



Opsware[®] SAS 6.5 User's Guide: Application Automation

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Opware Inc. Unpublished Confidential Information.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	23
<hr/>	
Contents of this Guide	23
Conventions in this Guide	25
Icons in this Guide	25
Guides in the Documentation Set and Associated Users	26
Opsware, Inc. Contact Information	27
Chapter 1: Visual Application Manager	29
<hr/>	
Overview of Visual Application Manager	29
New in This Release – VAM 2.0.1	30
Over of VAM Features	31
Discovering, Mapping, Managing Applications	33
VAM Usage Examples	33
Launching VAM	36
Launching VAM from Servers or Device Groups	37
Launching VAM from Search Results	38
Launching VAM from Generated Reports	39
VAM User Interface	39
VAM Toolbar	41
Menus and Menu Options	44
Link Selection	44
Understanding VAM	46

Data Collection and Display	46
VAM Application	47
VAM Maps: Visualizing Applications	53
Symbols Used in Maps	62
Property Pages	64
VAM Options	74
Virtualization Settings	74
Scan Time-Out Preference	75
Accessing Servers From VAM	76
Opening a Device Explorer	76
Opening a Remote Terminal	76
Opening a Global Shell	77
Creating Application Definitions	77

Application Templates	78
Application Tiers	79
Creating an Application Tier	79
Editing an Application Tier	80
Deleting an Application Tier	80
Cutting and Copying an Application Tier	80
Pasting an Application Tier	81
Application Component Signatures	81
Evaluation Order	82
Creating a Component Signature	84
Editing Component Signatures	85
Deleting Component Signatures	85
Cutting and Copying a Component Signature	85
Pasting a Component Signature	86
VAM File Management	87
Opening a .vam File	87
Saving a .vam File	87
Saving a .vam File as an Application Template	88
Comparing Scan Results	88
Source and Target Scan Results	89
Scan Results Comparison Types	89
Comparing Two Sets of Scan Results	93
Filtering VAM Data	94
Data Filtering	94
Filter Criteria	95
Error Messages	98
Scan Results Comparison Heuristics	100

Chapter 2: Audit and Remediation	103
Overview of Opsware Audit and Remediation	104
Audit and Remediation Examples	104
Audits	106
Audit Policies	106
Audits and the Compliance Dashboard	107
Audit Reports	107
Snapshots	107
Terms and Concepts	108
Understanding Audits	110
Audit Comparison Types	110
The Auditing Process	111
Audit Elements	112
Creating an Audit	113
Using Save As for Audit or Snapshot Specification	116
Viewing Server Audit and Snapshot Usage	116
Configuring an Audit	118
Opsware Audit and Remediation Rules	120
Server Objects Used in Audits and Snapshots	122
Configuring Opsware Audit and Remediation Rules	124
Configuration Rules: Expected (Target) and Remediation Values	124
Audit Sources: Server or Snapshot?	128
Configuring Specific Rules	131

Configuring Application Configuration Rules	131
Configuring COM+ Rules	138
Configuring Custom Scripts Rules	139
Configuring Event Logging Rules	141
Configuring File System Rules	143
Configuring Hardware Rules	144
Configuring IIS Metabase Rules	145
Configuring Operating System Rules	146
Configuring Software Rules	148
Configuring Users and Groups Rules	148
Configuring Windows Registry Rules	150
Configuring Windows Services Rules	151
Configuring The Opsware Network (TON) Rules	152
File System Inclusion and Exclusion Rules	155
Example: Including all .txt Files in a Snapshot or Audit	157
Example: Including Only File a in a Snapshot or Audit	158
Example: Including last temp.txt file and exclude all else	159
File System Rule Overlap	159
Audit Rule Exceptions	161
Rules That Cannot Have Exceptions	162
Considerations When Applying Exceptions to Device Groups	162
Adding a Rule Exception to an Audit	162
Editing or Deleting a Rule Exception	164
Audit Policies	165
Creating an Audit Policy	166
Linking and Importing Audit Policies	166
Running an Audit	169

Running an Audit from the Library	169
Running an Ad-Hoc Audit	170
Running an Audit on a Server from All Managed Servers	170
Re-running an Audit from Audit Results	171
Scheduling an Audit	172
Scheduling a Recurring Audit	173
Editing an Audit Schedule	174
Viewing a Completed Audit Job	175
Viewing and Remediating Audit Results	175
Viewing and Remediating Differences of Audit Results Objects	178
Viewing Audit Results with Exceptions	181
Searching for Audits	182
Understanding Snapshots	182
Snapshot Specification Elements	184
The Snapshot Process	186
Creating a Snapshot Specification	187
Creating a Snapshot Specification from a Server	187
Creating a Snapshot Specification from the Library	187
Configuring a Snapshot Specification	188
Configuring a Snapshot Specification	188
Configuring Snapshot Specification Rules	190
Saving a Snapshot as an Audit Policy	190
Deleting a Snapshot Specification	191
Running a Snapshot Specification	191
Scheduling Snapshot Jobs	192

Scheduling a Recurring Snapshot Job	193
Editing an Snapshot Job Schedule	194
Viewing a Snapshot Job Schedule	195
Deleting a Snapshot Job Schedule	195
Locating Snapshots	196
Searching for Snapshots	196
Viewing Snapshot Contents	197
Detaching a Snapshot From a Server	199
Deleting a Snapshot	200
Copying Objects from a Snapshot to a Server	200
Copying Objects to a Server from a Snapshot	201
Chapter 3: Compliance Dashboard	205
Overview of Compliance Dashboard	205
Compliance Dashboard Usage: Proactive and Reactive	206
Viewing the Compliance Dashboard	207
General Compliance Dashboard Categories	208
Compliance Dashboard Statuses	209
Refreshing to Get the Latest Compliance Information	210
Compliance Dashboard Terms and Concepts	210
Compliance Dashboard Remediation	212
Software Compliance	214
Application Configuration Compliance	215
Understanding Application Configuration Compliance Status	215
Application Configuration Compliance Remediate Options	216
Patch Compliance	216

Understanding Patch Compliance Status	216
Patch Remediate Options	217
Audit Compliance	217
Audit Compliance Status: All Scheduled and Individual	218
Showing Individual Audits	219
Audit Remediate Options	220
Duplex Compliance	220
Duplex Remediate Options	221
Filtering and Sorting Compliance Dashboard Information	221
Showing/Hiding Specific Compliance Tests	222
Exporting the Compliance Dashboard	223
Chapter 4: SAS Client Reports	225
Overview of SAS Client Reports	225
Reports Features	226
Opsware SAS Client Reports	226
User Permissions	230
Launching the Reports Feature	230
Reports Display	231
Running a Report	234
Report Results	235
Exporting a Report	237
Printing a Report	238
Chapter 5: Patch Management for Windows	239
Overview of Patch Management for Windows	239

Patch Management for Windows Features	240
Library	242
Patch Management for Windows Prerequisites	244
Microsoft Patch Database	245
Opware SAS Integration	246
Support for Windows Patch Testing and Installation Standardization .	246
Supported Windows Patch Types	247
Supporting Technologies for Patch Management	248
Windows Hotfixes	248
Searching for Patches and Policies	249
Patch Management Roles for Windows	250
Patch Management Process	251
Patch Properties	254
Patch Dependencies and Supersedence	256
Viewing Windows Patches	257
Editing Windows Patch Properties	257
Importing Custom Documentation for a Patch	258
Deleting Custom Documentation for a Patch	258
Finding Vendor-Recommended Windows Patches	259
Finding Servers That Have a Windows Patch Installed	259
Finding Servers That Do Not Have a Windows Patch Installed	259
Importing a Patch	260
Automatically Importing Windows Patches	260
Exporting a Windows Patch	263
Exporting Windows Patch Information	264
Deleting a Patch	265
Policy Management	266

Patch Policy	266
Patch Policy Exception	268
Precedence Rules for Applying Policies	269
Remediation Process.	270
Remediating Patch Policies	271
Setting Remediate Options	273
Setting Reboot Options for Remediation	273
Specifying Pre and Post Install Scripts for Remediation.	274
Scheduling a Patch Installation for Remediation	275
Setting Up Email Notifications for Remediation.	276
Previewing a Remediation	277
Verifying Patch Policy Compliance	278
Creating a Patch Policy	279
Deleting a Patch Policy	280
Adding a Patch to a Patch Policy	280
Removing a Patch from a Patch Policy	281
Attaching a Patch Policy to a Server	281
Detaching a Patch Policy from a Server	282
Setting a Patch Policy Exception.	282
Finding an Existing Patch Policy Exception.	283
Copying a Patch Policy Exception	283
Removing a Patch Policy Exception	284
Patch Compliance	285

Patch Compliance Scans	285
Patch Compliance Levels	286
Patch Compliance Rules	286
Patch Compliance Reports	287
Patch Administration for Windows	288
Setting the Patch Availability	288
Importing the Microsoft Patch Database	289
Selecting Windows Products to Track for Patching	290
Scheduling a Patch Policy Compliance Scan	290
Setting the Patch Policy Compliance Level	291
Importing Windows Patch Utilities	291
Exporting Windows Utility Files	292
Editing the Customized Patch Policy Compliance Level	292
Locales for Windows Patching	293
Supported Locales	293
Overview of Locale Configuration Tasks	293
Configuring the Opsware Core for Non-English Locales	294
Selecting the Locales of Patches to Import	294
End User Requirements for Non-English Locales	295
Patch Installation	296

Installation Flags	297
Application Patches	298
Service Packs, Update Rollups, and Hotfixes	299
Installing a Windows Patch	299
Setting Windows Install Options	300
Setting Reboot Options for a Windows Patch Installation	301
Specifying Install Scripts for a Windows Patch Installation	302
Scheduling a Windows Patch Installation	304
Setting Up Email Notifications for a Windows Patch Installation	304
Previewing a Windows Patch Installation	305
Viewing Job Progress of a Windows Patch Installation	306
Patch Uninstallation	307
Uninstallation Flags	308
Uninstalling a Windows Patch	309
Setting Uninstall Options	310
Setting Reboot Options for a Windows Patch Uninstallation	310
Specifying Install Scripts for a Windows Patch Uninstallation	312
Scheduling a Windows Patch Uninstallation	313
Setting Up Email Notifications for a Windows Patch Uninstallation	313
Previewing a Windows Patch Uninstallation	314
Viewing Job Progress of a Patch Uninstallation	314
Chapter 6: Patch Management for Unix	329
 Overview of Patch Management for Unix	329

Patch Management for Unix Features	330
Opware SAS Integration	332
Support for Unix Patch Testing and Installation Standardization	332
Library	334
Search Feature	335
Patch Management Roles for Unix	335
Patch Management for Specific Unix Operating Systems	336
Supported Unix Versions and Patch Types	336
Underlying Technologies for Patch Management on Unix	338
AIX Patches	339
Solaris Patches	340
HP-UX Patches	340
Patch Uploads for Unix	341
Patch Uploads for Specific Unix Versions	341
Patch Properties	342
Viewing Unix Patches	343
Editing Unix Patch Properties	343
Finding Servers That Have a Unix Patch Installed	344
Finding Servers That Do Not Have a Unix Patch Installed	344
Exporting a Patch	344
Deleting a Patch	345
Software Policies	345
Patch Compliance Reports	346
Patch Administration for Unix	346
Setting the Default Patch Availability	346
Patch Installation	347

Installation Flags	348
Application Patches	349
Installing a Unix Patch	349
Setting Unix Install Options	351
Setting Reboot Options for a Unix Patch Installation	351
Specifying Install Scripts for a Unix Patch Installation	352
Scheduling a Unix Patch Installation.	354
Setting Up Email Notifications for a Unix Patch Installation.	354
Previewing a Unix Patch Installation.	355
Viewing Job Progress of a Unix Patch Installation.	356
Patch Uninstallation	356
Uninstallation Flags	357
Uninstalling a Unix Patch.	358
Setting Uninstall Options	359
Setting Reboot Options for a Unix Patch Uninstallation.	359
Specifying Pre and Post Install Scripts for a Unix Patch Uninstallation	360
Scheduling a Unix Patch Uninstallation	361
Setting Up Email Notifications for a Unix Patch Uninstallation	362
Previewing a Unix Patch Uninstallation	362
Viewing Job Progress of a Patch Uninstallation.	363
Chapter 7: Software Management	365
Overview of Software Installation.	365
Software Installation Process	366
Ways to Install Software in Opware SAS	368
Installing Software Using a Software Policy	368

Attaching a Software Policy to a Server	368
Attaching a Server to a Software Policy	370
Overview of Software Policies Remediation	372
Ways to Open the Remediate Window	372
Remediating Software Policies	374
Installing Software	380
Uninstalling Software Using a Software Policy	382
Detaching a Software Policy from a Server	382
Overview of Software Template	383
Ways to Open the Install Software Templates Window	384
Installing Software Using a Software Template	385
Overview of Running ISM Controls	392
Ways to Open the Run ISM Control Window	392
Running ISM Controls	393
Software Policy Compliance	396
Performing a Software Compliance Scan	397
Software Policy Reports	397
Chapter 8: Application Configuration Management	399
Overview of Application Configuration Management (ACM)	399
Application Configuration Creation and Use	400
ACM Components	402
Configuration Template	403
Application Configuration	403
Value Set Editor	403
Configuration Markup Language (CML)	407
Application Configuration Inheritance	407

Application Configuration Default Values	408
Application Instance Values	409
Sequence Merging and Inheritance	412
Sequence Replace	413
Sequence Append	413
Sequence Prepend	415
Using ACM	416

Creating an Application Configuration	417
Creating a Configuration Template	419
Searching for Application Configurations	420
Viewing Application Configuration Template Sources	421
Adding or Removing Configuration Templates	421
Deleting Application Configurations	422
Loading a Template File	423
Setting a Configuration Template to Run as a Script	425
Specifying Template Order	426
Editing Default Values for an Application Configuration	427
Attaching an Application Configuration to a Server or Group	430
Setting Application Configuration Values on a Server or Group	431
Loading Existing Values into a Configuration Template	433
Pushing Changes to a Server or Group	434
Scheduling an Application Configuration Push	436
Comparing Two Configuration Templates	438
Comparing a Template Against an Actual Configuration File	438
Scanning Configuration Compliance	439
Scheduling a Configuration Compliance Scan	440
Rolling Back to a Previous State	442

Chapter 9: Operating System Provisioning **445**

Supported Operating Systems for OS Provisioning	446
OS Provisioning Overview	448
Server Lifecycle for OS Provisioning	448
OS Provisioning	450

Overview of OS Provisioning	450
Solaris OS Provisioning	452
Linux or VMware ESX OS Provisioning	452
Windows OS Provisioning	452
Hardware Preparation	453
New Server Booting	454
Booting New Servers with Different Operating Systems	454
OS Build Agent	455
Booting a Windows (DOS), Linux, or VMware ESX Server with PXE ..	455
Booting a Windows Server with PXE Using WinPE	458
Booting a Solaris Server Over the Network	459
Ways that the OS Build Agent Locates the Opware Build Manager .	460
Installing OS Build Agents	460
Verifying Installation of an OS Build Agent	461
Recovering when an OS Build Agent Fails to Install	462
OS Installation with the SAS Client	462
Creating an OS Installation Profile	463
Creating an OS Sequence	463
Selecting Servers in the Unprovisioned Servers List	468
Before Running an OS Sequence – Firewall Considerations	469
Running An OS Sequence	470
Reprovisioning a Managed Server	472
Appendix A: VAM Platform Support	475
Supported Platforms in VAM	476

Appendix B: Glossary	479
Index	497

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, managing patches, reconciling servers, script execution, configuration tracking, and deploying and rolling back code. 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: Visual Application Manager: Describes how to use the Visual Application Manager tool to draw detailed layout views of the operational architecture and behavior of distributed business applications in your IT environment. Provides instructions about how to create, edit, and export physical and logical drawings that can help you diagnose and resolve problems.

Chapter 2: Audit and Remediation: Describes how to define server configuration policies and make sure that servers in your facilities meet those policy standards. When servers are found to be 'out of compliance' (not configured the way you want them to be), you can remediate the differing server configurations.

Chapter 3: Compliance Dashboard: Describes how the Compliance Dashboard allows you to view at a glance the overall compliance levels for all the devices in you facility and helps you to remediate compliance problems. The Compliance Dashboard displays compliance tests for software policies, application configurations, audits, patches, and

duplex status. Each of these compliance tests is based upon an Opware Server Automation System (SAS) "policy" (user or system defined) which define a unique set up server or device configuration settings or values that help ensure your IT environment is configured the way you want it to be.

Chapter 4: Reports: Provides information about how to create reports in the SAS Client and how you can perform actions on objects within the reports. These reports include: Server Reports, Compliance Reports, Sarbanes-Oxley (SOX) Reports, Network Reports, User and Security Reports, and Custom Reports.

Chapter 5: Patch Management for Windows: Provides information about managing patches for the Windows 2000, Windows 2003, and Windows NT4.0 operating systems. These patches include Service Packs, Update Rollups, and hotfixes. It describes the user roles: a policy setter, a patch administrator, and a system administrator. It also describes reconciling, previewing (an install), installing, and uninstalling patches by using patch policies and patch policy exceptions.

Chapter 6: Patch Management for Unix: Provides information about managing patches for Unix operating systems by using software policies. It discusses patch types, testing, and installing and uninstalling patches. It review the roles of the patch administrator and system administrator in applying patches, and the permissions required for performing patch management.

Chapter 7: Software Management: Provides information about installing and uninstalling software using software policies, installing software using software policy template, running ISM Controls, and performing software compliance scans.

Chapter 8: Application Configuration Management: Provides information about managing application configurations through the SAS Client, and includes such topics as creating Application Configurations, Application Configuration inheritance, editing value sets, and applying Application Configurations to a server.

Chapter 9: Operating System Provisioning: Provides information about supported environments for OS provisioning and an overview of the permissions and server life cycles associated with OS provisioning. It also describes the process for provisioning, an overview of the hardware preparation, information about booting new servers, and using the SAS Client to install operating systems using OS sequences.

Appendix A: VAM Platform Support: Provides information about the operating system platforms and architecture that VAM supports scanning and displaying application (process families), server, and device information.

Appendix B: Glossary: Defines terminology and acronyms that are unique to Opware SAS.



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.
<code>Courier</code>	Identifies text that is entered or displayed at the command-line prompt, such as Unix commands, Opware 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.

ICON	DESCRIPTION
	<p>This icon represents a tip. It identifies information that can help simplify or clarify tasks.</p>
	<p>This icon represents a warning. It is used to identify significant information that must be read before proceeding.</p>

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 Planning 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 *Opware[®] SAS Policy Setter's Guide* is intended for system administrators responsible for setting up OS provisioning, configuration tracking, code deployment, and software management.
 - The *Opware[®] SAS Content Utilities Guide* is intended for advanced system administrators responsible for importing content such as software packages into Opware SAS. It documents the following command-line utilities: OCLI 1.0, IDK, and DET (CBT).
 - The *Opware[®] Automation Platform Developer's Guide* is intended for software developers responsible for customizing, extending, and integrating Opware SAS. It documents how to create Web Services, Java RMI, Python, and CLI clients that invoke methods on the Opware API.

Opware, Inc. Contact Information

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

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

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

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

For troubleshooting information, see the Opware Knowledge Base:

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

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

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

Chapter 1: Visual Application Manager

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Visual Application Manager
- Discovering, Mapping, Managing Applications
- Launching VAM
- VAM User Interface
- Understanding VAM
- VAM Options
- Accessing Servers From VAM
- Creating Application Definitions
- VAM File Management
- Comparing Scan Results
- Filtering VAM Data
- Error Messages
- Scan Results Comparison Heuristics

Overview of Visual Application Manager

The Visual Application Manager (VAM 2.0.1) helps you manage the operational architecture and behavior of distributed business applications in your IT environment by displaying detailed application information in physical and logical drawings.

VAM displays detailed application information in physical and logical drawings, and it shows the relationship between a collection of component signatures, process families, connections, and dependencies. VAM enables you to create models of your application tiers and hierarchies, which map to actual application processes. This provides a clearer picture of how all an application's components and processes interact – including all

related servers and network devices. When you better understand an application's processes and interrelationships, you can understand how an application's resources are distributed and more effectively troubleshoot errors when they occur.

VAM also enables you to compare scan results and see differences between an application's maps and definitions at a specific point in time. You can compare two scan results to see changes that have occurred, and remediate any differences as you see fit. You can also view SAS and NAS compliance information, so you can troubleshoot servers or devices that are out of compliance.

VAM is tightly integrated with other features in the Opware SAS Client and Opware NAS Client. This enables you to perform change management tasks, such as reconfiguring, patching, auditing, and remediating software and patch policies.

New in This Release – VAM 2.0.1

VAM 2.0.1 is released as part of the Opware Server Automation System (SAS) 6.5 and provides the following new features:

- The ability to visualize applications that run on virtual servers (Solaris 10 local zones and VMware ESX virtual machines) and display both hypervisor and guest virtual servers relationships for each technology. VAM also displays ESX 3 virtual switches (vSwitches) and port groups in both the Virtual Map and Device Map.
- The ability to create and compare VAM scan results – these results present a picture of the current state your applications (its visual maps, structures, component signatures and process families, and so on) as viewed through VAM. You can compare differences between past VAM scan results with current scan results and visualize meaningful differences between specific application or server components. A single .vam file allows you to save as many VAM scan results as you want.
- The ability to display compliance information from both SAS and NAS. You can troubleshoot and remediate any out of compliance configurations by launching the appropriate server or device inside of SAS or NAS.
- The ability to filter scan results based upon all existing criteria plus new compliance criteria.
- The ability to visualize Microsoft IIS 5.0 and 6.0 as an application inside of VAM.

VAM 2.0.1 supports internationalization; it can display application, server, or network information using a language other than English. It also supports Opware TON predefined VAM application definitions of common applications.



Opware's Visual Application Manager (VAM) is a separately licensed product that requires the Opware Server Automation System (SAS) in order to run. In order to visualize any networking information inside of VAM, you must have both a licensed version of NAS integrated with your Opware SAS core, plus an additional license to run VAM showing NAS data. If you have not purchased VAM or NAS and would like to, contact your Opware sales representative.

Over of VAM Features

VAM enables you to perform the following tasks:

- Discover, map, and visualize the components, connections, and dependencies of multi-tiered business applications.
- Visualize applications that run on virtual servers, showing virtual servers in relationship to their hypervisors, as well as **virtual switches and port groups** (VMware ESX only).
- Visualize application information in multiple physical and logical layouts, such as an application view, a server view, a network view (includes virtual devices), and a virtual view that displays virtual server and the applications that run on them.
- Organize recognized components into multi-tier applications to create a logical view that can be analyzed to verify correct operation.
- Map actual application process families to application component signatures and highlight them with user-defined custom colors.
- Manage VAM scan results in .vam (Visual Application Manager archive) files, which contain an application definition and one or more VAM scan results of an application.
- Create and share application templates that represent an ideal application definition.
- Troubleshoot and resolve problems by launching the SAS Client Device Explorer, Network Device Explorer, Global Shell, and Remote Terminal, to perform in-depth analysis or to perform actions on the systems under investigation.
- Save maps as .gif, .jpg, and .svg files.
- Print maps.

VAM Prerequisites

VAM requires an Opware Agent version 5.1 or higher on managed servers. This enables VAM to scan them.

Backwards Compatibility with Previous Versions

Both VAM 1.0 (SAS 6.0) and VAM 2.0.1 (SAS 6.1) files are supported in VAM 2.0.1.

Supported Operating Systems

VAM collects and displays data about managed servers that are running AIX, Linux, HP-UX, Solaris, VMware ESX, and Windows operating systems. If you are running non-standard kernels on a Linux operating system, VAM might depend on the kernel version, in addition to the operating system version.

