groups Option, perform the following steps:

1 From the navigation panel, click Servers ► Manage Servers. The Manage Servers page appears.

2 From the **Resource** menu, choose **New Group**. The New Private Static Group dialog appears.

3 To create a private dynamic group, navigate to the place in the group hierarchy where you want to create the group by clicking the group links displayed in the Name and Type field.

Or

To create a public dynamic group, click the Public Groups link in the Name and Type field, and navigate to the location in the group hierarchy where you want to create the group.



The Save In drop down list is populated according to the location you drill down to in order to create your group. You can use the Save In drop-down list to verify that you are in the correct location in the group hierarchy, and you can move to a different location in the hierarchy by selecting it from the Save In list.

- 4** Enter the name of the group in the Save As text box.
- 5** In the Options field, select Dynamic.
- 6** Click **Create Rules**. The Manage Servers Properties page appears, with the Rules tab selected, as Figure 3-7 shows.

Figure 3-7: Manage Servers Properties Page with Rules Tab Selected

Manage Servers / Dynamic Content Servers / **Properties** ?

Properties Rules History

▼ Rules

The following rules describe the server members of this group and do not apply to subgroups. Match: If any rules are met

Search for servers whose:

Agent Reporting Is Minimal
Not Reporting
Ok
Partial

Search

Save Cancel

- 7** Select the criteria that apply to the servers you would like the group to include. See “From the navigation panel, click Servers ► Search. The Search Rules page appears with the Servers tab displayed.” on page 118 in this chapter for more information.



Create as many lines of criteria as required to adequately describe your server group rules by clicking the plus button to add a new line. Conversely, to remove lines of criteria, click the minus button next to the line you want to remove.

- 8** From the Match drop down list select either “If all rules are met” or “If any rules are met.”
- 9** (Optional) Click **Search** to apply the rules to a server search to validate the results.
- 10** Click **Save** to save the rules that apply to your group.

Creating Groups by Using Search

To create groups by using search, perform the following steps:

- 1** From the navigation panel, click Servers ► Search. The Search Rules page appears with the Servers tab displayed.

To search for servers, use this tab. To search for groups, select the Groups tab.

- 2** Use the criteria to create the rules used for a server search and to identify servers for dynamic groups. The options in the user interface for specifying dynamic group rules and for using Search are the same. See “Criteria for Search and Dynamic Group Rules” on page 118 in this chapter for more information.



Create as many lines of criteria as required to adequately describe your server search or your server group rules by clicking the plus button to add a new line. Conversely, to remove lines of criteria, click minus button next to the line you want to remove.

- 3** If you are defining dynamic server group rules, click **Save** to save your rules. You can also click **Search** to use your rules to perform a server search.

Or

If you are doing a server search, click **Search** to perform the server search, or click **Save** to save the search as the rules for a dynamic group.

The New Dynamic Group dialog appears.

Criteria for Search and Dynamic Group Rules

The following table describes the rules that you can use to search for servers or to create dynamic server groups.

Note that anywhere you can enter text, you can enter a wildcard (*) character to broaden your results.

Table 3-2: Rules for Server Search and for Creating Dynamic Groups

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---------------------------|-----------|-------|
| PROPERTIES | | |

Table 3-2: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|--|---|
| Agent Discovery Date: The date that the Opware Agent was installed. | <ul style="list-style-type: none"> Is after Is before | Drop-down lists with the day, month, and year |
| | <ul style="list-style-type: none"> Is within the last | User-entered text |
| | <ul style="list-style-type: none"> Is today | N/A |
| Agent Reporting: Whether the Opware Agent is reporting to Opware SAS. | <ul style="list-style-type: none"> Is Is not | <ul style="list-style-type: none"> Has not reported OK Registration in progress Reporting error |
| Agent Status: Whether the Opware Agent is reachable by Opware SAS. | <ul style="list-style-type: none"> Is Is not | <ul style="list-style-type: none"> Not reachable OK |
| Agent Version: The version of the Opware Agent – such as 14.2.3b. | <ul style="list-style-type: none"> Contains Does not contain Is Is not | User-entered text |
| Custom Attribute (any): The name of a custom attribute that is associated with the server through attachment or inheritance. | <ul style="list-style-type: none"> Contains Is | User-entered text |
| Custom Attribute (local): The name of a custom attribute that is locally attached to the server, | <ul style="list-style-type: none"> Contains Is | User-entered text |
| Customer: The customer or account that the server is associated with. | <ul style="list-style-type: none"> Is Is not | Popup window of customers |

Table 3-2: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|--|--|---|
| Deployment Stage: The stage the server performs within a lifecycle environment. | <ul style="list-style-type: none"> • Is • Is not | <ul style="list-style-type: none"> • In Deployment • Live • Not Specified • Offline <hr/> <p>The values that appear in this list are customizable; in addition to the values above, values specific to your environment might appear.</p> <hr/> |
| Facility: The collection of servers managed by an Opsware SAS installation. | <ul style="list-style-type: none"> • Is • Is not | <p>Popup window of facilities</p> <hr/> <p>When Opsware SAS is running multimaster mode, the list can contain many facilities.</p> <hr/> |
| Group Membership: Whether the server belongs to a group. | <ul style="list-style-type: none"> • Is | <p>Popup window of groups</p> |
| Host name: The host name of the server – such as m004.company.com. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | <p>User-entered text</p> |
| Name (any): This enables searching for any name or IP address associated with a server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | <p>User-entered text</p> |

Table 3-2: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|--|--|
| Notes: The contents of the Notes field from the Properties tab for a server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| Opware Display Name: The user-configurable name for the server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text <hr/> By default, Opware SAS uses the configured host name of the server until a user edits it. <hr/> |
| Server Use: How the server is being used – such as Development, Staging, Production. | <ul style="list-style-type: none"> • Is • Is not | <ul style="list-style-type: none"> • Development • Not Specified • Production • Staging <hr/> You can customize values that appear in this list. In addition to the values above, values specific to your environment might appear. <hr/> |

Table 3-2: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|--|---|---|
| <p>Service Level: The user-defined category that can be used as an organizational tool.</p> | <ul style="list-style-type: none"> • Is • Is attached here or below | <p>Popup window of service levels</p> <hr/> <p>Servers can be associated with multiple service levels.</p> <hr/> <p>See the <i>Opware® SAS User's Guide</i> for more information about working with service levels.</p> |
| OPSWARE PROPERTIES | | |
| <p>Application Configuration: Whether the server uses the Application Configuration feature.</p> | <ul style="list-style-type: none"> • Is not used • Is used | N/A |
| <p>Code Deployment: Whether the server uses the Code Deployment feature.</p> | <ul style="list-style-type: none"> • Is not used • Is used | N/A |
| <p>Configuration Tracking: Whether the Configuration Tracking feature is monitoring or backing up specific files or configurations on a server.</p> | <ul style="list-style-type: none"> • Is off • Is on | N/A |
| <p>Lifecycle: The server states that are part of bringing a server into Opware SAS.</p> | <ul style="list-style-type: none"> • Is • Is not | <ul style="list-style-type: none"> • Available • Build Failed • Deactivated • Installing OS • Managed |

Table 3-2: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|--|--|
| Server ID: The internal ID Opware SAS uses to identify the server. | <ul style="list-style-type: none"> • Is • Is not | User-entered text <hr/> In most cases, the Server ID is the same as the MID. <hr/> |
| SOFTWARE | | |
| Attached Software: The software that is assigned or modeled through the Opware reconcile operation – the installation process. | <ul style="list-style-type: none"> • Is • Is attached here or below | Popup window of software |
| Installed Patches: Whether a patch is installed on the server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| Installed Software: The package reported installed on the server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text <hr/> A package does not have to be installed by Opware SAS to be reported as installed on a server. <hr/> |
| OS Version: The OS version defined by OS definitions in the OS Provisioning feature. | <ul style="list-style-type: none"> • Is • Is not | Popup window of operating systems |
| Reported OS: For Windows – the version reported by the OS, for Unix – the version returned by the uname command. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |

Table 3-2: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|--|-------------------|
| Windows Service: The names of the Windows services that are reported to Opsware SAS. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| NETWORK | | |
| DNS Search Domains: The domains configured to be searched in the server's network settings. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| DNS Servers: The IP addresses of the DNS servers configured in the server's network settings. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| Default Gateway: The IP address of the default router. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| IP Address: Any Internet Protocol address for the server. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| MAC Address: Any Media Access Control address, which is the network interface card's unique hardware number. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |

Table 3-2: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---|---|---|
| WINS Servers: The Windows Internet Naming Servers configured in the server's network settings. | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | User-entered text |
| HARDWARE | | |
| CPU Make and Model: The vendor name and CPU model for the server – such as GENUINEINTEL Intel(R) Pentium(R) 4 CPU 2.60GHz. | <ul style="list-style-type: none"> • Is • Is not | Popup window of CPU makes and models |
| CPU Speed: The Central Processing Unit speed in gigahertz [GHz]. | <ul style="list-style-type: none"> • Does not equal • Equals • Is greater than • Is less than | User-entered text <hr/> A 600 Mhz machine should be entered as 0 . 6. <hr/> |
| Make and Model: The vendor name and server model for the server – such as Compaq - DL360. | <ul style="list-style-type: none"> • Is • Is not | Popup window of server makes and models |
| Number of CPUs: The number of CPUs on the server. | <ul style="list-style-type: none"> • Does not equal • Equals • Is greater than • Is less than | User-entered text |
| RAM: The amount of RAM on the server in megabytes [MB]. | <ul style="list-style-type: none"> • Does not equal • Equals • Is greater than • Is less than | User-entered text <hr/> To enter 1 Gigabyte, type 1024. <hr/> |

Table 3-2: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|--|---|---|
| <p>Serial Number: The serial number of the server.</p> | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | <p>User-entered text</p> |
| <p>Storage Make and Model: The vendor name and storage model for the server – such as WDC - WD800BB-75DKA0.</p> | <ul style="list-style-type: none"> • Is • Is not | <p>Popup window of storage makes and models</p> |
| CUSTOM FIELDS | | |
| <p>A Numeric field</p> | <ul style="list-style-type: none"> • Does not equal • Equals • Is greater than • Is less than | <p>User-entered text</p> |
| <p>A String field</p> | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | <p>User-entered text</p> |
| <p>A URI field</p> | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | <p>User-entered text</p> |
| <p>A File field</p> | <ul style="list-style-type: none"> • Contains • Does not contain • Is • Is not | <p>User-entered text</p> |

Table 3-2: Rules for Server Search and for Creating Dynamic Groups (continued)

| RULE NAME AND DESCRIPTION | OPERATORS | VALUE |
|---------------------------|----------------------|---|
| A Date field | • Is after | Drop-down lists with the day, month, and year |
| | • Is before | |
| | • Is within the last | User-entered text |
| | • Is today | N/A |

Adding a Server to Static Groups, Both Public and Private

To add a server or a server group to a static group, perform the following steps:

- 1** From the navigation panel, click Servers ► Manage Servers. The Manage Servers page appears.
- 2** Navigate to the static group that you want to add servers and server groups to.
- 3** Click the check box next to the group and select **Edit ► Add Servers**. The Select Servers and Groups to Add to [group name] window opens. The window is populated with the same servers and groups visible on the Manage Servers page. You can use the Status, OS, and Customer filters to change the servers and groups that appear on the list.



When you select **Edit ► Add Servers** and then select a group, the servers in the group are added. The group itself is not added. If you want to add a group to a group, select **Edit ► Copy to Group** from the menu.

- 4** Click the check box next to the servers and groups you want to add, and click **Add** at the bottom of the window.

Adding a Server to Dynamic Groups, Both Public and Private

Servers are added to dynamic server groups automatically, based on the rules created for the group. To change the membership of a dynamic group, add, delete, or update the dynamic group rules.

To update the rules for a dynamic group, perform the following steps:

- 1** From the navigation panel, click Servers ► Manage Servers. The Manage Servers page appears.
- 2** Navigate to the Properties page of the dynamic group that you want to update rules for; navigate by clicking the group links in Manage Server list.
- 3** Select the Rules tab.
- 4** Click **Edit**.
- 5** To add criteria to the existing rules, click the plus (+) button next to existing criteria. The fields for the criteria appear. Enter the values for the rule. See “From the navigation panel, click Servers ► Search. The Search Rules page appears with the Servers tab displayed.” on page 118 in this chapter for descriptions of these values.
- 6** To delete criteria, click the minus (-) button next to the criteria you want to delete from the rules. The criteria are removed from the page.
- 7** Click **Save**.

Groups can be added manually to a dynamic group, in either of the following two ways:

- 1** Click the check box next to the name of the server group to which you want to add a server group, and select **Edit ► Copy to Group**. Then follow the steps for creating a new group, either static or dynamic. See “Ways to Create Server Groups” on page 113 in this chapter for more information.

Or

Navigate to the server group that you want to add a server group to.

- 2** Click the check box next to “This Group” and select **Resource ► New Group**. Then follow the steps for creating a new group, either static or dynamic. See “Ways to Create Server Groups” on page 113 in this chapter for more information.

Removing Servers from Static Groups, Both Public and Private

To remove a server from a static server group, perform the following steps:



A server does not need to be deactivated to be removed from a group. You can remove a server from a Static group at any time. The Membership tab for a server displays all the server groups of which the server is a member.

- 1 From the navigation panel, click Servers ► Manage Servers. The Manage Servers page appears.
- 2 Navigate to the static server group that you want to remove servers from until you reach the level where the servers in question are located.



Servers can belong to more than one group, so if you want to remove a server from each group it belongs to, you must locate and remove each instance of the server from all groups.

- 3 Click the check box next to the servers you want to remove and select **Edit ► Remove from Group**.

As soon as you select that menu option, the server is removed and the screen refreshes to display the current members of the group.

Removing Servers from Dynamic Groups, Both Public and Private

Servers are removed from dynamic server groups automatically, based on the rules created for the group. To remove servers, you can create or update rules. To add rules to an existing server group, perform the following steps:

- 1 From the navigation panel, click **Servers ► Manage Servers**. The Manage Servers page appears.
- 2 Navigate to the Properties page of the dynamic group that you want to add rules to, by drilling down into the group and clicking the “This Group” link.
- 3 Select the Rules tab.
- 4 Click **Edit**.
- 5 Add one or more lines of criteria to the existing rules by following the procedure described in the “Creating Dynamic Server Groups by Using the New Group Option” topic, step 3 on page 118 or see “From the navigation panel, click Servers ► Search. The Search Rules page appears with the Servers tab displayed.” on page 118 in this chapter for more information.

Moving Servers from One Static Group to Another

The process of moving servers from one group to another is similar to copying servers from one group to another with the exception that with a move, the servers do not remain in the original group.



The **Edit ► Move** menu option is disabled when you select servers in a dynamic group. To move servers in a dynamic group, choose the **Edit ► Copy to Group** menu command. This menu command is also disabled for public groups when you do not have permission to manage public groups.

To move servers from one group to another, perform the following steps:

- 1** From the navigation panel, click **Servers ► Manage Servers**. The Manage Servers page appears.
- 2** Navigate to the group containing the servers you want to move and select those servers.
- 3** From the **Edit** menu, choose **Move**. The Move [group name] Group dialog box appears.
- 4** Navigate to the place in the group hierarchy where you want to create the group by clicking the group links displayed in the Name and Type field.
- 5** Click **Move**.

Duplicating Server Groups

Duplicating a group is similar to copying a group. To duplicate an existing group, perform the following steps:

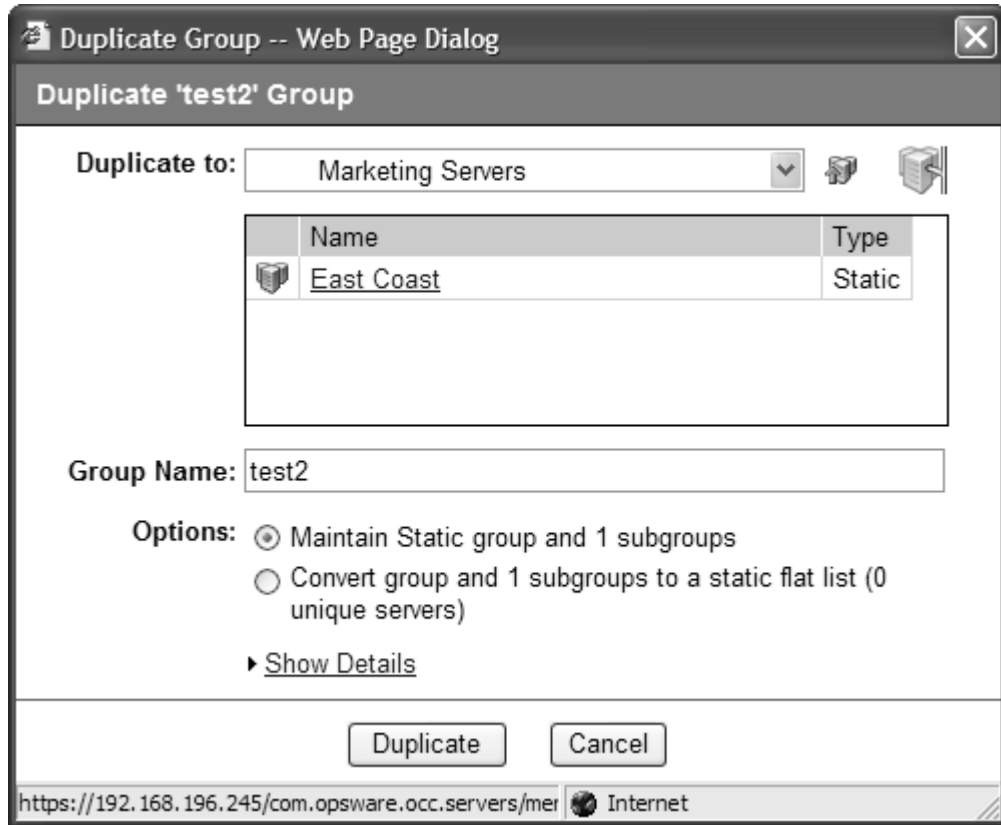
- 1** Navigate to the group that you want to duplicate.
- 2** Click the check box next to the group name.



You can only select one group at a time to duplicate.

- 3 From the **Edit** menu, choose **Duplicate Group**. The Duplicate [group name] Group dialog appears as shown in Figure 3-8.

Figure 3-8: Duplicate Group Dialog Box



- 4 Select the location for the newly-duplicated group using the Duplicate In drop down list.



The Duplicate In drop down list's default location is the location of the group you select to duplicate. Navigate through the hierarchy of groups to find the location where you want the newly-duplicated group to reside.

- 5 Enter the name of the new group.
- 6 The values on the Options field will vary depending on whether the group duplicated is static or dynamic.

- 7** If the group is static, select one of the following options:
 - **Maintain Static group and <number> subgroups:** Selecting this option copies the group and its hierarchy as is.
 - **Convert group and <number> subgroups to a static flat list:** Selecting this option copies the group and flattens the hierarchy.
- 8** If the group is dynamic, select one of the following options:
 - **Maintain rules & group hierarchy:** Selecting this option copies the group and leaves it as a dynamic group, it also copies any subgroups.
 - **Convert to static group, maintain hierarchy:** Selecting this option copies the group, but turns it into a static group with the current servers defined by the rules and also copies any subgroups.
 - **Convert group and <number> subgroups to a static flat list:** Selecting this option copies the group, and turns it into a static group and flattens any hierarchy.
- 9** Click **Duplicate**. The newly-duplicated group appears in the selected destination.

Rules for Deleting Server Groups

It's important to note the distinction between removing servers from a group, deleting a server, and deleting a group.

Removing servers from a group only removes the server from the selected static group, but the server itself remains in the global list of servers and is still managed by Opware SAS.

Deleting a server can only be performed on a server whose status is deactivated. Selecting the Delete Server option completely removes the server from within Opware SAS, although its history remains.

Deleting a server group removes the group, but the servers in the group still remain in the list of servers and in any other groups for which they are members.

A group cannot be deleted when any of the following conditions apply:

- Software is attached to the group or a subgroup of the group.
- Access control boundaries are attached to the groups or a subgroup of the group.
- Servers and groups are selected together for deletion.

Deleting a Server Group

To delete a server group, perform the following steps:

- 1** Click the check box next to the server group you want to delete.
- 2** From the **Edit** menu, choose **Delete Group**. A confirmation message appears, detailing the number of servers and server groups in the server group that you want to delete.
- 3** Click **OK** to complete the deletion of the server group.

The screen refreshes, showing the list of servers and groups without the deleted group of servers.

Server Life Cycle

This section provides information on the server life cycle within Opsware SAS and contains the following topics:

- OS Provisioning and the Server Life Cycle
- Server Properties
- Server Management Tasks Related to the Server Life Cycle
- Changing the Use and Stage Values for Servers
- Editing the Properties of a Server
- Tasks Associated with Deactivating a Server
- Deactivating a Server
- Deleting a Server from Opsware SAS
- Overview of Cloning a Server

OS Provisioning and the Server Life Cycle

Opsware SAS is designed to enable multiple teams to work together to provision servers. The OS Provisioning feature allows IT teams to separate the tasks of readying servers for provisioning (such as mounting servers in racks and connecting them to 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 Opware build network. Next, they boot the server for the first time by using an Opware 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 the target OS installed and might not have access to disk resources.

During OS provisioning, servers progress through the Opware SAS life cycle state changes:

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

Table 3-3 describes the Opsware SAS server life cycle values.

Table 3-3: Opsware SAS Life Cycle Values for Servers

| OPSWARE LIFE CYCLE VALUE | DESCRIPTION |
|--------------------------------|---|
| Server Pool Values | |
| Available | <p>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.</p> <p>The OS Build Agent is a small agent that can run in the memory of the bare-metal server.</p> <p>See “Operating System Provisioning” on page 331 in Chapter 8 for more information.</p> |
| Installing OS | <p>Indicates that a user is installing the target OS on the server.</p> <p>The server stays in the Server Pool list until the installation process finishes successfully. Then the server moves to the Manage Servers list.</p> <p>See “Installing an OS by Using a Template” on page 336 in Chapter 8 for more information. See “Installing an OS by Using a Custom Installation” on page 341 in Chapter 8 for more information.</p> |
| Build Failed | <p>Indicates a server on which the OS Build Agent was installed and is running, but the installation of the target OS failed.</p> <p>The server remains in the Server Pool list with this status for 7 days before Opsware SAS deletes the entry.</p> <p>See “Recovering when an OS Installation Fails” on page 343 in Chapter 8 for more information.</p> |

Table 3-3: Opsware SAS Life Cycle Values for Servers (continued)

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

Table 3-4 shows which server icons appear in Opsware SAS and explains what each icon indicates in regard to the server life cycle.

Table 3-4: Server Icons in Opsware SAS




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

Table 3-4: Server Icons in Opware SAS












| SERVER ICON | DESCRIPTION |
|---|--|
|  | <p>Indicates a server that the SAS Web Client is managing and that Opware SAS can communicate with. An Opware Agent is running on the server.</p> <p>Appears in the My Jobs panel in the home page, in the list in the My Jobs page, in the Manage Servers list, and in the server lists in the Opware wizards.</p> |
|  | <p>Indicates a server that is scheduled for an operation (install software, uninstall software, and so forth).</p> <p>Appears in the My Jobs panel in the home page and in the list in the My Jobs page.</p> |
|  | <p>Indicates a managed server that Opware SAS cannot communicate with (it is Not Reachable) because the Opware Agent on the server cannot connect to Opware SAS.</p> <p>If you want to discover reasons why the managed server is unreachable, you can run a Communication Test. See the <i>Opware[®] SAS User's Guide: Server Automation</i> for information.</p> <p>Appears in the Manage Servers list.</p> |
|  | <p>Indicates a managed server on which an error occurred while Opware SAS was installing or uninstalling software.</p> <p>Appears in the My Jobs panel in the home page and in the list in the My Jobs page.</p> |
|  | <p>Indicates a managed server on which a warning occurred while Opware SAS was installing or uninstalling software.</p> <p>Appears in the My Jobs panel in the home page and in the list in the My Jobs page.</p> |
|  | <p>Indicates a server that was deactivated in Opware SAS so that it is currently not managed and not reachable.</p> <p>Appears in the Manage Servers list and in the server lists in the Opware wizards (however, it is not selectable in the wizards).</p> |

Table 3-4: Server Icons in Opware SAS

| SERVER ICON | DESCRIPTION |
|---|--|
|  | <p>Indicates a managed server on which the configuration file on the server is out of sync with the Application Configuration Template (Opware model).</p> <p>Appears only in the Application Configuration feature and the server list in the SAS Client.</p> |
|  | <p>Indicates a static group of servers. The same states that apply to single servers apply to groups.</p> <p>See “Server Groups” on page 104 in this chapter for information about the different types of group of servers.</p> |
|  | <p>Indicates a dynamic group of servers. The same states that apply to single servers apply to groups.</p> |
|  | <p>Indicates a public and static group of servers. The same states that apply to single servers apply to groups.</p> |
|  | <p>Indicates a public and dynamic group of servers. The same states that apply to single servers apply to groups.</p> |

Server Properties

Figure 3-9 shows the Server Properties columns. Table 3-5, Table 3-6, and Table 3-7 describe the Status, Stage, and Use properties for managed servers.

Figure 3-9: Server Properties Columns in the Manage Servers List

Manage Servers (Summary View) ?

All Statuses ▾
All Operating Systems ▾
All Stages ▾
All Uses ▾

All Customers ▾
All Manufacturers ▾
All Models ▾

Managed & Deactivated ▾
Update

Resource Edit View Tasks Configuration Tracking

| | Name [▲] | Host Name / IP Address | OS Version | Stage | Use | Customer |
|--------------------------|----------------------|--------------------------------------|-------------------------------|---------------|---------------|--------------|
| Public Groups | | | | | | |
| <input type="checkbox"/> | core.tr3.opsware.com | core.tr3.opsware.com 172.16.36.18 | Red Hat Enterprise Linux AS 3 | Not Specified | Not Specified | Opsware |
| <input type="checkbox"/> | m201.tr3.opsware.com | mh03 172.16.36.34 | Windows 2000 | Not Specified | Not Specified | Not Assigned |

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

Status (short for Agent Status) is set automatically by Opware SAS.


Opware SAS toggles each server between OK and Not Reachable by reachability checks.

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

Table 3-5: Values for the Status Property for Managed Servers

| STATUS VALUE | DESCRIPTION |
|--------------|--|
| OK | Server is manageable by Opware SAS. Represented as text (OK) in the properties page for a server. Represented as an icon in the Manage Servers and Server Pool lists: <div style="text-align: center; margin-top: 10px;"> </div> |

Table 3-5: Values for the Status Property for Managed Servers (continued)

| STATUS VALUE | DESCRIPTION |
|---------------|---|
| Not Reachable | <p>Server is unreachable by Opware SAS due to an error (for example, it cannot connect to the Opware core); automatically set by Opware SAS. Represented as text (Not Reachable) in the properties page for a server. Represented as an icon in the Manage Servers list:</p> <div style="text-align: center;">  </div> <p>If you want to discover reasons why the managed server is unreachable, you can run a Communication Test. See the <i>Opware® SAS User's Guide: Server Automation</i> for information.</p> |

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 Opware administrator can change the values for the Stage property. By default, Opware SAS is installed with the following Stage values.