For more detailed information on VAM platform support, see Appendix A, "VAM Platform Support".

.vam Files

VAM manages all application information that is collected, displayed, and captured in a .vam file, which can be opened, closed, saved (as a file or template), and edited.

Each .vam file contains an application definition (which can be empty) and zero or more scan results, which contain the visual maps (Application, Server, Virtual, and Devices). The application definition specifies an application's logical construction in terms of tiers, subtiers and the application components contained in those tiers. Each set of scan results represents the state of a set of network devices and managed servers, the process families running on those servers, the connections among those process families, and any external clients and dependencies. You can refresh scan results to view the latest application data, and these results are saved in the .vam file.

A VAM application that does not have any scan results attached is useful for creating an application definition that can be saved as a templates. See "VAM File Management" on page 87 for instructions on how to save, open, and edit a .vam file. See "Application Templates" on page 78 for information on creating a VAM template.

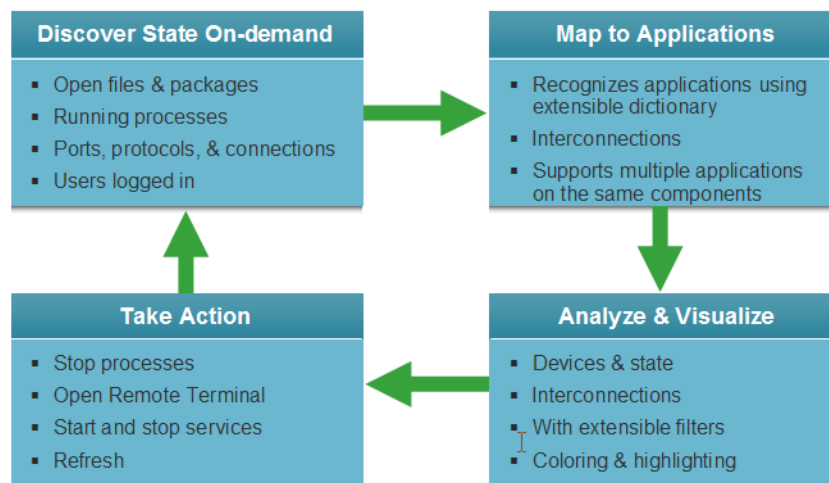
Discovering, Mapping, Managing Applications

When you launch VAM, it scans, collects, models, maps, and displays information about the operational architecture and behavior of applications in your IT environment. VAM manages this information with .vam files that contain application definitions and VAM scan results. VAM scan results capture the information gathered when you scan an application (and the servers and devices it runs on) as well as application definitions.

VAM uses information gathered from the Opware SAS and Opware NAS data models, leveraging the architecture to collect more data on-demand (such as processes that are running, open ports, and the number of users logged in). It also maps application data to enable you to visualize and analyze your operational environment.

Figure 1-1 shows the process you follow using VAM to display application data in your IT environment.

Figure 1-1: VAM Processes



VAM Usage Examples

Understanding how VAM functions within the context of a real data center is best illustrated with some general usage examples, such as:

- Launch VAM
- Discover and Map Applications on Servers
- View Related Networking Information
- Define and Customize an Application Definition

- Troubleshoot Problems and Take Action

These examples are discussed in greater detail in the following sections.

Launch VAM

An application administrator starts a new job at a company and one of his first tasks is to add a new feature to an application, which was maintained by a prior employee, but the former employee left very little documentation. The administrator was provided with the application's source code but does not understand how all pieces of the application work together from an operational standpoint.

To gain a better picture of the application, he opens the Opsware SAS Client, selects a group of servers that the application runs on, and launches VAM.

Discover and Map Applications on Servers

VAM performs a scan of the selected servers and discovers all applications, component signatures, processes and process families, and connections related to all the applications on the selected servers. VAM displays detailed “maps” of the application (processes and process families) and servers and connections associated with the servers selected (virtual and non-virtual), as well as any related network relationships.

The application administrator examines this information and sees a short list of items that contains two servers, one of which is his server, and two network devices. Looking at the Server Map, he selects the box that represents his server, and a properties pane opens to display more detailed information about the server. He notices that the server has virtual machine-related information, so he concludes that his application may be running on a virtual machine instance. He selects the Virtual Map, and now the map only shows a single server. He clicks that server and once it expands, he sees his server within it. He now understands that his application runs on a VMware virtual machine (VM), and has visibility into the physical host (hypervisor) on which the VM runs.

View Related Networking Information

The administrator then selects the Network Map and sees his server again, but notices that it has a green line connecting it to another box. By clicking on that box and examining the properties, he determines that the other box is a VMware vSwitch, which in turn is connected to a Cisco switch. He can see precisely which VLAN, port group, switch port, and network interfaces are involved when his application communicates over the network. He now understands how his application fits into the network, both physical and virtual.

He takes a closer look at the lines emanating from his server, and notices a prominent, thick black line pointing at some IP address, so he clicks on it. He sees that the line represents 64 connections to some another host on port 1433. It looks like the database that he knows his application uses. He goes back to the SAS Client, finds a server, which has that same IP address, selects it and his original server, and opens them in VAM, initiating another VAM session. Now he sees that the thick black line is pointing at the new server, and after drilling down, he discovers it pointing specifically at an SQL Server process. He continues this until he finds his application running across and depending on 10 separate servers.

Define and Customize an Application Definition

The application administrator naturally does not want to have to perform all of this manual mapping and discovery each time he wants to view and manage his application. He knows that the application vendor's documentation contains a logical architectural diagram of the application, so to make his job easier, he uses VAM to create an application template.

His first step is to create the logical tiers of the application. He selects the Application Tree tab, selects the Application node, and creates three main tiers for the application: Web, Application, and Database. He creates sub-tiers in the Application tier for authentication services and integration services. He then defines application component signatures to add to each tier, specifying which tier a recognized component signature should fall in.

In order to create reusable component signatures for each tier, he opens the component signature dialog by right-clicking on a tier. This allows him to specify, for each application component, the criteria used to recognize it, including the process's name, open files, listener ports, command line, and so on.

He continues to do this for each tier in the application, and then color codes the component signatures in each tier. When the application is visualized in the Application or Server maps, he will be able to see each tier of the application in different colors. Next time he launches VAM, the application will map and display according to his definition.

Finally, he saves his application definition as a template, so it can be reused by others who want to work with the same or similar application.

Troubleshoot Problems and Take Action

To help keep track of the state of an application at any given time, the application administrator continually uses the **Refresh The Scan Results** button in order to create new scan results of the application. Each scan result is saved inside of a .vam file that

contains his application definition, which can be used later to compare previous scans of the application with a current state to find any important differences and troubleshoot errors.

For example, if at some point his application malfunctions and stops working, the administrator can open his saved .vam file containing his application, select the Compare feature, and visualize the differences between the current state of the application and the last known good state. Comparing the scan results will show if specific devices are not communicating with other devices. For example, he can drill into the network map and see that an interface is missing from his VMware ESX hypervisor from the same diagram and select Open Remote Terminal to remedy to problem.

Launching VAM

To view applications on your servers with VAM, you need to select a server and launch VAM from the SAS Client. VAM will perform an extensive scan of the servers you selected – including all virtual servers and their hypervisors.

You can launch VAM by selecting a group of servers, one or more managed servers, or search or report results, as illustrated in the following sections.

For more information about how VAM scans a server, see “Data Collection and Display” on page 46.



The Allow Analyze permission is required to use VAM. You also need read access to each managed server that you plan to scan. Write access to each managed server is not required to run the VAM; however, write access is required to perform any actions on the servers, such as when you use a Remote Terminal.

To visualize virtualization dates inside of VAM, the View Virtual Server permission must be set to Yes for the user group that your user belongs to. (Without this permission, virtual servers will be displayed just like regular physical servers.) To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide*.

You can launch VAM from inside the SAS Client from the following different locations:

- Launching VAM from Servers or Device Groups
- Launching VAM from Search Results
- Launching VAM from Generated Reports

Launching VAM from Servers or Device Groups

To launch VAM individually or from multiple selected servers, groups of servers, virtual servers, or hypervisors, perform the following steps:

- 1** Launch the SAS Client from one of the following locations:
 - Click the SAS Client link in the Power Tools section of the SAS Web Client home page.
 - Double-click the SAS Client icon on your desktop (if you installed it on your desktop when you installed the SAS Client).
 - Select **Start** menu ► **All Programs** ► **Opware SAS Client**.
- 2** From the Navigation pane, select the Devices tab.
- 3** From the Device Groups, All Managed Servers list, or Virtual Servers list, select a server group, or one or more managed servers. To do so, perform one of the following actions:
 - From the **Actions** menu, select **Open with** ► **Visual Application Manager**.

Or

 - Right-click, and from the menu, **Open with** ► **Visual Application Manager**.

Or

- From the **Tools** menu, select **Visual Application Manager ► Open Selection**.


After scanning is completed, the VAM window appears containing the Device Tree, Application Tree, Property Page, Server Map, Network Map, Virtualization Map, and the Application Map.



If you have selected virtual servers or a virtual server's hypervisor to open with VAM, you will initially be asked if you want to scan virtualization relationships – in other words, if you want to scan any virtual and host servers related to the servers that you selected. This could increase the time it takes to complete the scan, depending on how many virtual servers or hypervisors are related to your selected servers.

Launching VAM from Search Results

To launch VAM from search results, perform the following steps:

- 1** Launch the SAS Client from one of the following locations:
 - Click the SAS Client link in the Power Tools section of the SAS Web Client home page.
 - Double-click the SAS Client icon on your desktop (if you installed it on your desktop when you installed the SAS Client).
 - Select **Start** menu ► **All Programs** ► **Opware SAS Client**.
- 2** From the Search panel, perform a search for servers. For example, from the top drop-down list, select Servers, and click the green search button .
- 3** In the search results, select one or more servers and then perform one of the following actions:
 - From the **Actions** menu, select **Visual Application Manager**.

Or

- From the **Tools** menu, select **Visual Application Manager ► Open Selection**.

After scanning is completed, the VAM window appears containing the Device Tree, the Application Tree, the Property Page, the Server Map, the Network Map, the Virtualization Map, and the Application Map.

Launching VAM from Generated Reports

To launch VAM from report results, perform the following steps:

- 1** Launch the SAS Client from one of the following locations:
 - Click the SAS Client link in the Power Tools section of the SAS Web Client home page.
 - Double-click the SAS Client icon on your desktop (if you installed it on your desktop when you installed the SAS Client).
 - Select **Start** menu ► **All Programs** ► **Opware SAS Client**.
- 2** From the Navigation pane, select the Reports tab.
- 3** Expand the Reports tab, and select a report that will display servers in its results. For example, expand the Server Reports folder and run one of the server reports.
- 4** From the report results, drill down and select an individual server or multiple servers, right-click, and select **Visual Application Manager**

After scanning is completed, the VAM window appears containing the Device Tree, the Application Tree, the Property Page, the Server Map, the Network Map, the Virtualization Map, and the Application Map.

If a VAM scanning process is taking too long, you can cancel. For more information on how to set the scan timeout value, see “Scan Time-Out Preference” on page 75.



When launching VAM on device groups or when refreshing a previous scan, the servers involved in that scan will consist of the members of those device groups at the time of the scan. Membership may change over time, so two scans of the same selection may produce a different set of scanned servers.

VAM User Interface

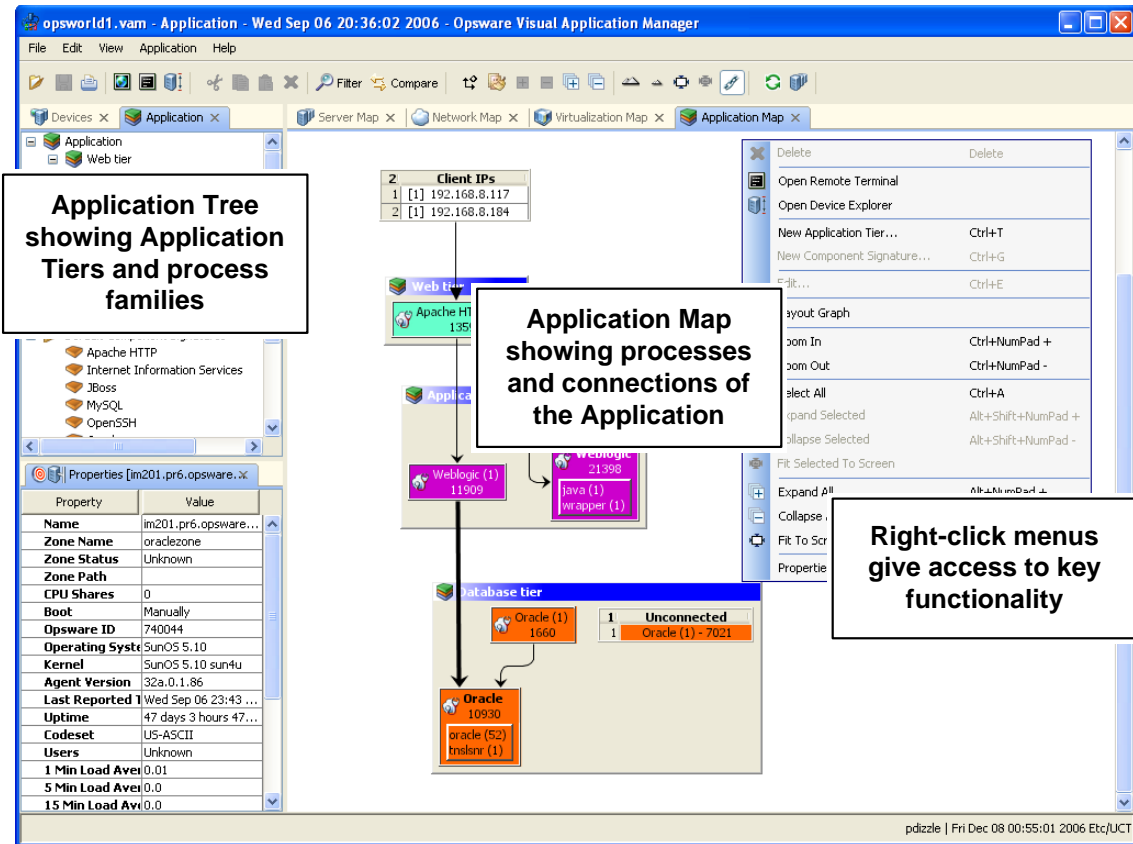
The VAM user interface shows an application and all its related processes, connections, and devices. It does so by providing the following user interface elements:

- Maps that display the physical layouts of applications – see “VAM Maps: Visualizing Applications” on page 53.

- Trees that display the logical layouts of applications – see “Application Tree” on page 47 and “Device Tree” on page 50.
- Property pages that provide very granular information about a selected object, component signature, process, or connection – see “Property Pages” on page 64.
- Dynamic tool bars and detailed tooltips to provide more information about tree and map objects.

Figure 1-2 shows the types of information that the VAM displays.

Figure 1-2: Opware Visual Application Manager User Interface



VAM Toolbar

The VAM toolbar allows you to open, close, resize, and organize different layout views and trees. Depending on the tree and view selected, certain toolbar icons will be unavailable. See Table 1-1 for a description of the toolbar icons.

Table 1-1: Toolbar Icons in VAM









TOOLBAR ICON	DESCRIPTION
	Opens a previously saved .vam or .vat file.
	Saves the current application (including maps) as a .vam or .vat file in your local file system or in the Opware Global File System (OGFS). If the application has not been previously saved, the Save As window displays.
	Prints the selected map. Displays the Print window where you specify page setup (including printing across multiple pages), a title for the printed map, and so on.
	Opens a Global Shell session.
	Opens the Open Remote Terminal window where you select a login ID for a Remote Terminal.
	Opens a Device Explorer for the selected managed server.
	Removes (cuts) a selected application component or a selected tier in the Application Tree and saves it to the clipboard.
	Copies a selected application component or a selected tier in the Application Tree and saves it to the clipboard.

Table 1-1: Toolbar Icons in VAM (continued)


















TOOLBAR ICON	DESCRIPTION
	Pastes a tier or an application component that has been previously cut or copied to the clipboard. See “Pasting an Application Tier” on page 81 and “Pasting a Component Signature” on page 86.
	Deletes a selected application tier or component signature in the Application Map or Application Tree.
	Allows you to filter the currently loaded scan results to find relevant data. For more information, see “Filtering VAM Data” on page 94.
	Allows you to compare to VAM snapshots. Clicking once will display the Compare pane at the bottom of VAM window. Click again to hide the Compare pane. For more information on comparing scan results, see “Comparing Scan Results” on page 88.
	Rotates the selected view, toggling it between a vertical and a horizontal orientation.
	Redraws all components in the selected view. Components that have been manually revised will retain their sizing.
	Expands selected tiers in the Application Tree or closed folders in the selected map. Tiers are expanded recursively down to the application component that they contain. Managed servers underneath the application components are not expanded.
	Collapses all tiers in the Application Tree or closed components in the selected view.
	Opens all tiers in the Application Tree or components in the selected map.
	Closes selected tiers in the Application Tree or folders in the selected map.

Table 1-1: Toolbar Icons in VAM (continued)

TOOLBAR ICON	DESCRIPTION
	Zooms into the selected view (enlarges the display size of).
	Zooms out of the selected view (reduces the display size of).
	Resizes the selected components in a currently active view to fit within the screen size.
	Resizes all components in the currently active map to fit within the screen size.
	Links trees and maps so that elements selected in a map will cause the corresponding element in a tree to also be selected. The Device Tree, Network Map, Server Map, and Virtualization Map can be linked together so that selecting an object in one causes it to be selected in all three. The Application Tree and Application Map can also be linked together so that selecting an object in one causes it to be selected in the other.
	Refreshes the scan results by collecting and displaying new information. Each time you click this button, VAM creates a new VAM snapshot which gets saved as part of the .vam file, and which can be used in a snapshot comparison.
	Toggles the maps to display the host server names in the title bar for virtual servers (Virtualization Map), virtual devices (Network Map), and application tiers (Application Map).

Menus and Menu Options

Most menus and menu options in the VAM are intuitive. This section discusses the menus and menu options that might not be self-explanatory.

File Menu

If you have made changes to the application definition and want to set this as the default, select **Set as Default Template** from the **File** menu.

If you have made changes to the application definition and want to restore the previously saved default application, select **Reset Default Template** from the **File** menu.


If you would like to import an application template that has already been saved, from the **File** menu, select **Import Template** and import the selected template.

For more information on application templates, see "Application Templates" on page 78.

View Menu

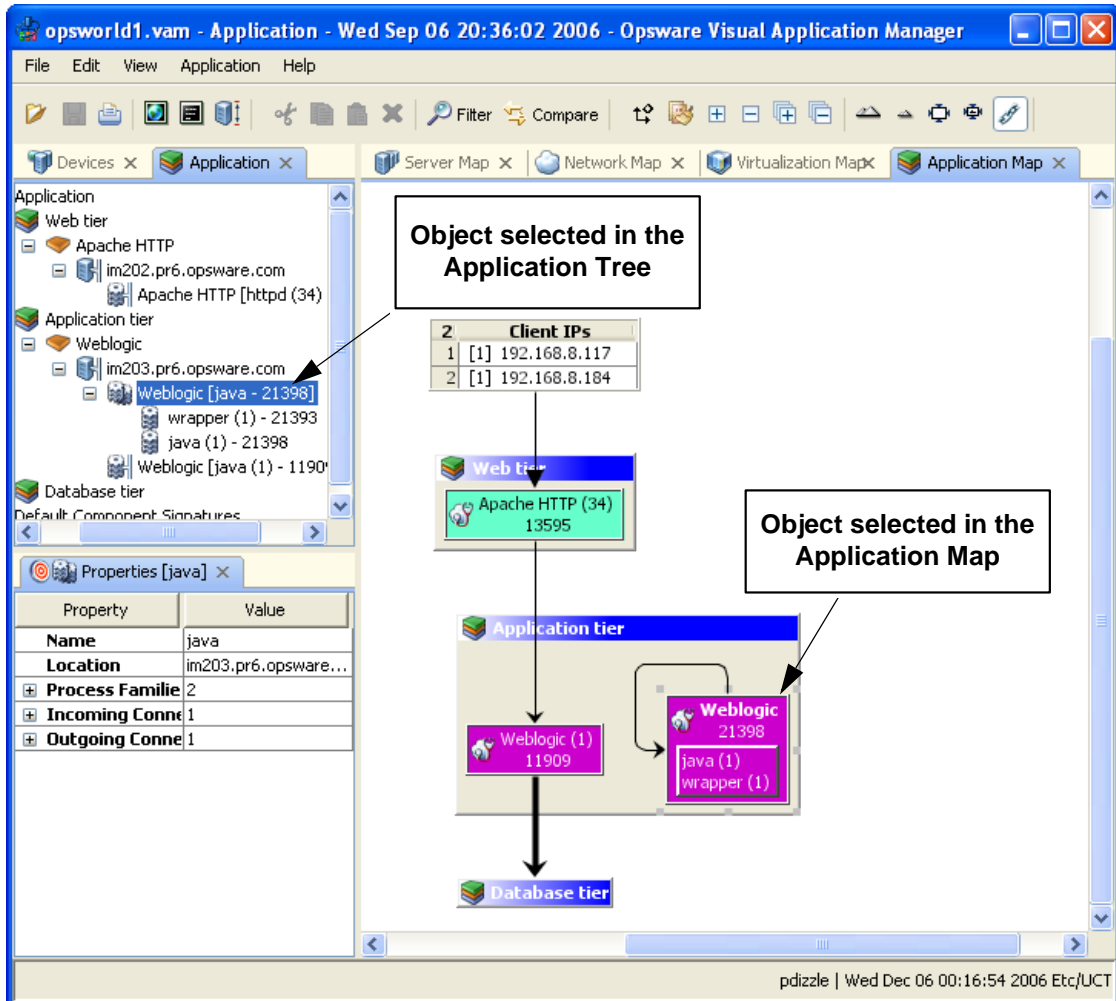
By default, the **Animate Layout** option is ON (preceded by a check mark). This causes the map to be animated (objects are displayed in motion) each time it is drawn, including a refresh. If the **Animate Layout** option is OFF (no check mark), the map will not be animated (objects are not displayed in motion) each time it is drawn.

Link Selection

When you toggle the Link Selection icon , selecting a server or process family in the Device Tree causes the corresponding element in the Network Map, Server Map, Virtual Map, Virtualization Map, or Application Map to also be selected. Conversely, selecting a node for a server in the Network Map causes all corresponding elements in the Device Tree to also be selected.

Also, you can select an application node in the Application Tree and the corresponding node will be selected in the Application Map. See Figure 1-3.

Figure 1-3: Application Tree with Link Selection



Understanding VAM

VAM's main function is to visualize applications in great detail, and to display the relationships among all their parts and processes and the servers and devices they depend on to function.

VAM scans a server (or multiple servers) you have selected to gather this information and displays it in the Maps and Application Tree, visualizing all processes and process families, connections, and devices related to the application. Each VAM session, which can be saved as a .vam file, allows you to create, visualize, analyze, define, share, and troubleshoot a VAM application. Clicking **Refresh the Scan Results** allows you to create a scan result of the current state of the VAM application and save it inside of the currently loaded .vam file. These scan results can be compared on a one to one basis using the compare feature (activated by the Compare toolbar button).

To provide a unique view into applications, VAM shows the following processes:

- Data Collection and Display
- VAM Application
- VAM Maps: Visualizing Applications
- Property Pages

Data Collection and Display

VAM scans a server or servers and draws layout maps based on data that is collected in real-time results of a scan. Server data is captured directly from servers and then recorded in scan results. Network device data is scanned and then recorded in scan results by NAS – where it is retrieved by the VAM from the Opsware NAS data model.

When you launch VAM, a set of programs runs on the selected managed servers and captures data. This scanning process collects data about processes running on those servers and the connections between them. It also collects detailed configuration information and current run-time state information about connections and processes. VAM then merges the server data with the network device data to show how servers, interfaces, switches and switch ports are connected together.

VAM collects and displays the following information about managed servers and network devices:

- Processes and process families (grouped into application component signatures) that are running on managed servers

- TCP and UDP connections between these processes
- Detailed configuration information
- Current runtime information about servers, connections, and processes
- Servers, interfaces, switches and vSwitches, and switch port connections

See “Processes, Process Families, and Extended Process Families” on page 50 for an explanation of how VAM interprets this data. See “Data Filtering” on page 94 for instructions on how to search the data that was collected by object type, such as by process family, network interface, and so on.

VAM Application

An application from the perspective of VAM is a complex collection of services that typically run across multiple servers and networking devices. An application in VAM consists of application definitions (tiers, component signatures, and property definitions) visible in the Application tree, and a collection of Maps that visualizes relationships between an application’s components, processes (and process families) and external clients and dependencies.

A VAM application maps to actual instances of applications that are running on servers that VAM has scanned and displayed. An application, as seen in the Application Map, is a collection of processes running on a managed server that maps to a VAM Application definition as specified in the Application Tree. The VAM Application is further explained in the following sections:

- Application Tree
- Application Tier
- Application Component Signature
- Processes, Process Families, and Extended Process Families

For information on how to create a VAM application, see “Creating Application Definitions” on page 77.

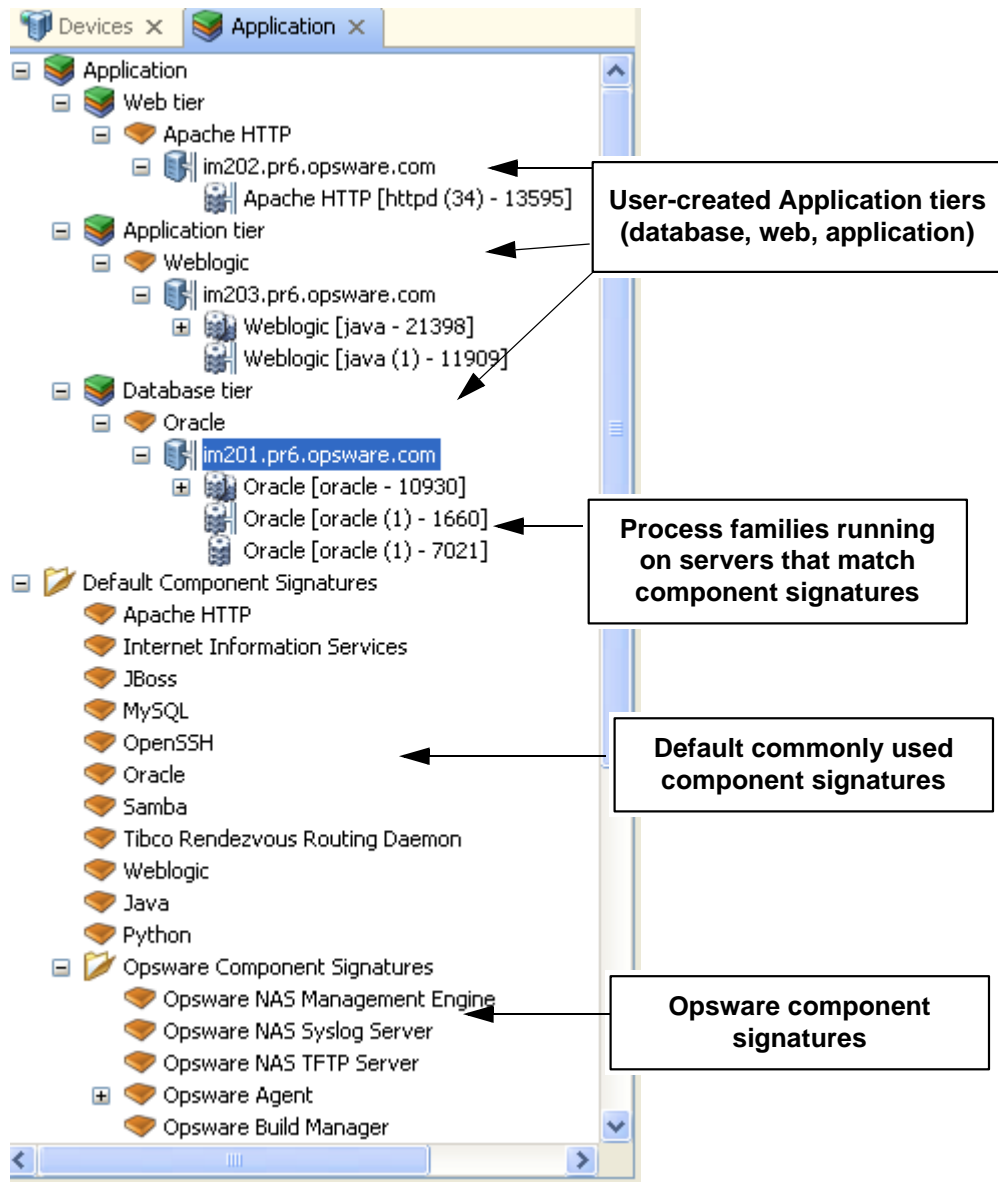
Application Tree

The Application Tree is a logical view of an application that provides a hierarchical representation of an application’s infrastructure. The Application Tree provides a different way of looking at what is visually displayed in the Application Map. The Application Tree

contains tiers and subtiers, which in turn contain their associated application component signatures. Inside of each component signature is the server or device that the process families run on, and inside of the server or device is the process family itself.


See Figure 1-4 for a picture of the Application Tree.


Figure 1-4: Application Tree





Application component signatures at the bottom of the Application Tree (such as ones that are associated with the root tier of an application) do not appear in the Application map – instead, they are highlighted in the Network Map, Server Map, and Virtualization Map.

When you select Link Selection from the toolbar  (or you can select it from the **View** menu), whenever you select an application tier or component signature in the Application Tree, the corresponding element in the Application map will also be selected (and vice-versa). See Figure 1-3 for an example of Link Selection.


If there are no matching process families for a component signature, a warning icon  appears next to it, and the tier that contains it, in the Application Tree.

Application Tier

An application tier enables you to create a logical structure of an application the way you want to understand it. You can define an application that helps you see all of its processes and process families as a diagram of elements that run across multiple servers, displaying the connections among them, clients connecting to them, and dependencies to which they connect.

The application tier definition contains a server filter, which restricts the servers whose process families will match the tier's component signatures.


Each application consists of a set of tiers and sub-tiers, such as a Web tier running Apache on Linux, an application tier running WebLogic on Windows, and a database tier running Oracle on Solaris.

An application tier is represented in the Application Tree by the  icon.

An application tier object consists of Application component signatures optional sub-tiers.

Application Component Signature

An application component signature is an object that represents a process or process family that comprise an application, such as Apache, Oracle, BEA WebLogic, Microsoft® SQL Server, and so on.

An application component signature is represented in the Application Tree by the  icon.

An application component signature object consists of a signature and preferences.


A signature is a set of rules that you provide and that VAM uses to identify a process family. This set of rules uses data such as process name, open files, command line, connected to port, modules, executable path, and listener port. If VAM discovers the process or process family during a scan according to the signature rule definition, then the process or process family will be added to the component signature and highlighted in the maps. See "Creating a Component Signature" on page 84 for more information.

Preferences specify the alias of the application component. These are displayed in the differing background and foreground text color of the different maps.

Processes, Process Families, and Extended Process Families


In VAM, a *process* is a running instance of a program in a Unix or Windows environment. A process is discovered and aggregated into process families and extended process families.

A *process family* is a collection of processes that are part of the same Unix session (same name and GID) or a collection of processes that are part of the same Windows session (same name and login session ID).

A process family is represented in the Application (or Device) Tree by this  icon. (Single processes are always grouped visually into process families, and so are also represented by the process family icon.) If the process family is connected to something else (another process family, for example), it is represented in the Application (or Device)

Tree by the  icon.

An *extended process family* is a set of processes that the VAM has heuristically computed to be related, but are not necessarily members of the same process hierarchy.

An extended process family is represented by the  icon.

Device Tree

The Device Tree is a logical, text-based view of top-level information about managed servers, process families, and network devices. This is a hierarchical display of the same top level information that is shown in the Network Map, Server Map, and Virtual Map.

The Device Tree contains servers and network devices (physical and virtual) as its top nodes. Below the servers are process families and extended process families. And, below the network devices are VLANs, ports, and port groups. The Device Tree will also show virtual devices that were scanned when you launched, including the following:

- All hypervisors for all scanned virtual servers are displayed, regardless of whether the hypervisor was scanned.
- Virtual servers are grouped as children of their hypervisors, whether or not the server is managed.
- VMware virtual switches (vSwitches) are grouped under each hypervisor. vSwitches can be expanded to view their port groups.
- Solaris Global zones can be expanded to list all running processes, but this list of processes includes processes on local zones not included in the current scan.

Figure 1-5 illustrates the Device Tree.

Figure 1-5: Device Tree

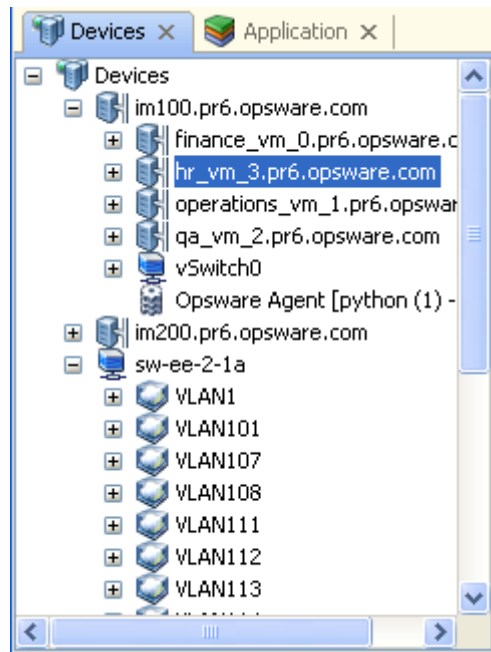
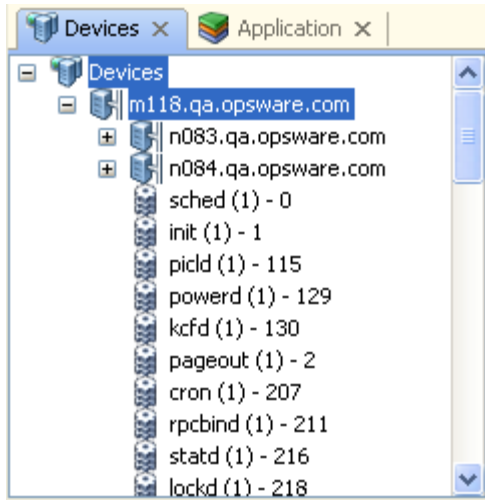


Figure 1-6 illustrates a virtual server (*m118.qa.opsware.com*) in the Device Tree with two virtual servers (*n083.opsware.com* and *n084.opsware.com*) being hosted on it.

Figure 1-6: Virtual Servers in the Device Tree




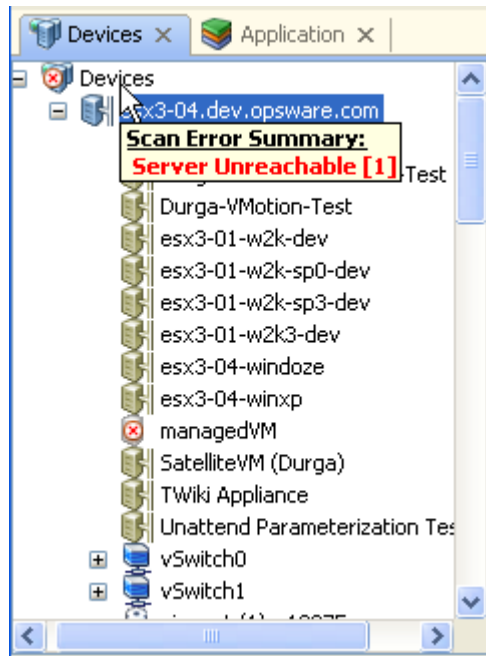
If a server device has an error associated with it, then it will appear in the Device Tree with an error icon on it , for example if the server is unreachable by Opsware SAS. When you move your mouse pointer over the device node in the tree, a tooltip message indicates the nature of the error, as shown in Figure 1-7.

Figure 1-7: Device Tree server node with tooltip indicating error



For more information on possible device errors, see “Error Messages” on page 98.

VAM Maps: Visualizing Applications

VAM provides four visual maps that display physical and logical drawings of managed servers, network devices, and connections in your environment: the Application Map, Server Map, Virtualization Map, and Network Map,

You can export a map to a .gif, .jpg, and an .svg file. See “Saving a Map to .gif, .jpg, or .svg” on page 63.

You can also print a map on single and multiple sheets of paper. See “Printing a Map” on page 64.

To enable you to see and understand how your application functions, VAM provides the following maps:

- Application Map
- Server Map
- Virtualization Map
- Network Map

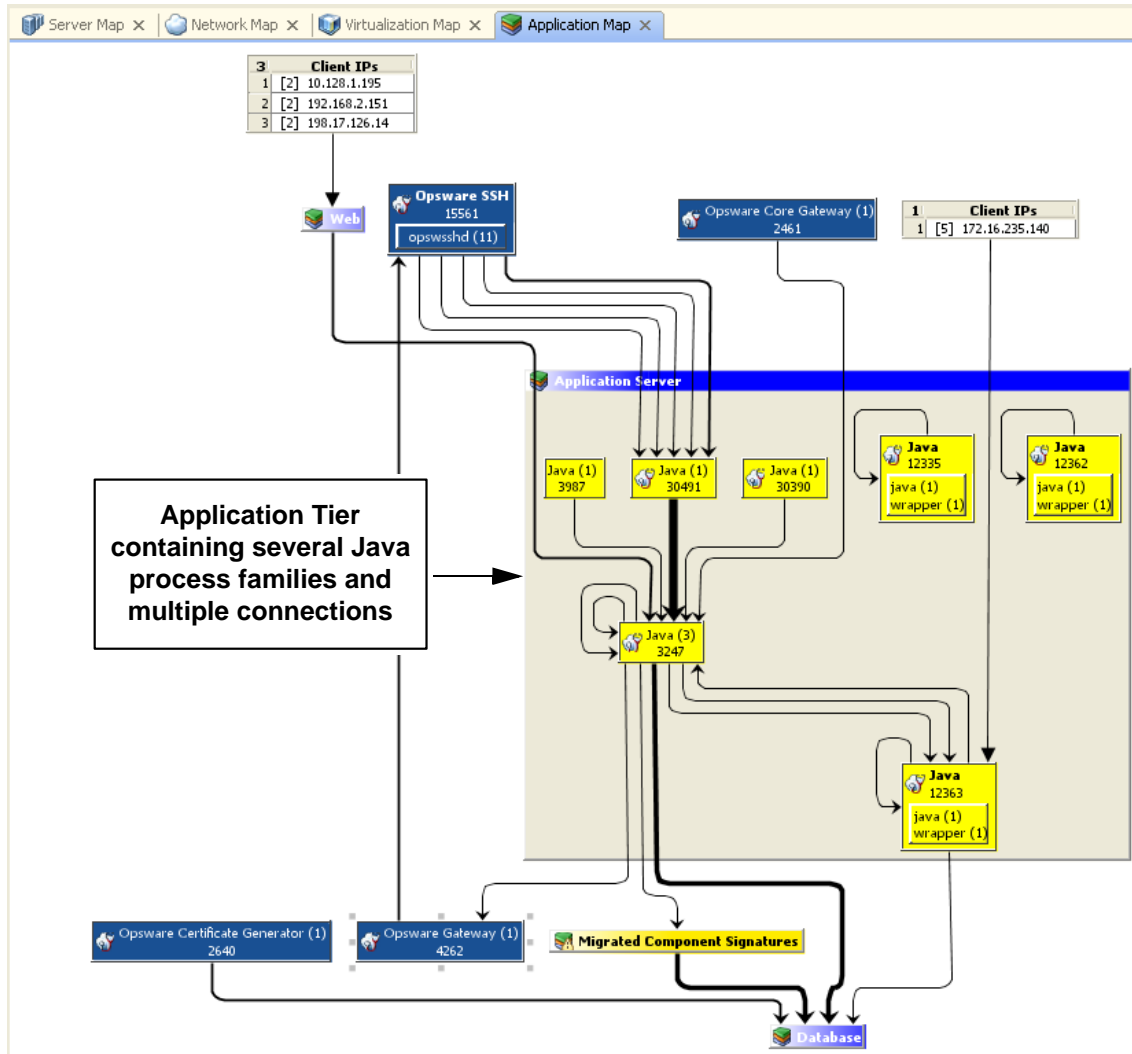
Application Map

The Application Map displays the logical structure of an application, including an application's tiers and component signatures and the connections between application component signatures, external clients, and other dependencies. By default, this map is initially empty until you create the tiers and define the component signatures that comprise an application. See "Creating an Application Tier" on page 79.



The Application Map also shows the external IP addresses (Client IPs) that are connected to the application and the external IP addresses (external dependencies) that the application connects to and depends on.

Component signatures can be attached to a tier and used to recognize process families as named elements of an application. Tiers that do not have process families are shown in the Application Map. You can group them within a tier and make them visually distinct by modifying their color and name. See Figure 1-8.

Figure 1-8: Application Map



If an application component signature does not have any process families associated with

it, then the tier object title bar will display a warning icon , for example,  .

Server Map

The Server Map displays the physical layout of how elements of an application map to a set of servers (virtual or non-virtual and hypervisors), including the process families that are running on servers and how those processes families are connected to one another.

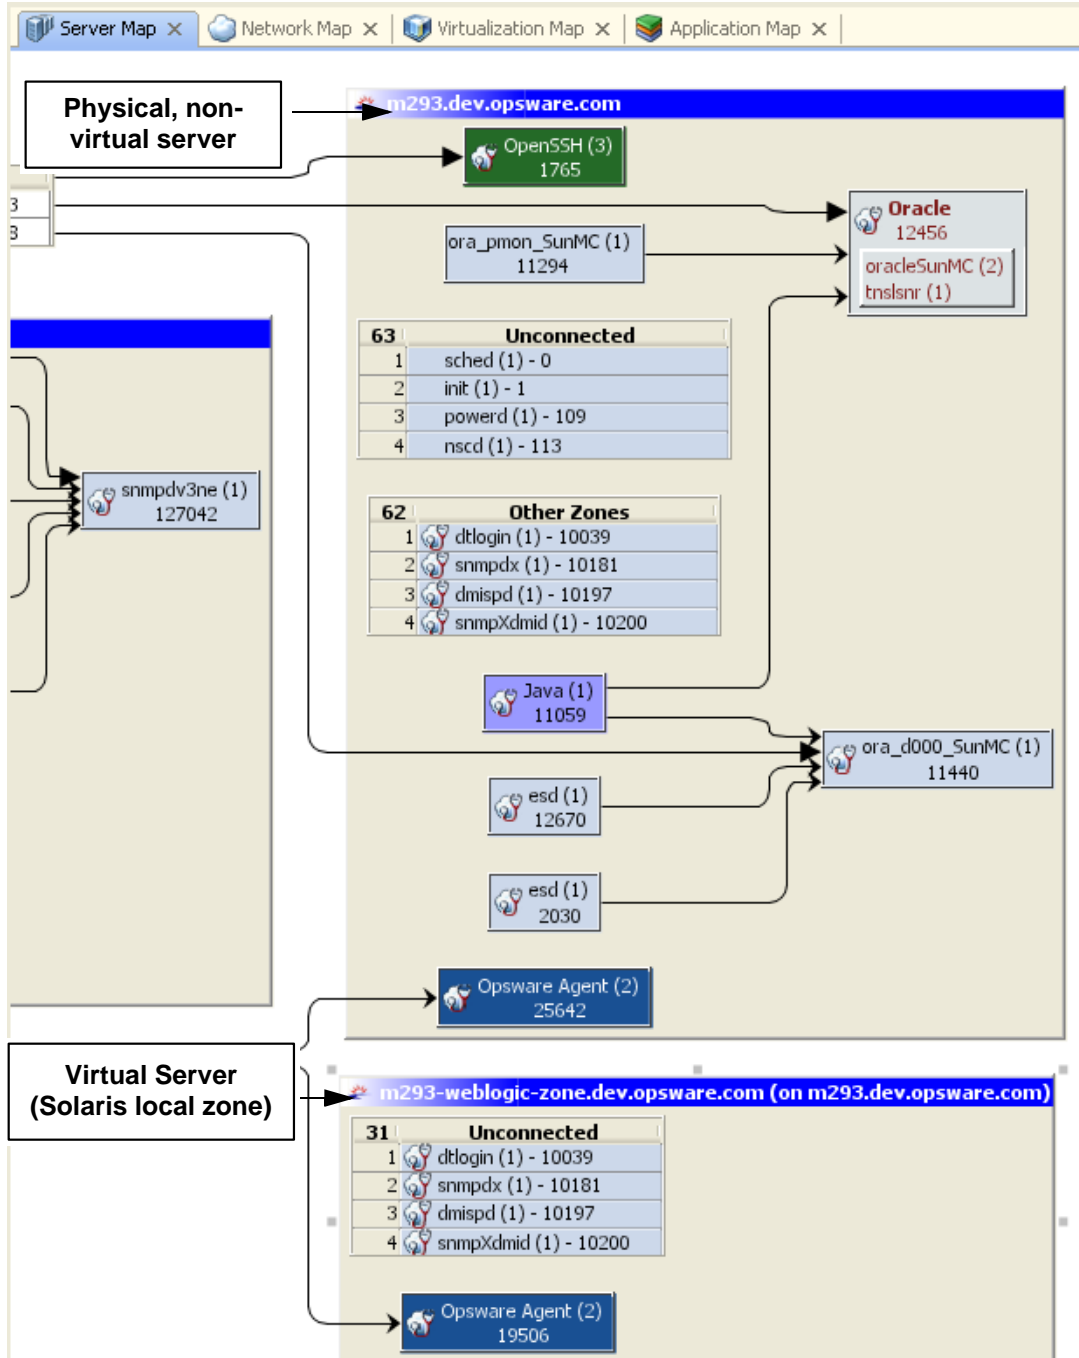
The Server Map also shows the external IP addresses (client IPs) that are connected to the application and the external IP addresses (external dependencies) that the application connects to and depends on.

Virtual servers are also displayed in the Server Map just as regular non-virtual servers are, but with some differences:

- Virtual servers' title bars provide the hypervisor name in parentheses. You can show or hide the hypervisor name using the Show Server Names button on the toolbar.
- Solaris global zones (hypervisors) contain a single list named "Other Zones" that represent all processes and process families from all local zones running on it (minus the global zone processes).
- Unmanaged virtual servers or hypervisors that were not included in the scan will not be shown in this map.

Figure 1-9 shows the Server Map displaying a regular non virtual server (top) and a virtual server (bottom) with the hypervisor name showing in the title bar.

Figure 1-9: Server Map



Virtualization Map

The Virtualization Map displays all virtual servers – hypervisors and VMware virtual machines and Solaris local zones – that you have scanned with VAM. The Virtual Map will show non-virtual servers as well.

If you open a hypervisor with VAM, the Virtual Map displays all the hypervisor's virtual servers inside the parent hypervisor's server box. (See Figure 1-10.)

If you choose to scan only the hypervisor but not its guests, then you will only see limited information about the virtual servers and will not be able to open them and view their contents.