Table 3-6: 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 Opware 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, Opware SAS is installed with the following Use values. Except for the Staging, Production, and Not Specified values, an Opware administrator can change the default values. The CDR feature depends on the Staging and Production values. Therefore, these default values cannot be changed or deleted.

Table 3-7: 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 Opware core servers). |
| Staging | A staging server for production. |

Server Management Tasks Related to the Server Life Cycle

Managing servers in Opware SAS involves the following standard tasks:

- Bringing a new server into Opware SAS so that it appears in the Server Pool
See the *Opware® SAS User's Guide: Application Automation* for information.
- Installing an operating system on a server
See “Ways to Install Operating Systems on Servers” on page 335 in Chapter 8 for more information.
- Installing a patch
See the *Opware® SAS User's Guide: Application Automation* for information.
- Installing software
See the *Opware® SAS User's Guide: Application Automation* for information.
- Reprovisioning a server with a new OS
See “Reprovisioning a Solaris or Linux Server” on page 347 in Chapter 8 for more information about reprovisioning a server by using the SAS Web Client.
See the *Opware® SAS User's Guide: Application Automation* for information using the SAS Client to reprovisioning a 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 Opware SAS and Opware-managed servers by using this feature.



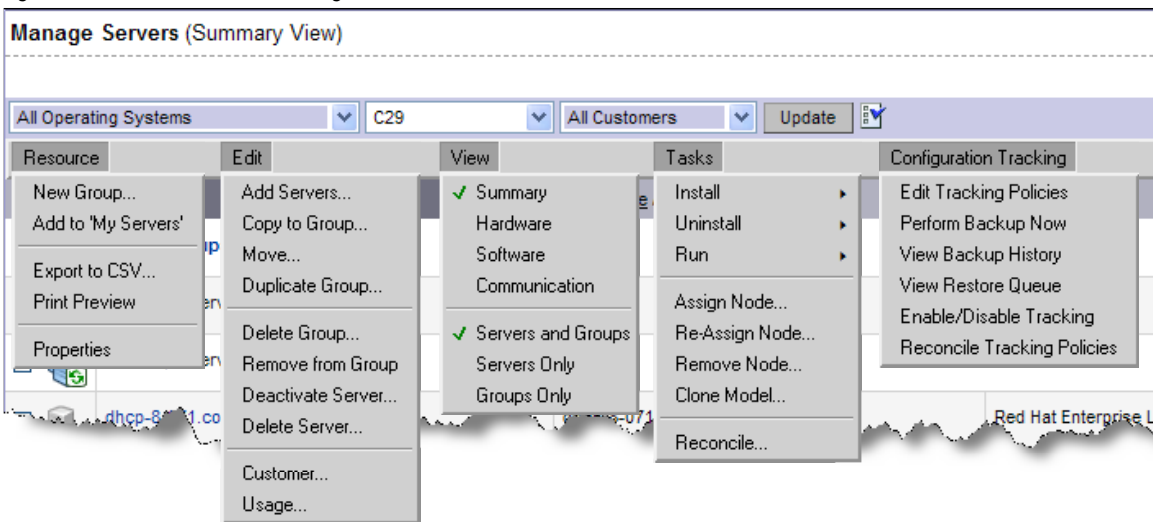
You cannot reprovision a Linux server so that it runs a Windows OS.

- Deactivating a managed server so that Opware SAS no longer manages it.

See “Tasks Associated with Deactivating a Server” on page 145 in this chapter for more information.

You accomplish server management tasks by using the menus in the Manage Servers list, as Figure 3-10 shows.

Figure 3-10: Menus in the Manage Servers List



See “Software Provisioning” on page 291 in Chapter 6 for more information.

See “Reconcile” on page 349 in Appendix A for more information.

See the *Opware® SAS User's Guide: Server Automation* for information about how to run a distributed script on a server.

See “Manage Servers List” on page 59 in Chapter 2 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 ► Manage Servers. The Manage 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 “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 for more information.

- 2 Select the servers that you want different Use or Stage values for.
- 3 Choose **Edit ► Usage** from the menu above the Manage Servers list. A window prompts you to select different values, as Figure 3-11 shows.

Figure 3-11: Edit Server Popup Window

Edit Servers

Select the new values

| Server | Facility | Stage | Use |
|-------------------------------------|-------------|---|---|
| 172.16.36.18 - core.tr3.opsware.com | DATACENTER1 | current: Not Specified Not Specified ▼ | current: Not Specified Not Specified ▼ |

- 4 Select the Use and Stage values from the lists.
- 5 Click **Save Changes**. The window closes and the Manage 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 SAS Web Client 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 Opsware SAS views it changes. Perform the following steps to edit the properties of a server:

- 1 From the navigation panel, click Servers ► Manage Servers. The Manage 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 “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 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 SAS Web Client, 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 Opware administrator defines the options for customer selections. Contact your Opware 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 the *Opware® SAS Configuration Guide* for more information.

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

See “Server Properties” on page 139 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 the *Opware® SAS User’s Guide: Server Automation* for information.

- 4** To save your changes, click **Save**.



To change the custom attributes of the server, select the Custom Attributes tab. The Manage Servers: Custom Attributes page appears. See “Managing Custom Attributes” on page 161 in this chapter for more information.

Tasks Associated with Deactivating a Server

You will want to deactivate a server when Opsware SAS 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 feature.

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 the *Opsware[®] SAS User's Guide: Server Automation* for 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 the *Opsware[®] SAS Configuration Guide* for more information.

Deactivating a Server

Perform the following steps to deactivate a server:

- 1** From the navigation panel, click Servers ► Manage Servers. The Manage 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 “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 for more information.

- 2** Select the servers that you want to deactivate.

- 3 Choose **Edit ► Deactivate Server** from the menu above the Manage Servers list. A confirmation dialog box prompts you to confirm the deactivation.
- 4 Click **OK**. The Manage Servers list refreshes and the server appears with a deactivated icon.

Deleting a Server from Opsware SAS

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



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

When you delete a server from Opsware SAS, 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 ► Manage Servers. The Manage 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 “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 for more information.
- 2 Select the servers that you want to delete.
- 3 Choose **Edit ► Delete Server** from the menu above the Manage Servers list. A confirmation dialog box prompts you to confirm the deletion.
- 4 Click **OK**. The Manage Servers list refreshes and the server disappears from the list.

Overview of Cloning a Server

The SAS Web Client includes a feature that allows users to copy a server. Copying (referred to as cloning in the SAS Web Client) 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 Opware SAS.

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.

Cloning a Server

Perform the following steps to clone a server:

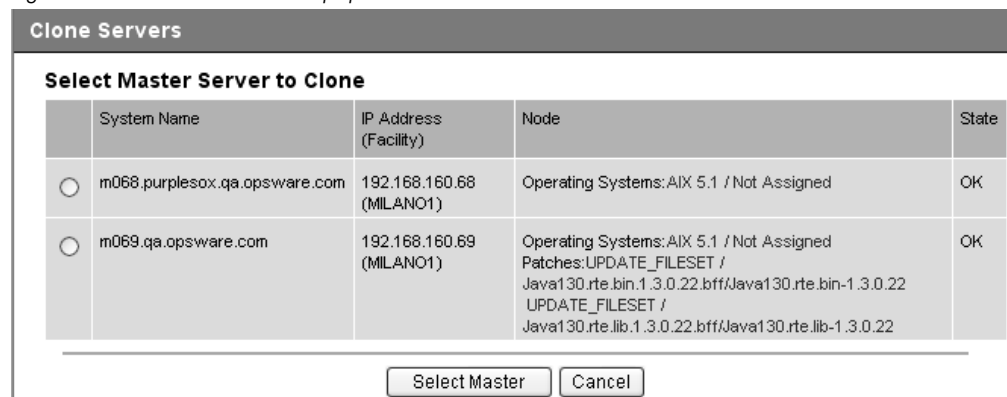
- 1** From the navigation panel, click Servers ► Manage Servers. The Manage Servers page appears. Browse the list to find the servers that you want to clone.
- 2** Select two or more servers. When copying servers, you must select at least two servers so that one server can be the master. The other servers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

Opware, Inc. recommends that you select servers with similar hardware architecture (based on the value in the Reported OS column in the Manage Servers list).

- 3** Choose **Tasks ► Clone** model from the menu above the Manage Servers list.

A window prompts you to select the master server, as Figure 3-12 shows.

Figure 3-12: Clone Servers Popup Window in the SAS Web Client



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

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

- 5 Click **Commit Clone**.

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 Opsware SAS to install the software on the target servers. See “Reconcile” on page 349 in Appendix A for information about reconciling servers.

Server Management Jobs

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

- My Jobs
- My Jobs Display Information
- My Jobs in the SAS Web Client
- My Jobs in SAS Client
- Viewing Job Details in the SAS Web Client
- Viewing My Job Details from inside the SAS Client

My Jobs

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 SAS Web Client:

- A panel on the SAS Web Client 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

Opsware SAS 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 SAS installation at your organization.)

My Jobs Display Information

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, select Job from the drop-down list and enter a Job ID and click **Go**.

My Jobs in the SAS Web Client

The My Jobs feature in the SAS Web Client 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

My Jobs in SAS Client

My Jobs feature in the SAS Client displays job logs about the following SAS Client jobs:

- Creating snapshots
- Pushing or Auditing an Application Configuration
- Auditing Servers
- Creating Server Snapshots
- Creating, Installing, and Reconciling Packages
- Any jobs scheduled to be run at a future date

My Jobs in the SAS Client appears in the Jobs Logs feature window.

Viewing Job Details in the SAS Web Client

Perform the following steps to view the job details:

- 1 From the SAS Web Client home page, click the link in the My Jobs panel for the job that you want to view, as Figure 3-13 shows.

Figure 3-13: My Jobs Panel in the SAS Web Client Home Page

| My Jobs | | | | | See All (22) |
|---|--------------------------|---------|--------|-------------------------|--------------|
| Name | Start Time | Servers | Groups | Status | |
|  Run Custom Extension | (not set) | 4 | 0 | Cancelled | |
|  Run Script | Wed Oct 31 21:45:00 2007 | 4 | 0 | Scheduled | |
|  Run Script | Tue May 31 21:22:06 2005 | 3 | 0 | Completed with errors | |
|  Audit Servers [Launch OCC Client] | Tue May 31 21:00:48 2005 | 2 | 0 | Completed | |
|  Create Snapshot [Launch OCC Client] | Tue May 31 20:57:05 2005 | 1 | 0 | Completed | |
|  Audit Servers [Launch OCC Client] | Tue May 31 20:50:09 2005 | 2 | 0 | Completed with warnings | |

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 3-14 shows.

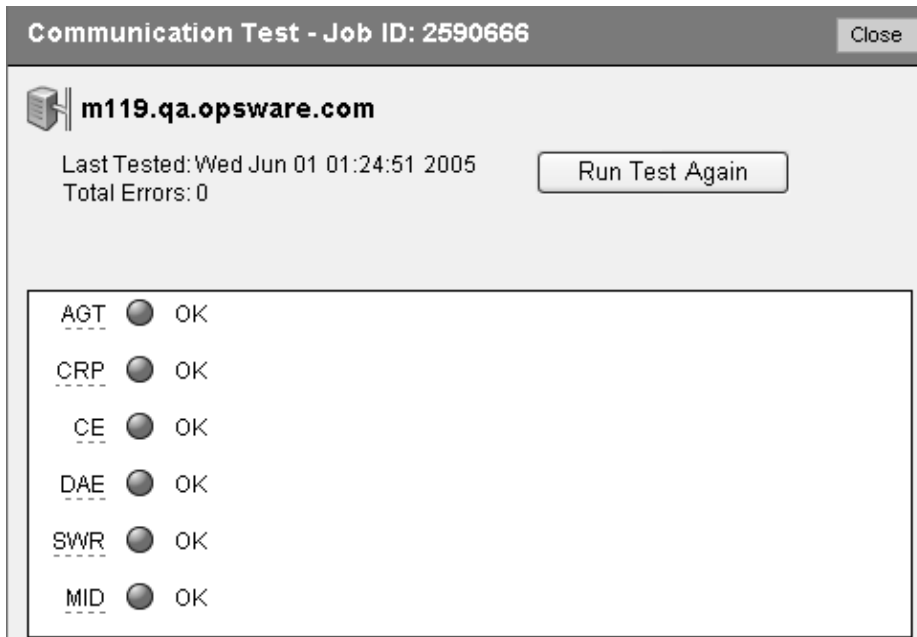
Figure 3-14: My Jobs Page Accessed from the Navigation Panel

| < Job ID > All Job Types No Time Restrictions All Job Status Update | | | | | | |
|---|---------|---------------------------------------|--------------------------|---------|--------|-----------------------|
| 30 Total | | | | | | |
| | Job ID | Job Type | Start Time ^ | Servers | Groups | Status |
| | 1440666 | Run Custom Extension | (not set) | 4 | 0 | Cancelled |
| | 1450666 | Run Script | Wed Oct 31 21:45:00 2007 | 4 | 0 | Scheduled |
| | 1700666 | Audit Servers [Launch OCC Client] | Tue May 31 22:29:54 2005 | 2 | 0 | Completed with errors |
| | 560667 | Audit Servers [Launch OCC Client] | Tue May 31 22:19:18 2005 | 2 | 0 | Completed with errors |
| | 550667 | Audit Servers [Launch OCC Client] | Tue May 31 22:16:50 2005 | 2 | 0 | Completed with errors |
| | 1630666 | Audit Servers [Launch OCC Client] | Tue May 31 22:12:26 2005 | 2 | 0 | Completed with errors |
| | 1580666 | Audit Servers [Launch OCC Client] | Tue May 31 22:04:27 2005 | 2 | 0 | Completed with errors |
| | 1570666 | Audit Servers [Launch OCC Client] | Tue May 31 22:00:17 2005 | 2 | 0 | Completed with errors |
| | 1560666 | Audit Servers [Launch OCC Client] | Tue May 31 21:58:35 2005 | 2 | 0 | Completed with errors |
| | 1550666 | Audit Servers [Launch OCC Client] | Tue May 31 21:57:33 2005 | 2 | 0 | Completed with errors |
| | 1460666 | Run Script | Tue May 31 21:22:06 2005 | 3 | 0 | Completed with errors |
| | 1240666 | Audit Servers [Launch OCC Client] | Tue May 31 21:00:48 2005 | 2 | 0 | Completed |
| | 1230666 | Create Snapshot [Launch OCC Client] | Tue May 31 20:57:05 2005 | 1 | 0 | Completed |

The My Jobs page displays the operations that you performed.

Click **View Details** 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 Opware SAS generates. See Figure 3-15.

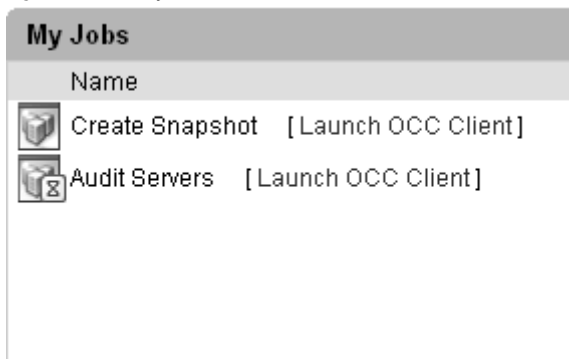
Figure 3-15: Communication Test Details Page in the My Job Window



Viewing My Job Details from inside the SAS Client

You can also view job details from inside the SAS Client. If the job was run from inside the SAS Client, you will see a link next to the job named Launch SAS Client, as shown in Figure 3-16.

Figure 3-16: My Jobs with link to SAS Client



To launch the SAS Client and view the job details, click the link.

Server Management Scheduling and Notification

This section contains the following topics:

- Scheduling a Job
- Sending Email Notification

The time used for the scheduled job is specified in the user's preferred time zone (which can be modified in My Profile). If the user does not have a preferred time zone set, the time zone is derived from the Opware SAS core server (usually UTC).

Scheduling a Job

Perform the following steps to schedule server management tasks:

- 1** In the Schedule and Notify page of an Opware Wizard, choose the Run Now option to execute the operation immediately or choose the Specify Time Option to schedule the operation at a later date and time. See Figure 3-17



The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opware SAS core server (usually UTC).

- 2** When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:
 - **Option 1:** Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.

- **Option 2:** Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

Figure 3-17: Scheduling a Job in an Opsware Wizard

Run Custom Extension ?

1 Select Extension
2 Select Servers
3 Specify Settings
4 Confirm Settings
5 **Schedule and Notify**
6 Run Extension

Schedule and Notify

Select the time and date for this job to run. You can choose to be notified after the job runs.

Schedule: Run Now
 Specify Time
 March 07 2005 23 45

Notify: None
 Condition Recipients
 on any result ssunderrajan@opsware.com

Options: Run the job on the servers from the Groups based on the time of Confirm Selections.
 At the scheduled time, update the list of servers from the groups (0) selected prior to running the job.

Cancel < Previous Run >

- 3** Additionally, you can view a scheduled job in the My Jobs page and change the date and 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.)

Sending 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 Opsware SAS.

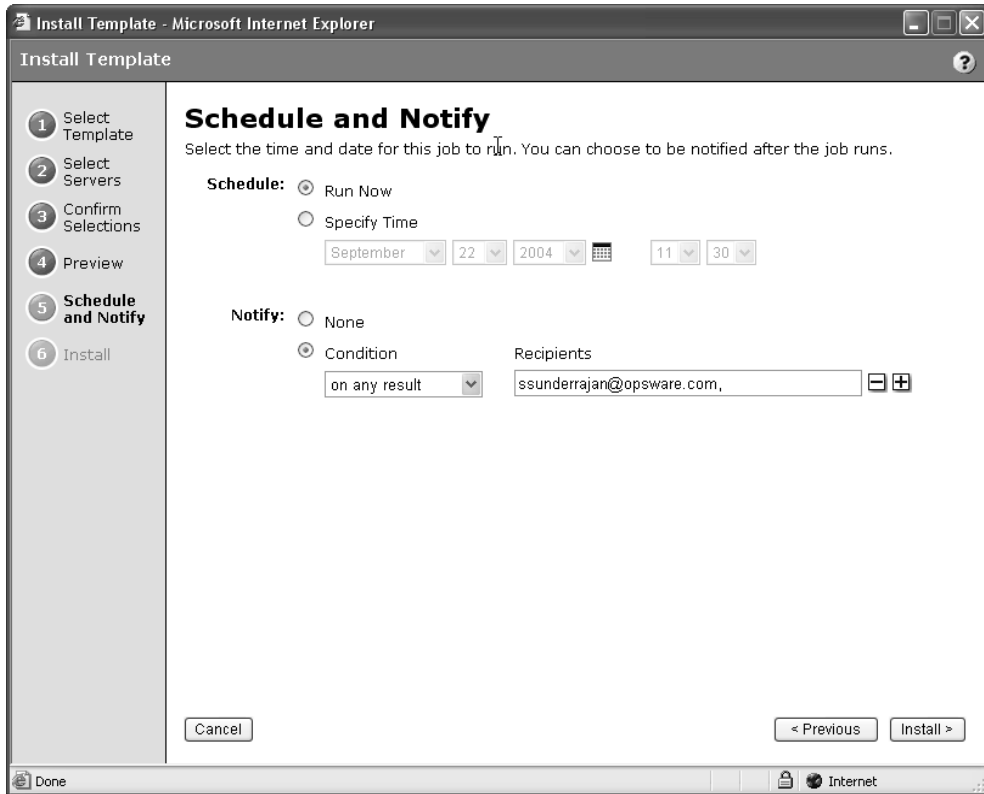
You have the option of sending the notification

- 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 (+) button next to the Recipients field. See Figure 3-18.

Figure 3-18: Notifying about a Job in the SAS Web Client



Time-Outs for Server Management Jobs

Table 3-8 describes the time-out values that apply to the server management operations in Opware SAS.

Table 3-8: Time-Outs in Opware SAS

| TIME OUT (MINUTES) | OPSWARE SAS OPERATION |
|--------------------|--|
| 420 (7 hours) | Reconciling software (installation and uninstallation). |
| 2 | Starting Command Engine sessions in response to a command. If the Opware Agent does not start executing the command within this time, the command will time out and the Command Engine script will continue. |
| 30 | <p>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 Opware Agent.</p> <p>If the Opware 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.</p> <p>Opware SAS polls the Opware Agent every 15 minutes and if the Opware Agent fails to respond two consecutive times, the command will time out.</p> |

Customizing the 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 Opware 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.

Custom Fields for Servers

This section provides information on custom fields for servers within Opsware SAS and discusses the following topics:

- Overview of Custom Fields for Servers
- Changing the Value for a Custom Field

Overview of Custom Fields for Servers

In Opsware SAS, you can store custom server data that is specific to your operational environment. These fields were created specifically for your Opsware SAS installation – all servers in an Opsware SAS installation contain the same number and type of custom fields. Custom fields can contain files, URLs, text strings, numbers, and dates.

In addition to adding files, URLs, text strings, numbers, and dates, you can use custom fields to search for servers based on a value stored in a custom field. You can also use a custom field as criteria to create a dynamic group of servers.



In order for the custom fields to appear in the Manage Server: Properties page, you will have to initially create a custom field. To create a custom field, you will need to install the `custFields.py` Custom Extension which is available only from the Content Starter Pack. Contact your Opsware Technical Support for assistance in installing the `custFields.py` Custom Extension in Opsware SAS.

Typed Data Supported for Custom Fields

The custom fields designated for your operational environment will vary. However, the custom fields created for your environment support values with the following data types:

- **Number:** The value provided must be a long integer.
- **Short String:** A text string that must be less than or equal to 4000 characters.
- **Long String:** A text string with no length restrictions. Custom fields with this type cannot be used in search or in the rules for dynamic group of servers.

- **Date:** The value will be verified for correctness.
- **File:** Indicates a file attachment.
- **URI:** A Uniform Resource Identifier string, which is validated as a URI.

Opware SAS validates the value you enter in a custom field based on the data type specified for the field.

Uses of Custom Fields

The custom fields designated for your operational environment will vary. However, you can use custom fields to accomplish any of the following goals:

- To store the date of patch installation
- To assign a severity rating between 1 and 10 to a Hotfix
- To store an ID from an internal bug tracking system with its associated patch
- To store a JPEG image of the back of a server and have that JPEG associated with that server
- To store a Microsoft Word document describing the disaster recovery steps for a server or group of servers
- To search for a server based on the value of a custom field

Changing the Value for a Custom Field

- 1** From the navigation panel, click Servers ► Manage Servers. The Manage Servers page appears. Browse the list to find the server whose custom fields you want to change.

Or

Search for the server whose custom fields you want to change.

- 2 Click the server name. The Manage Servers: Properties page appears. Scroll to the Custom Fields section of the page. The custom fields that appear in the properties page are specific to your operational environment. Figure 3-19 shows an example of the types of custom fields that can appear for a server.

Figure 3-19: Custom Fields That Appear for a Server

| CUSTOM FIELDS | |
|------------------------|---|
| Date to be installed | 04 October 2004 [Calendar Icon] 0 00 TVT Unset |
| External OSS system ID | CQ12345 ... Enter up to 4000 characters in this field |
| Microsoft bulletin | http://www.microsoft.com/bulletin/hotfixes/Q23909.html [Edit] Unset |
| Severity | 4 Please enter an integer with a maximum of 18 digits |
| Support document | support.doc This is the support document for this patch. Remove |
| Save Cancel | |

- 3 To change a value in a field requires a number or text, enter the value in the field.
- 4 To add a file, click **Browse** and select the file from the dialog box.
- 5 To remove a file, click **Remove**.
- 6 To specify a URL, enter the URL in the field or click **Edit** to change an existing URL.
- 7 To add or change a date, click the [Calendar Icon] icon and select the date or enter the date in the field by selecting the appropriate day, month, and year from the drop-down lists.
- 8 Click **Save**.

Custom Attributes for Servers

This section provides information on custom attributes for servers within Opware SAS and contains the following topics:

- Overview of Custom Attributes for Servers
- Managing Custom Attributes
- Adding Server Custom Attributes
- Editing Server Custom Attributes
- Deleting Server Custom Attributes

Overview of Custom Attributes for Servers

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 SAS Web Client 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 3-20 shows.

Figure 3-20: Custom Attributes Set for a Server

Return to Manage Servers

| Properties | Network | Membership | Attached Nodes | Installed Packages | Custom Attributes | Config Tracking | History |
|--|------------------------------|------------|----------------|--------------------|--|-----------------|---------|
| <input type="button" value="New"/> <input type="button" value="Edit"/> <input type="button" value="Delete"/> <input type="button" value="Revert"/> | | | | | | | |
| Name | Value | | | | Source | | |
| OPSW.reprovision_device_attributes_to_preserve | kernel_arguments reboot_c... | | | | Operating Systems / Red Hat Enterprise Linux AS 3 / Not Assigned (Operating Systems) | | |
| <input type="button" value="Save"/> <input type="button" value="Cancel"/> | | | | | | | |

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 Manage Servers list. After locating the server and displaying the server properties, select the Custom Attributes tab. The SAS Web Client displays the currently defined custom attributes for the selected server.

When you add or edit server custom attributes using the SAS Web Client, Opsware SAS removes leading and trailing whitespace characters from custom attribute values.

See “Adding Server Custom Attributes” on page 162 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, select 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 the *Opware® SAS Configuration Guide* 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. Select 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-specific 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 group of servers by viewing a group, and then selecting the Custom Attributes tab for that group. The procedure to add custom attributes to a group of servers is the same as adding custom attributes to individual servers.

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 ► Manage Servers. The Manage 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 Manage Servers: Server Properties page appears.
- 3** Select the Custom Attributes tab. The Manage Servers: Custom Attributes page appears.
- 4** Click **New**.

5 Enter the name and value for the custom attribute that you want to add.

6 Click **Save**.

See “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 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 ► Manage Servers or Server Pool. The Manage 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 Manage Servers: Server Properties page appears.

3 Select the Custom Attributes tab to change the custom attributes of the server. The Manage 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 **Save** to save your changes. The Manage Servers: Custom Attributes page reappears with the updated value.

See “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 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 ► Manage Servers or Server Pool. The Manage 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 “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 for more information.

- 2** Click the display name of the server. The Manage Servers: Server Properties page appears.
- 3** Select the Custom Attributes tab. The Manage Servers: Custom Attributes page appears.
- 4** Select the check box for the custom attribute that you want to delete.
- 5** Click **Delete**. The SAS Web Client displays a confirmation page.
- 6** Click **OK** to delete the custom attribute.

See “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 for more information.