For all Solaris global zones that are scanned, the Virtual Map displays a list named Other Zones. It contains all processes visible in the global zone that are actually running in a local zone, but that were not included in the scan.

For VMware ESX hypervisors, vSwitches are shown alongside virtual machines, in addition to the connections between the virtual machines and the vSwitches' port groups, which will always appear as a green matching-duplex ethernet connection.

The information that VAM is able to display is also dependent upon whether or not the hypervisor or virtual server has an agent installed on it.

For example:

- It is possible that a virtual server has an agent installed on it, but not its hypervisor. In this case, VAM will only display the scanned virtual server but not its hypervisor.
- It is possible to have a hypervisor that has an agent installed on it, but some or all of its guest virtual servers do not. In this case, VAM will show all guest virtual servers but only with limited virtual server information. In other words, you won't be able to open it (drill down into it).

Figure 1-10 shows a virtual server (a Solaris zone) contained inside a hypervisor in the Virtualization Map.




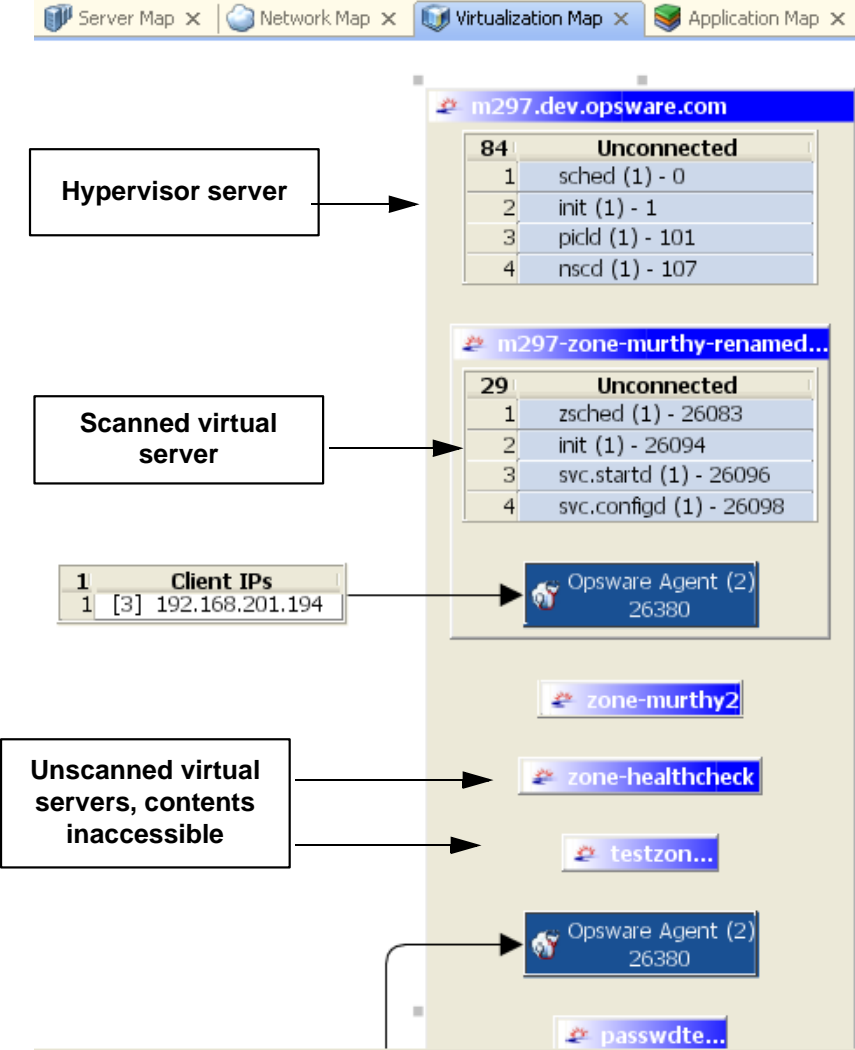
If you open a virtual server and the arrangement of server boxes is difficult to view inside the map, click the Rotate Layout icon  on the toolbar and VAM will rotate the layout for a different unique maps of the servers.

Figure 1-10: Virtual server and its hypervisor visualized in the Virtualization Map



Network Map

The Network Map displays a physical (and virtual) layout of how the elements of an application connect to each other within the network, including the network interfaces on a server and the devices (switches and vSwitches) to which the server is connected.

The map shows the process families that are connected over network interfaces on a server, and the ports and port groups, VLANs, and listeners that a server's network interfaces are connected to. All network elements are displayed in green.





The Network Map also shows external IP addresses (client IPs) that are connected to an application and the external IP addresses (external dependencies) that an application connects to and depends on.

The Network Map also highlights layer 1 connections that have speed or duplex mismatches between interfaces and network devices. The Network Map uses the following color scheme:

- Green lines and arrows indicate duplex and speed matches.
- Red lines and arrows indicates either a duplex or speed mismatch.
- Gray lines and arrows indicate that not enough information was gathered to determine the duplex or speed matches.

VMware vSwitches are shown alongside virtual machines. Connections between them will appear as a green matching-duplex ethernet connection. If VAM cannot tell from the information gathered in a scan the devices a group of processes are connected to, it groups them into a box named "Extraneous."

Network interfaces and devices use the following symbols:

-  **Network Device:** (As shown in the Device Tree) A switch or a vSwitch. In the Network Map, the network device will display a box with the device manufacturer's logo in the title bar. For example, if it is a Cisco switch, the icon in the Network Map will show the Cisco logo in the title bar ." data-bbox="241 774 668 791"/>
-  **Network Interface Card (NIC):** When available, a NIC is shown with its IP address and any processes connected to it.
-  **Listeners:** Process families that are listening on or connected to more than one network interface appear multiple times in the Network Map.

-  **Network Device Port**
-  **Virtual LAN**


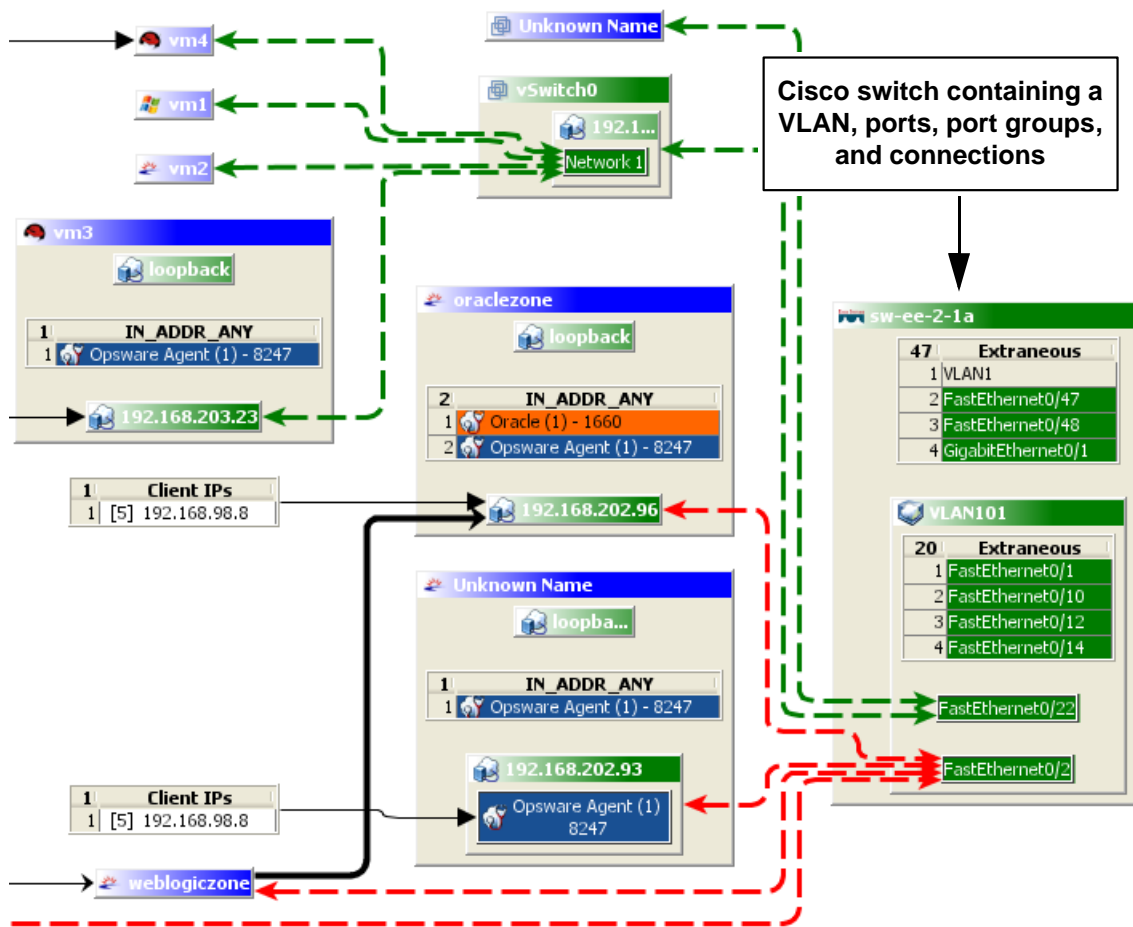
Network devices without connections are also shown. If it is unknown whether the network device is connected to a server or another network device, a warning icon  appears next to it in the Device Tree. See “Error Messages” on page 98.

Figure 1-11 illustrates network devices (green) as shown in the Network Map, with switches and virtual switches, ports and port groups, and network connections (green and red lines).

Figure 1-11: Network Map



Symbols Used in Maps





VAM uses a variety of symbols in the Network Map, Virtual Map, Server Map, and Application Map, such as lines, arrows, diamonds, and so on. This section contains the following topics:

- Process Connection Symbols
- Map Process and Network Connections
- Saving a Map to .gif, .jpg, or .svg
- Printing a Map
- Source and Target Scan Results

Process Connection Symbols

In the VAM Maps, lines and arrows represent connections between process families. In some cases, VAM will know both the source of the connection and the destination. In other cases, VAM may only know either the source or the destination of the connection.

To represent these process connection relationships, VAM uses the following lines and arrows:

- **Source Unknown – Diamond:**  An arrow with a diamond at its source indicates that VAM does not know the source of the connection. If there is a solid line from a process connection source (in other words, it doesn't show a diamond) then VAM knows the process source.
- **Destination Unknown – Hollow Arrow:**  A hollow arrow represents an inbound connection to a process family destination that is unknown.
- **From Remote IP – Solid Arrow:**  A solid arrow represents an inbound client connection from a remote IP, such as TCP or UDP, where the destination process family is known.
- **From Process – Lined Arrow:**  A lined arrow represents a connection from a process family where the destination process family is known.

Map Process and Network Connections


In the VAM Maps, lines represent the following connections between devices:

- **Client link:** An internal connection that is labeled by the client IP address.

- **Process link:** A collection of TCP or UDP connections between process families. This link displays processes that provide a network service, such as listening for network connections, and processes that have a connection to another process or server.
- **Layer 1 connection:** A physical link between a server's network interfaces and switch ports/switches. Layer 1 connections are indicated by colored, dashed lines in the map:






 A green dashed line indicates that there is no duplex mismatch.

 A red dashed line indicates that there is a duplex mismatch.

 A gray dashed line indicates that there may or may not be a duplex mismatch because at least one value is unknown.

- **Line Thickness:** The thinness or thickness of the line represents the number of connections associated with the link. A smaller number is indicated by a thinner line and a larger number is indicated by a thicker line, as illustrated in Figure 1-12

Figure 1-12: Line Thickness and Process Connection Relationship

	1-4
	5-16
	17-64
	65-256
	257+

Saving a Map to .gif, .jpg, or .svg

You can export a map to a .gif, .jpg, .svg file for use in other applications where you can annotate the drawing or map the exported file in a web browser.

To export a map to a .gif, .jpg, or .svg file, perform the following steps:

- 1** From the **File** menu, select **Save As Image**.
- 2** Select a directory where you want the file to be located.
- 3** Enter a file name that includes either .gif, .jpg, or .svg as the file name extension.
- 4** Click **Save As Image**.

Printing a Map


You can print a map on single and multiple sheets of paper, and you can also title the map for better presentation.

If the map you want to print is very dense and complex, you can make adjustments by zooming in and zooming out, and by creating rows and columns that will break the map up over several pages. Doing this enables you to print the map on multiple sheets of paper, thus increasing the map's readability.

To adjust the map before you print:

- Use the Zoom In, Zoom Out, or Zoom selector to increase or decrease the size of the map before you print.
- Click the Row and Column up or down arrows to add or remove rows and columns, which will break up the map over several pages
- Enter a title for the map

To print a map, perform the following steps:

- 1** From the **File** menu, select **Print** or select the  toolbar icon.
- 2** (Optional) In the Print window, specify page setup and printer options, including a title that you want to appear on the printed map.
- 3** Click **Print**.

Property Pages

VAM displays a Property Page for the component selected in the Device Tree, Application Tree, or any of the maps.

The information displayed in the Property Page varies depending on the component type, such as a server, network device, process family, tier, application component, and link. Depending on the type of object you select, this page includes the number of users

logged in, load average, swap usage, memory usage, application components, network devices, network ports, VLANs, tiers, links, and so on. It shows the MAC addresses for each network interface, as shown in Figure 1-13.

Figure 1-13: Property page for a server

Property	Value
Name	m208.dev.opsware.com
ID	420043
Operating System	Red Hat Enterprise Linux AS 3
Kernel	Linux 2.4.21-27.ELsmp i686
Code Set	UTF-8
Users	4
1 Min Load Average	0.23
5 Min Load Average	0.05
15 Min Load Average	0.02
Free Memory	19,357,696 bytes
Total Memory	2,104,594,432 bytes
Swap Used	0 bytes
Swap Total	2,089,209,856 bytes
Swap In	0 bytes/s
Swap Out	0 bytes/s
Interfaces	2
loopback	
MAC Address	
Broadcast	
Subnet Mask	255.0.0.0
Device	localhost
192.168.165.130	
MAC Address	00123F20BBBC
Broadcast	192.168.165.143
Subnet Mask	255.255.255.240
Device	eth0
File Systems	3
/boot	
Size	103,515,136 bytes

Server Property Page

The Property Page for a server displays the following information:

- **Name:** The host name of the server.
- **Opsware ID:** The Opsware unique identifier for the server.
- **OS:** The operating system of the server.
- **Kernel:** The kernel version of the operating system (when applicable).

- **Agent Version:** The version of the Opsware Agent that enables the server to be managed and scanned.
- **Last Reported Time:** The most recent time that the Opsware Agent communicated with the Opsware core.
- **Uptime:** The length of time the server has been powered on.
- **Codeset:** The character encoding for the server's locale.
- **Users:** The number of users that are currently logged in.
- **Load averages:** 1-minute, 5-minute, and 15-minute load averages. The load average for servers running a Windows operating system will display unknown because it is not supported by Microsoft.
- **Memory usage:** The total free memory.
- **Swap usage:** The total used swap and swap in/out activity.
- **Interfaces:** The number of network interfaces. For each interface on a server, the following information is displayed:
 - MAC address
 - Broadcast address
 - Subnet mask
 - Device
- **File systems:** The size of free space, percent used, and associated device for each file system.
- **Virtual Machines/Zones:** If the selected server is a hypervisor (Solaris Global Zone or VMware ESX server), you can expand the list and view all local zones (Solaris) virtual machines (VMware ESX). Each virtual machine or local zone will display its own server properties. For more information, see "Virtual Server Properties" on page 70.

Properties for Servers and Devices with Compliance Policies

For servers or network devices that have compliance policies associated with them (Software, AppConfig, Patch, Audit, Duplex), the server's properties will show a rollup compliance status for all attached policies. You can expand the compliance list to view each individual compliance policy attached to the server.




Each compliance category will display one of the following compliance statuses:

- **Compliant:** The compliance scan ran successfully and the actual server or device configuration matches the criteria defined in the policy.
- **Partial:** The compliance scan ran successfully, but the server or device configuration did not fully pass the compliance criteria defined in the policy.
- **Noncompliant:** The compliance scan ran and the actual server or device configuration did not match the criteria defined in the policy.
- **Scan Failure:** The compliance scan was unable to run.
- **Scan Needed:** The results are unavailable, perhaps because a compliance scan was never run (for example, on a new installation), or the configuration on the server changed since the last time information was reported to the Compliance Dashboard.
- **Scanning:** The compliance scan currently being run.

You can launch the Device Explorer or remote terminal in the SAS Client to view and remediate any compliance discrepancies by clicking on a compliance status link in the properties window. For NAS-enabled cores, clicking a compliance status link launches the NAS Web interface.

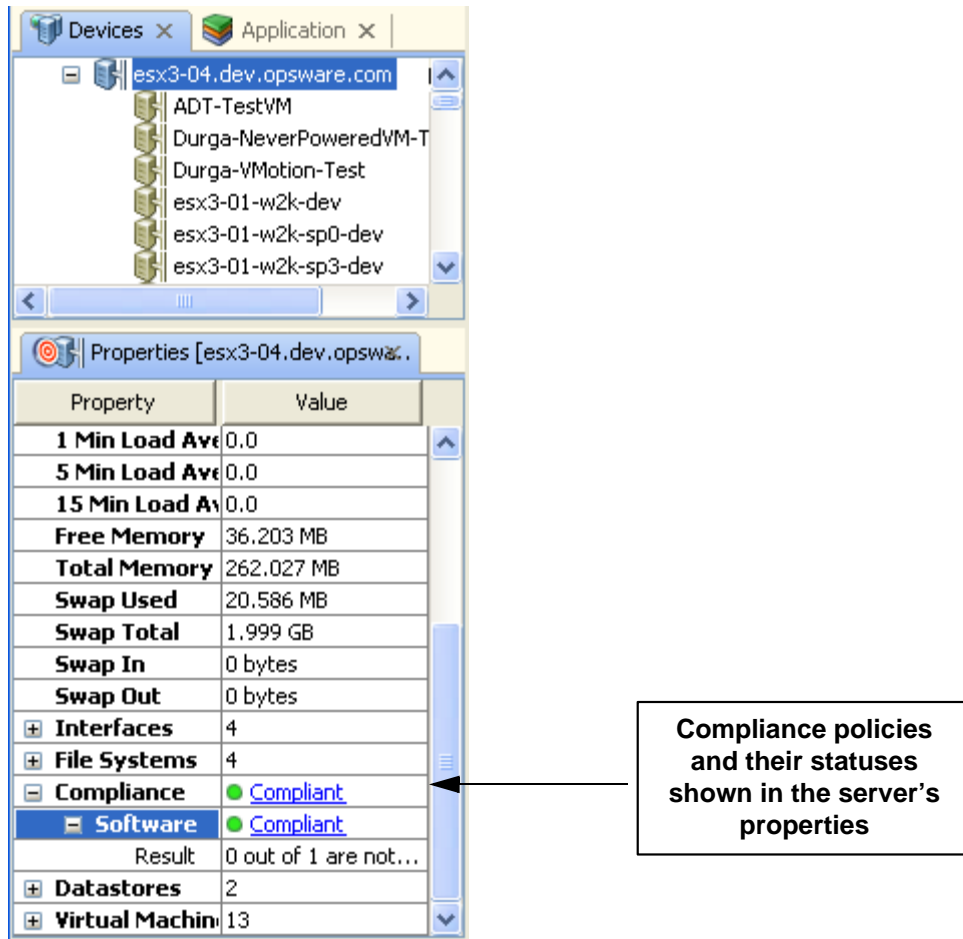
For information on launching a Device Explorer, remote terminal, or global shell, see “Accessing Servers From VAM” on page 76.

Figure 1-14 shows a server’s properties and lists compliance information about the server. In this case, the server has the following one compliance policy attached to it:

- **Audit:** Non-compliant (indicated by the  icon).
- **Software:** Compliant (indicated by the  icon).
- **Platform Test:** Non-compliant (indicated by  icon).

Note that when any compliance policy on the server is non-compliant, then the main compliance policies row shows a non-compliant status, as seen in Figure 1-14.

Figure 1-14: Server properties compliance Information



A VAM scan is not the same thing as a compliance scan, but they are related. A compliance scan runs in the SAS Client and checks a server or device's compliance status and reports this information to the Compliance Dashboard inside of the SAS Client (and by extension, a server's properties in VAM).

The actual compliance state you are viewing in VAM may have changed since you last scanned the server or device. To get the most current information, click **Refresh Scan**

Results. For more information on the Compliance feature, see the *Opware® SAS User's Guide: Application Automation*.

Link Property Page

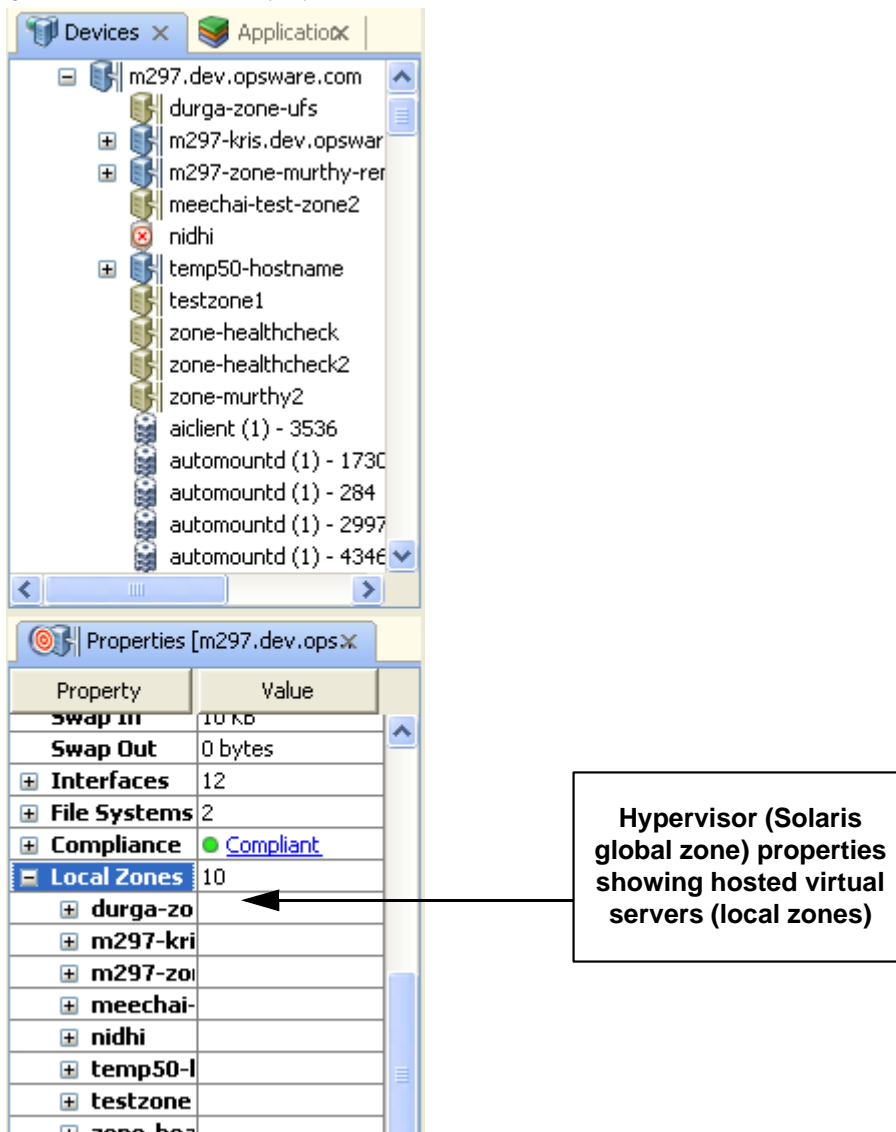
The Property Page for a link displays the following information:

- **Protocol:** The TCP or UDP.
- **Port:** The destination port that is associated with this link.
- **Connections:** The number of connections associated with this link. For each connection, the following information is displayed:
 - **End points:** The process families (if known). IP addresses (if unknown).
 - **Ephemeral port number:** A random port that is assigned by the operating system.

Virtual Server Properties

The properties of a virtual server display all the same information as a physical server except that VMware and Solaris 10 hypervisors will show all hosted virtual servers. You can expand each hosted virtual server and view its properties. Conversely, each virtual server will contain in its properties the hypervisor that is hosting it. Figure 1-15 illustrates this feature.

Figure 1-15: Virtual Server properties



Network Device Property Page

The Property Page for a network device displays the following information:

- **Name:** The name of the network device.
- **Last Reported Time:** The date of the last successful scan of the network device by Opsware NAS.
- **Manufacturer:** The Vendor that manufactures the network device.
- **Model:** The model number of the network device.
- **Operating System:** The operating system running on the network device.
- **Firmware Version:** The firmware version number for the device.
- **Asset Tag:** The assigned number used for tracking the network device.
- **VLANs:** The total number of VLANs that this network device has.
- **Ports:** The total number of ports that this network device has.
- **Compliance:** For network devices that have compliance policies associated with them (Duplex), the properties will display its compliance status. For information on Compliance statuses, see “Properties for Servers and Devices with Compliance Policies” on page 66.

Virtual Switch Properties

Virtual Switch properties will display the following information:

- **Port Groups:** These can be expanded to view port groups configured for the selected vSwitch.
- **Network Interfaces:** These can be expanded to view network interfaces assigned to the selected vSwitch.

Port Group Properties

Port group properties will display the following information:

- **Port Group Name:** The name of port group
- **VLAN ID:** The VLAN ID of port group. This is optional in the VMware management user interface.

Network Interface Property Page

The Property Page for a network interface displays the following information:

- **IP Address:** The IP address that is associated with a network interface.
- **MAC Address:** The Media Access Control ID that is associated with a network interface.
- **Subnet Mask:** The subnet that is associated with a network interface.
- **Broadcast:** The broadcast address that is associated with a network interface.
- **Device:** The device that is associated with a network interface.
- **Duplex:** The configured duplex (if it can be collected).
- **Negotiated Duplex:** The negotiated duplex (if it can be collected).
- **Speed:** The configured speed in Mbps (if it can be collected).
- **Negotiated Speed:** The negotiated speed in Mbps (if it can be collected).

Process Family Property Page

The Property Page for a process family displays the following information:

- **Name:** The name of the controlling process of the family.
- **Family ID:** The unique ID given to the process family.
- **Extended Family:** The name of the extended process family, if the selected process family belongs to an extended process family.
- **Max. Resident Memory:** The maximum permanent memory used by the process family, in bytes.
- **Max. Virtual Memory:** The maximum permanent memory used by the process family, in bytes.
- **Max. Run Time:** The length of time the process has been running.
- **Total CPU Time:** The total length of time the process used CPU resources.
- **Max CPU Utilization:** The total amount of CPU resources used by the process.
- **Group ID:** The group ID of the process family on Unix and the session ID on Windows.
- **Listeners:** The interface and port for each listener.
- **Incoming connections:** The connections incoming to the process family, grouped by process family (if known, the IP address otherwise) and interface.

- **Outgoing connections:** The connections outgoing from the process family, grouped by process family (if known, IP address otherwise) and interface.
- **Modules:** The shared libraries associated with the process family. These include DLLs on Windows and shared object files on Unix.
- **Open files:** The files that the process family currently has open.
- **Software Packages:** The packages associated with the files that the process family has open.
- **Processes:** The number of individual processes in the process family. For each process, the following information is displayed:
 - **PID:** The process ID.
 - **User:** The user ID the process is running as.
 - **Command line:** The command line used to start the process.
 - **Path:** The path to the process binary.
 - **Memory statistics:** The percentage of physical memory consumed by the process, the resident size (in bytes) of the process and the virtual size (in bytes) of the process.
 - **Run time:** The time (in milliseconds) that the process has been running.
 - **CPU Statistics:** The CPU time accumulated by the process and the percentage of CPU consumed by the process since it began.
 - **Environment:** The name and value of each environment variable in the process environment.

Application Tier Property Page

The Property Page for an application tier displays the following information:

- **Name:** The name of the application tier as displayed in the Application Tree.
- **Application Tiers:** The number of subtiers that are currently recognized in the tier.
- **Component signatures:** The number of application component signatures that are currently recognized in the tier.
- **Server Filter:** The servers that are associated with this tier. Only matching servers are filtered for matching application components.

Component Signature Property Page

The Property Page for an application component displays the following information:

- **Name:** The name of the application component as displayed in the Application Tree.
- **Alias** (Optional): The name of the application component as displayed in the different views. The name will be shown as an alias in not defined.
- **Families:** The number of process families recognized as this application component.
- **Process Name:** The process name filter used to recognize this application component.
- **Command Line:** The command line filter used to recognize this application component.
- **Connected To Port:** The port that the server is connected to.
- **Listener Port:** The listen port used to recognize this application component.
- **Executable Path:** The executable path filter used to recognize this application component.
- **Open Files:** The open file filter used to recognize this application component.
- **Modules:** Shared libraries that are associated with the process family, DLL files on Windows operating systems and shared object files on Unix operating systems.
- **Background:** The background color displayed in the different maps.
- **Foreground:** The foreground text color displayed in the different maps.

VAM Options

For VAM, you can specify the following options:

- Virtualization Settings
- Scan Time-Out Preference

Virtualization Settings

You can configure SAS Client options that allow you to choose whether or not you want to perform a scan on any virtual servers or hypervisors related to the virtual server you want to open in VAM.

For example, if you want to visualize a VMware virtual machine (VM) or Solaris local zone in VAM, by default you will be asked if you also want to scan any virtualization relationships – in other words, the system asks if you want VAM to also scan the hypervisor that is hosting the selected virtual server. Depending upon the virtual server you select, VAM might have to scan several related virtual servers in order to visualize a single virtual server in VAM.

Conversely, if you select a hypervisor to open in VAM, you will also be asked if you want to scan any virtualization relationships – in this case, VAM would need to scan all of the hosted virtual servers, which could take a long time to perform.

By default, VAM will always ask you if you want to scan virtual relationships, but you can set your own default behavior for scanning related virtual servers with the following virtualization options:

- Ask each time if you want to scan related virtual and host servers.
- Always scan related virtual and host servers.
- Never scan related virtual and hypervisor servers.

To change the virtualization settings, perform the following steps:

- 1** From the **Edit** menu, select **Options**.
- 2** In the Set Options window, in the Views pane, select Visual Application Manager.
- 3** Specify your desired Virtualization Settings, then click **OK** when you are finished.

Scan Time-Out Preference

VAM is optimized to scan a maximum of 50 servers. A number of factors affect the time it takes for a scan to complete, including the load on the scanned servers and the load on Opware. The default scan time-out is set to 300 seconds. You can reset this time-out value to a minimum of 30 seconds or to a maximum of 3600 seconds.

To change the scan time-out, perform the following steps:

- 1** From the **Edit** menu, select **Options**.
- 2** In the Set Options window, in the Views pane, select Visual Application Manager.
- 3** In the Scan Timeout section, move the slider to increase or decrease the number of seconds at which you want the scanning process to stop.
- 4** Click **OK** to save your changes or click **Cancel** to close the window without saving your changes.

You can also access these options from inside the SAS Client by selecting **Options** from the **Tools** menu.

Accessing Servers From VAM

As a means of helping you troubleshoot and take action for server and application errors, VAM gives you easy access to servers through the following methods:

- Opening a Device Explorer
- Opening a Remote Terminal
- Opening a Global Shell

Opening a Device Explorer

To view detailed information about a server using the Server Explorer, perform the following steps:

- 1** In a map, select one or more servers.
- 2** Right-click and then select **Open with Opsware SAS Client** to open a Server Explorer for each selected server.

See the *Opsware® SAS User's Guide: Server Automation* for information about how to use the Server Explorer.

Opening a Remote Terminal

The Remote Terminal enables you to log into devices (servers and network devices) and run native commands.

To open a Remote Terminal, perform one of the following tasks:

- 1** In a map, select one or more servers.
- 2** Right-click and then select **Open Remote Terminal** to open the Select Remote Login window.
- 3** In the Login column, select a login ID from the drop-down list, such as root or LocalSystem, or any of the user logins that might be configured.

- 4 Click **OK** to open a Remote Terminal for each selection.

See the *Opware® SAS User's Guide: Server Automation* for information about using utilities in a Remote Terminal.


Opening a Global Shell

You can use the Global Shell feature to navigate between servers and connected network devices by tracing their layer 1 connections in the `/opsw/Servers/@` and `/opsw/Network/@` directories in the OGFS.

In the OGFS, you can also run scripts to perform the following tasks:

- Find servers and network devices.
- Find all servers that are connected to a certain switch.
- Display the network interfaces of a certain server.
- Get the IP addresses of all devices.
- Compare two files to identify changes made, such as what changes were made to a device configuration (.conf) file.
- Change device details, such as the snmp-location.

To launch the Global Shell, perform one of the following tasks:

- From the **File** Menu, select **Global Shell**.
- Select the  toolbar icon.

See the *Opware® SAS User's Guide: Server Automation* for information about how to use Global Shell.

Creating Application Definitions

An application definition allows you to transform a data display that contains extraneous and hard-to-understand information into a focused and easy-to-understand view of the relevant data. Based on the application tiers and component signatures that you create, VAM recognizes actual application processes and displays them in the Application Map according to any visual customizations you have made to them.

You create application definitions in order to recognize processes by giving them meaningful names and appearances (colors). You also use application definitions to define the logical tiers of an application and display component signatures according to the tier in which they reside.

See “Evaluation Order” on page 82 for information on the order in which VAM scans component signatures and matches them to processes and process families on servers.

For information on understanding and creating application definitions, see the following topics:

- “Application Tiers” on page 79
- “Creating an Application Tier” on page 79
- “Application Component Signatures” on page 81
- “Creating a Component Signature” on page 84

Application Templates

When you first scan servers with VAM and visualize them, the Application Tier is empty – it has no application definitions until you create them. (There are, however, some predefined commonly used default applications built into the product, such as Apache, WebSphere, and so on, contained in the Default Component Signatures folder.) Once you create and define an application definition with tiers and component signatures, you can save the application definition as a template, which can be reused by yourself or others on your team.

You can also set an application definition to use as the default template, so that whenever you open VAM, it always opens using the application definitions saved in the default template. If you make changes to an application that is based upon a template, and do not wish to save the changes, you can restore the default template.

Setting a Default Application Template

If you have made changes to the application definition and want to set this as the default, select **Set as Default Template** from the **File** menu.

Resetting the Default Application Template

If you have made changes to the application definition and want to restore the previously saved default application, select **Reset Default Template** from the **File** menu.

Importing an Application Template

If you would like to import an application template that has already been saved, from the **File** menu, select **Import Template** and select the template to import.



Importing an application template will replace any existing application definitions in your current VAM session.

Application Tiers

Application tiers provide an architectural framework to organize and display application component signatures. You can add, edit, delete, cut, copy, and paste application tiers in the Application Tree. You can paste an application tier before or after a selected position in the Application Tree to rearrange the order. The order of application tiers (and the application components they contain) is significant because it affects the order that the process families are assigned to component signatures. For more information, see “Evaluation Order” on page 82.

Tiers that do not have any component signatures (including sub-tiers that do not contain any recognized process families) are not drawn in the map. If any tiers have application component signatures that do not recognize any process families, they and their ancestors are represented with warning icons in the tree and by yellow title bars in the view. This allows you to quickly identify component signatures that should be running but are not.

Creating an Application Tier

To create an application tier in the Application Tree, perform the following steps:

- 1** In the Application Tree, select a tier.
- 2** From the **Application** menu, select **New Tier** or right-click and then select **New Tier** to display the New Tier window.
- 3** Enter a name and server filter.
- 4** (Optional) Select the check box to enable case sensitivity for the server filter.
- 5** Click **Apply** to save your changes or click **OK** to save your changes and close the window.

Or

Click **Cancel** to close the window without saving your changes.

Editing an Application Tier

To edit an application tier in the Application Tree, perform the following steps:


- 1** In the Application Tree, select a tier.
- 2** From the **Application** menu, select **Edit** or right-click and then select **Edit** to display the Edit Tier window.
- 3** Make your changes.
- 4** (Optional) Select the check box to enable case sensitivity for the server filter.
- 5** Click **Apply** to save your changes or click **OK** to save your changes and close the window.

Or

Click **Cancel** to close the window without saving your changes.

Deleting an Application Tier



To delete an application tier from the Application Tree, perform the following steps:

- 1** In the Application Tree, select a tier.
- 2** From the **Edit** menu, select **Delete** or right-click and then select **Delete** (or, click the delete toolbar button ).

Cutting and Copying an Application Tier


You can cut and copy an application tier to the clipboard. After you do this, you can paste the application tier before or after a selected position in the Application Tree to rearrange the order. The order of application tiers (and the component signatures they contain) is significant because it affects the order that the process families are assigned to component signatures.

To cut and copy an application tier in the Application Tree, perform the following steps:

- 1** In the Application Tree, select a tier.
- 2** From the toolbar select either the  icon or the  icon, or right-click and select **Cut** or **Copy**.

Pasting an Application Tier

To paste an application tier in the Device Tree, perform one of the following tasks:

- Select a tier in the Device Tree and then select the Paste icon . The tiers that you cut or copied to the clipboard will be appended to the selected tier's children. When you select a component signature in the Device Tree, the Paste icon will be disabled.
- Select a tier in the Device Tree and then select **Paste Before** from the **Edit** menu. The tiers cut or copied to the clipboard will be inserted into the selected tier's parent tier that is *before* (above) the selected tier. When you select an application component or the root tier in the Device Tree, **Paste Before** will be disabled.

Application Component Signatures

Application component signatures are organized and displayed in application tiers and contain processes and process families. You can add, edit, delete, cut, copy, and paste components signatures in the Application Tree. You can paste a component signature

before or after a selected position in the Application Tree to rearrange the order. The order of component signatures (and the tiers that contain them) is significant because it affects the order that the process families are assigned to component signatures.

VAM comes with a set of predefined default component signatures that recognize a variety of commonly used component signatures, such as Apache HTTP, Microsoft IIS, JBoss, Oracle, and so on. So, if your server has any of these applications installed, VAM will be able to recognize and display them in the Application Tree and the maps.

VAM also includes a set of Opsware component signatures, such as the Opsware Agent, Opsware Build Manager, NAS Syslog Server, Opsware Command Engine, and so on. Many of these component signatures will appear only if you use VAM to scan the server or servers that the Opsware core is installed on, while others, like the Opsware Agent, will appear on all reachable managed servers.

Evaluation Order

The order that component signatures are recognized in VAM is important because a process family is associated with the first component signature that it matches in the Application Tree. Evaluation order is significant especially when the recognition criteria for a component signature matches the same process family found in multiple component signatures.

Component signatures are evaluated in a depth-first, top to bottom order: component signatures in a tier's sub-tiers are evaluated before the tier's component signatures (depth-first), and tiers in the Application Tree are evaluated from top to bottom. Components signatures are applied in the order in which they appear in each tier.

After all user-created application tiers and component signature hierarchies are evaluated, then all of the default Opsware component signatures are evaluated; for example, Opsware NAS Management System, Opsware NAS Syslog Server, and so on. After the Opsware component signatures are evaluated, then all of the default predefined component signatures are evaluated; for example, Apache HTTP, Internet Information Server, and so on.


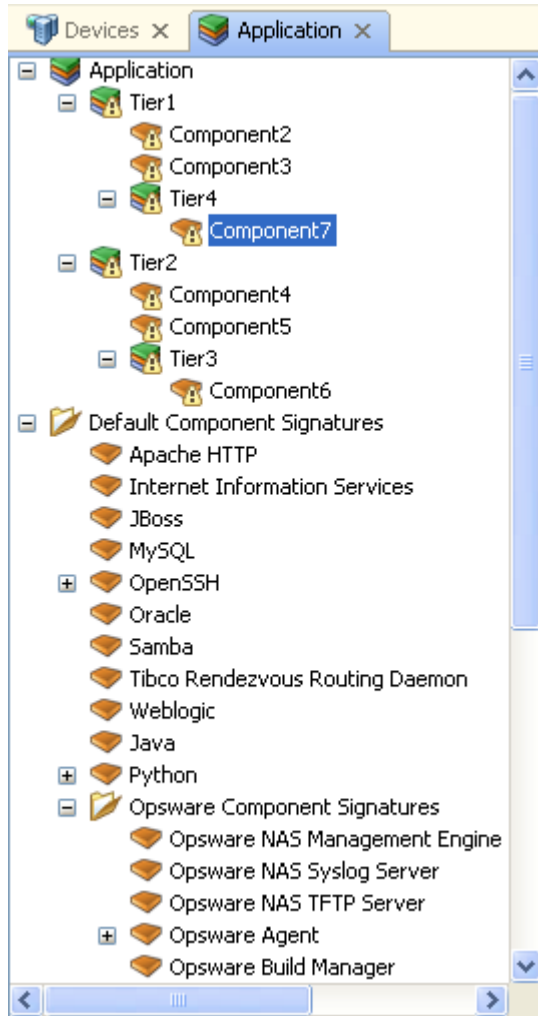
Consider an application definition that has the structure shown in Figure 1-16. In this image, no processes or process families match the component signature definitions, and so the signatures are represented with the  icon.

Figure 1-16: Application Tier Component Signature Evaluation Order



In this application definition example, the component signatures are evaluated in the following order:

1. Component7
2. Component2
3. Component3

4. Component6
5. Component4
6. Component5
7. Opsware NAS Management System
8. Opsware NAS Syslog Server
9. <All remaining Opsware component signatures, top to bottom>
10. Apache HTTP
11. Internet Information Services
12. <All remaining default component signatures, top to bottom>

Creating a Component Signature

To create a component signature in the Application Tree, perform the following steps:

- 1** Select the Application Tree.
- 2** From the **Application** menu, select **New Component Signature**. The New Component signature window displays.
- 3** Enter a name for the new component signature.
- 4** In the Signature section, enter any of the following information (Optional):
 - **Process Name**: The name of the process family.
 - **Command Line**: The command line that a component signature was started with.
 - **Executable Path**: The path to the executable file of this application component.
 - **Open Files**: The name of an open file.
 - **Modules**: The shared libraries associated with the process family. These include DLLs on Windows and shared object files on Unix.
 - **Connected to Port**: The port the component signature is connected to.
 - **Listener Port**: The port on which the component signatures is listening.
- 5** In the Preferences section, enter the following optional information:
 - **Alias**: The name of the application component as displayed in the different views.
 - **Background color**: The background color displayed in the different views.

- **Foreground color:** The foreground text color displayed in the different views.

6 Click **Apply** to save your changes or click **OK** to save your changes and close the window.

Or

Click **Cancel** to close the window without saving your changes.

Editing Component Signatures

To edit a component signature in the Application Tree, perform the following steps:

- 1** In the Application Tree, select a component signature.
- 2** From the **Application** menu, select **Edit** or right-click and then select **Edit** to display the Component Signature window.
- 3** Make your changes.
- 4** Click **Apply** to save your changes or click **OK** to save your changes and close the window.

Or

Click **Cancel** to close the window without saving your changes.

Deleting Component Signatures

To delete a component signature from the Application Tree, perform the following steps:

- 1** In the Application Tree, select an application component.
- 2** From the **Edit** menu, select **Delete** or right-click and then select **Delete**.



Cutting and Copying a Component Signature

You can cut and copy a component signature to the clipboard. After you do this, you can paste the component signature before or after a selected position in the Application Tree to rearrange the order.





Default component signatures and Opware component signatures can be copied and pasted into user-created application tiers, but cannot be deleted or overwritten.

To cut and copy an application tier in the Application Tree, perform the following steps:

- 1** In the Application Tree, select a tier.
- 2** From the toolbar, select the  icon or the  icon, or right-click and select **Cut** or **Copy**.

Pasting a Component Signature

You can perform the following paste actions if one or more component signatures have been cut or copied to the clipboard:

- Select a component signature in the Device Tree and then select the Paste icon . The application components that you cut or copied to the clipboard will be appended to the selected tier's component signatures. When you select a component signature in the Device Tree, the **Paste Before** menu action will be disabled.
- Select a component signature in the Device Tree and then select the **Paste Before** action from the **Edit** menu. The component signatures that you cut or copied to the clipboard will be inserted into the selected component signature's parent tier *before* the selected component signature.
- Select an component signature in the Device Tree and then select the Paste icon . The components that you cut or copied to the clipboard will be inserted into the selected component signature's parent tier *after* (below) the selected component signature. When you select a tier in the Device Tree, the **Paste Before** menu action will be disabled.



If you select a combination of tiers and component signatures in the Device Tree, or if you do not select any tiers or applications in the Device Tree, the Paste icon and **Paste Before** menu action will be disabled.



For default component Opware component signatures, you can copy and paste them into user-created application tiers, but you cannot delete or overwrite them.

VAM File Management


In VAM, scan results represent the state of a set of network devices and managed servers, the process families running on those servers, the connections among those process families, and any external clients and dependencies. Scan results are saved as part of a .vam file. A .vam file can contain any number scan results, each of which can be loaded into VAM, and which can be used for scan results comparisons.

You manage scan results by managing .vam files. You can open, edit, and save a .vam file. You can also delete scan results associated with a .vam file, which enables you to keep the size of .vam files manageable and discard information that you no longer need.

Opening a .vam File

After you have launched the VAM, you can open a previously saved .vam file.

To open a .vam file, perform the following steps:

- 1** In the Opsware VAM window, select the  toolbar icon or select the **File** menu and then select **Open** to display the Open window.
- 2** In the Look in drop-down list, select the directory where the .vam file was saved.
- 3** In the left pane, double-click on the .vam file to open it. In the Application Tiers pane, you can preview the basic tier structure of the application defined.
- 4** In the Scan Results pane, select one or more scan or scan results in the .vam file that you want to open. The default scan result is the last one saved.
- 5** Click **Open**.

Saving a .vam File

To save a .vam file, perform the following steps:

- 1** From the **File** menu, select **Save** or **Save As** to open the Save window. (Note that by default, all scan results are selected to be saved.)
- 2** If you chose Save As, in the Save in drop-down list, select Opsware Global File System or Desktop to indicate where you want to save the .vam file.
- 3** In the Scan Results pane, select the scan you want to save in the .vam file. Scan results that are not selected will be deleted from the .vam file.

A Scan Result is identified by the time stamp that was generated when it was recorded. If the last scan result is deleted from a .vam file, only the application definition is saved.

- 4** In the File name field, enter the name of the .vam file.
- 5** In the Files of type drop-down list, select VAM Archives (*.vam).
- 6** (Optional) You can save a copy of this .vam file as an application template, which captures the application definition you created and can be reused later.
- 7** Click **Save**.



If you exit VAM before saving your changes (either application definition changes or topology changes), you will be prompted to choose whether you want to save your changes and then exit or exit without saving your changes.

Saving a .vam File as an Application Template

If you would like to save the current application definition as a template, so it can be reused or set to open VAM using that definition, see "Application Templates" on page 78.

Comparing Scan Results

VAM provides the ability to compare the results of two VAM scans, to help you determine if any changes have occurred between the current state of the application and its state as captured in a previous scan.

When you click **Refresh Scan Results** and then click **Save**, VAM captures and saves all the information related to your application in the currently open .vam file, including all servers and processes associated with the application, the current state of all running processes, and all values and definitions you have created in the Application Tree.

Each time you refresh scan results and then save the .vam file, scan results are saved within the .vam file. Each set of scan results can be used in a one to one comparison between the currently loaded scan results and a saved one. You can compare scan results from the currently loaded .vam file or from another .vam file.