Service Levels

This section provides information about service levels within Opsware SAS and contains the following topics:

- Overview of Service Levels
- Adding a Service Level to the SAS Web Client
- 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

Overview of Service Levels

Service levels are user-defined categories that enable you to group servers in an arbitrary way and design your own organizational schemes. For example, you can organize your servers by functionality (finance, engineering, and so forth) or tier (Web, application, and database) or by ontogeny (development, staging, and production).

You can also 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.

Please note that assigning servers to service levels does not cause Opware SAS to operate any differently with respect to those servers. When you first use service levels, the categories will be fairly empty and by default, when an Opware Agent is installed on a server in the operational environment, the server will be added to the UNKNOWN Service Level.

Adding a Service Level to the SAS Web Client

Perform the following steps to add a service level to the SAS Web Client:

- 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 3-21 shows.

Figure 3-21: Service Level Hierarchy

Service Levels > OTHER

Details Add

No Sub-Nodes

| Properties | Custom Attributes 0 | Members 0 | Config Tracking | Templates 0 | History |
|---|----------------------|-----------|-----------------|-------------|---------|
| Cannot edit or delete this Node. This Node is reserved. | | | | | |
| Name: | OTHER | | | | |
| Description: | Other service level | | | | |
| Customer: | Customer Independent | | | | |
| Operating System: | OS Independent | | | | |
| Locked: | No | | | | |
| Allow Servers: | Yes | | | | |
| ID: | 50003 | | | | |

- 3** Click **Add**. 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 **Save**. 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.

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 **Edit** 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 **Save**.

Ways to View the Service Level for Servers

Find the server whose service levels you want to view by searching or browsing the Manage Servers list.

- If you are browsing the Manage Server list, you can find the service level for a server by locating the value in the Environment column, as Figure 3-22 shows.

Figure 3-22: Service Level Node Appearing in the Software Tab of the Server List

| Managed Servers: Software View | | | | |
|---|-------------------------------|---|--|---|
| Ok All Stages All Uses C06 Not Assigned Update | | | | |
| Server | Software | Configuration Tracking | View | |
| <input type="checkbox"/> | Name | Hostname / IP Address | Software | Environment |
| <input type="checkbox"/> | ventoux.snv1.corp.opsware.com | ventoux.snv1.corp.opsware.com 192.168.9.79 | OS: Red Hat Linux 7.3 / Not Assigned | Cust: Not Assigned Fac: C06 HW: DELL COMPUTER CORPORATION / PRECISION WORKSTATION 360 Rlm: Svl: Platinum Service |

- If you searched for the server, click the server name and then select the Attached Nodes tab. You can find the service levels, as Figure 3-23 shows.


Figure 3-23: Nodes Tab That Shows the Service Level to Which a Server is Assigned

| Properties | Network | Membership | Attached Nodes | Installed Packages | Custom Attributes |
|---|---------|------------|----------------|--------------------|-------------------|
| Expand All Contract All | | | | | |
| <div style="border: 1px solid gray; padding: 5px;"> <p>Directly Attached to Server</p> <ul style="list-style-type: none"> Software Service Levels <ul style="list-style-type: none"> (HA) from Service Levels / Opsware / gateway adh from Service Levels / Opsware buildmgr from Service Levels / Opsware cogbot from Service Levels / Opsware hub from Service Levels / Opsware occ from Service Levels / Opsware spin from Service Levels / Opsware twist from Service Levels / Opsware way from Service Levels / Opsware word from Service Levels / Opsware Opsware </div> | | | | | |

- 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. Select the Members tab, as Figure 3-24 shows.

Figure 3-24: Manage 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:

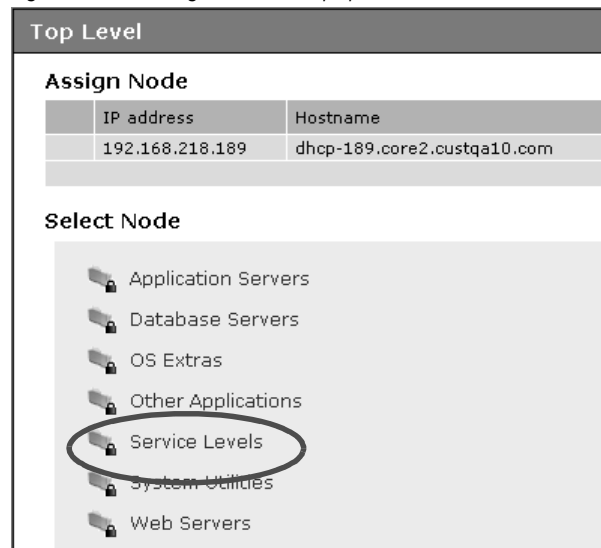
- 1 From the navigation panel, click Servers ► Manage Servers. The Manage Servers page appears. Browse the list to find the server you want to assign to a service level.

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 **Tasks ► Assign Node** from the menu above the Manage Servers list. A window displays the categories of nodes, as Figure 3-25 shows.

Figure 3-25: 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 **Assign**. The window closes and you are returned to the Manage Servers list.

See “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 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 ► Manage Servers. The Manage Servers page appears. Browse the list to find the server that you want to remove from a service level.

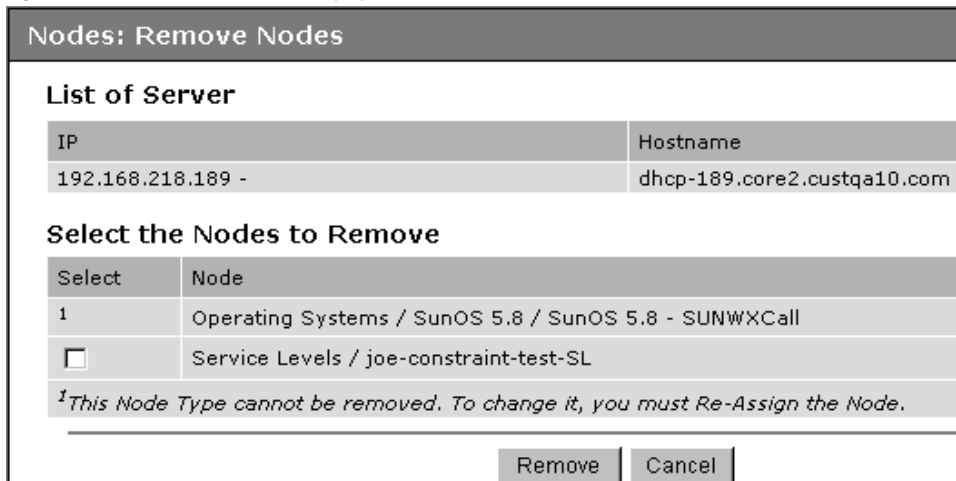
Or

Search for the server that you want to remove from a service level.

See “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 for more information.

- 2 Select the server that you want to remove from a service level.
- 3 Choose **Tasks ► Remove Node** from the menu above the Manage Servers list. A window displays the nodes to which the server is assigned, as Figure 3-26 shows.

Figure 3-26: Remove Nodes Popup Window



- 4** Select the service level node from which you want to remove the server and click **Remove**. You are prompted to confirm that you want to remove the server from the service level.
- 5** Click **Confirm Remove**. The window closes and you are returned to the Manage Servers list.

Network Configuration

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

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

Overview of the Server Network Configuration

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

The OS Provisioning feature 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. Opsware SAS enables system administrators to do this through the SAS Web Client rather than logging onto the server manually after OS provisioning is complete.



Opsware SAS does not support managed servers that have IPv6 addresses.

The Server Network Configuration feature 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 Opsware SAS should use when managing the server)

See “Locating the Management IP Address of a Managed Server” on page 100 in Chapter 3 for more information.

- Gateway (the IP address of the default router)
- 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, host name, and subnet mask

You can alter any of these options and then apply the settings to the managed server. Opsware SAS 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 ► Manage Servers. The Manage 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 “Using the Search Feature” on page 65 in Chapter 2 for more information. See “Server Searching by IP Address” on page 81 in Chapter 2 for more information.

- 2** Click the name of the server that you want to configure networking for. The Manage Servers: Properties page appears for that server.
- 3** Select 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.

- **Host Name:** The host name 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 Opsware SAS 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 Opsware SAS 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 host names
- **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, host name, and subnet mask) are not relevant for this interface, and the SAS Web Client 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).
 - **Host Name:** The local host name 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, Opsware SAS displays the management IP address and MAC address for the server; however, you cannot change these values reported by the Opsware Agent.

- 5 Click **Update Server** at the bottom of the page.

(If you click **Revert**, 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 SAS Web Client 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 SAS Web Client. 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.

Opware SAS 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 3-27.

Figure 3-27: DNS Domain Field Displays in the Network Tab for a Server

[Return to Manage Servers](#)

| Properties | Network | Membership | Attached Nodes | Installed Packages | Custom Attributes | Config Tracking | History |
|---|---------|--------------------------------|----------------|--------------------|-------------------|-----------------|---------|
| NETWORK INFORMATION (as of Tue May 24 19:21:06 2005) | | | | | | | |
| Hostname: | | core.tr3.opsware.com | | | | | |
| Management IP: | | 172.16.36.18 | | | | | |
| Facility: | | DATACENTER1 | | | | | |
| Management Interface: | | eth0 | | | | | |
| Gateway: | | 172.16.36.17 | | | | | |
| DNS Servers: | | 66.54.32.78 66.54.0.78 | | | | | |
| Search Domains: | | tr3.opsware.com opsware.com | | | | | |

Chapter 4: Software Provisioning Setup

IN THIS CHAPTER

This section discusses how to prepare to provision applications by setting up nodes in the Software Tree and creating templates. This section discusses the following topics:

- Software Tree
- Node Management in 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 SAS Web Client. Contact your Opware administrator to obtain the necessary access rights.

This section assumes that users have already packaged software applications and uploaded the packages to the Software Repository. See “Overview of Uploading a Package” on page 275 in Chapter 5 for information about how to upload packages to the Software Repository.

See *Planning Deployments for Opware® SAS* for more information about Software Tree best practices.

Software Tree

This section provides information on the Software Tree within Opware SAS and contains the following topics:

- Overview of Software Tree
- Example: Software Tree
- Guidelines for Setting Up the Software Tree
- Software Tree Hierarchy and Navigation
- Software Reconciled onto Servers
- The Time to Reconcile

Overview of Software Tree

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.

Opware technology enables IT organizations to implement their policies by using Opware'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 the *Opware® SAS User's Guide* for more information.

The SAS Web Client 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 4-1 shows. You can access the Applications page from the navigation panel by clicking Software ► Applications.

Figure 4-1: Top Level of Applications Page

Applications

| | Name | Modified |
|---|-------------------------|------------|
|  | Application Servers | 04-01-2005 |
|  | Database Servers | 04-01-2005 |
|  | Operating System Extras | 04-01-2005 |
|  | Other Applications | 04-01-2005 |
|  | System Utilities | 04-01-2005 |
|  | Web Servers | 04-01-2005 |

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.



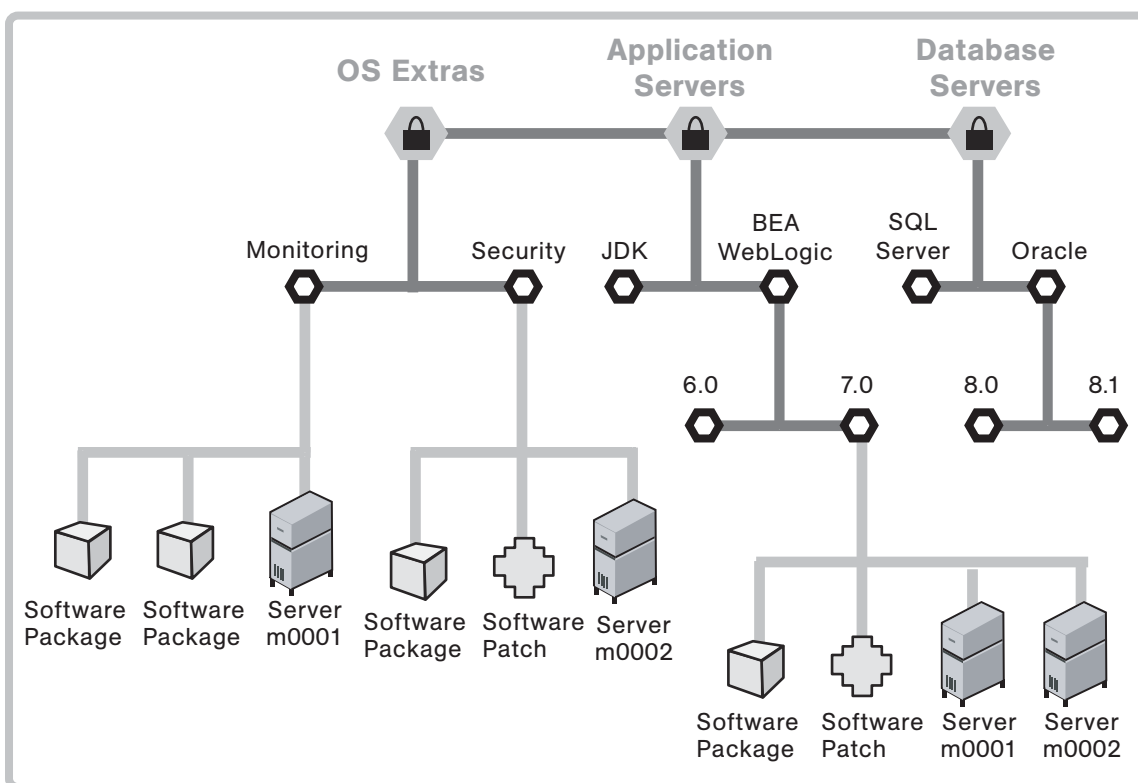
The Software Tree has the following characteristics:

- 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 199 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: Software Tree

Figure 4-2 shows how three categories and the nodes within them might be defined at a site.

Figure 4-2: 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, Opware 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 ► Solaris 7 ► 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.
- Opware 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.

See *Planning Deployments for Opware® SAS* for more information about Software Tree best practices.

Software Tree Hierarchy and Navigation

Figure 4-3 demonstrates the node hierarchy as it goes from the general to the specific. System Utilities is the category, Opsware is a node in that category, Agent Deployment Helper is a subnode of Opsware, Windows is a subnode of Agent Deployment Helper, and 2000 is a subnode of Windows. The tabs used for all functions in Applications are also shown.

Figure 4-3: Application Hierarchy



When you navigate through the Software Tree, you can perform actions on a node by using the tabs. Table 4-1 gives descriptions of the available tabs.

Table 4-1: 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. |
| 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. |

Table 4-1: Tabs Available for Applications

| TAB PAGE | DESCRIPTION |
|-----------------|--|
| Members | Manage all servers and group of servers associated with a node. Manage servers by reconciling servers between nodes, installing applications, installing patches, and so forth. See the <i>Opsware® SAS User's Guide</i> for information about how to use these functions. |
| Config Tracking | View and add entries in the configuration tracking policy for a node. See "Configuration Tracking Using the Software Tree" on page 309 in Chapter 7 for more information. |
| Templates | Lists the templates attached to a node. |
| 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. |

Software Reconciled onto Servers

A user installs software on a server by using the packages that are stored in the Software Repository. In the SAS Web Client, the user adds a package to a node, and then attaches a server to that node. When the user reconciles the server, Opsware SAS installs the software on the server.

See *Opsware® SAS User's Guide* for more information about how to reconcile.

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. Opsware SAS 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 SAS Web Client displays a preview of the actions about to be performed so that you can make any necessary changes before you proceed.

The Time 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.

Node Management in the Software Tree

This section provides information on how to manage nodes in the Software Tree and contains the following topics:

- Software Provisioning Setup Tasks
- Node Adding in the Software Tree
- Node Editing in the Software Tree
- Node Deleting in the Software Tree
- Node Copying in the Software Tree
- Node Moving in the Software Tree
- Node Configuration Tracking Policy Management
- Node History

Software Provisioning Setup Tasks

It is important to manage the nodes in 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 "Node-Based Tracking Policies" on page 320 in Chapter 7 for more information.
- View the changes to nodes.
- Create, edit, and delete templates.

Node Adding in the Software Tree

Within each of the categories, users can create subordinate nodes that extend a node. Nodes can inherit software packages or other attributes from their parent nodes, and new software, updates, or other information can be added as well.

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 the *Opware® SAS User's Guide* for more information about reconcile.

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 in 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.

Adding 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 **Add**.
- 4 If you are adding a node to an existing branch in 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 **Add**.

The Add Sub-Node form appears, as Figure 4-4 shows. (If a lot of nodes display, you might need to scroll down to see the form.)

Figure 4-4: Adding a Sub-Node to an Existing Node or Category

| ADD SUB-NODE TO System Utilities / Winzip / Windows 2000 | |
|--|--|
| Name: | <input type="text"/> Example: Patch 4 |
| Description: | <input type="text"/> Example: Save for IT Testing |
| Notes: | <div style="border: 1px solid gray; height: 60px; width: 100%;"></div> |

- 5 Complete the entries that define the node:
 - Enter a name (required).
The name is limited to 255 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 ► Solaris 6
 - (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 **Save** 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 section 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.

Node Editing 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 appear 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 (Sun Solaris node) cannot be attached to a node associated with a Linux operating system.

Editing 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 **Edit**. The Edit Attributes for [node] page appears, as Figure 4-5 shows.

Figure 4-5: Edit Attributes Page

| EDIT ATTRIBUTES FOR Winzip | |
|----------------------------|--|
| Name: | <input type="text" value="Winzip"/> |
| Description: | <input type="text"/> |
| Notes: | <input type="text"/> |
| Locked: | Yes |
| Allow servers: | <input checked="" type="checkbox"/> |
| Change customers: | <div style="border: 1px solid black; padding: 2px;"> Customer Independent </div> |
| Change operating systems: | <div style="border: 1px solid black; padding: 2px;"> OS Independent Windows NT 4.0 </div> |

- 4** Table 4-2 shows the fields that you can edit to update the node. Click **Save** when you are done.

Table 4-2: Field Data to Update for the Node

| FIELD | DESCRIPTION |
|---------------|---|
| Name | Name, limited to 225 characters, that is used in the display of the node navigation path. For example: Application Servers ► WebLogic ► Soalris 8 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: <ul style="list-style-type: none"> • 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 are 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 4-2: Field Data to Update for the Node

| FIELD | DESCRIPTION |
|------------------|---|
| Change Customers | <p>The customer that the node is associated with. This attribute determines who can view and edit the node (with the appropriate Opware 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.</p> <p>The list of customers that displays in the selection box varies. Only valid choices display. The following conditions affect what displays in the list:</p> <ul style="list-style-type: none"> • 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 4-2: Field Data to Update for the Node

| FIELD | DESCRIPTION |
|--------------------------|--|
| Change Operating Systems | <p>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.</p> <p>The list of operating systems that display in the selection box vary. Only valid choices display. The following conditions affect what the list displays:</p> <ul style="list-style-type: none"> • 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. |

Node Deleting 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 for 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).

Deleting 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 **Delete**.



The **Delete** button does not appear if one or more of the above restrictions is true.

The SAS Web Client 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 4-6.

Figure 4-6: Applications: Delete Node Page

| Applications: Delete Node System Utilities / WinZip / Windows 2000 / 8.0 | |
|---|---|
| Return to Browse Applications | |
| DELETE WinZip / Windows 2000 / 8.0 | |
| 0 sub-Node, 0 servers attached, and is attached to 0 templates. | You can delete it. |
| This Node is not derived data. | |
| This Node is not Locked and you are in the node type owner's group. | |
| 0 Install Order Dependencies and 0 reverse Install Order Dependencies. | [Install Orders Dependencies do not affect your ability to delete Nodes.] |
| Attached software: | [none] |
| Are you sure you want to delete this node? | |
| <input type="button" value="Confirm Delete"/> <input type="button" value="Cancel"/> | |

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 **Confirm Delete**.

The SAS Web Client returns you to the Properties page for the parent node.

Node Copying 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 for Copying Nodes

You can copy a node at any level, except the top-level categories.

Copying 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 ► Solaris 8).
- 3** Click **Copy**. The Copy Node page appears as Figure 4-7 shows.

Figure 4-7: Copy Node Page

Copy Node

Specify the name of the new node

From : SunOS 5.8

To :

Hide Options

Copy Subnode

yes

no

- 4** Enter a name for this node. Remember that this node name must be unique. The limit is 225 characters.
- 5** Select Yes to copy any subnodes. Select No if you only want to copy this node and no subnodes.

- 6** Click **Save**. 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.

Node Moving 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 servers attached to 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 servers 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 list describes the restrictions for moving a node:

- You can only move a node to a location in 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 in 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).

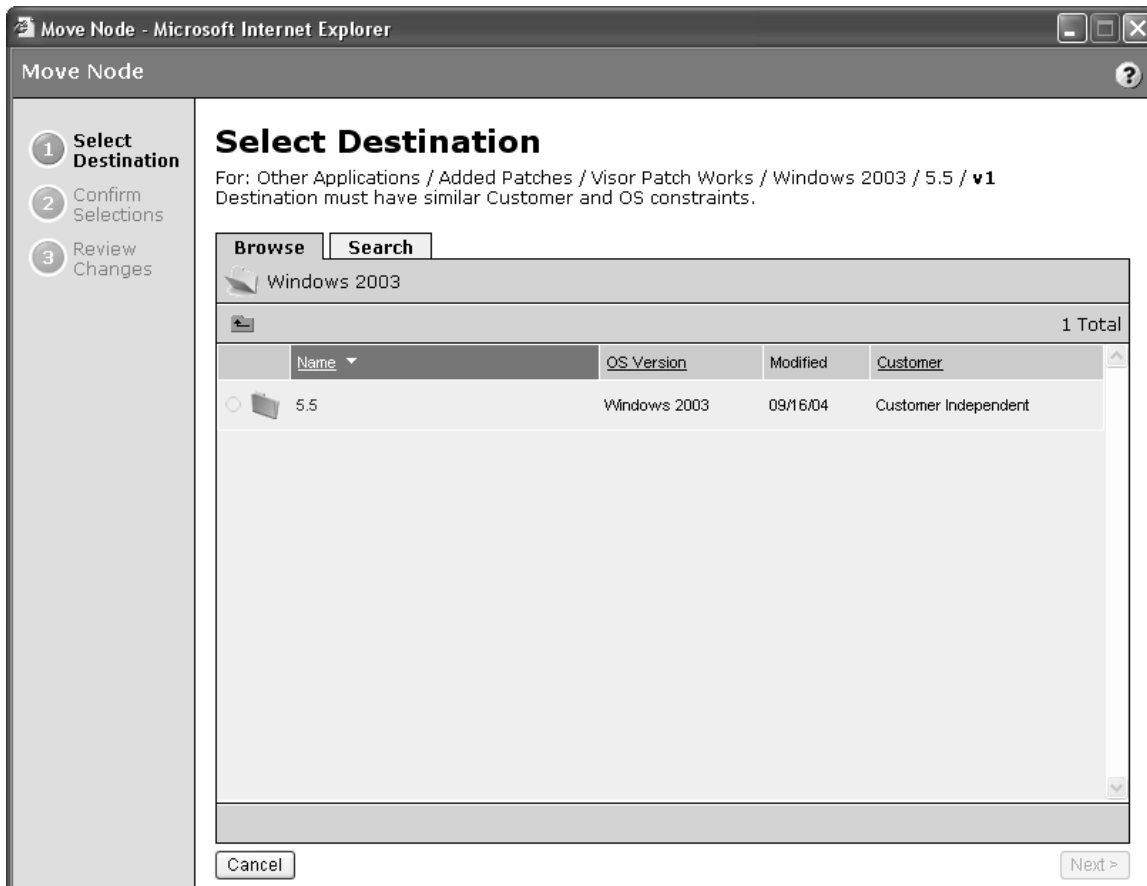
- You cannot move nodes within the Hardware and Opware Software features in the navigational panel.

Moving 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 ► Solaris 8).
- 3** Click **Move**. The Move Node Wizard appears.

Figure 4-8: 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. Click **Close** when you finish.

Node Configuration Tracking Policy Management

The 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 Software Tree nodes. Using nodes allows you to take advantage of the node's model-based architecture. Through nodes, you can create configuration policies and deploy them to the appropriate servers based on what nodes the servers are attached to.

See "Configuration Tracking Using the Software Tree" on page 309 in Chapter 7 for more information.

Node History

Using the SAS Web Client, 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 SAS Web Client 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 the 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 the node

Each history entry contains three pieces of information, as Table 4-3 shows.

Table 4-3: 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.

Viewing Node History

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 ► Solaris 8).

- 3** Select the History tab. The History for Node page appears, as Figure 4-9 shows.

Figure 4-9: History for Node Page

| HISTORY FOR NODE: | | |
|---|--|-------------------------|
| Winzip / Windows 2000 / 8.0 | Show Last: Week Two Weeks Month Quarter | |
| Event Description | Modified By | Date Modified |
| Updated ic_certified field to "true" | vhsiehadmin | 21-Aug-2003 04:46:36 PM |
| Created child node : "Addons" (ID: 86660007) | vhsiehadmin | 21-Aug-2003 04:45:27 PM |
| Deleted child node : "8.1" (ID: 86650007) | vhsiehadmin | 21-Aug-2003 04:45:03 PM |
| Created child node : "8.1" (ID: 86650007) | vhsiehadmin | 21-Aug-2003 04:44:40 PM |
| Created child node : "Winzip Patch" (ID: 86220007) | admin | 18-Aug-2003 08:55:53 PM |
| Updated allow_dvc field to "true" | admin | 18-Aug-2003 08:54:33 PM |

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 SAS Web Client 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 Opware SAS and contains the following topics:

- Overview of Software Attached to Nodes
- Modeling Software in Nodes
- Software Configuration Settings
- Viewing Software Attached to a Node
- Adding a Package to a Node
- Removing a Package from a Node
- Changing the Installation Order of Software
- Software Inheritance 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

Overview of Software Attached to Nodes

Packages are attached to nodes. Using nodes to install packages provides an automated and consistent mechanism for software installation based on the best practices that your organization defines.

See “Overview of Package Management” on page 251 in Chapter 5 for information about how to manage (upload and download) packages through the SAS Web Client.

A user installs software on a server by using packages stored in the Software Repository. In the SAS Web Client, the user adds a package to a node and then attaches a server to that node. When the user uses the Install Software Wizard, Opware SAS reconciles the server and installs the software on the server. See the *Opware® SAS User's Guide* for more information about Reconcile.

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 the packages.

When you modify software that belongs to a node, remember to click **Save Edits** or the changes will not be 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 **Save Edits** to commit all the changes at the same time. See “Software Inheritance from Other Nodes” on page 210 in this chapter for more information.

Modeling Software in Nodes

Opware SAS 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 “Software Provisioning Setup” on page 175 in this chapter for more information about the full list of package types that Opware SAS manages.

Table 4-4 shows only the package types that you can attach to nodes.