When you compare scan results, VAM evaluates certain key objects and their attributes on a one to one basis, and displays any differences in value between those objects. The results of the comparison are displayed in the Compare Results tab at the bottom of the VAM window.

Source and Target Scan Results

The currently loaded scan results in VAM is referred to as the *source* scan results, while the set of scan results you are comparing against the currently loaded scan result is called the *target*. When you compare scan results, you are always comparing the currently loaded scan result (*source*) with another saved scan result (*target*).

Scan Results Comparison Types

VAM will display comparison results based on the following criteria:

- Object Existence Comparison
- Object Attribute Difference
- “Interesting” Object Attribute Differences

Object Existence Comparison

Comparing two sets of scan results helps determine whether or not an object exists between the them. If an object exists in one scan result, but does not exist in the other, the comparison results display the object with an attribute named "existence" and describes it as either "found" or "missing" on either the source or target. Figure 7-8 describes the types of objects compared during a VAM scan results comparison.

Table 1-2: Objects Checked for Existence in a VAM Scan Results Comparison

OBJECT	UNIQUE IDENTIFIER
Servers	Opsware server ID
File systems on servers	Opsware server ID plus device
Network interfaces on servers	Opsware server ID plus device
Processes on servers	Opsware server ID plus process id plus process start time (~one hour)
Process families on servers	Must contain at least one process in common

A process family could be running (and thus, "exist") when you refresh scan results and save the .vam file. But if the process is no longer running, and you refresh the scan results again and save the .vam file, the results change. When you compare the saved scan results (target) with the currently loaded scan results (source), the results of the comparison will display as shown in Figure 1-17.

Figure 1-17: Scan Results Comparison Showing Python Process Difference

Server Name	Name	Attribute	Source Value	Target Value	Difference	% Diff...
aix43-ppc.build.opsware.com	inetd [F740044_51...	Process Listener	127.0.0.1 : / (...			
aix43-ppc.build.opsware.com	inetd [F740044_51...	Process Listener	192.168.50.1...			
aix43-ppc.build.opsware.com	rpc.statd [F74004...	Process Listener	192.168.50.1...			
aix43-ppc.build.opsware.com	inetd [F740044_51...	Process Listener		127.0.0....		
aix43-ppc.build.opsware.com	sshd [F740044_90...	Process Listener		192.168....		
aix43-ppc.build.opsware.com	python [F740044_...	Existence	Missing	Found		
aix43-ppc.build.opsware.com	rpc.statd [F74004...	Process Listener	127.0.0.1 : 95...			
aix43-ppc.build.opsware.com	inetd [F740044_51...	Process Listener		192.168....		
aix43-ppc.build.opsware.com	inetd [F740044_51...	Process Listener		192.168....		
aix43-ppc.build.opsware.com	inetd [F740044_51...	Process Listener		192.168....		
aix43-ppc.build.opsware.com	inetd [F740044_51...	Process Listener		127.0.0....		

pdizzle | Mon Dec 04 14:32:26 2006 America/Dawson

The results of the selected row show that on the AIX server, a python process was running during the target scan results (the earlier saved scan results), but that it was not running on the currently loaded scan results.

Object Attribute Difference

The VAM compare feature also evaluates two scan results to determine any differences in the value of an object attribute. If the same attribute does not match between the two sets of scan results, then it will be marked as a “difference” and displayed in the comparison results. For attribute values with numerical differences, the comparison results will display both the numerical difference and percentage of change.

If you scan a server in VAM and the server shows that it has 2 gigabytes of RAM, and then at a later point in time, one gigabyte of RAM is removed from the server, then when you perform a comparison, the results will show the server's total memory as having a difference of one gigabyte. The results will also indicate that the target (the earlier saved scan results) had a value of two gigabytes, and the source (currently loaded scan results) has a value of one gigabyte.

Table 1-3 lists all object attributes evaluated during a scan results comparison.

Table 1-3: Object Attributes Checked for Difference in a VAM Scan Results Comparison

OBJECT CATEGORY	OBJECT ATTRIBUTES COMPARED
Server	boot time, OS, kernel, host name, codeset, total memory, total swap
Interface	MAC address, IP address, broadcast address, subnet mask, switch port (if NAS enabled), VLAN, configured speed, negotiated speed, configured duplex, negotiated duplex
File system	size, device
Process family	user (on any member process), listener (on any member process)

“Interesting” Object Attribute Differences

VAM will also compare a set of attributes by using special heuristics specific to certain attributes, so that differences VAM considers “interesting” will be shown.

If an attribute value in one of the scan results exceeds a minimum (or maximum) threshold and its value changes by at least a certain percentage between the scan results being compared, then VAM will present this in the comparison results.

Table 1-4 shows special object attribute differences.

Table 1-4: “Interesting” Object Attribute Differences Calculated by VAM Heuristics

OBJECT	OBJECT ATTRIBUTES COMPARED
Server	Load average, percentage of free memory on a server, percentage of free swap memory
Server's file system	Percentage of free space
Process families	Number of open files, total number of all related connections, total count of process family member connections.

For more information on the heuristics used to calculate what is considered a significant difference, see “Scan Results Comparison Heuristics” on page 100.

Comparing Two Sets of Scan Results

In order to compare two sets of scan results, you must have at least one saved scan result. If you select **Save** from the **File** menu, this will save the currently loaded scan results. If you click **Refresh Scan Results** on the VAM toolbar, and then save again, this will create and save a new set of scan results.

To compare scan results in VAM, perform the following steps:


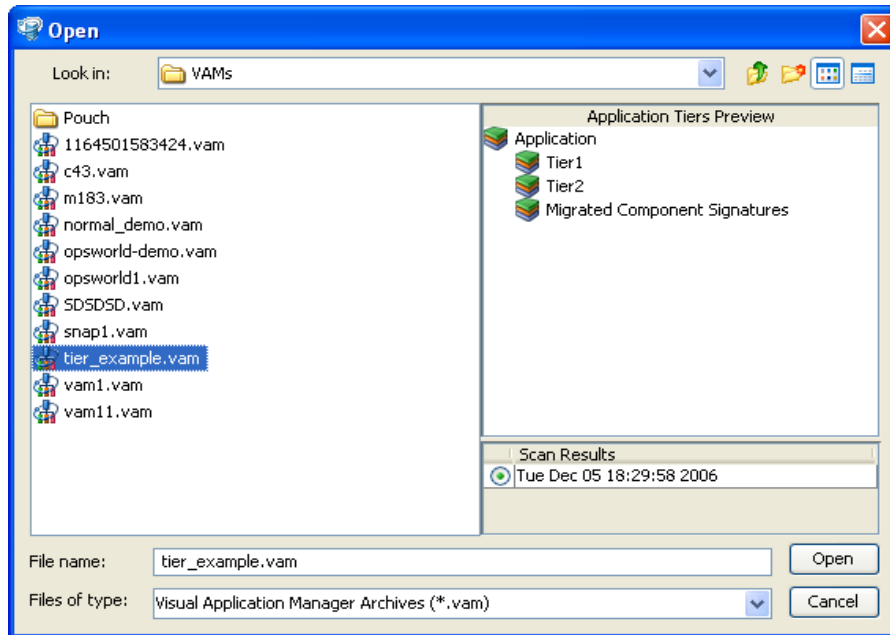
- 1 From the **Edit** menu, select **Compare**. (Or, click the Compare  **Compare** button on the VAM toolbar). The Compare pane appears at the bottom of the VAM application window.
- 2 From the Target Scan Results drop down list, select a scan results item from the currently loaded .vam file, and then click **Compare**.
- 3 If you want to select scan results from a different .vam file, from the Target Scan Results drop down list, select More. The Open window, as shown in Figure 1-18, allows you to browse the current or another .vam file and all the scan results contained in the selected .vam file.

Figure 1-18: Scan Results Selected from a .vam File



- 4 After you select a .vam file, click **Select**.

- 5** From the Target Scan Results drop down list, select scan results to compare with the currently loaded scan results.
- 6** Click **Compare**.
- 7** In the Comparison Results pane, from the Display drop down list, select Server, Interface, File System, or Process Family.
- 8** To close the Compare pane, from the Edit menu, select **Compare** again.

Filtering VAM Data

From VAM, you can filter data that was collected during a VAM scan, and from the results, launch other Opsware SAS Client features, such as Server Explorer, Network Device Explorer, Global Shell, and Remote Terminal.

You can use these features to stop processes, to start and stop services, to perform other troubleshooting or automation tasks, and to help with troubleshooting and remediation tasks.

Data Filtering

You can filter the data that VAM has collected to find the data most interesting to you, such as specific information on VAM in the following object categories:

- **Server:** Name, Opsware ID, Opsware agent version, total memory, users logged in at the time of the scan, Compliance status, and so on.
- **File Systems:** Mount Point, device name, percentage used, free space, and so on.
- **Networks Interfaces:** Device, MAC address, IP address, and so on.
- **Process Families:** Name, command line argument, number of open files, and so on.
- **Network Ports:** Name, MAC address, speed, duplex, and so on.

For more information on filtering criteria and regular expressions, see “Filter Criteria” on page 95.

To filter data that was collected in the currently loaded VAM scan results, perform the following steps:

- 1** From the **Edit** menu, select **Filter**. (Or, from the VAM toolbar, click the



button.) The Filter pane appears at the bottom of the VAM window.

- 2** In the Filter pane, from the drop-down list near the top of the pane, select the type of object you want to filter for, such as servers, file systems, network interfaces, process families, and network ports.
- 3** Enter criteria in any field to narrow the filter results. See “Filter Criteria,” and “Examples of Regular Expressions” on page 96.

If you are looking for servers and would like to filter according to Compliance, keep in mind the following rules:

- If you choose Compliance = None, then the filter will ignore all compliance information.
- If you choose Any, the filter will look for any compliance policy on all servers that match the selected compliance status.

- 4** Click **Filter** to display the Filter results. A new pane opens at the bottom of the VAM window called Filter Results.
- 5** From any of the results related to server, you can select the server on which the results were found, right-click, and select **Open Remote Terminal** or **Open Device Explorer** to browse the server.
- 6** (Optional) To filter within the results, perform the following steps:
 1. Select the Filter pane after you have performed a filter operation.
 2. Enter criteria in any field to narrow the results.
 3. Click **Filter In Result**. The Filter Results pane is refreshed and displays information based on the criteria you entered.
- 7** (Optional) To navigate within the results, perform the following steps:
 1. From the View As drop-down list, select an object type to find objects that are related to the current results. See “How Filter Results are Related” on page 96.
 2. Click the directional arrows to view previous results.

Filter Criteria

In the standalone text boxes in the Filter pane, enter Perl 5 compatible regular expressions as filtering criteria. You can filter by using standard text matching and also by adding any regular expression patterns.

By default, all filtering is performed in case-sensitive mode. To filter in a case-insensitive mode, select the corresponding check box for the filter criteria.

In paired text boxes ([_] to [_]), enter numbers only to filter within a numeric range. If either the beginning or ending number in the range is left blank, then that part of the range will not be included in the filter. For example, to find all file systems that are almost full, you would filter for file systems that are using at least 80% of their capacity. Enter 80 to specify the beginning of the range and leave the end of the range blank, such as [80] to [_]. To filter for an exact numerical match, enter the same number in the beginning and ending range positions, such as [80] to [80].

The units of measure for filtered items should match what is shown in the Filter Results and Property Page, such as:

- **Memory:** Bytes
- **Uptime:** Days
- **Percentages:** A number from 0 to 100, (such as disk space used and CPU utilization)
- **Disk space:** Bytes

Examples of Regular Expressions

The following examples show how to use regular expressions in filter text boxes:

- **Operating System:** To find all servers that are not running a Windows operating system, look for servers whose operating system does not begin with an "M" (for Microsoft Windows). For example, enter `^[^M]` in this text box.
- **Kernel:** To find servers whose kernel is one of 2.6.5, 2.6.6 or 2.6.7, enter `2.6.[5-7]` in this text box.
- **Mount Point:** To find all mounted Unix file systems other than `/`, enter `/.+` in this text box.

How Filter Results are Related

To find objects that are related to the current filter results, the VAM filter feature allows you to filter results related to the original filter by browsing other object categories.

For example, to find all servers that start with the letter "m", in the Filter pane, select Servers from the drop down list, and in the Name field enter the following expression, `m.*`

Then, click **Filter**. The filter results return all servers that start with the letter "m".

To find data related to these results, from the View As drop down list at the top of the Filter Results pane, select File System. All the file systems on servers whose names start with “m” will display. Select Network Interfaces from the View As drop down list, and all network interfaces on servers that start with “m” will display.

To further filter results, you can select from the following categories:

- If you filtered for **Servers**, then you can view the following related information from the drop-down list:
 - **Network Interfaces**: Displays all interfaces on the listed servers.
 - **File Systems**: Displays the file systems in the listed servers.
 - **Process Families**: Displays the process families running on the listed servers.
 - **Network Ports**: Displays the network ports connected to interfaces on the listed servers.
- If you filtered for **File Systems**, then you can view the following related information from the drop-down list:
 - **Servers**: Displays the servers that contain the listed file systems.
 - **Network Interfaces**: Displays the interfaces in the servers that contain the listed file systems.
 - **Process Families**: Displays the process families (on the server that contains the file system) that have open files in the listed file systems.
 - **Network Ports**: Displays the network ports connected to interfaces that are in the servers that contain the listed file systems.
- If you filtered for **Network Interface**, then you can view the following related information from the drop-down list:
 - **Servers**: Displays the servers that contain the listed interfaces.
 - **File Systems**: Displays the file systems in the servers that contain the listed interfaces.
 - **Process Families**: Displays the process families (on the server that contains the interface) that have open connections through the listed interfaces.
 - **Network Port**: Displays the network ports connected to the listed interfaces.
- If you filtered for **Process Families**, then you can view the following related information from the drop-down list:




- **Servers:** Displays the servers on which the listed process families are running.
- **Network Interface:** Displays the interfaces through which the listed process families have open connections.
- **File Systems:** Displays the file systems on which the listed process families have open files.

Network Ports: Displays the network ports that are connected to interfaces through which the listed process families have open connections.

- If you filtered for **Network Ports**, then you can view the following related information from the drop-down list:
 - **Servers:** Displays servers that have interfaces connected to the listed network ports.
 - **Network Interfaces:** Displays interfaces that are connected to the listed network ports.
 - **File Systems:** Displays all file systems on servers that have interfaces connected to the listed network ports.
 - **Process Families:** Displays all process families that have open connections on interfaces that are connected to the listed network ports.

Error Messages

VAM indicates when an error occurred on a managed server by displaying the following server icons when you move your mouse pointer over the icon:

- **Server Error Icon** : There was an error in gathering information from the server when VAM scanned it (see Table 1-5 for possible causes for the error).
- **Server Unreachable Error Icon** : The Opware core was not able to communicate with the Opware Agent installed on the server.
- **Server Unknown** : VAM is unable to scan the server at all, possibly because the server is no longer in the core and under Opware management.

It will show these icons before the server name in the Device Tree, Network Map, Virtualization Map, and Server Map. You can move your cursor over the server name to display the detailed error message.



Scan failures and scan time-outs typically occur when the Opware managed server is very busy, or when network traffic is very heavy or running over a low bandwidth connection. If these types of errors occur too frequently, please contact your Opware administrator for assistance.

Table 1-5 describes these errors and recommended actions.

Table 1-5: Error Messages in VAM

ERROR	DESCRIPTION	ACTION
Not Enough Disk Space	A selected managed server does not have enough disk space to perform a scan.	Free up disk space.
Scan Timed Out	The scan process has exceeded the time-out limit.	See “Scan Time-Out Preference” on page 75.
Server Access Denied	By using the OGFS, you are unable to access the server's file system as root (on a Unix server) or as LocalSystem (on a Windows server).	Contact your Opware administrator for the required permissions.
Server Capture Failed	The remote capture of data or the transfer of data back to the Opware core failed.	Review the log file that is in /tmp/.sitemap/<number> for details in your global shell session.
Server ID Invalid	The server's directory was not found in the OGFS, which means that Opware SAS does not know the server exists.	
Server Scan Agent Failed	The driver used to collect data could not be correctly copied to the managed server. This could be caused by a checksum mismatch.	Contact Opware Support and provide the log file.

Table 1-5: Error Messages in VAM (continued)

ERROR	DESCRIPTION	ACTION
Server Unreachable	The managed server is unreachable by Opware SAS. This could be caused if the Opware core cannot communicate with the server's Opware Agent.	Try again later. If this condition persists, contact your Opware administrator.
Unknown Scan Error	An unknown error occurred during the scanning process.	Try again later. If this condition persists, contact your Opware administrator.
Unsupported Agent for Scan	The VAM does not support the Opware Agent version running on a selected managed server.	Opware Agent 5.1 or higher is required.
Unsupported OS for Scan	The VAM does not support the operating system running on a selected managed server.	See "Supported Operating Systems" on page 32.

Scan Results Comparison Heuristics

When you compare scan results in VAM, specific objects and their attributes are evaluated between each scan result and any differences (and non-existence of objects) will be displayed in the comparison results. (For information on the basic set of objects and attributes compared in a scan results comparison, see "Comparing Scan Results" on page 88.)

In addition to the basic set of attributes evaluated in a comparison, VAM also applies a certain set of heuristics to some attributes in order to discover unique differences that VAM has determined to be interesting or useful.

Specifically, if an attribute value in one of the scan results exceeds a minimum (or maximum) threshold and its value changes by at least a certain percentage between the scan results, then VAM will present this in the comparison results.

The following special object attribute differences are compared:

- **Server:** Load average, percentage of free memory on a server, percentage of free swap memory.

- **Server's File System:** Percentage of free space.
- **Process Families:** Number of open files, total number of all related connections, total count of process family member connections.

The heuristics applied to certain attributes in scan results during a comparison are listed in Table 1-6.

The following variables are used in the expressions:

- X = The maximum value of the attribute between the two scan results.
- N = The minimum value of the attribute between the two scan results.
- P = The percentage change in value of the attribute between scan results.

Table 1-6: Scan Results Comparison Heuristics

OBJECT ATTRIBUTE	EQUATION
server – 15 minutes load average	$X > 0.8 * \text{cpu count AND } P > 20\%$ OR $X > \max(1, 0.25 * \text{cpu count}) \text{ AND } P > 100\%$
server – percentage memory free (%)	$N < 0.1 * \text{total mem AND } P > 25\%$
filesystem – percentage free (%)	$N < 0.2 * \text{size AND } P > 10\%$
process family – open file count (on any member process)	$X > 50 \text{ AND } P > 50\%$
process family – connection count (aggregate across all member processes)	$X > 50 * \text{process count AND } P > 30\%$
process family - connection count (on any member process)	$X > 50 \text{ AND } P > 50\%$

Chapter 2: Audit and Remediation

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Opsware Audit and Remediation
- Understanding Audits
- Creating an Audit
- Viewing Server Audit and Snapshot Usage
- Configuring an Audit
- Opsware Audit and Remediation Rules
- Configuring Opsware Audit and Remediation Rules
- Audit Rule Exceptions
- Audit Policies
- Running an Audit
- Scheduling an Audit
- Viewing and Remediating Audit Results
- Understanding Snapshots
- Creating a Snapshot Specification
- Configuring a Snapshot Specification
- Scheduling Snapshot Jobs
- Locating Snapshots
- Viewing Snapshot Contents
- Copying Objects from a Snapshot to a Server

Overview of Opware Audit and Remediation

Opware Audit and Remediation allows you to define server configuration policies and ensure that servers in your facilities meet those policy standards. When servers are found to be out of compliance (not configured the way you want them to be), you can remediate the differing server configurations.

With Audit and Remediation, you can audit server configuration values based on a live server (or server snapshot), or based on your own custom values (or both). Opware Audit and Remediation also allows you to take snapshots of a server to capture the current state of a system, so you can perform server comparisons against a baseline, or use the snapshot inside of an audit. You can also create custom audit policies that define company or industry server configuration compliance standards, which can be used inside of audits or snapshot specifications.

If you subscribe to The Opware Network (TON), you can be kept up to date on the latest industry compliance standards based on the needs of your data center. For example, subscribing to TON Essential Content gives you access to regularly updated security best practices, such as the Center for Internet Security (CIS), NSA, and so on, as well as the Opware patch supplement for Microsoft Windows. The TON Subscription Service enables you to access the most current regulatory compliance policies (FISMA, Sarbanes-Oxley, etc.) and daily vulnerability alerts. You can also join the TON content developer communities to share and access custom-created audit policies and rules. And much more.

For information about subscribing to TON, contact your Opware sales representative.



Audit and Remediation Examples

The following examples illustrate ways the Opware Audit and Remediation feature helps you manage server configurations in your facility:

- Capturing Golden Server Configurations
- Enforcing Security Policies
- Creating Your Own Ad-Hoc Audit

Capturing Golden Server Configurations

Sometimes a server becomes configured in such a way that it represents the ideal state of server configuration for some purpose in your facility. For example, if you want to set up a collection of servers that handle web traffic, you might configure a single server that represents a perfect configuration – a golden server configuration – for a group of Web servers. After you configure this golden server, you can duplicate the golden server configuration across a group of servers.

For example, you have a Red Hat Linux server with a unique configuration of Apache Web Servers, and you want to duplicate this exact configuration across several other servers. With Opsware Audit and Remediation, you can create an audit that uses the golden server as the source. In the audit, you select those configurations to use to audit other servers, such as an application policy and specific application configuration rules.

Then, select those servers as the target of the audit to be configured like the golden server. After you run the audit, you can remediate any target server's configurations that do not match the golden source. Then, you can schedule the audit to run on a regular basis, so if any of the servers become non-compliant, you can remediate them when they deviate from the golden standard.

Enforcing Security Policies

Your IT organization likely has security policies you want to enforce, to make sure servers are configured properly and are safe from security attacks. Your organization can use Opsware Audit and Remediation to enforce these policies.

For example, your organization wants to ensure that a collection of Windows 2003 servers has a recent Microsoft security patch, regardless of the applications installed on the servers. By being subscribed to The Opsware Network (TON), you can access the rules for this patch in an audit policy to define this security configuration. This policy would let systems administrators who directly configure and manage those servers know that this policy exists. You can create an audit and link it to the audit policy that contains the patch, and then set the Windows 2003 servers as targets of the audit. The audit can be scheduled to run regularly. If the audit results show that any of the target servers do not contain the new security patch, those servers can be remediated to have the patch.

installed. If new patches come out and need to be installed on the target servers, you can update the audit policy with the new patch, and the audit that runs against the target servers is automatically updated to reflect the new patch definition.

Creating Your Own Ad-Hoc Audit

As a system administrator, you might monitor a class of servers that run a home grown application built by your team, such as a database server or middle ware application. As you spend time configuring and monitoring the servers that run the application, you keep a list that tracks the ideal state of the configuration. Such a list might include, file, disk, partition permissions, application configuration definitions, unique registry permissions, HBA card configurations, RAID levels and configurations, and so on.

You can create an audit that defines these configurations, audit the servers after the application gets installed, and the audit results would confirm whether or not the application installed and has been installed and configured successfully according to your criteria. If something goes wrong, you can create an ad hoc audit to troubleshoot something related to the problem. When the audit results indicate the error, you can remediate the server to match your ideal configuration. To ensure that the configuration change actually works in production, you can schedule the audit to run hourly, or daily, and have an email sent when the results are finished. If this configuration proves to work well, you can save the audit as an audit policy and it can be used by others on your team.

Audits

An audit is the tool you use to define the desired configuration values for a server, compare expected configurations against live servers, and remediate any differences found by the audit. Using audit rules, you can define the audit to look for such configurations as IIS Metabase, Windows Services, file system checks, hardware configurations, application configurations, event logging, COM+, and so on. You can define what the audit should look for, what values you expect to find on the server, and what value to use to fix when differences are found.

For more information on audits, see “Understanding Audits” on page 110.

Audit Policies

An audit policy is used to define rules for checking the configuration of a server and can be reused by other people in your organization. An audit policy contains a set of ideal server configuration rules that help define compliance best practices for others to use for running audits. Audit policies can be linked to audits or snapshot specifications, which maintain the latest changes made to the audit policy.

For more information on audit policies, see “Audit Policies” on page 165.

Audits and the Compliance Dashboard

The Compliance Dashboard allows users to view the overall compliance levels for servers in their facility and helps them remediate compliance problems. The Compliance Dashboard also displays the following types of compliance audits:

- **All Scheduled Audits:** By default, a rollup of scheduled audits will appear in a single column named Audit in the Compliance Dashboard. This status enables you to view, at a glance, the compliance status of all audits that you have scheduled to run on a regular basis. You will only see the Compliance status for those audits that have been scheduled on servers that your user has access to. Any servers that you do not have access to will not be represented in the Compliance Dashboard in the audit roll up.
- **Individual Audits:** Individually scheduled audits can be displayed on a per-audit basis. These audits will not appear by default and must be activated to display in the Compliance Dashboard. You must have access to view the server (facility, customer, or group) where the audit is running in order to see its compliance displayed in the Compliance Dashboard. Servers you do not have access to will not be included in the audit compliance category.

For more information, see “Compliance Dashboard” on page 205

Audit Reports

The SAS Client enables you to generate the several types of Opsware Audit and Remediation reports. For more information on how to run and view reports in the SAS Client, see “SAS Client Reports” on page 225.

Snapshots

Snapshots differ from audits in that snapshots allow you to take a picture of the current state of configuration of a server. Snapshots are useful for capturing the configuration of a golden or baseline server that you would like to compare against other servers in your facility. You can use the snapshot as the source of an audit if any servers do not match the configuration captured in the snapshot, then you can remediate those servers after the audit has run from the Audit Results window.

For more information on snapshots, see “Understanding Snapshots” on page 182.

Terms and Concepts

The following list defines key Audit and Remediation terms and concepts:

- **Audit:** A set of rules that expresses the desired state of a managed server's configuration objects – for example, a server's file system directory structure or files, a server's Windows Registry, application configuration, and so on. An audit's rules can be linked to an audit policy. An audit can be run to compare server configuration object values against a baseline server, a server snapshot, or user-defined values, to determine how values differ. When an audit reveals a difference between servers or user-entered values, the user can install software and server objects to remediate the variance.
- **Audit Job:** The process that occurs when you run an audit. An audit job can be run immediately onetime, or on a recurring basis by scheduling the job. When an audit job is finished, it produces an Audit Result.
- **Audit rule types:** An audit can contain both types of the following rules:
 - **Server comparison:** An audit that compares a server's or snapshot's configurations of a server with other servers or snapshots.
 - **Value-based (user-specified):** An audit that compares one or more servers against a set of user-defined values. This type of audit includes an audit that links to an audit policy.
- **Audit policy:** A collection of rules that defines a desired state of configuration for a server. A policy can be used by an audit in the following ways:
 - **Link:** A linked policy maintains a persistent connection between the audit and the policy. This means that the rules in the audit are exactly those of the audit policy, and if any updates are made to the policy, then the latest changes are also reflected in the audit to which the policy is linked.
 - **Import (replace, non-linked):** When a user imports a policy into an audit, then the connection between the audit and the audit policy is no longer maintained, and the user can make changes to the audit without affecting the policy. Conversely, any changes or updates made to the policy will not be reflected in the audit.
 - **Import (merge):** When an audit policy is imported and merged into an audit, the audit policy's rules are added to the rules already present in the audit. No persistent link between the audit and the audit policy is maintained. During the merge, if rules are found to conflict, the newly imported rules from the audit policy

will replace the rules in the audit policy.

- **Audit Result:** The results of running an audit. This shows how a target server or a group of servers' configuration object values match or mismatch the values as defined in the audit.
- **Compliance:** Denotes the degree to which a server object conforms to a test. Compliance in Opsware Audit and Remediation is defined by the audit's or snapshot's rules, which specify the values expected of the target servers. If the values are different than specified, then the server is out of compliance.
- **Policy Setter:** A person in an organization who is responsible for defining server configuration compliance standards – the way a server should be configured – and who defines audit policies.
- **Rule:** A check on a particular server configuration object along with a desired value, and optional remediation value. Rules come in two types: server-based, which derive directly from a source server, and user-defined, which are created by a user.

If you are subscribed to The Opsware Network (TON), you can access pre-created rules that define a wide range of industry compliance standards, such as the latest patch supplement for Microsoft Windows, current regulatory compliance policies (for example, FISMA, Sarbanes-Oxley), user-created rules from the TON developer community, daily vulnerability content updates, and so on.

- **Server Object:** An object from a server to which an audit or snapshot specification rule can be applied. This can be a value (such as minimum password length) or an object, such as a file or directory, registry entry, Windows Services hardware configuration, and so on. For more information on servers objects used in audits and snapshot specifications, see “Server Objects Used in Audits and Snapshots” on page 122.
- **Snapshot:** Shows a picture of how an Opsware-managed server is configured at a certain point in time. A snapshot is the result of a snapshot specification job that has been run.
- **Snapshot Specification Job:** The process that occurs when you run a snapshot specification. A snapshot job can be run once, or on a recurring basis by scheduling the job. When a snapshot specification job is completed, it produces a snapshot.
- **Snapshot Specification:** An object window that allows you to define and create a snapshot. In other words, you can define the rules and servers to take a snapshot of.

- **Target:** The server or servers that you run an audit against or take a snapshot of. The target for an audit can be a server, several servers, a group of servers, or a snapshot. The target for a snapshot can also be other servers.

Understanding Audits

An audit consists of a collection of rules that enable you to define what should be or what should not be for a server's configuration. An audit contains rules, a source, target servers, and a schedule that defines when and how often the audit will run.

Audit rules allow you to define and check the state of various objects on a server, such as the state of server's file system, registry settings, installed patches and packages, events, software, application configurations, operating system settings, and so on. If the configuration of the object on the target server is different than the state you defined in the audit, you can remediate the object configuration to make sure the target server's configuration is in compliance with the desired configuration.

You can audit server configuration values for a single server, groups of servers, or another server snapshot. You can also schedule audits to run immediately, or on a recurring schedule, and send email notifications when the audit has finished.

Audit Comparison Types

In general, an audit can contain the two following types of comparisons, based on the source of the audit:

- **Server Comparison:** An audit based on configuration values from a source server or source snapshot specified at the time the audit is created. The source server or server snapshot is also known as a "golden" or reference server. For example, you might want to compare file directories or file contents, registry structures, IIS Metabase entries, or user group settings among servers. Using a snapshot as the source of an audit, you can compare the snapshot with other servers in your facility.
- **User-Defined Value Comparison:** An audit based on custom, user-defined values for each server object (file system, windows services, IIS Metabase, users and groups, and so on). These values can be derived from a source server, or from Opsware attributes or custom attributes. This type of audit includes those based on an audit policy. In an audit policy, a user (known as a "policy setter") pre-defines values for each configuration object based on company or industry compliance standards.

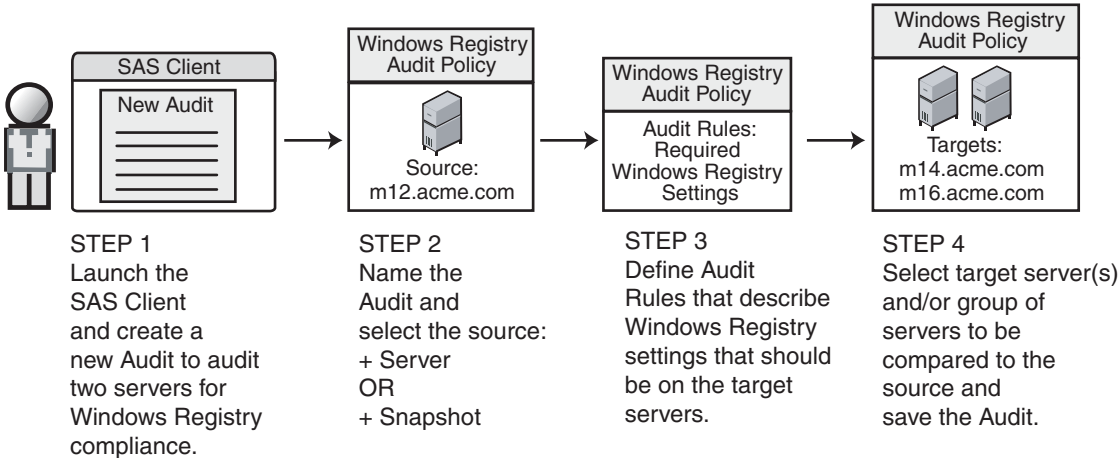
The Auditing Process

The following diagram illustrates a basic example of creating and running an audit.

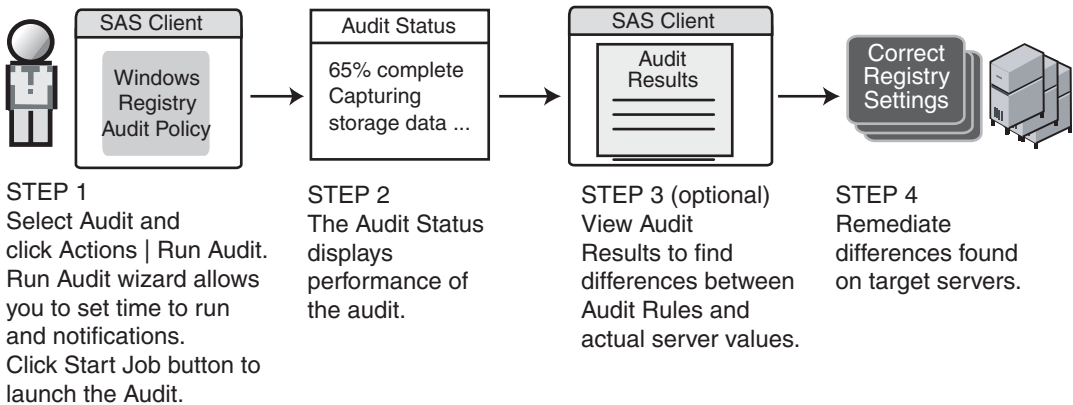
Figure 2-1: The Auditing Process

AUDITING PROCESS

Part A: Create Audit of Windows Registry Settings



Part B: Run Audit and View Results



Audit Elements

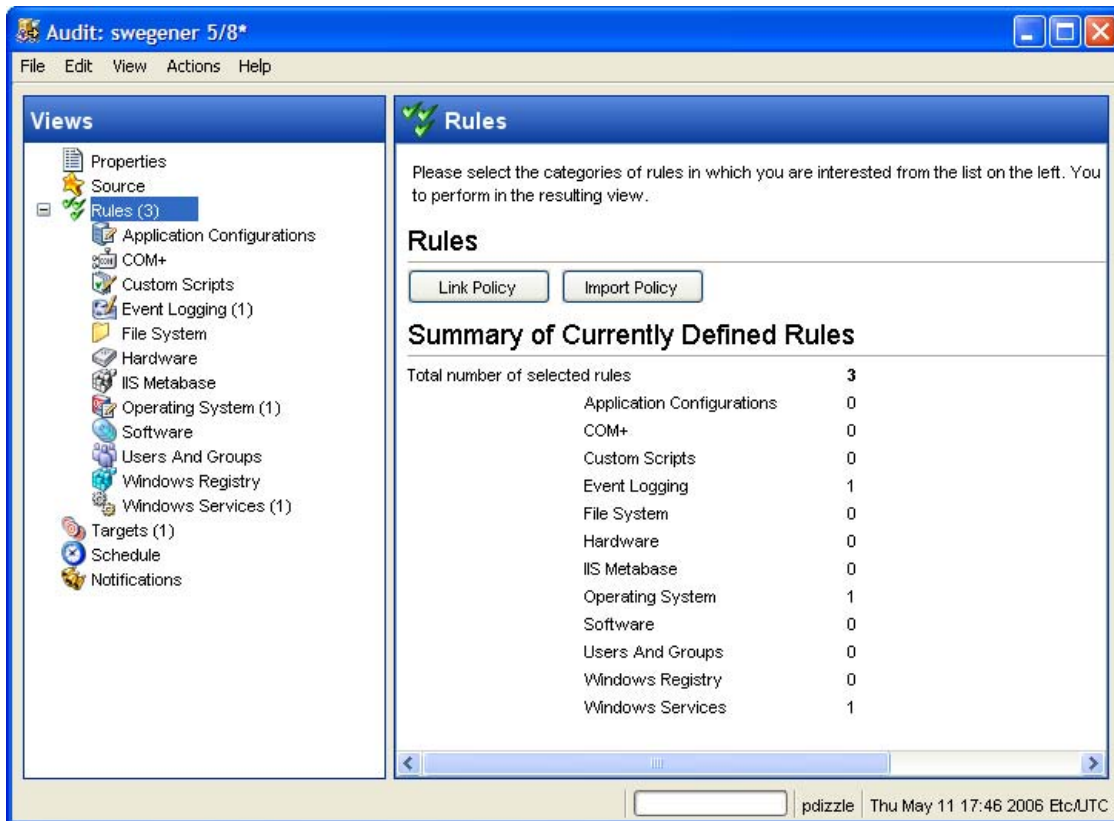
An audit consists of the following elements:

- **Properties:** The name and description of the audit.
- **Source:** The source of an audit can be a server, a snapshot, or no source at all. (However, some rules require a source.) Choosing a server as the source for an audit allows you to select server objects from that server as the basis of your audit. Choosing a snapshot as the source of an audit allows you to use the configuration values of the snapshot. If you choose no source, then you can define only your own custom values for the audit or snapshot.
- **Rules:** A check on a particular server object with a desired value and an optional remediation value. For example, you might check to see if this server contains a specific Windows Service, and if found, determine if the service is turned off. For a description of server objects and rules, see "Server Objects Used in Audits and Snapshots" on page 122.
- **Targets:** The servers that you select to audit. You can choose as many servers and groups of servers as needed for an audit or snapshot.
- **Schedule:** You can run an audit on a onetime basis, or on a recurring schedule. Audits that run on a recurring schedule appear as a single compliance column in the Compliance Dashboard.
- **Notifications:** You can send emails when the audit has finished running, and base the notification on the success, failure, or the completion of an audit job.

To configure an audit, select server configuration objects and then apply rules to those objects in order to define their desired configuration state. F

or example, Figure 2-2 shows an audit that has defined three rules. These rules will determine if any target server matches the rules set for event logging, operating system, and windows services.

Figure 2-2: Audit Window Showing Elements of an Audit



Creating an Audit

You can create an audit from several locations inside the SAS Client. You can choose to audit a specific server by selecting it from the server list, you can audit a group of servers, you can an audit from a snapshot, and so on.

You can create an audit from the following locations inside the SAS Client:

- From the Managed Server list, using the selected server as the source of the audit. You can choose to run the audit on a single server or a group of servers.
- From the Device Groups list, choosing a group of servers as the target at the audit.

- From the Library, by creating a new audit.
- From a snapshot, by creating an audit based on the snapshot.
- From an audit policy, by creating an audit based on the audit policy.

Creating an Audit from a Server

When you create a new audit from a managed server, the audit will use the selected server as the source of the audit. You can choose another server or snapshot for the audit source, if you want, or choose no source at all and define your own custom rules.



To audit a managed server, the server must be reachable and you must have access to the server.

To create an audit from a server, perform the following steps:

- 1** From the Navigation pane, select **Devices** and then select **All Managed Servers**.
- 2** Select a server, and then from the **Actions** menu, select **Create Audit**.

Creating an Audit from a Group of Servers

If you create an audit from a group of servers, then the audit will evaluate all the servers in that group. However, the audit will only evaluate those servers in a group to which you have access.

To audit a group of servers, perform the following steps:

- 1** From the Navigation pane, select **Devices** and then select **Device Groups**.
- 2** In the Navigation pane, browse until you see the group of servers that you want to audit.
- 3** Select the group of servers from inside the Content pane, right-click, and select **Create audit**.
- 4** When you perform an audit by selecting a group of servers, the group of servers becomes the target. If the audit rule requires a source, you must supply one.

Creating an Audit from the Library

To create a new audit and set all your own parameters, from the SAS Client Library, perform the following steps:

- 1** From the Navigation pane, select Library and then select Audit and Remediation.
- 2** In the Navigation pane, select Audits, and then Windows or Unix.
- 3** Right-click inside the Content pane and from the **Actions** menu, select **New**.

Creating an Audit from a Snapshot

You can select any snapshot in the Library and create an audit based on the server configuration captured in the snapshot. The snapshot will serve as the source of the audit, but you can also select another snapshot or server as the source after you create the new audit from the snapshot.

- 1** From the Navigation pane, select Library and then select Audit and Remediation.
- 2** In the Navigation pane, select Snapshots, then Windows or Unix.
- 3** From the Content pane, select a snapshot to create an audit from, right-click, and select **Create Audit**.

Creating an Audit from an Audit Policy

Audit policies are designed to be used by audits. When you create an audit from an audit policy, the audit policy is linked to the audit. So, if any updates are made to the audit policy, those changes are automatically reflected in the audit.

- 1** From the Navigation pane, select Library and then select Audit and Remediation.
- 2** In the Navigation pane, select audits, and then Windows or Unix.
- 3** From the **Actions** menu, select **Create Audit**.

Using Save As for Audit or Snapshot Specification

You can save an audit or snapshot specification using the “Save As” function, which will create a new audit or snapshot specification with a new name. Or, you can choose to save the audit or snapshot specification as an audit policy, which will save only the rules from the audit or snapshot specification and create a new audit policy.



For more information on audit policies, see “Audit Policies” on page 106.

To use save as to create a new audit or snapshot specification from an existing one, or to save an audit or snapshot specification as an audit policy, perform the following steps:

- 1** From inside the Audit or Snapshot Specification window, from the File menu, select Save As.
- 2** In the Save As window, enter a name. If you are renaming an audit or snapshot specification, you must use a unique name,
- 3** (Optional) Enter a description.
- 4** From the Type drop-down list, select either Audit/Snapshot Specification or Audit Policy.
- 5** Click **OK**.

Viewing Server Audit and Snapshot Usage

After you have created and run audits, you can view from the All Managed Servers list or from the Server Explorer, all the audits that are associated with a specific server.

Viewing a Server's Audit and Snapshot Usages from All Managed Servers

To view a server's audit usage from the All Managed Servers list, perform the following steps:

- 1** From the Navigation pane, select Devices and then select All Managed Servers.
- 2** In the Content pane, select a server.
- 3** From the View drop-down list, select Audit and Remediation. Notice that the lower Details pane shows information about audit and snapshot usage.
- 4** In the Details pane, select one of the following options:

- **Audit - Server is Source:** Shows all audits where the selected server is used as the source of the audit.
 - **Audit - Server is Target:** Shows all audits where the selected server is the target of the audit.
 - **Audit Results - Server is Source:** Shows the results of all audits where the selected server was used as the source of the audit.
 - **Audit Results - Server is Target:** Shows the results of all audits where the selected server was used as the target of the audit.
- 5** From any one of these views, you can select an audit or audit results, and perform actions from the Actions menu. For example, you can open an audit, re-run an audit, and so on.

Viewing a Server's Audit Usage from the Server Explorer

To view a server's audit usage from the Server Explorer, perform the following steps:

- 1** From the Navigation pane, select Devices and then select All Managed Servers.
- 2** In the Content pane, select a server, right-click, and select **Open**.
- 3** In the Server Explorer, from the Views pane, select Audit and Remediation.
- 4** In the Content pane, from the Show drop-down list, select one of the following options:
 - **Audit - Server is Source:** Shows all audits where the selected server is used as the source of the audit.
 - **Audit - Server is Target:** Shows all audits where the selected server is the target of the audit.
 - **Audit Results - Server is Source:** Shows the results of all audits where the selected server was used as the source of the audit.
 - **Audit Results - Server is Target:** Shows the results of all audits where the selected server was used as the target of the audit.
- 5** From any one of these views, you can select an audit or audit results, and perform actions from the Actions menu. For example, you can open an audit, re-run an audit, and so on.

Configuring an Audit

Configuring an audit consists of performing the following general steps:

- Name and describe the audit
- Select a source for the audit: a server, a snapshot, or none
- Configure the audit rules
- Choose a target server, group of servers, or snapshot to audit
- Add audit rule exceptions (optional)
- Schedule the audit
- Set the Email Notification (optional)
- Save the audit

To configure an audit, perform the following steps:

- 1** Create the new audit from one of the methods described in “Creating an Audit” on page 113. The Audit window opens.
- 2** Enter the following information for the audit:
 - **Properties:** Enter a name and description for the audit.
 - **Source:** Every audit can use a server or snapshot as its source. (Or, you can choose no source and define your own rules.) If you use a server as the source, you can browse the server for values to define the audit’s rules. If you choose a snapshot, you will be limited to the rules in the snapshot and the snapshot results when you define the audit rules. If you choose no source, you must define your own rules, or choose to link to an audit policy in the rules section. Some rules, however, require a source in order to be defined.
 - **Rules:** Choose a rule category from the list to begin configuring your audit’s rules. Each audit rule is unique and requires its own instructions. For information on how to configure individual audit rules, see “Configuring Opsware Audit and Remediation Rules” on page 124.

If you want to use an audit policy to define the rules of your audit, click either Link Policy or Import Policy. When you link an audit policy, the audit maintains a direct connection with the audit policy. So if any changes are made to the policy, the

audit will update with the new changes. If you import an audit policy, the audit will use all the rules defined in the policy but will not maintain a link to the audit policy. For information about audit policies, see “Audit Policies” on page 165.

- **Targets:** Choose the Targets of the audit. These are servers, groups of servers, or snapshots that you want the configured audit rules to evaluate and compare. To add a server or group of servers, click **Add**. To add a snapshot target, in the Snapshot Targets section, click **Add**.
- **Exceptions:** Click **Add** to add exceptions to the rules in your audit. In the Add Exception window, select a server or multiple servers (or device groups), and then select one or more rules you want to except from the chosen servers. You can except any of the rules in the audit from any of the target servers or snapshots. You can optionally add an explanation, a ticket ID, and an expiration date for the exception.
- **Schedule (Optional):** Choose whether you want to run the audit once, daily, weekly, monthly, or on a custom schedule. Parameters include:
 - **None:** No schedule will be set. If you want to run the audit immediately, or on a onetime basis, you have to select the audit, right-click, and select **Run Audit**.
 - **Daily:** Choose this option to run the audit on a daily basis.
 - **Weekly:** Choose the day of the week that you want the audit to run.
 - **Monthly:** Choose the months that you want the audit run.
 - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

```
0 0 * * 1-5
```

An asterisk (*) in any of these fields represent all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.
- **Time and Duration:** For each type of schedule, specify the hour, minute, day of the week, and month for the schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose an end date, select End. From the calendar selector, choose an end date. The Time Zone is set according to the time zone set in your user profile.

- **Notifications:** Enter email addresses to notify people when the audit job finishes running. You can choose to send the email on both the success and the failure of the audit job (not the success of the audit rules). To add an email address, click Add Notification rule. (This is only relevant if the audit is set to run on a recurring schedule.)

3 When you have finished configuring the audit, from the **File** menu, select **Save**.

Opsware Audit and Remediation Rules

An audit enables you to determine how your servers are configured, and whether or not those servers are configured correctly – that is, as defined in an audit or audit policy. You achieve this goal by creating rules about server objects. You can gather information about the server objects listed in Table 2-1 and either take a picture of their current state – in a snapshot – or define the desired configuration state for these objects – in an audit or an audit policy. (For a list of all server objects you can configure for an audit, see “Server Objects Used in Audits and Snapshots” on page 122.)