Table 4-4: Package Types Organized by Operating System

| PACKAGE TYPE | FILE FORMATS | NOTES |
|---------------------|--------------|---|
| AIX | | |
| AIX Update Fileset | N/A | None |
| AIX Base Fileset | N/A | None |
| AIX APAR | N/A | <p>Users can attach complete and incomplete APARs to nodes. An APAR is incomplete when Opware SAS does not have all required update filesets for the APAR.</p> <p>Users can install incomplete APARs on servers when the Software Repository contains the update filesets corresponding to the installed base filesets.</p> <p>Possibly, the missing update filesets are not applicable because their base filesets have not been installed. Opware SAS warns users when they try to reconcile servers when necessary update filesets have not been uploaded.</p> |
| 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 |

Table 4-4: Package Types Organized by Operating System

| PACKAGE TYPE | FILE FORMATS | NOTES |
|--------------------------|--|---|
| 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 | Opware SAS supports interactive (requires a response file) and non-interactive Solaris packages. |
| 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. See "Editing Package Properties" on page 281 in Chapter 5 for information about adding a response file for a Solaris package. |
| RPM | .rpm | None |
| Windows | | |
| Windows Hotfix | .exe | None |
| Windows MSI | .msi | None |
| Windows OS Service Pack | .exe | None |
| Windows ZIP File | .zip | None |

Software Configuration Settings

Each package should contain the configuration settings necessary to run the application on a server. If a configuration change is required in an operational environment, Opware Inc. recommends that you repackage the software and add the updated package to the Software Repository.

See “Package Management” on page 251 in Chapter 5 for information about how to manage (upload and download) packages with the SAS Web Client.

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, Opware SAS might be unable to rebuild a server to its latest configuration. Be sure to backup any manual configuration changes.

Even though the SAS Web Client 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 ► Solaris 8).
- 3** Select the Packages tab. The Packages Attached to a Specific Node page appears, as Figure 4-10 shows.

Figure 4-10: 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 Package to a Node



Each 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 (Sun Solaris 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.

Opware SAS provides two methods to add a software package to a node. You can use any one method to add a software package to a node and both the methods yield the same results.

Method 1: Adding a Package to a Node

Perform the following steps to add a 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 ► Solaris 8).
- 3** Select the Packages tab. A page that lists the software attached to the node appears.
- 4** Click **Edit Packages**. A page appears that shows the directly attached packages and the inherited packages.



When a node contains a large number of packages, the page might take a few seconds to load.

- 5 Click **Add Packages**. The Search for Packages page appears that allows you to search for the software that you want to add, as Figure 4-11 shows.

Figure 4-11: 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 **Find Software**.



The time it takes the SAS Web Client 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. (Packages already attached to the node do not have a check box.)
- 9 Click **Add to Node**. 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 209 in this chapter for more information.

11 Click **Save Edits**.



Additions and changes are not stored in Opware SAS until you click **Save Edits**.

Method 2: Adding a Package to a Node

- 1** From the navigation panel, click Software ► Packages. The Packages: Browse Packages page appears.
- 2** Select the package. The Packages: Edit Properties page appears.
- 3** Select the Nodes tab. The Package: View Nodes page appears, as Figure 4-12 shows.

Figure 4-12: The Package:View Nodes Page

Return to Browse Packages

Properties

Nodes

The following nodes currently utilize this package. To delete a package, it cannot be utilized by any nodes. You may be able to click on a node to edit it.

| NODE NAME | NODE TYPE | DESCRIPTION |
|--|------------------|-------------|
| System Utilities Opware Tools ISMtool Windows 2003 | System Utilities | |

- 4** Click **Add**. The Add Node to Package window appears.
- 5** Navigate to a node (for example, Application Servers ► WebLogic ► Solaris 8). After you have selected the application, click **Select**. The Add Node to Package window displays the package added to the node.
- 6** Click **Close** to finish.

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 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 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 **Find Software**. 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 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 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 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 Package from a Node

Users can delete one, many, or all of the 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 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 ► Solaris 8).
- 3** Select the Packages tab. A page appears that shows the list of packages attached to the node.

- 4 Click **Edit Packages**. 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 4-13 shows.

Figure 4-13: 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 **Select All** located below the list. (You can also reset the selection by clicking **Deselect All**.) You can select multiple packages to delete by holding down the CTRL key while selecting packages.

6 Click **Remove Packages**.

The selected packages disappear from the select box but the packages are not yet detached from the node.

7 Click **Save Edits** at the bottom of the page to complete detachment of the software.



Additions and changes are not stored in Opsware SAS until you click **Save Edits**.



To apply changes that you made to a node to servers that already have the node attached, you need to reconcile those servers. See the *Opsware[®] SAS User's Guide* for more information about Reconcile.

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 ► Solaris 8).
- 3** Select the Packages tab. A page appears that shows the list of packages attached to the node.
- 4** Click **Edit Packages**. 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 or the down arrow.

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 or move to bottom arrow.

Highlighted packages are moved to the top or bottom of the installation order.

- 7** When you finish making changes to the installation order, click **Save Edits** at the bottom of the page.



Installation order changes are not stored in Opware SAS until you click **Save Edits**.

Software Inheritance 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.

Example: 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 (+) to the left. These packages are referred to as plus override. Any software package with a plus (+) by its name is inherited and installed on servers attached to this node.

If a software package has a minus (-) 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** ► **Solaris 8**).
- 3** Select the **Packages** tab. A page appears that shows the list of packages attached to the node.
- 4** Click **Edit Packages**. 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 **Save Edits**.
Selected packages that were prefaced with a plus (+) are prefaced with a minus (-), and packages that were prefaced with a minus (-) are prefaced with a plus (+).
- 7** Click **Save Edits**.



The changes made to override values are not committed to Opware SAS until you click **Save Edits**.

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 packages allows those packages to be installed in the correct order when a server is reconciled.

Opware SAS 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 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 selecting 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 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 SAS Web Client 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 ► Solaris 8).
- 3** Select the Install Order tab. The Install Current Node After list and Install Current Node Before list display, as Figure 4-14 shows.

Figure 4-14: Install Order Tab

| |
|-------------------------------------|
| INSTALL CURRENT NODE AFTER: |
| Add |
| [none] |
| INSTALL CURRENT NODE BEFORE: |
| [none] |



The SAS Web Client displays [none] when there are no installation order dependencies.

Adding Software Installation Dependencies

When you create software installation order dependencies, Opware 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 SAS Web Client 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** Select the Install Order tab. The Install After list and Install Before list display.
- 4** Click **Add**. 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 **Add**. The SAS Web Client closes the window to display the new software installation order dependency, as Figure 4-15 shows.

Figure 4-15: 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 ► Solaris 8).
- 3** Select the Install Order tab.
- 4** Select the check box next to the specific dependency that you want to delete.
- 5** Click **Remove Dependency**.

Custom Attributes Set for the Environment

This section provides information on custom attributes set for the environment within Opware SAS and contains the following topics:

- Overview of Custom Attributes and Folders
- Managing Custom Attributes
- Adding Custom Attributes for a Node
- Editing Custom Attributes for a Node
- Deleting Custom Attributes for a Node

Overview of Custom Attributes and Folders

Users might need to store specific miscellaneous information in the Opware Model Repository, for example, to facilitate server or application installation and configuration or scripting.

The SAS Web Client 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 SAS Web Client 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 the *Opware® SAS User's Guide: Server Automation* for information about setting custom attributes for a server.

To set custom attributes that affect every server for a customer, you can set custom attributes for that customer.

See the *Opware® SAS Policy Setter's Guide* for information about setting custom attributes for customers.

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 ► Solaris 8).
- 3** Select 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 **Edit Custom Attributes**, and then **Add Custom Attributes**.

- 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 ellipsis (...) button to open a window with a larger text box.

- 5** Click **Save**. The View Custom Attributes page appears.
- 6** Click **Edit Custom Attributes**.
- 7** Click **Save Edits**.

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 ► Soalris 8).
- 3** Select the Custom Attributes tab.
- 4** Click **Edit Custom Attributes**. A new form appears, which shows three columns of information: Name, Inherited Value, and Local Value. See Figure 4-16.

Figure 4-16: Edit Custom Attributes Page

| Properties Packages 2 Custom Attributes 4 Install Order 1 Servers 0 Config Tracking History | | | | | | |
|---|--|--|--|--|--|------------------------|
| The following Custom Attributes are for this Node | | | | | | Edit Custom Attributes |
| Name | | | | | | Value |
| weblogic.httpd.enableLogFile | | | | | | true |
| weblogic.httpd.logFileFormat | | | | | | extended |
| weblogic.httpd.logFileName | | | | | | wl_access.log |
| weblogic.httpd.logFileFlushSecs | | | | | | 60 |

- 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 **Save Edits** 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 ellipsis (...) button to open a window with a larger text entry box.
- 8** Click **Save Edits** to commit the changes.



The changes made to custom attributes are not committed to Opware SAS until you click **Save Edits**.

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** ► **Solaris 8**).
- 3** Select the **Custom Attributes** tab. A page appears that shows the custom attributes for the selected node.

- 4** Click **Edit Custom Attributes**. The three-column form appears with a check box labeled Delete available next to the Local Value field, as Figure 4-17 shows.

Figure 4-17: Delete Custom Attributes

| The following Custom Attributes are for this Node | | | Add Custom Attributes | |
|---|-----------------|----------------------|-------------------------------------|--------|
| Name | Inherited Value | Local Value | | |
| weblogic.httpd.enableLogFile | | <input type="text"/> | <input checked="" type="checkbox"/> | Delete |
| weblogic.httpd.logFileFlushSecs | | 60 | <input type="checkbox"/> | Delete |
| weblogic.httpd.logFileFormat | | <input type="text"/> | <input checked="" type="checkbox"/> | Delete |
| weblogic.httpd.logFileName | | wl_access.log | <input type="checkbox"/> | Delete |

- 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 Opware SAS and contains the following topics:

- Overview of Working with Templates
- Templates, Folders, and Inheritance
- Template Inheritance
- Attachments: Local, Inherited, and Blocked
- Blocking and Reattaching Inherited Attachments

Overview of Working with Templates

Opware templates allow you to create groups of logical relationships between nodes and a server. With templates, you can arrange related sets of packages together so that you can apply them to a server in a single operation by using the Opware 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 feature of Opsware SAS. 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 templates and 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 4-18 shows that Solaris 5.8 and Basic Veritas VFS Software v1 are attached locally to Folder 1.

Figure 4-18: Folder with Nodes Attached Locally



This form of attachment is the simplest. Understanding local attachments is important, as more complex examples are shown.

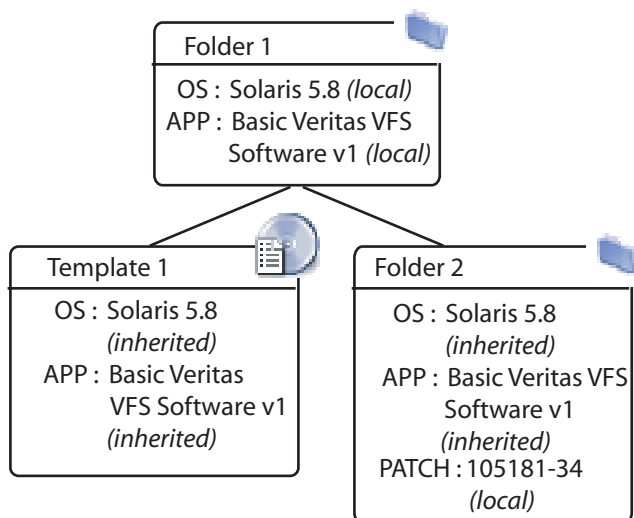
Inherited

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 4-19 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 4-19 shows.

Figure 4-19: 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.

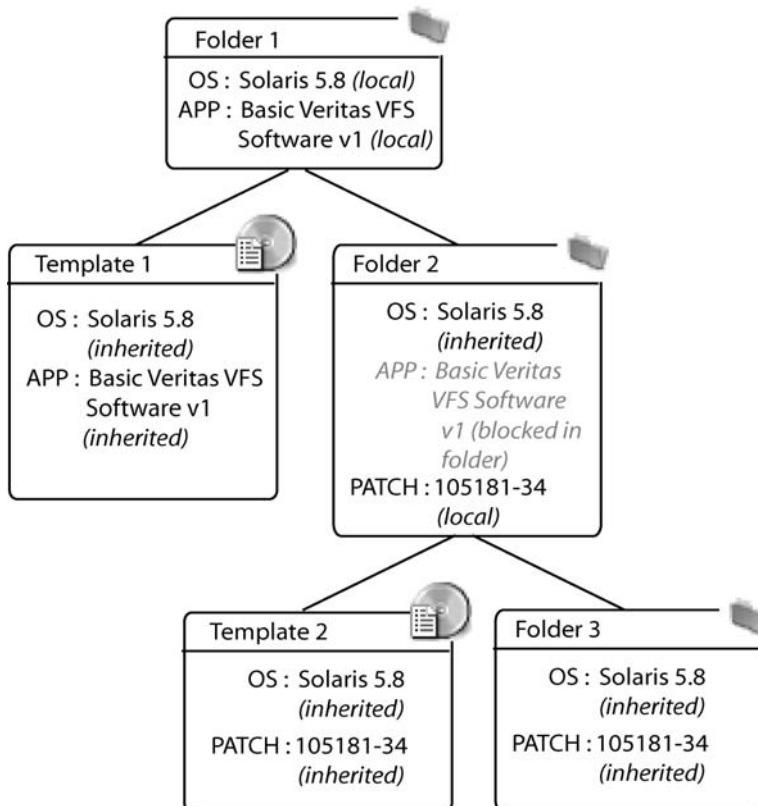
Blocked

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 4-20 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 4-20: Folders with Child Folders and Templates That Show Attachments Blocked from Inheritance



Blocking and Reattaching Inherited Attachments

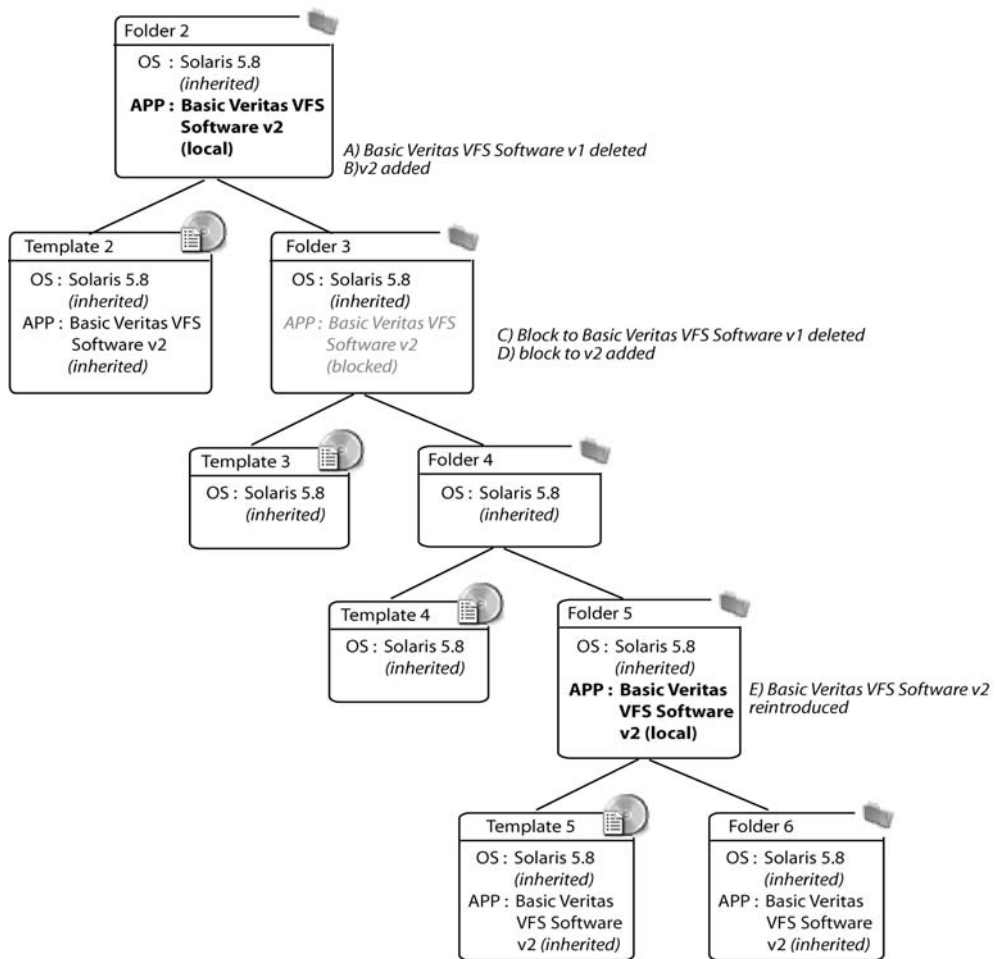
Blocking and then reattaching an attachment lower in the tree is possible, but this action can be complex and should be done with care.

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 4-21, 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 4-21 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 4-21: 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 Opware SAS and contains the following topics:

- Overview of Folders and Templates
- Creating Templates
- Creating Folders
- Editing or Removing Service Levels in Templates or Folders
- Copying Templates and Folders
- Deleting Templates or Folders
- Folder and Template Inheritance
- Blocking vs. Removing Attachments

Overview of Folders and Templates

Opware SAS 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 Opware SAS 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.

Templates and Folders Application

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 SAS Web Client, 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 **New Template**. The Templates: Create Template page appears, as Figure 4-22 shows.

Figure 4-22: Templates: Create Template Page

| Templates: Create Template | |
|---|---|
| Return to Templates: Manage Templates | |
| Name: | <input type="text"/> |
| Description: | <div style="border: 1px solid gray; height: 100px; width: 100%;"></div> |
| Notes: | <div style="border: 1px solid gray; height: 100px; width: 100%;"></div> |
| OS Version: | <input type="text" value="OS Independent"/> |
| Customer: | <input type="text" value="Customer Independent"/> |
| Assign Customer: | Assign customer to server upon template application? <input type="radio"/> Yes <input checked="" type="radio"/> No |
| <input type="button" value="Save"/> <input type="button" value="Cancel"/> | |

- 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 | <Template Name> page appears, as Figure 4-23 shows.

Figure 4-23: Templates: Edit Folder | Template Name Page

| Templates: Edit Folder My New Template | |
|---|----------------------|
| Return to Templates: Manage Templates | |
| Summary | Properties |
| Operating System | Patches |
| Applications | Service Levels |
| History | |
| Properties Edit | |
| Name: | My New Template |
| Location: | Templates |
| Description: | (not set) |
| Notes: | (not set) |
| OS Version: | OS Independent |
| Times Used: | 0 |
| Last Modified: | 03/30/04 18:52:28 |
| Customer: | Customer Independent |
| ID: | 23960010 |
| Operating System | |
| OS does not apply to OS Independent. | |
| Patches | |
| Patches do not apply to OS Independent. | |
| Applications Add... | |
| None. | |
| Service Levels Add... | |
| None. | |

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 SAS Web Client, 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 **New Folder**. The Templates: Create Folder page appears, as Figure 4-24 shows.

Figure 4-24: Templates: Create Folder Page

| Templates: Create Folder | |
|---|--|
| Return to Templates: Manage Templates | |
| Name: | <input type="text"/> |
| Description: | <div style="border: 1px solid #ccc; height: 100px;"></div> |
| Notes: | <div style="border: 1px solid #ccc; height: 100px;"></div> |
| OS Version: | OS Independent ▼ |
| Customer: | Customer Independent ▼ |
| <input type="button" value="Save"/> <input type="button" value="Cancel"/> | |

- 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 | <Folder Name> page appears, as Figure 4-25 shows.

Figure 4-25: Templates: Edit Folder | Folder Name Page

Templates: Edit Folder | My New Folder

Return to Templates: Manage Templates

[Summary](#)
[Properties](#)
[Operating System](#)
[Patches](#)
[Applications](#)
[Service Levels](#)
[History](#)

Properties Edit

| | |
|----------------|----------------------|
| Name: | My New Folder |
| Location: | Templates |
| Description: | Test |
| Notes: | test |
| OS Version: | OS Independent |
| Times Used: | 0 |
| Last Modified: | 03/30/04 18:30:49 |
| Customer: | Customer Independent |
| ID: | 23950010 |

Operating System

OS does not apply to OS Independent.

Patches

Patches do not apply to OS Independent.

Applications Add...

None.

Service Levels Add...

None.

Operating System Management for 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 SAS Web Client, 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 4-26 shows.

Figure 4-26: Templates: Edit Template Page for Adding an Operating System to a Template

Templates: Edit Template | TM Windows NT 4.0 (2306)

[Return to Templates: Manage Templates](#)

[Summary](#) | [Properties](#) | [Operating System](#) | [Patches](#) | [Applications](#) | [Service Levels](#) | [History](#)

Properties [Edit](#)

| | |
|------------------|--------------------------|
| Name: | TM Windows NT 4.0 (2306) |
| Location: | Templates |
| Description: | Used for TM Testing |
| Notes: | (not set) |
| OS Version: | Windows NT 4.0 |
| Times Used: | 0 |
| Last Modified: | 03/10/04 08:00:21 |
| Customer: | Customer Independent |
| Assign Customer: | No |
| ID: | 19390010 |

Operating System [Remove...](#) [Change...](#)

| | Name | Location | Description | Attachment |
|---|--------------------------|------------------------------------|------------------------------|------------|
|  | TM Windows NT 4.0 (6564) | Operating Systems / Windows NT 4.0 | Used to test OS Provisioning | Local |

Patches [Add...](#)

None.

Applications [Add...](#)

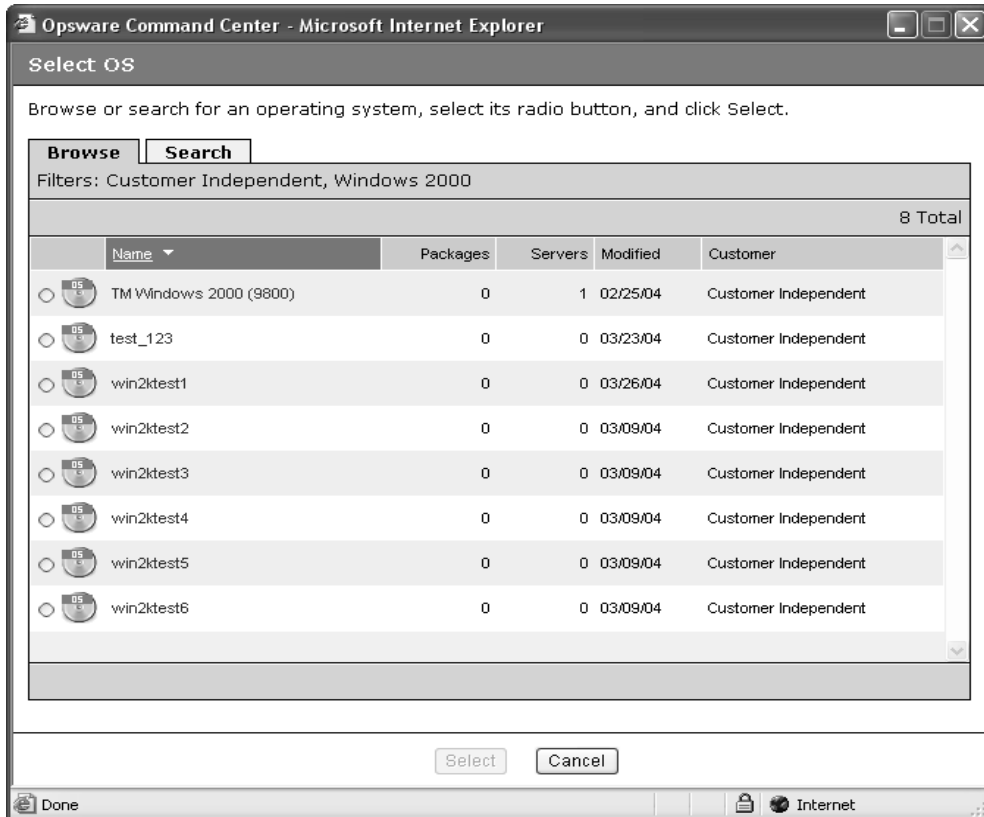
None.

Service Levels [Add...](#)

None.

- Click **Select** in the Operating System field. The Select OS window appears, as Figure 4-27 shows. (If the Template already has an operating system, click **Change** to change the operating system.)

Figure 4-27: Select OS Window



- Browse or search for the operating system that you want to add to the template and select it by selecting its radio button.
- Click **Select** 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:

- In the navigation panel of the SAS Web Client, 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.) The Edit Templates page appears, as Figure 4-28 shows.

Figure 4-28: Templates: Edit Template Page for Removing an Operating System from a Template

Templates: Edit Template | TM Windows NT 4.0 (2306)

[Return to Templates: Manage Templates](#)

Summary Properties **Operating System** Patches Applications Service Levels History

Properties Edit

Name: TM Windows NT 4.0 (2306)

Location: Templates

Description: Used for TM Testing

Notes: (not set)

OS Version: Windows NT 4.0

Times Used: 0

Last Modified: 03/10/04 08:00:21

Customer: Customer Independent

Assign Customer: No

ID: 19390010

Operating System Remove... Change...

| | Name | Location | Description | Attachment |
|---|--------------------------|------------------------------------|------------------------------|------------|
|  | TM Windows NT 4.0 (6564) | Operating Systems / Windows NT 4.0 | Used to test OS Provisioning | Local |

Patches Add...

None.

Applications Add...

None.

Service Levels Add...

None.

- 3 In the Operating System field, click **Remove**.
- 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 SAS Web Client, 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 4-29 shows.

Figure 4-29: Folder Contents

Templates: Manage Templates

The following templates are available for use.

Templates /

My New Template Details

Copy... Delete... | 📁 | New Template New Folder 2 Total

| <input type="checkbox"/> | Name ▼ | OS Version | <u>Times Used</u> | <u>Modified</u> | Customer |
|--------------------------|-------------------------|------------|-------------------|-----------------|----------------------|
| <input type="checkbox"/> | Best Template Practices | SunOS 5.8 | 0 | 03/30/04 | Opware |
| <input type="checkbox"/> | My Newer Template | SunOS 5.8 | 0 | 03/30/04 | Customer Independent |

- 3** Click **Details**. The folder's property summary page appears, as Figure 4-30 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 4-30: Folder Properties

Templates: Edit Folder | test

[Return to Templates: Manage Templates](#)

Summary | **Properties** | Operating System | Patches | Applications | Service Levels | History

Properties Edit

| | |
|----------------|-------------------------|
| Name: | test |
| Location: | Templates / test folder |
| Description: | (not set) |
| Notes: | (not set) |
| OS Version: | Red Hat Linux 7.3 |
| Times Used: | 0 |
| Last Modified: | 03/25/04 20:03:10 |
| Customer: | Customer Independent |
| ID: | 23670010 |

Operating System Select...

None.

Patches

Patches do not apply to Linux operating systems.

Applications Add... Edit

| | Name | Location | Description | Attachment |
|--|------|-------------------------------|-------------|--------------|
| | 1.1 | Other Applications / DCI / en | (not set) | Inherited... |

Service Levels Add...

None.

- 4** To add an operating system, from the Operating System section, click **Select**.
- 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 **Select** 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 SAS Web Client, 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 **Details**.

- 4** In the Summary page for the selected folder, in the Operating System field, click **Remove**, as Figure 4-31 shows.

Figure 4-31: Remove Operating System from a Folder

Templates: Edit Folder | Folder with OS

[Return to Templates: Manage Templates](#)

[Summary](#) | [Properties](#) | [Operating System](#) | [Patches](#) | [Applications](#) | [Service Levels](#) | [History](#)

Properties [Edit](#)

Name: Folder with OS

Location: Templates

Description: Patrick testing.

Notes: (not set)

OS Version: Red Hat Linux 6.2


Times Used: 0

Last Modified: 03/30/04 23:27:41

Customer: Customer Independent

ID: 24000010

Operating System [Remove...](#) [Change...](#)

| | Name | Location | Description | Attachment |
|---|-----------------------------|---------------------------------------|------------------------------|------------|
|  | TM Red Hat Linux 6.2 (6352) | Operating Systems / Red Hat Linux 6.2 | Used to test OS Provisioning | Local |

Patches

Patches do not apply to Linux operating systems.

Applications [Add...](#)

None.

Service Levels [Add...](#)

None.

- 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.



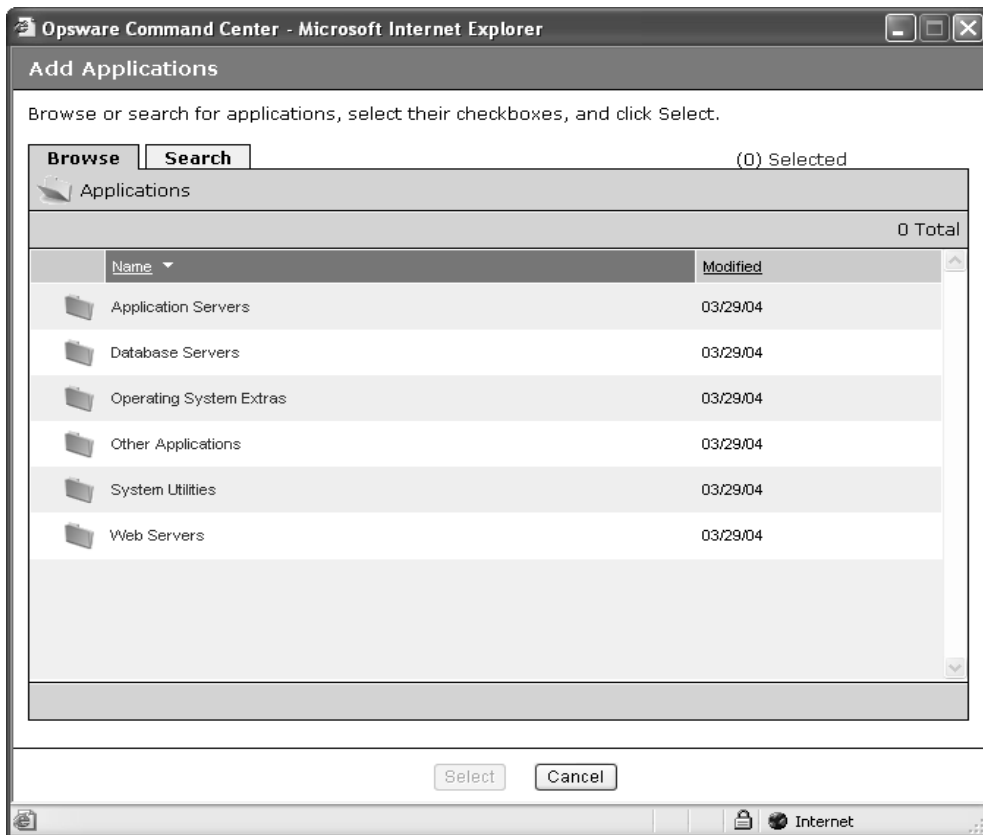
If you remove an operating system from a folder that has child folders, you are not able to add another operating system at the same level. Instead, use the **Change** button in the Operating System field.

Adding Applications to Templates or Folders

Perform the following steps to add an application to a template or to a folder:

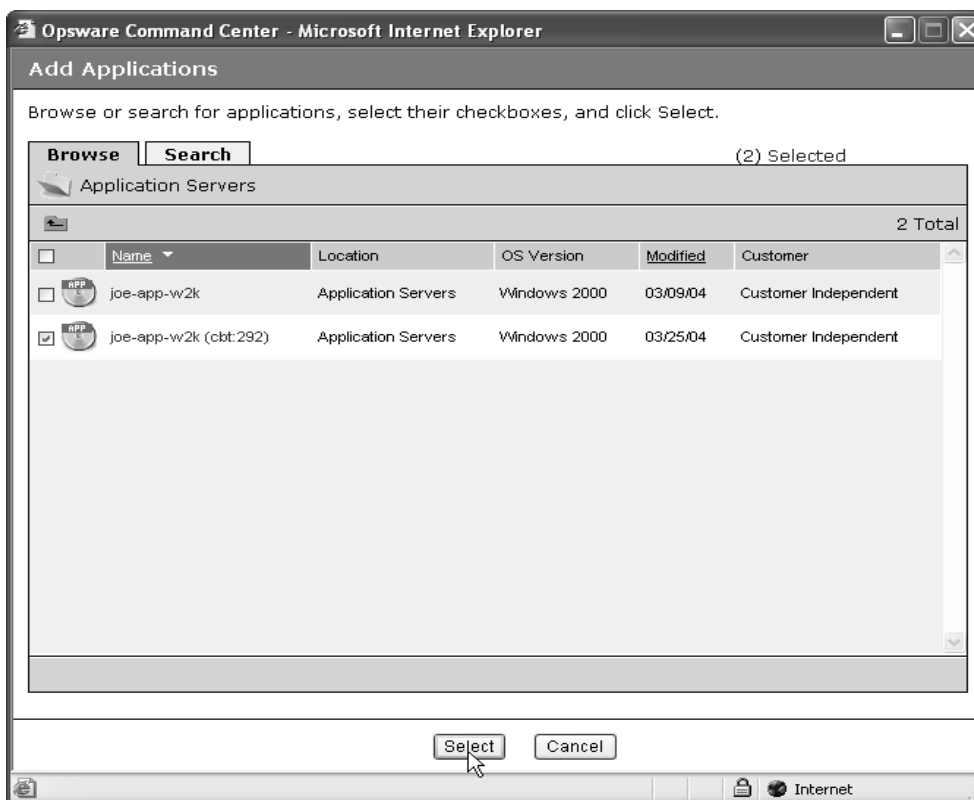
- 1** From the navigation panel of the SAS Web Client, click Software ► Templates. The Templates: Manage Templates page appears.
- 2** Select the template or folder where you want to add an application. (If a template or folder is inside another folder, you need to navigate to its location before you can select it.)
- 3** If you are adding an application to a folder, click **Details**. If you are adding an application to a template, go to the next step.
- 4** To include applications in your template, click **Add** in the applications field. The Add Applications page appears, as Figure 4-32 shows.

Figure 4-32: Add Applications Page



- 5 Browse or search for the applications that you want to add. Select the applications by clicking their check boxes, and then click **Select** to add them to your template or folder, as Figure 4-33 shows.

Figure 4-33: 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 4-34 shows.

Figure 4-34: Applications Added to a Template

| Applications | | | | | |
|--------------|-----------------------|--|-------------|------------|------|
| | | | | Add... | Edit |
| | Name | Location | Description | Attachment | |
| | AllInternalZip | Other Applications / DCI / en / 1.1 / Windows 2000 | (not set) | Local | |
| | joe-app-w2k (cbl:292) | Application Servers | (not set) | Local | |

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 SAS Web Client, 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 **Details**. If you are adding an application to a template, go to the next step.
- 4** Click **Edit** in the Applications field. The Templates: Edit Template page appears with the Applications tab active, as Figure 4-35 shows.

Figure 4-35: Edit Applications Page

| Templates: Edit Template DCI Test | | | | | | | |
|---------------------------------------|-----------------------|--|------------------|----------|----------------------|----------------|---------|
| Return to Templates: Manage Templates | | | | | | | |
| Summary | | Properties | Operating System | Patches | Applications | Service Levels | History |
| Remove... | | Add Applications... | | | | 2 Total | |
| <input type="checkbox"/> | Name | Location | OS Version | Modified | Customer | Attachment | |
| <input type="checkbox"/> | AllInternalZip | Other Applications / DCI / en / 1.1 / Windows 2000 | Windows 2000 | 03/09/04 | Customer Independent | Local | |
| <input type="checkbox"/> | joe-app-w2k (cvt:292) | Application Servers | Windows 2000 | 03/25/04 | Customer Independent | Local | |

- 5** To remove applications, click the check box next to the name of the application that you want to remove.



Inherited and blocked applications do not have a selectable check box so you cannot remove them here. You have to remove them from the folder where they are attached locally.

- 6** Click **Remove**.
- 7** The Remove Applications confirmation page appears. Click **Remove** 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

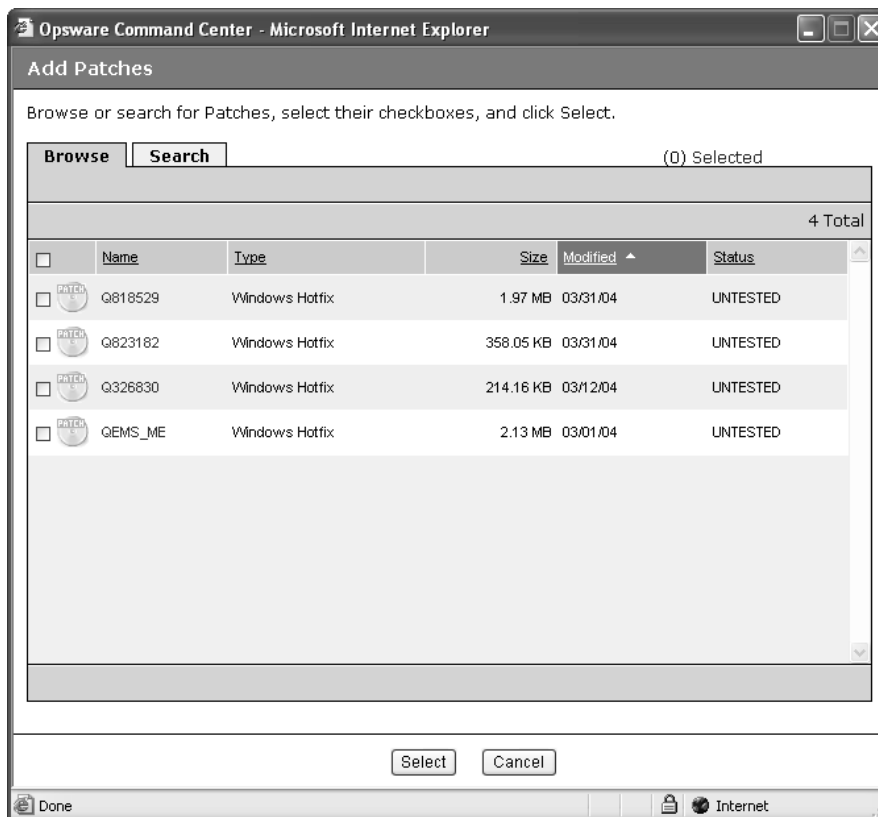
In the Patches area of the Edit Template page, you have the option of either adding or editing patches. If you choose to edit patch attachments, you can both remove existing patches and add new patches.

Perform the following steps to add patches to a template or a folder:

- 1** In the navigation panel of the SAS Web Client, 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 **Details**. If you are adding a patch to a template, go to the next step.

- 4 Click **Add** in the Patches field. The Add Patches window appears, as Figure 4-36 shows.

Figure 4-36: Add Patches Window



- 5 Browse or search for the patches that you want to add to the template or folder. Select the check boxes for the desired patches, and then click **Select** 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 SAS Web Client, 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 **Details**. If you are adding a patch to a template or folder, go to the next step.
- 4 Click **Edit** in the Patches field. The Templates: Edit Template page appears with the Add Patches tab active, as Figure 4-37 shows.

Figure 4-37: Templates: Edit Template Page - Edit Patches

| Templates: Edit Template TM Windows 2003 (1107) | | | | | | | | |
|---|------------|-----------------------------|----------------|--------------|----------------|----------|------------|--|
| Return to Templates: Manage Templates | | | | | | | | |
| Summary | Properties | Operating System | Patches | Applications | Service Levels | History | | |
| Remove... | | Add Patches... | | | | | 1 Total | |
| <input type="checkbox"/> | Name | Location | Type | OS Version | Size | Modified | Attachment | |
| <input type="checkbox"/> | Q823980 | Patches / NT / 5.2 / HOTFIX | Windows Hotfix | Windows 2003 | 1.42 MB | 03/31/04 | Local | |

- 5 To remove patches, click the check box next to the name of the patches that you want to remove.



Inherited and blocked patches do not have a selectable check box so you cannot remove them here. You must remove them from the folder where they are attached locally.

- 6 Click **Remove**.
- 7 The Remove Patches confirmation page appears. Click **Remove** to remove the patches from the template or folder.
- 8 To add patches to the template or folder, click **Add**. The Add Patches page appears.
- 9 Browse or search for the patches that you want to add. Select them by clicking their check boxes and then click **Select**.

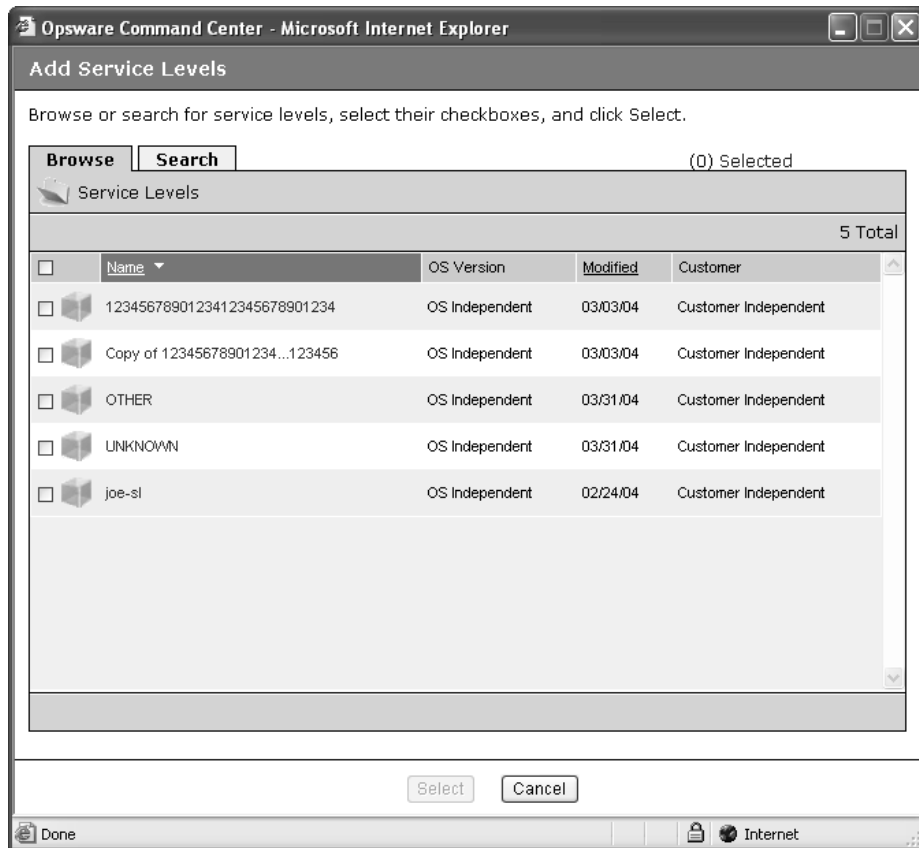
Adding Service Levels in Templates or Folders

Perform the following steps to add service levels in templates or folders:

- 1 In the navigation panel of the SAS Web Client, 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 **Details**. If you are adding a service level to a template, move to the next step.
- 4 Click **Add** in the Service Levels field. The Add Service Levels page appears, as Figure 4-38 shows.

Figure 4-38: 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 check boxes, and then click **Select** to add them to your template or folder.

See the *Opware® SAS User's Guide* for more 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 **Edit** in the Service Levels field. The Edit Template page appears with the Service Levels tab active.
- 2** To remove service levels, click the check box next to the name of the service levels that you want to remove.



Inherited and blocked service levels do not have a selectable check box so you cannot removed them from here. You have to remove them from the folder where they are attached locally.

- 3** Click **Remove**.
- 4** The Remove Service Levels confirmation page appears. Click **Yes** 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.

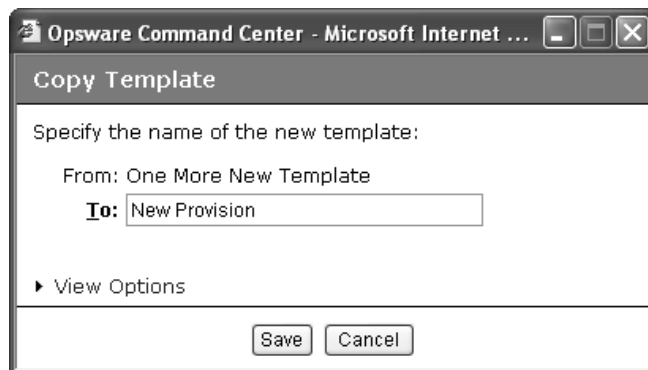
When you copy a folder, the copy has the same attachments and the same children with all of the attachments as the original folder.

Perform the following steps to copy a template or a folder:

- 1** In the navigation panel of the SAS Web Client, 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 **Details**. If you are copying a template, move to the next step.
- 4** Click **Copy**.

- 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 4-39 shows.

Figure 4-39: 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 **Save**. 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 SAS Web Client, 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 **Delete**.

- 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.

Folder and Template Inheritance

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.

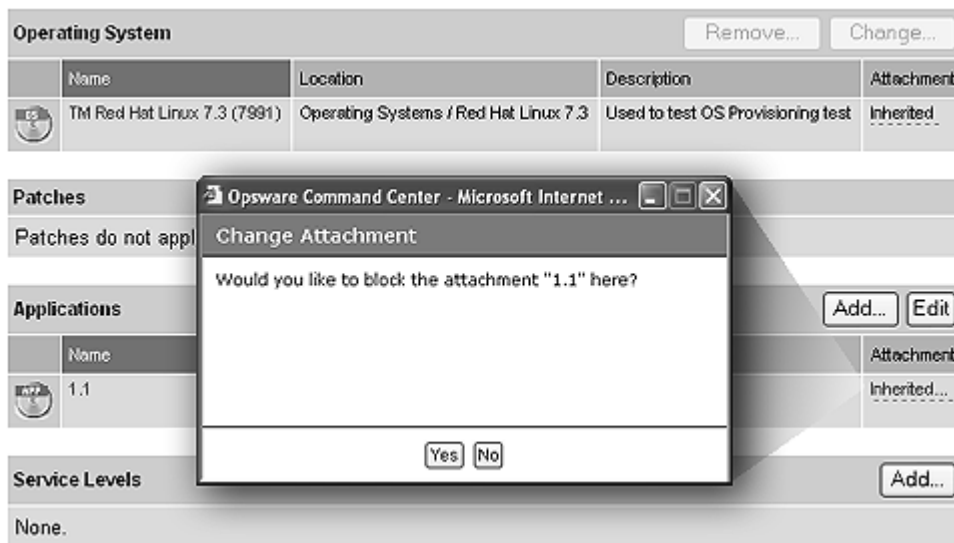
Blocking an Attachment from Being Inherited

Perform the following steps to block an attachment from being inherited:

- 1 In the navigation panel of the SAS Web Client, 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 4-40 shows.

Figure 4-40: 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 5: Package Management

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Package Management
- AIX Packages
- HP-UX Packages
- Linux Packages
- Solaris Packages
- Windows Packages
- ZIP Packages
- Package Management Tasks

Overview of Package Management

Packages are made available in Opware SAS by uploading the packages to the Software Repository with the Opware SAS Web Client or by using the Opware Command Line Interface. See the *Opware® SAS Content Utilities Guide* for information about the Opware Command Line Interface (OCLI) for package management.

The SAS Web Client provides options that enable you to perform the following tasks:

- Upload packages to the Software Repository.

The Software Repository provides a data store for all software that Opware SAS 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. Opware

SAS reconciles the differences between how a server is configured by assigned nodes and what software is actually installed. The Opware Agent installed on each Opware-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.



Opware SAS 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 Opware SAS, 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

Opware SAS will manage servers that run the following operating systems.

Each operating system that Opware SAS supports has a list of package types that you can upload. Opware SAS supports these package types on the supported operating systems, as the following table shows.*

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 |

Table 5-1: Supported Operating Systems and Package Types

| OPERATING SYSTEM | PACKAGE TYPE | FILE FORMATS | ADDITIONAL METADATA* |
|------------------|--|--|----------------------|
| 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 |

*For certain package types, the SAS Web Client requires that you provide additional meta-data for the package.

The Build Customization Script feature is available for Linux, Solaris, and Windows.



Opware SAS 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 Packages

This section provides information on AIX package management within Opsware SAS and contains the following topics:

- Overview of AIX Packages
- LPP Metadata

Overview of AIX Packages

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, Opsware SAS 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. Opsware SAS automatically creates package entries for all the filesets that the LPPs contain. When viewing an LPP in the SAS Web Client, 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 SAS Web Client shows filesets and APARs in the Installed Packages list for a server.

See the *Opsware[®] SAS User's Guide: Application Automation* for more information about how Opsware SAS manages AIX APARs.

LPP Metadata

Opsware SAS 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 updated fileset, a list of APARs addressed by the fileset
- For each APAR listed, the list of filesets that make up that APAR



Opware SAS 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, Opware SAS 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 Packages

This section contains the following topics:

- Overview of HP-UX Package
- Depot Metadata
- Preparing for 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

Overview of HP-UX Package

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 Opsware SAS.



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.
- Opware SAS supports products, filesets, and patch products as installable software.



Opware SAS 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. Opware SAS does not support using HP-UX codewords.

When a user uploads a depot, Opware SAS 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 SAS Web Client 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. Opware SAS attempts to install it anyway. An error results if a newer version is already installed.



Opware SAS 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 281 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, Opware SAS does overwrite newer versions of filesets and products with older versions. Opware SAS does not support relocating packages for HP-UX.

Depot Metadata

Opware SAS 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

Preparing for 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 260 in this chapter for more information. See “Example: File - Script to Split a Depot by Bundle” on page 260 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 Opware SAS.

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 Opware  
# 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 Opware 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 Packages

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

Opware SAS 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, Opware SAS performs the following actions:

- Opens the RPM and parses its metadata.
- Registers the RPM as an installable package.

Solaris Packages

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:
 - **File system format:** The format for packages stored in a directory structure.
 - **Data stream format:** The format for standalone package files. This format is required for uploading Solaris packages into Opware SAS.

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. Opware SAS adds package instances to and removes package instances from servers by using the reconcile function. See the *Opware[®] SAS Content Migration Guide* for more information about Reconcile.

In the SAS Web Client, you can upload, view, download, and delete Solaris packages, and you can view, deprecate, and attach to nodes the instances that they contain.

Opware SAS supports Solaris packages in the following ways:

- Users upload Solaris packages in the uncompressed data stream file format.
- Opware SAS can install interactive and non-interactive Solaris package instances. Interactive Solaris package instances require response files.

- Opware SAS 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 Opware Agent on managed servers includes its own admin file that it uses when installing Solaris package instances. The admin file that the Opware Agent uses is only used by Opware SAS 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.

Opware SAS supports response files in the following ways:

- Users create response files outside of Opware SAS by using the `pkgask` Solaris utility.
- By using the Package Properties page in the SAS Web Client, 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 Opware SAS 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, Opware SAS 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 SAS Web Client, Opware SAS might report an error when any server is reconciled with that Solaris package instance.

When a user uploads a Solaris package, Opware SAS 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

Opware SAS 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 data stream format before you can upload it to the Opware 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 Packages

Opware SAS 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

Opware SAS supports .msi files as installable software.

MSI Package Metadata

Opware SAS 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 Opware SAS, you are required to provide the 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.

To discover the ProductName and ProductVersion of an MSI package, perform the following steps:

- 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

Opware SAS 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 <http://www.microsoft.com> or by including the `--withmsi` option on the Opware Agent Installer command line.

See the *Opware® SAS User's Guide: Server Automation* for more information about the steps to install an Opware 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, Opware requires the user to provide the version of the service pack. When you upload Hotfixes and Security Patches, Opware requires the user to provide the operating system version and the patch type.

ZIP Packages

This section provides information on ZIP package management within Opware SAS and contains the following topics:

- ZIP Package Support
- ZIP Packaging
- ZIP Packages Creation
- ZIP Packages Upload
- Package Installation and Remove Scripts Definition
- Editing Properties for ZIP Packages
- Info-Zip Compatible ZIP Packages

- Windows Performance for Uploading Packages

ZIP Package Support

The SAS Web Client adds support for ZIP packages on the following operating systems:

- Windows NT4
- Windows 2000
- Windows 2003

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. Opware SAS 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.

Opware SAS supports ZIP encapsulation for application package files that were built using other standalone installation programs, for example, InstallShield.

Opware SAS requires silent install operation for programs designed for interactive installation. When you package these program files to upload to Opware SAS, 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.

ZIP Packages Creation

Opware SAS 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.

ZIP Packages Upload

You can upload, re-upload, download, delete, deprecate, and attach ZIP packages to nodes with the SAS Web Client. You can also use the Opaware 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 the *Opaware[®] SAS Content Utilities Guide* for information.

Package Installation and Remove Scripts Definition

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 Software Repository.

Editing Properties for ZIP Packages

After you upload a ZIP package to the Software Repository, the SAS Web Client 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 Opware SAS:

- 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

Opware SAS offers package management support for Info-Zip compatible.zip packages. The files that are archived within Info-Zip are installable files on Opware SAS. You can download the.zip package creation tool from www.info-zip.org.

Info-Zip Compatible Package Metadata

Opware SAS uses the ZIP package file name 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 Opware 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

The Package Management feature of Opware SAS provides tools for defining package types and uploading packages into Opware SAS.

The tasks that you can perform when you use this feature include:

- Displaying Packages
- Searching for Packages
- Viewing Packages Assigned to Nodes
- Overview of Uploading a Package
- Overview of Overwriting a Package
- Editing Package Properties
- Overview of Deleting a Package
- Overview of Deprecating a Package
- Downloading a Package

Displaying Packages

By default, the SAS Web Client 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 page displays only the packages that the user has permission to view, as Figure 5-1 shows. Packages are associated with customers and users are given permission to manage resources for specific customers in Opware SAS.

Figure 5-1: Packages: Browse Packages Page

Packages: Browse Packages

Browse Search Upload

All Types SunOS 5.9 All States Customer Independent Update

Delete Deprecate 1-10 of 1931 | Show All

| <input type="checkbox"/> | Name | Type | Size | Last Modified | Description |
|--------------------------|-----------|---------------|-----------|---------------|-------------|
| <input type="checkbox"/> | 107038-02 | Solaris Patch | 73 KB | 08/06/2003 | (not set) |
| <input type="checkbox"/> | 108434-08 | Solaris Patch | 916.63 KB | 07/30/2003 | (not set) |
| <input type="checkbox"/> | 108434-09 | Solaris Patch | 930.21 KB | 06/19/2003 | (not set) |
| <input type="checkbox"/> | 108435-06 | Solaris Patch | 773.55 KB | 07/30/2003 | (not set) |
| <input type="checkbox"/> | 108435-09 | Solaris Patch | 774.69 KB | 06/19/2003 | (not set) |
| <input type="checkbox"/> | 108528-14 | Solaris Patch | 16.36 MB | 07/30/2003 | (not set) |
| <input type="checkbox"/> | 108528-17 | Solaris Patch | 20.86 MB | 06/19/2003 | (not set) |
| <input type="checkbox"/> | 108652-47 | Solaris Patch | 7.58 MB | 07/30/2003 | (not set) |
| <input type="checkbox"/> | 108652-61 | Solaris Patch | 7.66 MB | 06/19/2003 | (not set) |
| <input type="checkbox"/> | 108725-08 | Solaris Patch | 208.54 KB | 07/30/2003 | (not set) |

Page 1 2 3 4 5 6 7 8 9 10 ...

For each package available for a selected customer, operating system, type, and state, the SAS Web Client displays the package's name, size, last modified date, and description, in addition to the selected operating system and customer.

- To display the properties of the package, click the package name.

See "Editing Package Properties" on page 281 in this chapter for more information.

Filters for Displaying Packages

The top panel of the page displays four filters that you can specify to qualify the packages that the SAS Web Client 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.
- **Package State:** Specifies whether the display lists all packages or only those that are available or deprecated. By default, the SAS Web Client 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 Opware administrator defines by using the Opware Administration feature from the navigation panel.

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 Select 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

Packages: Search Packages

Browse Search Upload

To search for a package, specify search parameters below and click Search

| | |
|---------------------------------------|---|
| Search for: | <input type="text"/> |
| Look for match in: | <input checked="" type="radio"/> Name <input type="radio"/> Description |
| Type: | All Types |
| Operating System: | Windows 2003 |
| State: | All States |
| Customer: | Customer Independent |
| <input type="button" value="Search"/> | |

- 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 SAS Web Client 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 a specific operating system.

- **State:** Specifies whether the display lists all packages or only those that are available or deprecated. By default, the SAS Web Client 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 Opware administrator defines by using the Opware Administration channel.
- 4** Click **Search**. The SAS Web Client 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 270 in this chapter for more information) or by searching (See “Searching for Packages” on page 272 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** Select 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

Package: View Nodes | 106300-21

[Return to Browse Packages](#)

Properties Nodes

The following nodes currently utilize this package. To delete a package, it cannot be utilized by any nodes. You may be able to click on a node to edit it.

| NODE NAME | NODE TYPE | DESCRIPTION |
|--|-----------|-------------------------------------|
| Patches SUNOS 5.7 SOL_PATCH 106300-21.tar.gz | Patches | 64-Bit Shared library patch for C++ |

- 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.

Overview of Uploading a Package

Each operating system supported by Opware SAS supports certain package types. See “Overview of Package Management” on page 251 in this chapter for more information.

If a package that is being uploaded already exists in the Software Repository, Opware SAS overwrites the package. See “Overview of Overwriting a Package” on page 280 in this chapter for more information.

If you upload Solaris patch clusters that contain patches that already exist in the Software Repository, the patches are overwritten. However, Opware SAS preserves any reboot options or flags set for the patches in the SAS Web Client.

If an update fileset is part of many different APARs, Opware SAS 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 SAS Web Client. However, Opware SAS 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 Opware SAS 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 264 in this chapter for information about how to obtain and use Orca.

Uploading a Package

Perform the following steps to upload a package:

- 1** From the navigation panel, click Software ► Packages. The Packages: Browse Packages page appears.
- 2** Select the Upload tab. The Packages: Upload Package page appears, as Figure 5-5 shows.

Figure 5-5: Packages: Upload Package, Specify Basic Properties Page

| Packages: Upload Package Specify Basic Properties | |
|--|--|
| <div style="display: flex; justify-content: space-around;"> Browse Search Upload </div> | |
| To upload a package, specify the following fields then click Next to proceed | |
| Customer: | <input type="text" value="MRXYZ"/> |
| Operating System: | <input type="text" value="HP-UX 11.11"/> |
| Type: | <input type="text" value="HP-UX Depot"/> |
| <input type="button" value="Next >>"/> | |

- 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 Opware administrator defines by using the Opware Administration feature from the navigation panel.

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.

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.



Opware SAS verifies that RPMs, Solaris patch clusters, AIX LPPs, Solaris packages, and HP-UX depots uploaded are the correct package type.

- 4** Click **Next**. 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

| Packages: Upload Package Specify Location | |
|--|---|
| Return to Upload Packages | |
| <input type="button" value="Browse"/> <input type="button" value="Search"/> <input type="button" value="Upload"/> | |
| To upload a package, specify the location of the package then click Upload when ready | |
| Customer: | MRXYZ |
| Operating System: | HP-UX 11.11 |
| Type: | HP-UX Depot |
| Encoding of metainformation in package. | ASCII |
| Local Path To Package: | <input type="text"/> <input type="button" value="Browse..."/> |
| <input type="button" value=" << Back"/> <input type="button" value=" Upload"/> | |
| After clicking Upload, the Upload Status window appears. This window closes after the package uploads. Do not click other links while this window is open. | |

- 5** If necessary, enter additional required information for the package and click **Next**. 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 **Browse** to locate and select the package to upload.
- In the “Encoding of metainformation in package” field, select the character encoding to be used by the package.

You need to specify an encoding so that Opware SAS can extract the metadata contained in the package and correctly display the information in non-ASCII characters in the SAS Web Client (for example, in the Package Properties pages). Package metadata includes comments, READMEs, scripts, descriptions, and content lists.

6 Click **Upload**.



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 SAS Web Client overwrites the existing file after it prompts you to confirm your choice. See “Overview of Overwriting a Package” on page 280 in this chapter for more information.

After you upload the package, the SAS Web Client 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 281 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.

Character Encoding for Package Metadata and Scripts

In Opware SAS, you can specify the character encoding for package metadata and scripts in the following ways:

- Specify the encoding for package metadata when uploading packages in the SAS Web Client (in the Packages feature in the navigation panel and Software Install Wizard) or by using the Opware Command Line Interface (OCLI).

When the encoding is specified, the SAS Web Client 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 for scripts when uploading them in the SAS Web Client (in the Run Distributed Script Wizard and Scripts channel).

Opware SAS converts the script contents from the UTF-8 encoding to the encoding that you select. Internally, Opware SAS stores the script in the UTF-8 encoding.

After a script runs, you can download a ZIP file that contains the results encoded in UTF-8 format. For example, on Unix you can use the `iconv` program to interpret the downloaded results of the script execution.

The SAS Web Client includes the following selections for character encodings:

- Arabic (ISO-8859-6)
- Baltic (Cp1257)
- Baltic (ISO-8859-13)
- Baltic (ISO-8859-4)
- Central European (Cp1250)
- Central European (ISO-8859-2)
- Chinese Hong Kong, Taiwan (Cp950)
- Chinese Simplified (EUC-CN)
- Chinese Simplified (GB18030)
- Chinese Simplified (GBK)
- Chinese Traditional (Big5)
- Chinese Traditional (Big5-HKSCS)
- Chinese Traditional (EUC-TW)
- Cyrillic (Cp1251)
- Cyrillic (ISO-8859-5)
- Cyrillic (KOI8-R)
- English (US-ASCII)
- Greek (Cp1253)

- Greek (ISO-8859-7)
- Hebrew (Cp1255)
- Hebrew Visual (ISO-8859-8)
- Japanese (EUC-JP)
- Japanese (ISO-2022-JP)
- Japanese (Shift_JIS)
- Korean (Cp949)
- Korean (EUC-KR)
- Korean (JOHAB)
- South European (ISO-8859-3)
- Thai (TIS-620)
- Turkish (Cp1254)
- Turkish (ISO-8859-9)
- Unicode (UTF-8)
- Vietnamese (Cp1258)
- Western (Cp1252)
- Western (ISO-8859-1)
- Western (ISO-8859-15)

Overview of 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 SAS Web Client, 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, Opsware SAS 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. Opsware SAS 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 SAS Web Client are not affected.

See “Editing Package Properties” on page 281 in this chapter for more information.

Overwriting a Package

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 270 in this chapter for more information. Or locate the package by searching. See “Searching for Packages” on page 272 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 **Replace** 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 **Browse** to locate and select the package to upload.

- 5** Click **Upload**.

Editing Package Properties

After you upload a new file or select an existing package in the Software Repository, the SAS Web Client 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 270 in this chapter for more information) or by searching (See “Searching for Packages” on page 272 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 the RPM to be installed with the -U option, which removes the old version of the RPM and installs the new version in a single step. By default, the Upgradable check box is selected. The -U option is used in the following two cases. First, the RPM is upgradable, a version of the RPM has already been installed, but not with Opware SAS, and the installed version has not been adopted by Opware SAS. Second, the RPM is upgradable, one or more versions of the RPM have already been installed with Opware SAS, these versions are upgradable, and these versions are slated to be removed during this reconcile operation. In other cases, even if the Upgradable check box is selected, the RPM is installed with the -i option.
 - **Install Flags:** Specify to run these optional arguments 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:** Specify to run these optional arguments 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 **Upload** 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. Specify the character encoding for the response file.
 4. Click **Upload**. The popup window closes and the response file name appears in the page.

See “Solaris Packages” on page 261 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 **Delete** 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 **Replace** and specify the path to the new file. The file can have the same name or a different file name as the existing file. Opware SAS replaces the response file with the new file.
- 7** Click **Save**.



You can also replace (overwrite), delete, or download a package directly from the Edit Properties page.

Overview of 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.



Opware, Inc. recommends that you do not delete packages that have been installed on the servers by Opware SAS. Deleting such a package that is in use can cause problems in Opware SAS.

Deleting a Package


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 **Delete**. A confirmation page appears that prompts you to confirm the deletion.
The SAS Web Client 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

Packages: Delete

Return to Search Packages

| Packages That Cannot Be Deleted | | | |
|--|---------------------|---------------------|---|
| | Package Name | Description | Reason |
|  | Apache HTTPD Server | Apache HTTPD Server | You will need to first remove this package from nodes: <ul style="list-style-type: none"> • Web Servers Apache Server 1.3.28 |

- 4** Click **Delete** to remove the packages from the Software Repository.

Overview of 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, Opware SAS installs the deprecated fileset. Additionally, users can install available update filesets for deprecated APARs. See “AIX Packages” on page 255 in this chapter for more information. See the *Opware® SAS User’s Guide: Application Automation* for information about how Opware SAS 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 281 in this chapter for more information.

Deprecating a Package

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 **Deprecate**. The Packages: Deprecate page appears, as Figure 5-8 shows.

Figure 5-8: Packages: Deprecate Page

Packages: Deprecate
[Return to Search Packages](#)

Packages That Can Be Deprecated

To deprecate the selected packages below, click Deprecate.

| | Package Name | Size | Last Modified | Description |
|-------------------------------------|--------------|-----------|---------------|-------------|
| <input checked="" type="checkbox"/> | 106300-21 | 719.85 KB | 08/20/2003 | |

Optionally create a common notes entry for the above packages

Notes:

Deprecate **Cancel**

- 4** To record the version of a new package that replaces the deprecated version, enter the information in the Notes field.
- 5** Click **Deprecate**.

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 270 in this chapter for more information. Locate the package by searching. See “Searching for Packages” on page 272 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. Click **Download** at the bottom of the page.
Or

In the File Name field, click on the link to download the package.

Figure 5-9: Downloading a Package

| Properties | |
|---|--|
| Enter appropriate changes below then click Save | |
| Name: | ISMismto0-2.0.1-1.pkg |
| Version: | 2.0.1-1 |
| File Name: | ISMismto0-2.0.1-1.pkg |
| File Type: | DATASTREAM |
| Size: | 14.99 MB |
| Last Modified: | Tue Mar 29 18:30:41 2005 by (not set) |
| Created: | Tue Mar 29 18:30:55 2005 by ocli |
| Description: | (I)ntelligent (S)oftware (M)odule tool (ISM Application Package) <small>Enter up to 2,000 characters in this field</small> |
| Customer: | Customer Independent |
| Operating System: | SunOS 5.9 |
| Type: | Solaris Package |
| Solaris Package Instance: | ISMismto0-2.0.1-1 |
| Notes: | |
| Package ID: | 2390003 |
| <input type="button" value="Save"/> <input type="button" value="Delete"/> <input type="button" value="Download"/> <input type="button" value="Cancel"/> | |

Chapter 6: Software Provisioning

IN THIS CHAPTER

This section discusses the following topics:

- Installing and Uninstalling Software
- Software Installation and Uninstallation Options
- Template Installation

Installing and Uninstalling Software

This section provides information on how to install and uninstall software by using Opsware SAS and contains the following topics:

- Overview of Installing and Uninstalling Software
- Software Package Management in the Model Repository
- Ways to Install Software
- Opsware SAS Categories for Software
- Software Provisioning Functionality
- Platform-Specific Reconcile

Overview of Installing and Uninstalling Software

Opsware SAS automates the time-consuming process of installing and uninstalling software on servers. Using Opsware SAS, you can quickly deploy applications across a large number of servers with a minimum amount of downtime. Opsware SAS allows you to select the 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.

Opware SAS 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. Opware SAS also provides the same degree of control for uninstalling software.

You can uninstall any software that you installed through Opware SAS. You can also preview the results of an installation and then schedule the installation for some later time. If, for example, Opware SAS 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 Package Management in the Model Repository

In Opware SAS, packages reside in a central Software Repository. Opware 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 packages that help control the way that the software is installed.

See the *Opware® SAS Configuration Guide* for information about how to upload packages to and manage packages in the Software Repository.

Opware SAS maintains detailed information about the state of every server under management in a central database called the Model Repository.

Ways to Install Software

Opware SAS provides several ways to install software. You can:

- Select a single package and install it.
- Select a template that includes a number of different (and usually related) software that you can install in a single operation.
- Use Visual Packager to prepare installable software packages. See the *Opware® SAS Policy Setter's Guide* for information about the Visual Packager.

Opware users create and test the templates, and they provide the same level of control and consistency that is available when you install individual software. Templates, for example, allow you to install a set of applications, such as a Web server, an Oracle database, and related software that allow you to quickly deploy a server in a fully operational state.

Opware SAS Categories for Software

The Model Repository usually stores information about hundreds of different types of packages. To make the selection of software easier, Opware SAS 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 Opware SAS supports. You can browse through the list of software or use advanced search capabilities in Opware SAS to locate the packages or templates that you want to apply.

Software Provisioning Functionality

The Software Provisioning feature has the following functionality:

- A central location (the Software Repository) for 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 packages quickly across a large number of servers
- The ability to quickly apply 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

The Reconcile step is the last step that you perform when installing software. During that step, Opware SAS 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. See Appendix A for information about reconcile.

Table 6-1 gives details about reconcile operations for various operating systems. This table shows the utilities used during the reconcile process for all supported operating systems.

Table 6-1: Reconcile Operation by OS

| OPERATING SYSTEM | SEE THIS PAGE |
|------------------|--------------------------------------|
| AIX | See "AIX Reconcile" on page 357. |
| HP-UX | See "HP-UX Reconcile" on page 357. |
| Linux | See "Linux Reconcile" on page 358. |
| Solaris | See "Solaris Reconcile" on page 358. |

Software Installation and Uninstallation Options

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
- Package Installation on Servers with Low Disk Space
- Uninstalling Software with the Uninstall Software Wizard

Script Error Conditions

A 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 Opware SAS 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 the *Opware[®] SAS Configuration Guide* for information about how to set up inheritance for software nodes.

Ordinarily, a branch in the tree is set up for one 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 the 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



The wizards are flexible; for example, when you select multiple servers and applications, Opware SAS will install the correct applications on the servers even if you selected different OS versions for the servers and applications in the respective steps. Opware SAS also matches the customer association for applications, operating systems, and templates with the customer association of servers. If Opware SAS cannot find a match, an error message appears at the end of the wizard; therefore, use caution when modifying these default values.

Perform the following steps to install software with the Install Software Wizard:

- 1** From the Opware SAS Web Client home page, click the Install Software link in the Tasks panel.

Or

From the navigation panel, click Software ► Applications. The Application page appears. Select the application and click **Install**.

- 2** The Select Software page appears, as Figure 6-1 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. Select one or more software to install. Figure 6-1 shows an example of possible software selections.

Figure 6-1: Select Software Page

Select Software
Select the software application(s) you wish to install.

Browse | **Search** (1) Software Selected

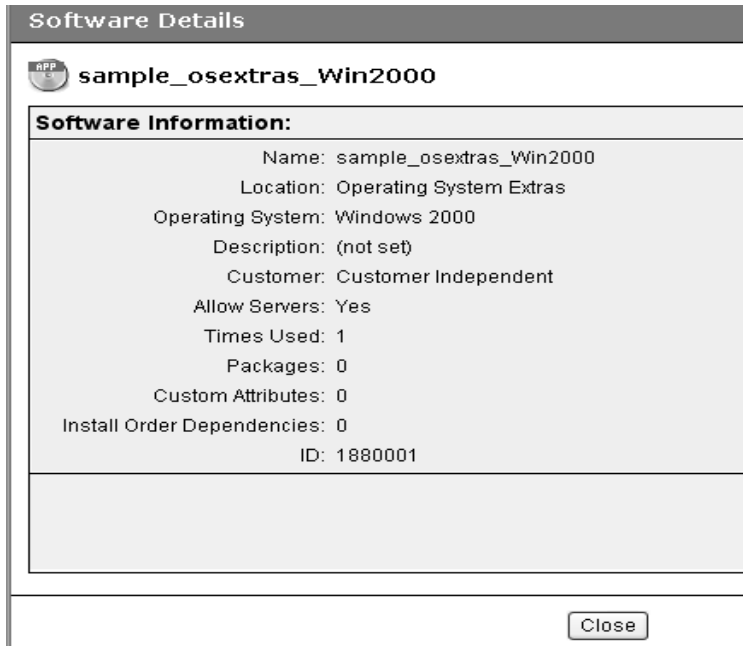
Windows

2 Total

| | Name | OS Version | Modified Date | Customer |
|-------------------------------------|------|--------------|---------------|----------------------|
| <input type="checkbox"/> | 2000 | Windows 2000 | 05-04-2005 | Customer Independent |
| <input checked="" type="checkbox"/> | 2003 | Windows 2003 | 05-04-2005 | Customer Independent |

- 4 (Optional) Click the name of any of the software to display additional information about the software, as Figure 6-2 shows.

Figure 6-2: Software Details Page



- 5** Click **Next** to continue. The Select Servers page appears, as Figure 6-3 shows.

Figure 6-3: Select Servers Page

The screenshot shows the 'Select Servers' page within the 'Install Software' wizard. On the left, a navigation pane lists steps: 1. Select Software, 2. Select Servers (highlighted), 3. Confirm Selections, 4. Preview, 5. Schedule and Notify, and 6. Install. The main area is titled 'Select Servers' and contains a search bar with 'Browse' and 'Search' tabs. Below the search bar, there are dropdown menus for 'All Customers' and 'Select OS Version', along with 'OK' and 'Update' buttons. A status bar indicates '(0) Items Selected' and '2 Total'. The main content is a table with the following data:

| | Name | Host Name / IP Address | OS Version | Customer |
|--------------------------|----------------------|--------------------------------------|-------------------------------|--------------|
| Public Groups | | | | |
| <input type="checkbox"/> | core.tr3.opsware.com | core.tr3.opsware.com 172.16.36.18 | Red Hat Enterprise Linux AS 3 | Opware |
| <input type="checkbox"/> | m201.tr3.opsware.com | mh03 172.16.36.34 | Windows 2000 | Not Assigned |

At the bottom of the page, there are 'Cancel', '< Previous', and 'Next >' buttons.

The selection of servers is defaults to common operating system and customer of the software that you selected.

- 6** Select the servers or groups 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** (Optional) 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.

Or

Click **Skip Preview** to skip the preview process.

- 8** On the Schedule and Notify page, you have the following options:

- **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 (+) button next to the Recipients field.
- **Schedule:** Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

- **Option 1:** Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2:** Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.



The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opware SAS core server (usually UTC).

- 9** (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.
- 10** Click **Close** 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 “Server Management Scheduling and Notification” on page 154 in Chapter 3 for information about the My Jobs feature.

Package Installation 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 Opware SAS should look for packages to install.



To accomplish this task you need to specify the location where Opware SAS should look for packages to install by defining a new custom attribute. See “Custom Attributes for Servers” on page 160 in Chapter 3 for more information.

You can enter any number of paths that Opware SAS attempts to use to find the package. Opware SAS tries each path, one after the other, until the package is located. If the package is not found in any of the specified locations, Opware SAS looks in the Software Repository. In that case, Opware SAS attempts to download the package if there is enough disk space. If Opware SAS 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 Opware SAS has read access.

Table 6-2 shows the operating systems and package types that you can use with the Low Disk Space feature.

Table 6-2: 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 install packages are in `/tmp/sys/inst.images`, run the command as follows:

```
cd /tmp/sys/inst.images ; inutoc
Or
inutoc /tmp/sys/inst.images
```

Specifying Paths for Package Installation

Perform the following steps to specify paths for package installation:

- 1** Log into the Opware SAS Web Client.
- 2** From the navigation panel, select Servers ► Manage Servers.
- 3** Navigate to the servers where you are defining a new custom attribute, and select the check box next to the name of each server that you are defining a new custom attribute for.
- 4** Select the Custom Attributes tab. The Manage Servers: Custom Attributes | [server name] page appears.
- 5** Click **New** 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 Opware SAS should look for the package to install. For example:

```
/mnt/cdrom
```

Or

```
/networkshare/packages/SunOS/5.6/
```

/mnt/cd0



You can specify the Software Repository as a path by using `opware_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.

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** From the Opware SAS Web Client home page, click the Uninstall Software link in the Tasks panel.

Or

From the navigation panel, click Software ► Applications. The Application page appears. Select the application and click **Uninstall**.

- 2 The Select Servers Page appears. 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 6-4 shows.

Figure 6-4: OS Versions Page

The screenshot shows the 'Uninstall Software' application interface. The title bar reads 'Uninstall Software' with a help icon. On the left is a sidebar with a progress indicator showing six steps: 1. Select Servers (highlighted), 2. Select Software, 3. Confirm Selections, 4. Preview, 5. Schedule and Notify, and 6. Uninstall. The main content area is titled 'Select Servers' and includes the instruction 'Select the server or servers you wish to modify.' Below this are 'Browse' and 'Search' tabs, and a filter set to 'OS Version'. A status bar indicates '(0) Servers Selected' and '34 Total'. A table lists the following OS versions:

| Name |
|-------------|
| AIX 4.3 |
| AIX 5.1 |
| AIX 5.2 |
| AIX 5.3 |
| HP-UX 10.20 |
| HP-UX 11.00 |

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.

- 3 Select servers or groups 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 6-5 shows.

Figure 6-5: 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.

If you selected a group of servers that has no software nodes in common, this list is empty.

- 4** 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.
- 5** 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.

Or

Click **Skip Preview** to skip the preview process.

- 6** (Optional) 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.

- 7** On the Schedule and Notify page, you have the following options:
 - **Schedule:** Click either **Run Now** to execute the operation immediately, or click **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 (+) button next to the Recipients field.

If you chose to uninstall immediately, progress bars appear that indicate the status of the uninstallation process. You can watch the progress bars to see when the software is uninstalled.

- 8** (Optional) When the process is complete, click **View Details** to see what packages were uninstalled or removed, and the output, if any, for the uninstallation scripts.
- 9** Click **Close** to end the wizard.

Closing the wizard does not stop the uninstallation if it is still running. After you close the wizard, you can view the progress of the uninstallation by viewing My Jobs.

See “Server Management Scheduling and Notification” on page 154 in Chapter 3 for information about the My Jobs feature.

Template Installation

This section provides information on how to install templates within Opsware SAS and contains the following topics:

- Overview of Installing Templates
- Installing Templates with the Install Templates Wizard

Overview of Installing Templates

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 the *Opware® SAS Configuration Guide* for information about how to create templates in Opware SAS.

You cannot uninstall the packages in a template as a group through an Opware wizard. You can, however, use the Opware 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 From the Opware SAS Web Client home page, click the Install Template link in the Tasks panel.

Or

From the navigation panel, click Servers ► Manage Servers. The Manage Server page appears. Select the servers and click **Tasks ► Install ► By Template**.

- 2 The Select Template page appears, as Figure 6-6 shows.

Figure 6-6: Install Templates Wizard Page

| Browse | | Search | | |
|-----------------------|--------------|----------------|------------|----------------------|
| Template | | | | |
| | | | | 2 Total |
| | Name | OS Version | Times Used | Customer |
| <input type="radio"/> | Opware | OS Independent | 0 | Opware |
| <input type="radio"/> | Opware Tools | OS Independent | 0 | Customer Independent |

- 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.
- 4 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. Select the servers or groups that you want to apply the template to by browsing or searching, and click **Next** to continue.
- 5 The Confirm Selection page appears. The Confirm Selection page displays the servers and template that you selected to install.
- 6 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.

Or

Click **Skip Preview** to skip the preview process.

- 7 (Optional) 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.
- 8 On the Schedule and Notify page, you have the following options:
 - **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 (+) button next to the Recipients field.
 - **Schedule:** Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:

- **Option 1:** Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.

- **Option 2:** Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.



The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opware SAS core server (usually UTC).

If you choose to install immediately, progress bars appear that indicate the status of the installation process on each server that you selected.

9 (Optional) When the process is complete, click **View Details** to see what templates were installed and the output, if any, from the installation scripts.

10 Click **Close** 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 “Server Management Scheduling and Notification” on page 154 in Chapter 3 for information about the My Jobs feature.

Chapter 7: Configuration Tracking Using the Software Tree

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Configuration Tracking
- Change Detection
- Configuration Tracking Policies
- Node-Based Tracking Policies

Overview of Configuration Tracking

The Configuration Tracking feature of Opsware SAS allows you to monitor critical configuration files and configuration databases. When Opsware SAS 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. A configuration tracking policy consists of one or more configuration tracking policy entries that specifies the configuration file, the directory of configuration files, or the configuration database that you want to track.

The Configuration Tracking feature is designed for flexibility. For example, you can set configuration tracking defaults for a particular software application, and all attached servers 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.

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 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 Configuration Tracking feature is not a general-purpose backup solution. 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.

File Types Supported

You can use the 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

Change Detection

All servers that Opware SAS manages have an Opware Agent installed on them. On servers that use Configuration Tracking, every four hours the Opware Agent inspects the configuration files and databases that you select to track.

The Opware 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 Opware 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 317 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.

Actions Performed

Opsware SAS 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).

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 Opware Agent on a server checks to see if any files have changed since the last full backup. If any files have changed, Opware SAS performs a new full backup. If no files have changed, the full backup does not take place.

You can also force Opware SAS to perform a full backup on a server by selecting the Perform Manual Backup option. (See the *Opware® SAS User's Guide* for more information about Performing Manual Backups.)

See the *Opware® SAS User's Guide* for more 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 the *Opware® SAS User's Guide* for more information about the procedure for deleting backups. See the *Opware® SAS Administration Guide* for information about mass deletion of backup files.

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/file111
```

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.

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. Opsware SAS 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
```

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` attributes 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.

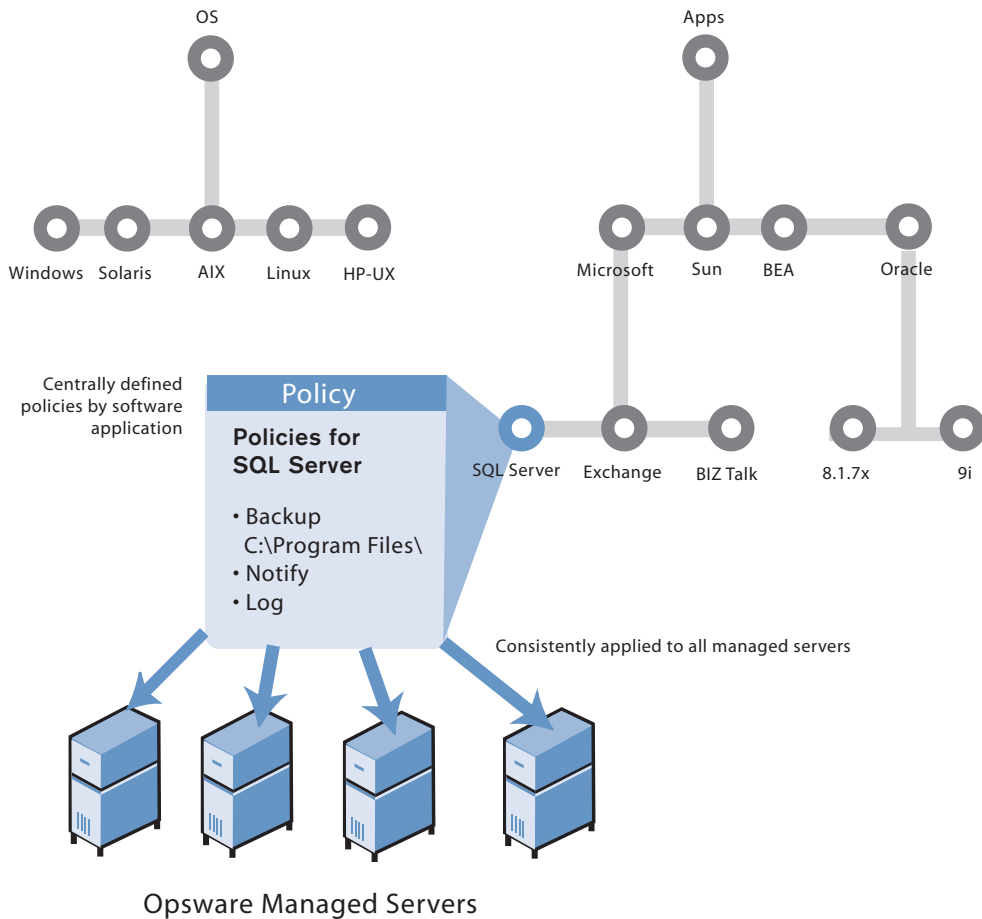
Configuration Tracking Policies

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, the directory of configuration files, or the 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. See Figure 7-1.

Figure 7-1: Configuration Tracking Policy



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).

Ways to Create Tracking Policies

You can create tracking policies in two ways:

- By using Opware nodes. (See “Node-Based Tracking Policies” on page 320 in this chapter for more information.)
- By creating custom tracking policies for a selected server or set of servers. (See the *Opware® SAS User’s Guide* for more information.)

The preferred way to create a tracking policy is by using Opware nodes. Among other functions, an Opware 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.

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, Opware SAS searches for files that match the pattern that you specify.
- If you use wildcards in a directory target, Opware SAS 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 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 file name 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, Opsware SAS is able to keep a more complete record of the file; Opsware SAS 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.

Conditions for 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 Configuration Tracking feature to delete files that were created in monitored directories. Because Opsware SAS was not tracking the file specifically, Opsware SAS 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.

Conditions for 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 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 following differences apply:

- 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.

Configuration Tracking Policy Limits

Because Configuration Tracking is not a general-purpose backup solution, Opware SAS 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 the *Opware® SAS User's Guide* for more information about Deleting Backups. See the *Opware® SAS Administration Guide* for information about mass deletion of backup files.

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 2500 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 "Tracking Policy Deployment" on page 324 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.

See the *Opware® SAS Administration Guide* for information about how to use the Opware SAS configuration parameters.

If the server does not have a backup notification list assigned to it, the warning message is sent only to the error email address.

Node-Based Tracking Policies

This section provides information on node-based tracking policies within Opware SAS and contains the following topics:

- Node-Based Tracking Policies
- Creating Node-Based Policy Entries
- Tracking Policy Deployment
- Reconciling a Node's Configuration Tracking Policy
- 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

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.

See "Node Management in the Software Tree" on page 183 in Chapter 4 for more information.

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 node:

- 1 From the navigation panel, click Software and then select the relevant type of software (for example, patches, applications, and so forth), as Figure 7-2 shows.

Figure 7-2: Selecting Node Types

The screenshot shows the Opsware interface. On the left is a navigation panel with a tree structure. The 'Software' node is expanded, and 'Applications' is selected. At the top right, there is a search bar with a dropdown menu set to 'Applications' and a 'Go' button. The main content area is titled 'Applications' and contains a table with the following data:

| | Name | Modified |
|--|-------------------------|------------|
| | Application Servers | 05-06-2005 |
| | Database Servers | 05-06-2005 |
| | Operating System Extras | 05-06-2005 |
| | Other Applications | 05-06-2005 |
| | System Utilities | 05-06-2005 |
| | Web Servers | 05-06-2005 |

- 2 Navigate to the node that you want to create a tracking policy for.
- 3 Select the Config Tracking tab.

- 4 Click **Add Entry**. The Track Configurations: Add Entry page appears, as Figure 7-3 shows.

Figure 7-3: Track Configurations: Add Entry

Return to Config Tracking

Add Entry to Other Applications Sample Application

Please complete the following form to add an entry:

| | |
|---|---|
| Type: | File <input type="text"/> |
| Target: | <input type="text"/> |
| | <input type="checkbox"/> Include sub-directories |
| Action(s): | <input checked="" type="checkbox"/> Backup <input type="checkbox"/> Email Backup Notification List for Server <input type="checkbox"/> Email <input type="text"/> <input type="checkbox"/> Log <input type="text" value="Info"/> |
| <input type="button" value="Save"/> <input type="button" value="Cancel"/> | |

- 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 7-1 describes each of the selections that you must make.

Table 7-1: Add Entry: Configuration Tracking

| FIELD | DESCRIPTION |
|-------|--|
| Type | <p>File: Monitor the file specified in the target field.</p> <p>Directory: Monitor all the files in the directory specified in the target field. The following types are available only for Windows servers.</p> <p>Windows Registry: Specify key in target field.</p> <p>IIS Metabase: The entire Metabase is monitored; do not specify target.</p> <p>COM + Registration Database: The entire Registry is monitored; do not specify target.</p> |

Table 7-1: Add Entry: Configuration Tracking

| FIELD | DESCRIPTION |
|---|---|
| Target | <p>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.</p> <p>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.)</p> <p>If you select the file or directory type, you can use wildcards in the target. (See “Configuration Tracking Policy Targets and Wildcards” on page 316 in this chapter for more information.)</p> <p>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, <code>HKEY_LOCAL_MACHINE\SOFTWARE</code>). If you selected the IIS Metabase or COM + Registration Database type, do not specify the target.</p> |
| Actions (you can select multiple options) | <p>Backup: Back up the specified file, directory, COM + Registration Database, Windows Registry keys, or IIS Metabase.</p> <p>Email Backup Notification List for Server: Send an email to the backup notification list for the selected server. See “Email Notification List” on page 313 in this chapter for more information.</p> <p>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.)</p> <p>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:</p> <ul style="list-style-type: none"> 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.

See “Reconciling a Node’s Configuration Tracking Policy” on page 324 in this chapter for more information.

Tracking Policy Deployment

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 324 in this chapter for information about the procedures for performing configuration tracking reconciles.

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.

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 (for example, patches, applications, and so forth).
- 2 Navigate to the node whose tracking policy you want to reconcile.
- 3 Select the Members 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 check box at the top of the list to select all the servers in the list.)

- Under the **Configuration Tracking** menu, select **Reconcile Tracking Policies** as Figure 7-4 shows.

Figure 7-4: Reconcile Tracking Polices

Managed Servers: Summary View

All Status ▾ All Stages ▾ All Uses ▾ All Facilities ▾ Intel Corporation ▾ Update

| Server | Software | Configuration Tracking | View | Name / IP Address | OS Version | Stage | Use | Facility |
|-------------------------------------|--|--|------|---|------------|---------------|------------|----------------------------|
| <input type="checkbox"/> | m0178white: Padma's - C... | <ul style="list-style-type: none"> Edit Tracking Policies Perform Backup Now View Backup History View Restore Queue Enable/Disable Tracking Reconcile Tracking Policies | | m0178white.cust.custqa4.com 88.160.71 | AIX 4.3 | Not Specified | Production | Folsom Data Center (core0) |
| <input checked="" type="checkbox"/> | m072.goldsox.qa.opsware.com Padma's - Curie.b | | | m072.goldsox.qa.opsware.com 192.168.160.72 | AIX 5.1 | Not Specified | Staging | Folsom Data Center (core0) |
| 2 Total | | | | | | | | |

At this point, Opware SAS performs a test reconcile to ensure that the reconcile can be performed without errors. The progress of the test reconcile process displays. Click **View Details** if you want to see information about what changes will be made when the reconcile is performed.

- 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.

Viewing a Node's Tracking Policy

Perform the following steps to view a node's tracking policy:

- From the navigation panel in the SAS Web Client, click Software then select the type of software.
- Navigate to the node whose tracking policy you want to examine.

- 3 Select 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 7-5 shows.

Figure 7-5: Viewing a Node's Configuration Policy

| Properties | Packages 0 | Custom Attributes 0 | Install Order 0 | Servers 1 | Config Tracking | History |
|--|-----------------------|---------------------|---|------------------|-----------------|------------|
| View: Enabled Entries Update | | | | | | |
| Add Entry | | | | | | |
| <input type="checkbox"/> | Target | Type | Action(s) | | | Inherited? |
| <input type="checkbox"/> | c:\configfiles\config | File Object | Backup, Email Backup Notification List for Server, Log Info | | | No |
| Disable | | Delete | | selected entries | | |

Tracking policy entries are identified by target. If a node inherits two or more policy entries from its parent nodes that have the same target, these entries display as a single entry when you view the child's and node'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 (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, select **Enabled Entries** from the **View** 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 7-6 shows.

Figure 7-6: Track Configurations: Edit Entry Page

| Track Configurations: Edit Entry c:\configfiles\config | |
|---|---|
| Return to Config Tracking | |
| Edit Entry in Other Applications Sample Application | |
| Please complete the following form to edit this entry: | |
| Source: | Other Applications Sample Application |
| Type: | File <input type="button" value="v"/> |
| Target: | c:\configfiles\config <input type="checkbox"/> Include sub-directories |
| Action(s): | <input checked="" type="checkbox"/> Backup <input checked="" type="checkbox"/> Email Backup Notification List for Server <input type="checkbox"/> Email <input type="text"/> <input checked="" type="checkbox"/> Log <input type="button" value="v"/> Info |
| <input type="button" value="Save"/> <input type="button" value="Cancel"/> | |

- 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 (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** 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 selecting the first check box in the list.
- 5 Click **Disable**.
- 6 Click **Disabled Entries** 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 (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** 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 selecting the first check box in the list.
- 5** Click **Delete**.
- 6** Click **Delete Entries** 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 (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 selecting the first check box in the list.
- 5** Click **Enable**.
- 6** Click **Enabled Entries** to commit the changes.



You must perform a configuration tracking reconcile to deploy the changes to the appropriate servers.

Chapter 8: Operating System Provisioning

IN THIS CHAPTER

This section discusses the following topics:

- Supported Operating Systems for OS Provisioning
- OS Installation with the SAS Web Client



Before users can install operating systems on servers with the OS Provisioning feature, the operating systems must be defined in Opware SAS. The OS media must be made available in Opware SAS. Additionally, OS Installation Profiles must be created in Opware SAS.

This section shows you how to install an OS using the OS Provisioning feature using the Opware SAS Web Client.

For overview info on OS provisioning and information on how to perform OS provisioning from the SAS Client, see the OS Provisioning chapter *Opware[®] SAS User's Guide: Application Automation*.

Supported Operating Systems for OS Provisioning

The OS Provisioning feature supports installation of the following versions of Linux, Sun Solaris, VMware ESX, and Microsoft Windows operating systems (architecture support is listed in parentheses):

- Red Hat Linux 6.2 (x86)
- Red Hat Linux 7.1 (x86)
- Red Hat Linux 7.2 (x86)
- Red Hat Linux 7.3 (x86)
- Red Hat Linux 8.0 (x86)

- Red Hat Linux Enterprise Linux 2.1 AS (x86)
- Red Hat Linux Enterprise Linux 2.1 ES (x86)
- Red Hat Linux Enterprise Linux 2.1 WS (x86)
- Red Hat Linux Enterprise Linux 3 AS (x86, x86_64, and ia64)
- Red Hat Linux Enterprise Linux 3 WS (x86, x86_64, and ia64)
- Red Hat Linux Enterprise Linux 3 ES (x86, x86_64, and ia64)
- Red Hat Linux Enterprise Linux 4 AS (x86 and x86_64)
- Red Hat Linux Enterprise Linux 4 WS (x86 and x86_64)
- Red Hat Linux Enterprise Linux 4 ES (x86 and x86_64)
- Sun Solaris 2.6 (SPARC)
- Sun Solaris 7 (SPARC)
- Sun Solaris 8 (SPARC)
- Sun Solaris 9 (SPARC)
- Sun Solaris 10 (SPARC x86 and x86_64)
- Fujitsu Solaris 8 (SPARC)
- Fujitsu Solaris 9 (SPARC)
- Fujitsu Solaris 10 (SPARC)
- Fujitsu Solaris 10 Update 2 (SPARC and Niagara)
- SUSE Linux Standard Server 8 (x86)
- SUSE Linux Enterprise Server 8 (x86)
- SUSE Linux Enterprise Server 9 (x86 and x86_64)
- SUSE Linux Enterprise Server 10 Update 1 and Update 2
- SUSE Linux Enterprise Server 10 (x86_64)
- VMware ESX 3
- VMware ESX 3 (x64)
- Windows NT 4.0
- Windows 2000

- Windows 2003
- Windows 2003 x64
- Windows XP Professional
- Windows XP Professional x64



Opware OS Provisioning does not support Windows OS provisioning on Itanium-based systems.



In order to provision a bare metal SPARC sun4u server with any of the supported SPARC Solaris versions, the server must support Solaris 10 U3 (06/06).

The OS Provisioning feature works with a floppy disk for Windows via the DOS preinstallation environment, a CD-ROM for Windows via the WinPE preinstallation environment, a CD-ROM for Linux, or network booting for all supported operating systems. Non-network booting is not supported for Sun Solaris (SPARC and x86).



The OS Provisioning feature does not provision HP-UX or AIX operating systems. However, you can integrate Opware SAS with Network Installation Management (NIM) to provision AIX and Ignite-UX to provision HP-UX. See the *Opware® SAS Administration Guide* for more information on how to integrate the Opware SAS with HP-UX and AIX OS provisioning systems.

OS Installation with the SAS Web Client

This section provides information on OS installation with the SAS Web Client and contains the following topics:

- Servers and the Server Pool
- 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

Servers and the Server Pool

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 Opware SAS but do not have the target OS installed. From there, you can install an OS by clicking **Install OS**. See Figure 8-1.

Figure 8-1: Server Pool List in the SAS Web Client



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

- The host name set by booting the server the first time over the network or by using an Opware Boot Floppy or CD
- The MAC address or chassis ID
- The manufacturer 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 life cycle associated with the server
- The customer and facility association

- Additional hardware information. (Clicking the server name opens a window that displays specific hardware information, as Figure 8-2 shows.)

Figure 8-2: Information Displayed in the Edit Server Page for a Server in the Server Pool



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.

Figure 8-3: Example of an OS Provisioning Baseline



- Performing a custom installation, which includes defining the installation on-the-fly by selecting an OS installation profile, which includes options for 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 Opware wizard itself or through the My Jobs interface.

See the *Opware® SAS Content Migration Guide* for more 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 SAS Web Client 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 **Install OS**.

Or

From the navigation panel, click Software ► Operating Systems. the Operating System page appears. Click the Install OS link.

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 8-4.

Figure 8-4: Overview Page in the Install Operating System Wizard



- 2 If necessary, select the Use Template option and click **Start**. If you did not select servers in Step 1, the Select Servers page appears.
- 3 If prompted, select the servers or groups that you want to provision and click **Next**. You can find the servers that you want to provision by browsing the list or by searching.

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 **Next**. The Assign Customer page appears.

By default, the OS Provisioning feature entered the server host name 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 Opware SAS Web Client UI.

- 5** Select a customer for the servers and click **Next**. 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 *Opware® SAS User's Guide: Server Automation* for information about the distinction between the customers Customer Independent and Not Assigned.

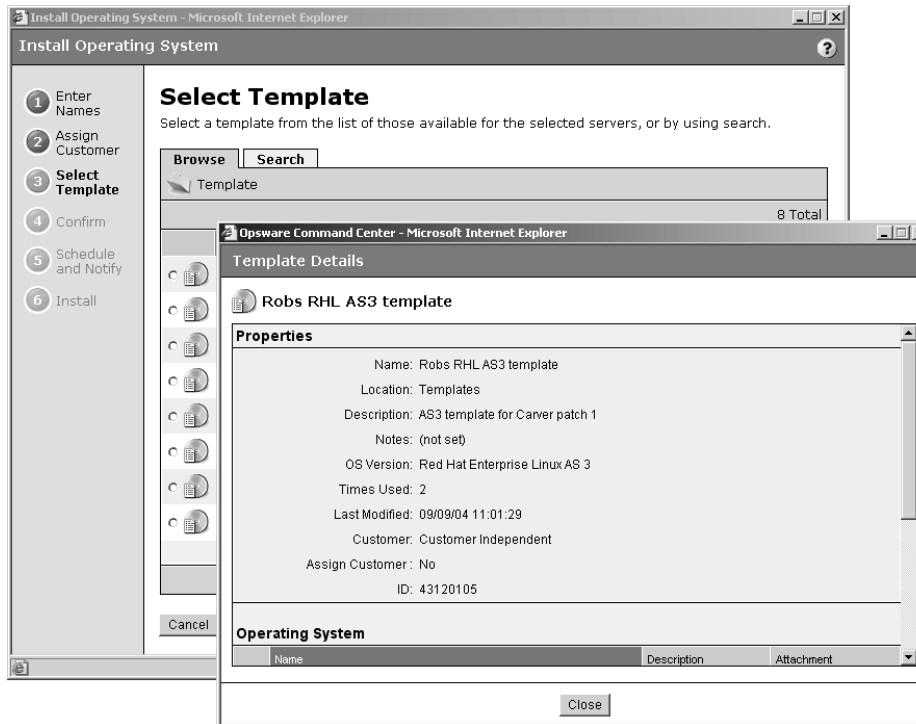
- 6** Select the template to use to provision the servers and click **Next**. Templates appear in the list only if they meet these requirements:

- 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 the *Opware® SAS Content Migration Guide* for more 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 8-5 shows.

Figure 8-5: 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 select 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 **Next**, 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:

- **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 (+) button next to the Recipients field.
- **Schedule:** Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

- **Option 1:** Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2:** Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.



The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opware SAS core server (usually UTC).

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 **View Details** to see progress or the results of the installation.
- 9** Click **Close** 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 *Opware® SAS User's Guide: Server Automation* for more information about 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 Opware SAS Web Client 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 **Install OS**.

Or

From the navigation panel, click Software ► Operating Systems. the Operating System page appears. Click the Install OS link.

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 8-6 shows.

Figure 8-6: Overview Page in the Install Operating System Wizard



- 2 Select the Custom Installation option and click **Start**. If you did not select servers in Step 1, the Select Servers page appears.
- 3 If necessary, select the servers or groups that you want to provision and click **Next**. 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 **Next**. The Assign Customer page appears.

By default, the OS Provisioning feature entered the server host name 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 Opware SAS Web Client UI.

- 5 Select a customer for the servers and click **Next**.

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 **Next**.

- 7 Select the OS patches that you want to apply to the servers and click **Next**. The Select Applications page appears.

See the *Opware® SAS User's Guide: Application Automation* for information on how Opware SAS performs patch management for Unix and Windows.

- 8 Select any applications that you want to install on the servers and click **Next**.

- 9 On the Schedule and Notify page, you have the following options:

- **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 (+) button next to the Recipients field.
- **Schedule:** Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:

- **Option 1:** Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.

- **Option 2:** Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.



The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opware SAS core server (usually UTC).

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 **View Details** to see progress or the results of the installation.

11 Click **Close** 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 *Opware[®] SAS User's Guide: Server Automation* for more information about 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 Opware SAS Web Client 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 Opware SAS, as Figure 8-7 shows.

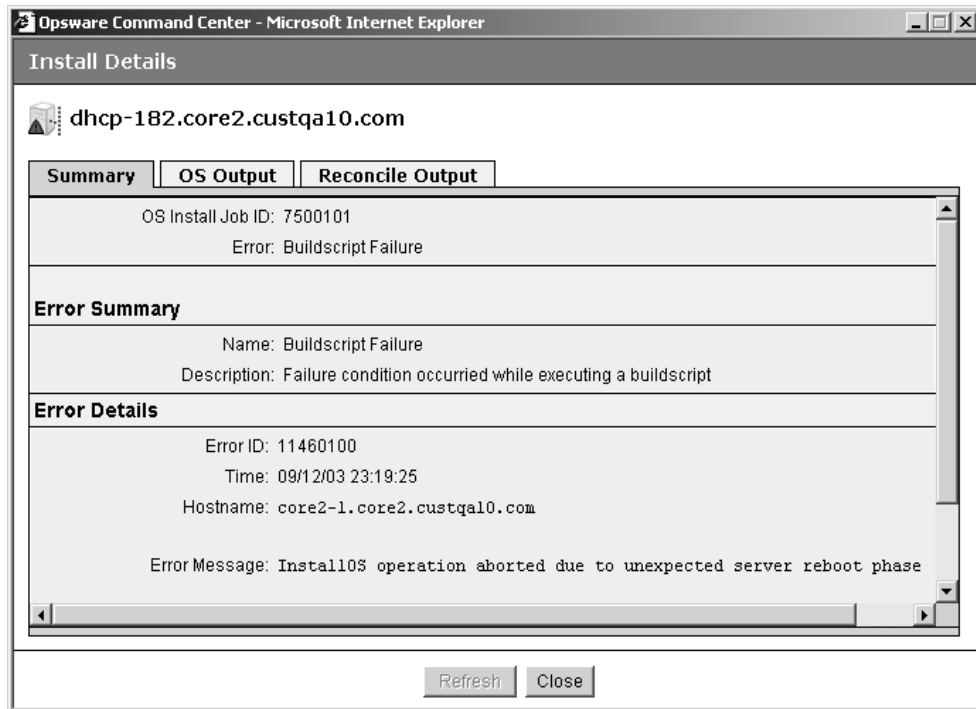
Figure 8-7: My Jobs Panel in the SAS Web Client Home Page

| My Jobs | | | | See All (29) |
|---|-------------------|---------|-----------------------|--------------|
| Name | Start Time | Servers | Status | |
|  Install OS | 09/09/03 03:56:41 | 1 | Completed | |
|  Install OS | 09/09/03 03:33:53 | 1 | Completed | |
|  Install OS | 09/09/03 03:14:31 | 1 | Completed with errors | |
|  Update Network Settings | 09/09/03 03:01:54 | 1 | Completed | |
|  Update Network Settings | 09/09/03 02:59:44 | 1 | Completed | |
|  Install OS | 09/09/03 02:43:42 | 1 | Completed | |

- 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.

- 3 Click **View Details** 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 feature generated, as Figure 8-8 shows.

Figure 8-8: 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 **Delete**.

You can leave the entry for the failed OS installation in the Server Pool list because Opware SAS automatically replaces the entry when you reboot the server.

Alternatively, Opware SAS removes the entry from the list after 7 days and all information about the server is removed from Opware SAS.

- 6 Reboot the server by using the network (PXE for Windows and Linux servers or DHCP for Solaris servers) or by using an Opware Boot Floppy (for Windows) or by using an Opware 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 336 in this chapter for more information.

See “Installing an OS by Using a Custom Installation” on page 341 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 Manage Server list.

See *Opware® SAS User's Guide: Server Automation* for more information about using My Jobs to obtain a history of your operations.

Network Configuration for Servers after OS Provisioning

The OS Provisioning feature provisions servers with an OS by using DHCP addresses. By using the SAS Web Client, users can configure network settings, including a static IP address, host name, 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. Opware SAS enables system administrators to do this through the SAS Web Client rather than logging onto the server manually after OS provisioning is complete.

See *Opware® SAS User's Guide: Server Automation* for information on how Opware SAS configures networking for servers.

Requirements for Reprovisioning Solaris and Linux Servers

When you choose to preserve network configuration for a server you are about to re-provision, 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 servers built or managed by Opsware SAS with this feature. You can reprovision Solaris and Linux servers so that they are running another version of the same OS if the hardware supports that new version of the OS.



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 ► Manage Servers. The Manage 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 **Tasks** 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 8-9.

Figure 8-9: Overview Page for the Install Operating System Wizard When Re-Provisioning

Enter Names

- 1 Enter Names
- 2 Assign Customer
- 3 Select Template
- 4 Confirm
- 5 Schedule and Notify
- 6 Install

Overview

This wizard helps you install an operating system and an Opsware Agent on a server that has registered with an OS Build Manager.

You can provision the OS by using an existing template, or by specifying the operating system, patches, and applications separately. You can also assign a customer to the servers and modify their names by using this wizard.

Use Template - Use a predefined template to specify the OS, patches, and applications.
 Custom Installation - Define the OS, patches, and applications separately.

Please preserve the network configuration for the selected servers

⚠ Reprovision Warning

All data on the selected servers will be erased when the operating system is installed.

Yes, I understand the OS installation process will erase all data on the selected servers.

- 4** Select the check box in the warning (indicating that you understand that Opsware SAS erases all data from the servers).
- 5** (Optional) Select the check box to preserve the network settings for the server.
Opsware® SAS User's Guide: Server Automation for more information about how Opsware SAS configures networking for servers.
- 6** Select the Use Template option or the Custom Installation option and click **Start**.
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 336 in this chapter for more information.

See “Installing an OS by Using a Custom Installation” on page 341 in this chapter for more information.

Boot Arguments for Reprovisioning Solaris Servers

When you reprovision a Solaris server, the OS Provisioning feature 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 “Boot Arguments for Reprovisioning Solaris Servers” on page 348 in this chapter for information about how to use interface settings when you boot a Solaris server.

Appendix A: Reconcile

IN THIS APPENDIX

This section discusses the following topics:

- Overview of Reconcile
- Ways to Perform Reconcile
- Reconcile Operations and Guidelines
- Reconcile on Supported Operating Systems
- Specifying Scripts During Reconcile
- Reconcile Output
- Assigning to and Removing Servers from Nodes
- Reconcile Software Wizard

Overview of Reconcile

This section explains how reconcile and test reconcile work for software provisioning and installing Unix patches, 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 or uninstallation wizard.

This section 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 templates.

Opware SAS uses a model-based approach that 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 Opware SAS 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 Opware SAS manages.

In Opware SAS, 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 Opware SAS to update the model of a server by adding or removing software nodes from the model. See the *Opware® SAS Configuration Guide* for more information about software provisioning setup.

During the process of using a wizard, if you decide to install or uninstall software immediately, Opware SAS begins a reconcile session. You can also schedule the installation or uninstallation for a later time, and Opware SAS 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 “Software Provisioning” on page 291 in Chapter 6 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 Opware SAS 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 331 in Chapter 8 for more information.)
- Install/Uninstall Unix Patch (See the *Opware® SAS User’s Guide: Application Automation* for information about remediating patch policies.)
- Install/Uninstall Windows Patch (See the *Opware® SAS User’s Guide: Application Automation* for information about remediating patch policies.)

- Install/Uninstall Software (See “Software Provisioning” on page 291 in Chapter 6 for more information.)
- Install Template (See “Software Provisioning” on page 291 in Chapter 6 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.

Reconcile Operations and Guidelines

This section provides information on how reconcile works within Opware SAS and contains the following topics:

- Reconcile Operations
- Reconcile and Package Metadata
- Installation and Uninstallation Order
- Software Installation Order for Adopted Software
- Patches and Reconcile
- Preview Reconcile
- Types of Reconcile

Reconcile Operations

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. Opware SAS then determines what operations are required to make the server conform to its model.

To make this determination, Opware SAS queries the Opware Agent on the server and examines the server’s model to assemble the following data:

- A list of all software (including patches) that is installed on the server. Although Opware SAS records this information in the Model Repository, Opware SAS uses the Opware Agent on the server to compile a list of all installed software. Querying the

Opware Agent is necessary in case any software was installed manually, without using Opware SAS.



Opware, Inc. recommends that users install all software through Opware SAS. Manually installed software, however, can be adopted by Opware SAS. See “Software Installation Order for Adopted Software” on page 353 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. Opware SAS obtains this information from the Model Repository.
- A list of software on the server that was installed through Opware SAS. Opware SAS requires this information because it cannot uninstall software that was not installed through Opware SAS. 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.

Opware SAS then calls the server’s native utilities to carry out the installation or uninstallation of software. (For example, on a Solaris server, Opware SAS uses the Solaris utility `pkgadd` to install software. On an AIX server, Opware SAS uses `installp`. See “Reconcile on Supported Operating Systems” on page 355 in this chapter for more information. If the reconcile operation requires both uninstallation and installation of software, the uninstallation occurs before the installation.

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 Software Repository (either through the command line interface, the Upload Patch Wizard, or through package management). Metadata includes information such as name and version, and this information varies depending on the package type. (The file name alone is not sufficient to identify the package in Opware SAS.)

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 the metadata must be entered manually, because the Windows utilities that obtain that type of information do not run on Unix.

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, Opware SAS 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 nodes, 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 Opware SAS 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 Opware SAS 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

Opware SAS does not uninstall software that was not installed through Opware SAS. Software that was manually installed, however, can be adopted by Opware SAS in a node. If a user attaches a package that was manually installed on a server, assigns the server to the node, and then performs a reconcile, the software becomes adopted by Opware SAS. (The software is not actually reinstalled during this operation.)



Opware SAS does *not* adopt Solaris patches. For example, if you uninstall a Solaris patch that was adopted into Opware SAS, the patch will not be uninstalled.

When software is adopted by Opware SAS, 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 Opware SAS.
- 2** A user installs three packages on this server using Opware SAS.
- 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.

Patches and Reconcile

All software in Opware SAS is associated with nodes. Ordinarily, these nodes are created explicitly by users during software 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.

Preview Reconcile

Before any changes are committed to a server, you have the option to perform 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 uninstall patch wizards, Opware SAS does not perform a preview reconcile.

Types of Reconcile

The two types of reconcile are partial and full.

During a partial reconcile, Opware SAS 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 Manage 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:

- AIX Reconcile
- HP-UX Reconcile
- Solaris Reconcile
- Linux Reconcile

After Opware SAS 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. The following table shows the utilities used during the reconcile session.

Table A-1: Utilities in Reconcile

| SOLARIS | LINUX | AIX | HP-UX |
|-------------------------------------|-------------------------------------|--|--|
| patchadd (installs patches) | RPM (installs and removes software) | installp (installs software) | swinstall (installs software) |
| patchrm (removes patches) | | installp -u (removes software) | swremove (removes software) |
| pkgadd (installs software) | | inutoc (generates a table of contents of packages to be installed) | swlist (copies individual packages into one large depot) |
| pkgrm (removes software) | | | swmodify (used to convert older format packages to a newer package format) |
| RPM (installs and removes software) | | RPM (installs and removes software) | |

See the *Opware® SAS Configuration Guide* for more information about the package types that Opware SAS supports.

AIX Reconcile

AIX software is delivered in LPPs, which are collections of filesets. When a server is reconciled and Opware SAS determines that the reconcile requires filesets to be installed, Opware SAS 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.

In AIX when you install a base fileset, it is always in the committed state and when you install an update fileset it is always in an applied state. When a fileset is in an applied state, the previous version of the fileset is saved. When a fileset is in a committed state, the previous version of the fileset is deleted. You can add the `-c` option in the install flag to install an update fileset in the committed state.

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 Opware SAS determines that the reconcile requires filesets to be installed, Opware SAS 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), Opware SAS 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, Opware SAS records the fact that the patches installed belong to a given patch cluster. This allows Opware SAS to identify patch clusters that are installed on the servers. Opware SAS can use this information to uninstall patch clusters.

Solaris patch clusters cannot be adopted by Opware SAS. If the patch clusters were not installed through Opware SAS, it is not possible to determine if a patch on a server originated from a patch cluster. Opware SAS can, however, adopt individual patches.

Linux Reconcile

RPM is the only package type that Opware SAS 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.

Specifying Scripts During Reconcile

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, Opware SAS does not rollback 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 Opware SAS 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. Opware SAS 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-install 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 Opsware SAS. 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 for 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 SAS Web Client, 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 **Tasks** menu, select **Assign Node**. The Assign Node Wizard appears.
- 5** Navigate to and select the node to which you want to assign the server.
- 6** Click **Assign**.

Removing Servers from Nodes

Perform the following steps to remove servers from nodes:

- 1** From the navigation panel in the SAS Web Client, 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 **Tasks** menu, select **Remove Node**.
- 5** Select the check boxes from the nodes that you want to remove and click **Remove**.

Reconcile Software Wizard

Servers are usually reconciled as a result of running any of the OS provisioning, patch management, or software 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



The Reconcile Software wizard is flexible when you run it on a group of servers. For example, when you select group of servers, Opware SAS will install the correct applications on the servers in the group, even if OS versions for the servers and the patches or applications do not all match. Opware SAS also matches the customer

association for applications, operating systems, and templates with the customer association of servers. If Opsware SAS cannot find a match, a “Valid devices not found” error message appears at the end of the wizard. Therefore, use caution when modifying these default values.

To directly reconcile servers, perform the following steps:

- 1** From the navigation panel in the SAS Web Client, 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 or groups that you want to reconcile.
- 4** From the **Tasks** 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.
- 8** Click **Preview** to continue. A preview reconcile occurs. Review the results of the preview reconcile and click **Next**.

Or

Click **Skip Preview** to skip the preview process.

- 9** On the Schedule and Notify page, you have the following options:
 - **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 (+) button next to the Recipients field.

- **Schedule:** Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

- **Option 1:** Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2:** Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.



The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opware SAS core server (usually UTC).

10 Click **Reconcile** to complete the process now.

Index

A

| | |
|--|---------------|
| adding | |
| application to a template | 240 |
| custom attributes for servers | 162 |
| custom attributes to nodes | 216 |
| nodes | |
| reasons to | 184 |
| restrictions | 184 |
| operating system to a template | 232 |
| packages to nodes | 204 |
| patches to a template | 243 |
| servers and server groups | 127 |
| servers to My Servers | 62 |
| service levels to a template | 245 |
| service levels, SAS Web Client to | 165 |
| software installation dependencies | 214 |
| adopting software, how to | 353 |
| AIX | |
| APARs | |
| about | 256 |
| deprecating | 286 |
| LPPs, about | 252, 255, 357 |
| package management | 255 |
| package types, supported | 252 |
| reconcile on | 357 |
| applications, package files | 251 |
| asset tracking, overview | 53 |
| assigning, servers to nodes | 359 |
| attachments, for folders and templates | |
| about | 221 |
| blocked | 223 |
| blocking inherited attachments | 224 |
| inherited | 222 |
| local | 222 |

B

| | |
|--------------------|-----|
| backups | |
| full backup | 311 |
| incremental backup | 311 |
| blocking | |
| attachments | 223 |

| | |
|---|-----|
| inherited attachments | 224 |
| vs. removing attachments with inheritance | 249 |

C

| | |
|-------------------------------------|--------|
| changing | |
| custom field values | 159 |
| use and stage values, servers for | 143 |
| cloning, servers | 146 |
| code deployment | |
| static NAT, use with | 102 |
| configuration | |
| settings for software | 203 |
| configuration tracking | |
| changes, detecting | 310 |
| email and logging actions, overview | 312 |
| email notification lists, creating | 313 |
| Opware SAS 6.0.1, availability in | 31 |
| overview | 309 |
| policies | |
| creating | 316 |
| overview | 314 |
| policy entries, re-enabling | 329 |
| policy limits | 310 |
| policy targets and wildcards | 316 |
| container packages, described | 252 |
| content migration | |
| best practices of | 33 |
| complete migration, explained | 30 |
| example of | 35 |
| full migration, explained | 26 |
| inheritance and negative overrides | 41 |
| optimizing | 35 |
| preview only mode | 32, 34 |
| process for | 25 |
| rolling back | 34 |
| service levels | 41 |
| software installation dependencies | 30 |
| Software Tree permissions | 43 |
| sync permissions for | 34, 44 |
| template inheritance | 42 |
| templates | 41 |

Content Migration Guide, purpose of 23

conventions used in the guide 18

copying

- nodes

 - in Software Tree 193
 - reasons for 193
 - restrictions 194

- servers, See cloning.

creating

- configuration tracking email notification list . . . 313
- configuration tracking policies, ways to 316
- node-based policies for configuration tracking 321
- server groups 113
 - dynamic groups 116
 - static groups 113, 115
- templates 227
- ZIP packages 266

custom attributes

- adding, for servers 162
- adding, to nodes 216
- deleting, for servers 163
- deleting, from nodes 218
- editing, for nodes 217
- editing, for servers 163
- managing, overview 161
- migrating 41
- nodes, managing in 216

custom fields

- changing value 159

customers

- association with servers 88
- Customer Independent, definition of 87
- Not Assigned customer, definition of 87
- overview 86

D

Data Center Intelligence reporting

- overview 90

deleting

- custom attributes from nodes 218
- custom attributes, for servers 163
- nodes
 - in Software Tree 191
 - reasons for 191
 - restrictions 191
- operating system from a template 234
- packages 284
- policies for node-based configuration tracking 328
- servers 146

dependencies

- between nodes for software installation 212

depots

- converting 260, 356
- metadata for 259
- package management for HP-UX 253
- reconciling 357
- script to split
 - by bundle 260
 - by product 260

deprecating, packages 286, 287

DHCP

- addresses for servers, usage 337, 346
- configuring for servers 172
- re-provisioning servers, requirements for 346

directories

- tracked configuration by wildcard targets 318

disabling, nodes' configuration tracking policies . 328

displaying, packages 270

DNS, configuring for servers 172

domains, changing for Windows servers 173

downloading, packages locally 288

duplicating, server groups 130

dynamic, server groups 108

E

editing

- application in a template 242
- custom attributes, for nodes 217
- custom attributes, for server 163
- nodes
 - configuration tracking policies 326
 - in Software Tree 186
 - reasons for 186
 - restrictions 186
- package file properties 281
- patch in a template 244
- server properties 143
- service levels 166
- services levels in templates 247
- ZIP package properties 267

email

- configuration tracking and logging actions 312
- configuration tracking notification list, format of 314
- notification list attribute, search order for 313

encoding

- available 279
- package metadata 278
- scripts 278
- Solaris response file 284

- errors
 - in scripts to install software294
- examples
 - commands to convert depots260
 - content migration35
 - inheritance for software, ways to use210
 - script to split depots
 - by bundle260
 - by product260
 - searching, ways to82
 - Software Migration Tool, command syntax35
 - Software Tree, nodes in179
 - swapping servers' disks86
- F**
- file formats, supported for packages252
- filters
 - packages to display272
- folders
 - adding
 - an application to240
 - an operating system to236
 - patches to243
 - service levels to245
 - blocking inheritance249
 - creating230
 - deleting248
 - editing/deleting, service levels247
 - editing/removing, applications from242
 - removing, an operating system from238
- G**
- guidelines, setting up Software Tree180
- H**
- hard disks, swapping for servers86
- histories
 - viewing, for servers92
- HP-UX
 - depots
 - converting260, 356
 - metadata for259
 - package management253, 258
 - reconciling357
 - package management257
 - reconcile on357
- I**
- icons, servers icons in the SAS Web Client136
- Info-Zip
 - compatible package metadata269
 - compatible Zip packages268
- inheritance
 - attachments for templates and folders222
 - blocking vs. removing249
 - for templates and folders, about221
 - installation wizards295
 - override values, changing211
 - overview295
 - software from other nodes210
- Install Operating System Wizard
 - custom installation341
 - template installation336
- Install Software Wizard295
- installable packages252
- installation
 - changing order for software209
 - dependencies
 - adding214
 - for software212
 - for software, viewing213
 - removing215
 - installing
 - issues for software294
 - operating systems
 - with custom installation341
 - with templates336
 - software
 - with Install Software Wizard295
 - templates
 - with Install Templates Wizard306
- IP addresses
 - configuring for servers171
 - management IP
 - definition of99
 - viewing100
 - primary for servers, setting102
 - primary IP, definition of100
 - searching for servers by81
 - setting for servers102
- IP ranges
 - overview170
- J**
- jobs
 - scheduling for servers154
 - time outs for157

L

| | |
|--|-----|
| Linux | |
| reconcile on | 358 |
| re-provisioning servers, details of | 346 |
| local attachments, for templates and folders | 222 |
| LPPs | |
| package metadata for | 255 |
| <i>See also</i> AIX LPPs. | |

M

| | |
|---|-----|
| management IP | |
| definition of | 99 |
| managing | |
| custom attributes | 161 |
| Microsoft | |
| Hotfixes, Security Patches, Service Pack packages | 265 |
| Microsoft Installer Packages. <i>See</i> MSI. | |
| Microsoft patches, package metadata for | 265 |
| model-based approach | |
| modeling software in nodes | 200 |
| Software Tree usage | 181 |
| modeling, public server groups | 108 |
| MSI | |
| package management, prerequisites | 264 |
| package metadata for | 264 |
| My Servers | |
| servers, adding to | 62 |
| servers, removing from | 63 |

N

| | |
|-------------------------------------|-----|
| NAT tables | |
| changing, affects of | 103 |
| <i>See also</i> static NAT. | |
| network | |
| configuring for servers | |
| after OS provisioning | 346 |
| overview | 171 |
| nodes | |
| adding | |
| custom attributes for | 216 |
| packages to | 204 |
| packages to, overview | 200 |
| reasons for | 184 |
| restrictions | 184 |
| to Software Tree | 184 |
| configuration tracking policies for | 310 |
| copying | |
| in Software Tree | 194 |

| | |
|---|-----|
| reasons for | 193 |
| restrictions | 194 |
| custom attributes | |
| deleting from | 218 |
| editing in | 217 |
| managing in | 216 |
| deleting | |
| in Software Tree | 192 |
| reasons for | 191 |
| restrictions | 191 |
| editing | |
| in Software Tree | 187 |
| reasons for | 186 |
| restrictions | 186 |
| installation dependencies | |
| overview | 212 |
| removing | 215 |
| migrating | 27 |
| removing, packages from | 207 |
| searching for packages to add to | 206 |
| servers, assigning and removing, overview | 359 |
| software attached, overview | 200 |
| software inheritance, explained | 210 |
| software, modeling in | 200 |
| viewing, software attached to a node | 203 |
| notifying, jobs | 155 |

O

| | |
|--|--------|
| operating systems | |
| custom installation | 341 |
| failure to install OS, recovering from | 343 |
| installing, with SAS Web Client | 334 |
| provisioning | 331 |
| re-provisioning, requirements for | 346 |
| supported types for packages | 252 |
| template installation | 336 |
| ways to install on servers | 335 |
| Opware guides | |
| contents of | 17 |
| conventions used | 18 |
| documentation set | 20 |
| icons in guide, explained | 19 |
| Opware SAS | |
| documentation set | 20 |
| related documentation | 20 |
| software in, overview | 292 |
| supported operating systems | 55, 57 |
| OS provisioning | |
| custom OS installation | 341 |
| failure to install OS, recovering from | 343 |

re-provisioning Linux, details of346
 template OS installation336
 override values of inherited software, changing ..211
 overwriting, packages280

P

package management
 overview251
 tasks, overview265

package types
 AIX APAR256
 HP-UX depots 253, 257, 258, 357
 HP-UX depots, reconciling356
 LPP 252, 255, 357
 RPM252, 356
 Windows Hotfix254
 Windows Service Packs265
 ZIP266

packages
 adding to nodes
 overview200
 process204
 AIX package management255
 AIX, managing for255
 container, described252
 deleting, restrictions285
 deprecating
 about286
 restrictions287
 displaying270
 downloading288
 editing, file properties281
 file formats252
 filtering in SAS Web Client272
 HP-UX, managing for257
 Info-Zip compatible268
 installable252
 Microsoft Hotfixes, Security Patches, and Service
 Packs265
 migrating29
 operating systems, supported252
 overwriting existing280
 searching for272
 viewing, assigned nodes274
 ZIP package management265

permissions
 server groups, required for105
 Software Tree, migrating43
 synchronizing after content migration44

prerequisites

Info-Zip compatible package management ... 269
 MSI package management264
 Solaris package management263
 primary IP, definition of100
 private, server groups107
 properties
 editing for servers143
 servers, overview139
 public, server groups107

R

reconcile
 adopting software353
 AIX357
 content migration, affect on29
 HP-UX357
 installation order for adopted software353
 Linux358
 operating systems, supported355
 output
 about358
 performing350
 process, overview351
 Reconcile Software Wizard360
 servers
 assigning to nodes359
 reconciling directly360
 removing from nodes359
 software reconciled onto servers, explained ..182
 Solaris358
 ways to perform350
 Reconcile Software Wizard360

removing
 application from template242
 operating system from a template234
 packages from nodes207
 patches from a template244
 servers
 from My Servers63
 from nodes359
 from nodes, overview359
 from server groups128, 129
 from service levels169
 service levels from templates247
 software installation dependencies from nodes ..
 215

reports
 overview90

re-provisioning
 Linux servers

| | | | |
|---|----------|---------------------------------------|-----|
| about | 347 | static | 107 |
| details of | 346 | using, server search | 118 |
| requirements for | 346 | server life cycle, Opware SAS | 133 |
| Solaris servers | | server lists | |
| about | 347 | filtering | 60 |
| requirements for | 346 | overview | 57 |
| RPM | | server management | |
| package metadata for | 261 | functions | 97 |
| package type | 252 | notifying, jobs | 155 |
| reconciling RPM packages | 356 | scheduling, tasks | 154 |
| | | server life cycle tasks | 141 |
| | | Server Pool | |
| | | overview | 57 |
| | | server search, using | 118 |
| | | servers | |
| | | asset tracking, overview | 53 |
| | | assigning to service levels | 168 |
| | | association with customers | 88 |
| | | changing value, custom field for | 159 |
| | | changing, use and stage values | 143 |
| | | cloning | 146 |
| | | configuration tracking policy | 310 |
| | | custom attributes, deleting | 163 |
| | | deactivating | 145 |
| | | deleting from the Opware SAS | 146 |
| | | DHCP, configuring | 172 |
| | | domains for Windows servers, changing | 173 |
| | | filtering in server lists | 60 |
| | | histories, viewing | 90 |
| | | identifying | 84 |
| | | jobs | |
| | | scheduling | 154 |
| | | management of | 58 |
| | | My Servers, removing from | 63 |
| | | network configuration | |
| | | configuring | 171 |
| | | overview | 170 |
| | | nodes, overview of assigning/removing | 359 |
| | | primary IP addresses, setting | 102 |
| | | properties | |
| | | editing | 143 |
| | | overview | 139 |
| | | reconciling directly | 360 |
| | | re-provisioning | |
| | | Linux details of | 346 |
| | | requirements for | 346 |
| | | Search box, using | 64 |
| | | search feature, using | 65 |
| | | searching by IP address | 81 |
| | | searching for | 64 |
| | | See <i>also</i> server groups. | |
| S | | | |
| SAS Client Library | | | |
| Migrated folder | 28 | | |
| SAS Web Client | | | |
| icons for servers, defined | 136 | | |
| OS provisioning | 334 | | |
| packages, filtering | 272 | | |
| server lists, filtering | 60 | | |
| ways to search | 65 | | |
| scheduling | | | |
| server jobs | 154 | | |
| time outs affecting | 157 | | |
| scripts | | | |
| error conditions when installing software | 294 | | |
| searching | | | |
| details | 79 | | |
| examples of | 82 | | |
| for packages to add to nodes | 206 | | |
| IP addresses | 81 | | |
| packages | 272 | | |
| Search box | 64 | | |
| search feature | 65, 79 | | |
| ways to | 65 | | |
| server groups | | | |
| adding, servers and server groups to | 127 | | |
| characteristics | 106 | | |
| creating | 113 | | |
| dynamic groups | 116 | | |
| static groups | 113, 115 | | |
| criteria for search | 118 | | |
| duplicating | 130 | | |
| modeling, public groups | 108 | | |
| overview | 104 | | |
| permissions, required | 105 | | |
| removing, servers and server groups to | 128, 129 | | |
| types | | | |
| dynamic | 108 | | |
| private | 107 | | |
| public | 107 | | |

- service levels, removing from 169
 - swapping disks, examples 86
 - ways to install operating systems 335
 - service levels
 - adding individual 165
 - assigning servers to 168
 - editing 166
 - migrating 27, 41
 - overview 164
 - removing, servers from 169
 - ways to view for servers 166
 - setting
 - primary IP address, server of 102
 - setup for servers
 - software provisioning 183
 - Software Tree, guidelines 180
 - software
 - adopted, installation order for 353
 - attached to nodes, overview 200
 - configuration settings for 203
 - inheritance from other nodes, explained 210
 - inherited override values, changing 211
 - Install Software Wizard 295
 - installation dependencies
 - between nodes 212
 - removing from nodes 215
 - installation issues 294
 - modeling in nodes 200
 - overview in Opware SAS 292
 - types for installation 293
 - Uninstall Software Wizard 300
 - uninstallation issues 294
 - viewing, installation dependencies 213
 - ways to install 292
 - software installation dependencies
 - migrating 30
 - Software Migration Tool
 - command syntax examples 34
 - command syntax for 33
 - running 32
 - software provisioning
 - features, summary of 293
 - modeling software in nodes 200
 - overview 291
 - overview, of setup tasks 183
 - servers, selecting 298
 - Software Tree
 - examples 179
 - inheritance, migrating 41
 - migrating nodes 27
 - overview 176
 - setup guidelines 180
 - synchronizing permissions, migration 44
 - usage explained 181
 - Software Tree nodes
 - adding, packages to 204
 - attaching, software to 200
 - copying 191
 - deleting 191
 - deleting, custom attributes from 218
 - editing 186
 - Opware-specific, migrating 27
 - removing, packages from 207
 - viewing, attached software 203
 - Solaris
 - package metadata for 263
 - reconcile 358
 - servers, re-provisioning 346
 - supported package types 252
 - stage values, overview for servers 143
 - static NAT
 - changing NAT tables, affects of 103
 - code deployment and rollback, use with 102
 - static, server groups 107
 - supported operating systems
 - for managed servers 55
 - for SAS Client 57
 - swmgmt-migrate.sh script, command syntax 33
- ## T
- templates
 - adding
 - an application to 240
 - an operating system to 232
 - patches to 243
 - service levels to 245
 - attachments
 - blocked 223
 - inherited 222
 - local 222
 - creating 227
 - deleting 248
 - editing/deleting, services levels 247
 - editing/removing, application 242
 - inheritance, migrating 42
 - inheritance, overview 221
 - Install Templates Wizard 306
 - migrating 27, 41
 - operating systems, installing with 336
 - removing, an operating system from 234
 - removing, patches from 244

- types of attachments221
- time outs
 - server management jobs157
- troubleshooting
 - operating systems, installation failure343
- types of, server groups107

U

- Uninstall Software Wizard300
- uninstalling
 - issues for software294
 - software with Uninstall Software Wizard300
- uploading
 - enhanced performance for269
 - packages275
- use values, overview for servers143
- using, server search118

V

- viewing
 - history
 - server92
 - job details in the SAS Client153
 - job details in the SAS Web Client151
 - nodes assigned to packages274
 - nodes' configuration tracking policies325
 - service levels, ways to166
 - software
 - attached to nodes203
 - installation dependencies213

W

- Windows servers
 - changing, domains173
 - Hotfixes, package type254
 - Service Packs
 - package management265
 - supported package types for252
 - uploading packages, enhanced performance .269
- wizards
 - Install and Uninstall Software, overviews295
 - Install Operating System
 - custom install341
 - template install336
 - Install Software295
 - Install Template306
 - Reconcile Software360
 - Uninstall Software300

Z

- ZIP Package Management265