In an audit and audit policy, you can also define what, if any, remediation value you would like the object to have. This is used only if a server object is found to be different than the desired state. The remediation value is not implemented automatically, but rather manually after the audit has been run.

An audit rule consists of a server object (file system, IIS Metabase entry, and so on), the specific thing about the object that you want to check (the specific files or directories you want to check), the desired state of the object, and a remediation value should the server configuration differ from the audit rule (optional).

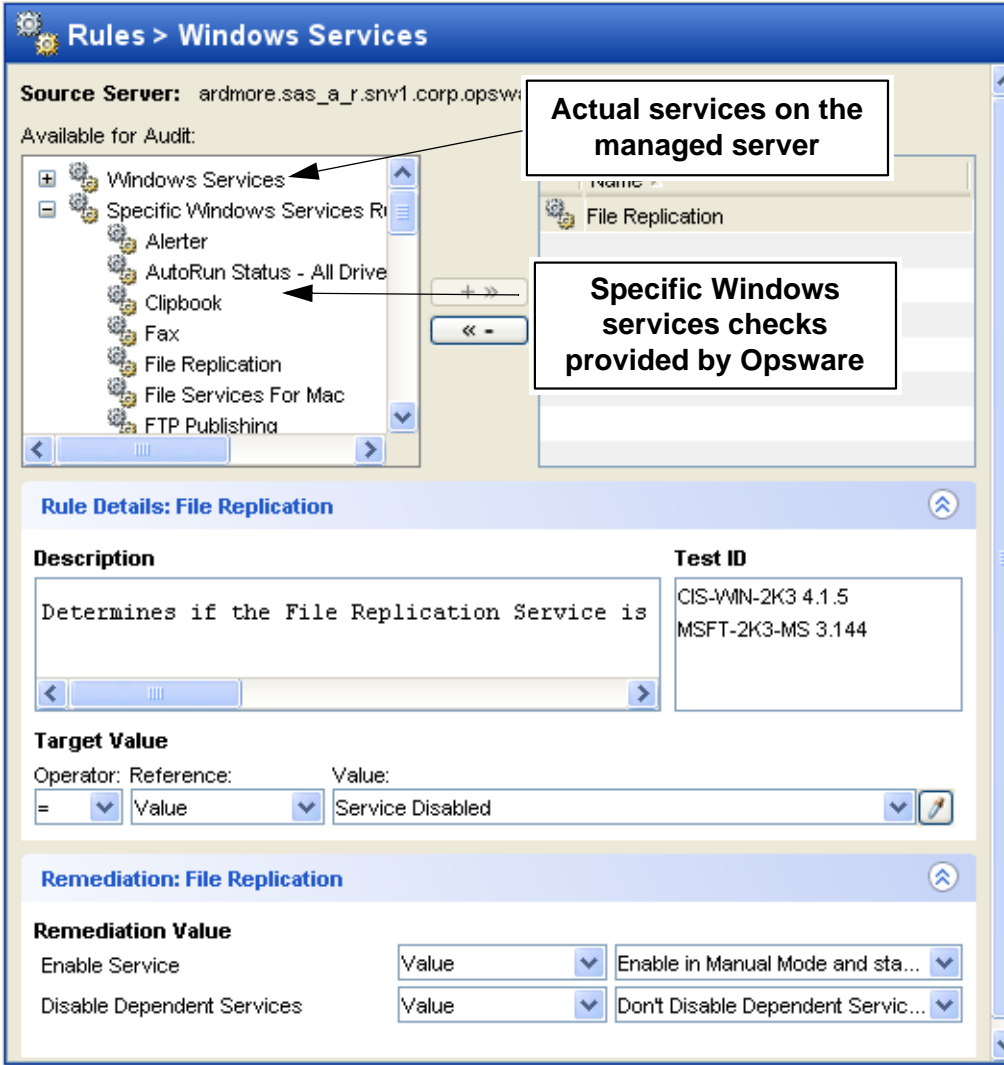
An audit rule consists of the following components:

- **Server Object:** This is a specific server configuration category that an audit can evaluate, such as a server's file system, application configurations, hardware, software (installed patches and packages), Windows Registry entries, and so on. A server object usually consists of several other things that you can check. For example, for the windows services server object, you might want to know if a specific service exists on target servers and whether or not the service is enabled or disabled.
- **Target Value:** This is the specific server configuration object element. For example, you might want to determine if a specific directory exists on a server, or if an application is configured properly, and so on.

- **Remediation Value:** This is the value that you want to change for the server object, if it is found to be different than desired. This value is not implemented automatically; you can make the remediation change after the audit has run.

Figure 2-3 illustrates an audit rule defined for a Windows Service named File Replication.

Figure 2-3: Windows Services Audit Rule



In this example, the audit rule has been configured in the following manner:

- **Available for Audit:** Lists all services from the source server available to be added to the audit, plus specific Windows services rules provided by Opsware.

- **Selected for Audit:** The service name File Replication has been chosen.
- **Description:** Describes what is being checked on the target server. In this case, the audit will check to see if the service is enabled or disabled.
- **Target Value:** This is the value compared against the target server. In this example, the user has set Service Disabled. This means that the audit will check to see if this service is disabled. If the service is in fact enabled, the audit results will indicate the variance, and the configuration would be considered out of compliance with CIS standards.

Depending on the type of check being done on a server, the target value can contain an operator (equals, greater than, and so on), a reference (use from the source of the audit), your own or a preset list of Values (for predefined rules, these values are built in), or a custom attribute that exists on the target server.

- **Remediation Value:** The remediation value determines the action to take if the service on the target server does not match the value you defined in the audit. Remediation values can be derived from a prebuilt or user-entered Value, a Server Attribute on the target server, or a custom attribute that exists on the target server.

Server Objects Used in Audits and Snapshots

Table 2-1 lists all server objects that you can create rules for inside an audit or a snapshot specification. Some server object values are captured and audited live and some objects are captured from the Model Repository.

Table 2-1: Audit and Remediation Server Objects

SERVER OBJECT	DESCRIPTION	CAPTURED LIVE AND/OR FROM MODEL REPOSITORY
Application Configurations	Contents of application configuration files and their values.	Live
COM+	COM+ objects and component categories.	Live
Custom Scripts	Write your own custom scripts to retrieve information from a server and compare contents. For example, you can run a script to gather output from a custom application and evaluate returned output against values set in the audit. (Python 1.5.2 only for python scripts.)	Live

Table 2-1: Audit and Remediation Server Objects (continued)

SERVER OBJECT	DESCRIPTION	CAPTURED LIVE AND/OR FROM MODEL REPOSITORY
Event Logging	Security, application, and system log files.	Live
File System	Contents of files and directories (and subdirectories), user and group access, checksum for files, file modification date, and Windows ACLs (Windows only).	Live
Hardware	CPU, storage devices, and memory.	Model Repository
IIS Metabase	IIS Metabase objects and configuration values to snapshot or audit.	Live
Windows Registry	Select Windows Registry directories or registry key values to capture and compare.	Live
Operating System	Operating system settings such as domain controller settings, numerous network settings (IP Source Routing Protection Level), among others.	Live
Software	Installed packages or patches.	Model Repository
Users and Groups	Compare information about users and groups on servers, such as user name for last login, whether or not CTRL + ALT + DELETE is enabled, and so on.	Live
Windows Services	Select Windows services.	Live
The Opsware Network (TON) Rules	If you are subscribed to TON, you have access to many different types of audit rules (called “pluggable checks”. The exact kind of rules you have access to depend on your subscription, but can include such rules as the latest patch supplements for Microsoft Windows, current regulatory compliance policies (for example, FISMA, Sarbanes-Oxley), user-created rules from the TON developer community, daily updated vulnerability content, and so on.	Live



A Windows COM+ category (folder) that does not have any objects will not be included in a snapshot or audit, even though Opware SAS will display an empty COM+ folder in the Server Explorer.



Opware Audit and Remediation does not support device files or sockets.

Configuring Opware Audit and Remediation Rules

Creating an audit (or snapshot specification) requires configuring Opware Audit and Remediation rules, which define:

- The type of server object to snapshot or audit and compare – objects such as the server's file system, hardware information, application configurations, installed patches or software, and so on.
- Information about that object to audit or snapshot. For example, for a server's file system, you can capture Windows NT file's Access Level Controls. For an application, you can capture the application configuration values you want to snapshot or audit, plus any remediation values to specify if differences are discovered between the rule and the actual value on the target server.

A rule can contain a custom script that seeks to determine if all the passwords stored in a file match a certain character length, or a rule can include a check to determine if a particular Windows Service is running or disabled on a server. For some rules, you can also specify the remediation value for the server object if the value defined in the audit or snapshot differs from the server's value after the audit has run. For example, if a Windows Service is disabled, you can specify that the Remediation value should restart the service.

Remediation values are implemented manually, after the audit has run, from the Audit Results window. For more information on how to remediate audit results, see "Viewing and Remediating Audit Results" on page 175.

Configuration Rules: Expected (Target) and Remediation Values

Some rules are a very simple to configure and define and do not require anything more than selecting the server objects that you want to snapshot or audit. Some rules might check to determine if a value or property exists on a configuration file on a server, without the need for any advanced parameters. For example, Opware Audit and Remediation

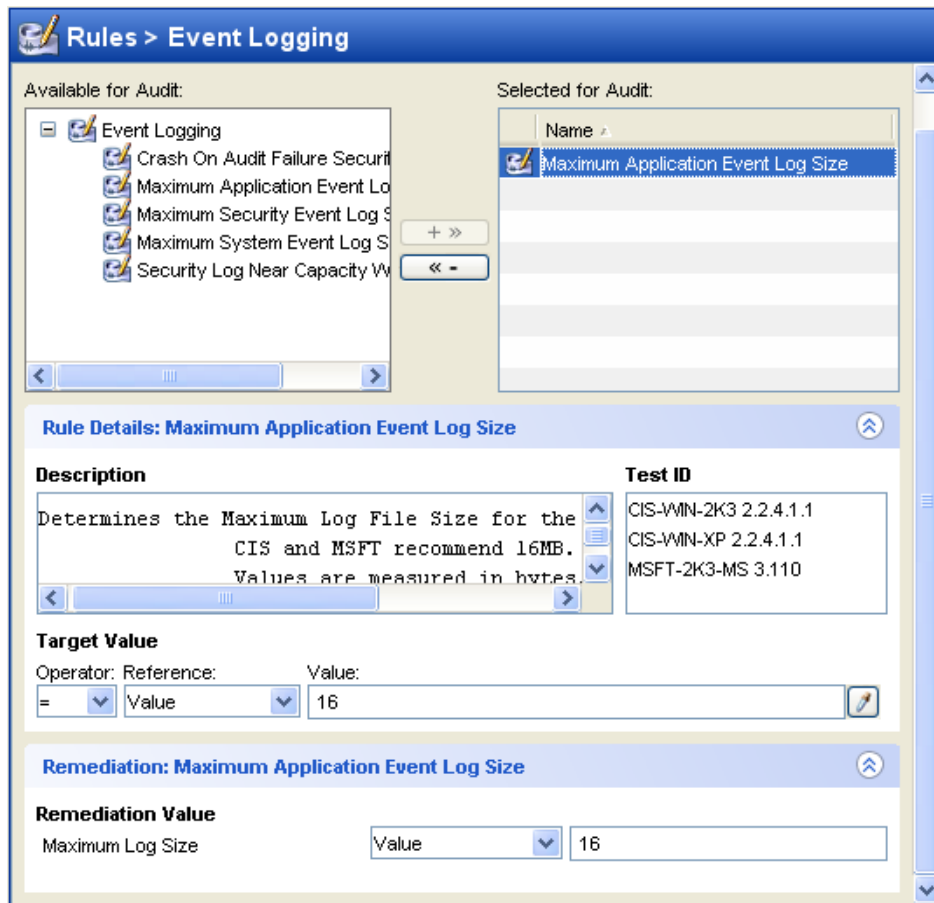
rules for the Software server object evaluate the patches or packages that are installed on the target servers. The Hardware rule allows you to check the CPU, memory, or storage values that exist on target servers. In this case, no extra rule parameters are necessary. Other rules are more complex and require more advanced configuration, such as specifying an expression that looks for a range of values and specifies a remediation that replaces undesired values.

Example Rule: Event Logging

Event Logging requires that operators and reference values (user-entered values, custom attributes from the source server, or server attributes) be defined. For example, you can choose to configure an Event Logging rule that will check the maximum application event log size. CIS standards and Microsoft recommends that this value to 16MB. You can

define the audit rule to determine if this value is no more than 16MB on your target servers. You can also set the remediation value to be 16MB, if the value found on a target server is greater than 16MB.

Figure 2-4: Example Audit Rule for Event Logging Server Object



In the example shown in Figure 2-4, the user has chosen to audit the Event Logging setting of “Maximum Application Event Log Size.” (This audit rule is one of the many predefined rules that come as part of the SAS Client product distribution.)

The top left side of the rules pane, Available for Audit, shows all Event logging objects available from the source server to add to the audit. The top right, Selected for Audit, shows all Event Logging rules that have been selected for the audit.

This rule consists of the following parameters:

- **Rule Details:** Describes this setting and the CIS and Microsoft recommended value, which is 16MB.
- **Target Value:** Allows you to define a target value, which is the value you expect to find on the server.
- **Operator:** Uses an operator to set the expression. Operators include equals (=), less than (<), greater than (>), and so on.
- **Reference:** Choose the source of the script output. You can choose from the following options:
 - **Source:** Takes the value of this setting from the source of the audit, either a server or a snapshot, or from the source of the snapshot specification, a server.
If you choose a server as source for an audit or snapshot specification, then you can select from all the objects available on that server.
If you choose a snapshot as your source for an audit, then you will only be able to select the snapshot rules and snapshot results for the audit. (You can only choose a server as the source for a snapshot specification.)
 - **Value:** Allows you to enter your own value.
 - **Server Attributes:** Common server attributes from the Opsware model.
 - **Custom Attributes:** Derives from the target server. (For the application configuration and custom script rule, if you choose a custom attribute for the rule definition, this custom attribute must also exist on the target servers.)
- **Value:** Either a user-entered value, a server attribute from the Opsware model, or a custom attribute from a target server.
- **Remediation Value:** The value that will replace those found on the target server that do not match the target value specified. The remediation value will not be implemented automatically. Rather, you must manually choose to remediate the value from the Audit Results window after the audit has run.

After you select parameters for the rule, the Value field will show the desired value for the selected configuration file. If the value set in the rule does not match the value on the target of the audit, then you can specify in the Remediate section.

Audit Sources: Server or Snapshot?

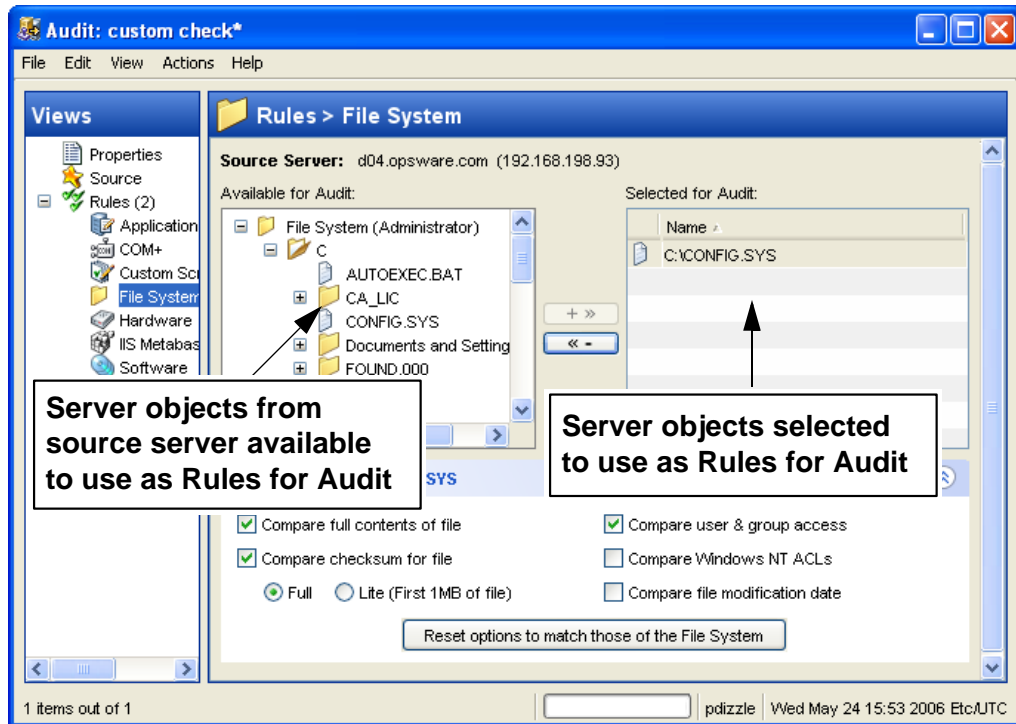
You have two options for choosing a source for an audit or snapshot specification: a server or a snapshot. The source of an audit determines what rules you are able to select from and configure in your audit or snapshot specification. Choosing a source depends on the purpose of your audit or snapshot specification:

- **Server as Source for an Audit or Snapshot Specification:** Choose a server as the source of an audit if you know that specific server contains the desired servers objects that you want to add to the audit or snapshot specification. For example, if you are interested in auditing or taking a snapshot of application configuration files for an Apache Web Server (for example, httpd.conf) on some target servers, choose as the source of your audit – a server that you know has Apache installed on it and that is configured correctly.

Remember that you can choose several different source servers as you build your audit or snapshot specification rules. In fact, you can choose a different source for each server object rule.

When you choose a server as the source for an audit, Figure 2-5 shows what you see in the audit or snapshot specification window's Content pane (right side of window):

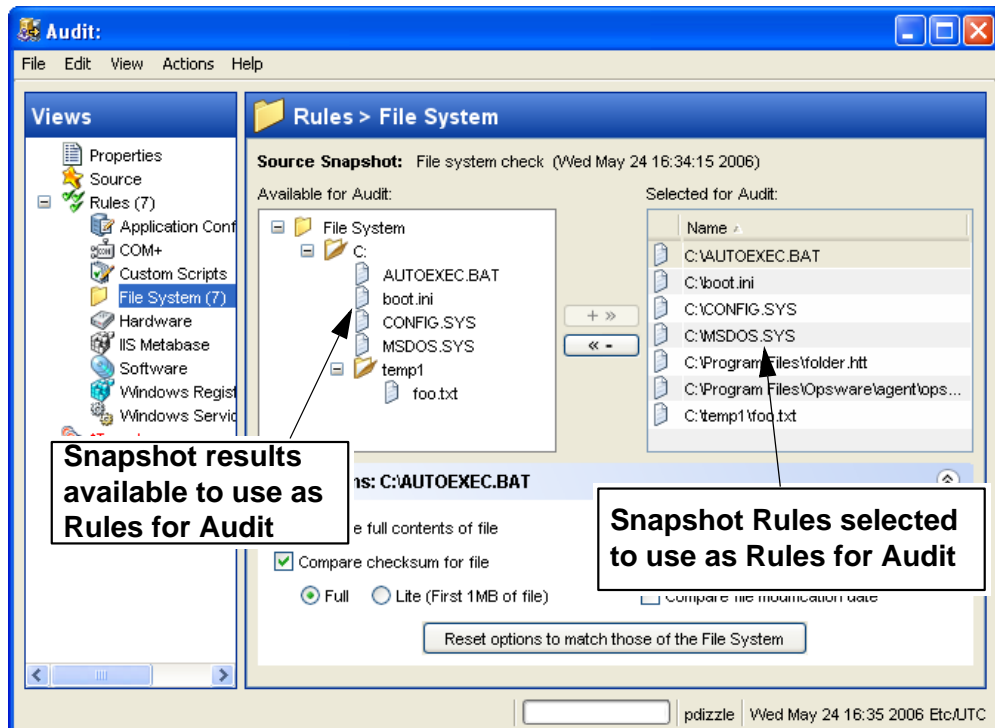
Figure 2-5: Server as Source of Audit: Available Server Objects to Build Audit Rules



- **Snapshot as Source for an Audit or Snapshot Specification:** Choose this option if you have a snapshot of a server that was in a known good state (a “golden” server configuration), and you would like to compare that snapshot with other servers in an audit. Or, choose this option to use the captured server values to take a snapshot of another server. Using a snapshot as the source for an audit or snapshot specification allows you to choose both the results and the rules of the original snapshot specification that the snapshot was based on.

Figure 2-6 displays the choices you have for building audit or snapshot specification rules when you use a snapshot as the source. You can choose from the snapshot's results and the snapshot's rules.

Figure 2-6: Snapshot as Source of Audit: Available Server Objects to Build Audit Rules



Rules That Use a Source Value From Source Server

Most rules require a source in order to define them, except the following rules:

- Any of the prebuilt rules that you do not set the value to derive from Source
- Custom Scripts rules that you do not set the compare value to derive from Source

You cannot save rule without giving a source if the rules specified require a source. You must select a source for all comparison checks and for rules that compare against a source value.

Configuring Specific Rules

For information on rules you can set for each type of server object, see the section for the specific server object that you want to configure a rule for, listed below:

- Configuring Application Configuration Rules
- Configuring COM+ Rules
- Configuring Custom Scripts Rules
- Configuring Event Logging Rules
- Configuring File System Rules
- Configuring Hardware Rules
- Configuring IIS Metabase Rules
- Configuring Operating System Rules
- Configuring Software Rules
- Configuring Users and Groups Rules
- Configuring Windows Registry Rules
- Configuring Windows Services Rules
- Configuring The Opsware Network (TON) Rules



You must have permissions to create and configure Opsware Audit and Remediation rules. To obtain these permissions, contact your Opsware administrator. See the *Opsware[®] SAS Policy Setter's Guide* for more information.

Configuring Application Configuration Rules

The application configuration rule inside an audit, snapshot specification, or audit policy allows you to configure values for application configuration files on a target server. You can define rules you want to capture and the configuration file values that you want to check.

You can choose from a list of predefined application configurations for the configuration file you would like to audit or take a snapshot of. You can also choose from custom application configurations that a user in your organization has created and made available for usage in an audit, snapshot specification, or audit policy.

An application configuration is a template (or collection of templates) that models the information found inside of a configuration file for an application. When you choose an application configuration inside an audit, snapshot specification, or audit policy and click **View**, you will see the contents of the configuration file from the source of the audit. All key-value pairs that you are able to add to the audit rule will display.

The information displayed inside an audit windows depends on the source of the audit or audit policy (or the target for a snapshot specification):

- If you choose a server as the source of the audit or audit policy, then the application configuration values displayed in the audit rule will be those of the configuration file on the server, as filtered through the application configuration template.
- If you choose a snapshot as the source of the audit or audit policy, then you will only be able to modify the values that were captured at the time the snapshot was taken.
- If you do not choose any source, then you will not be able to configure a rule for the application configuration file.
- If you choose to configure an application configuration in a snapshot specification, then the values of the configuration will derive from the target server.



You will only see values of the source configuration file that have been modelled in the application configuration. If the application configuration is customized and has no name-value pair defined (but the value exists in the source configuration file), you will not see it in the audit or audit policy.

After you view the contents of the source application configuration file, you can select values and define target values – the values that the audit should look for on the target server. You can also define remediation values should the audit find a difference.

Creating an Application Configuration Rule

A useful way to understand how to configure an application configuration rule is to look at an example. Let's say you want to create an audit rule for a UNIX hosts file (/tmp/hosts). You know that the UNIX hosts file on a specific server represents the ideal state of the

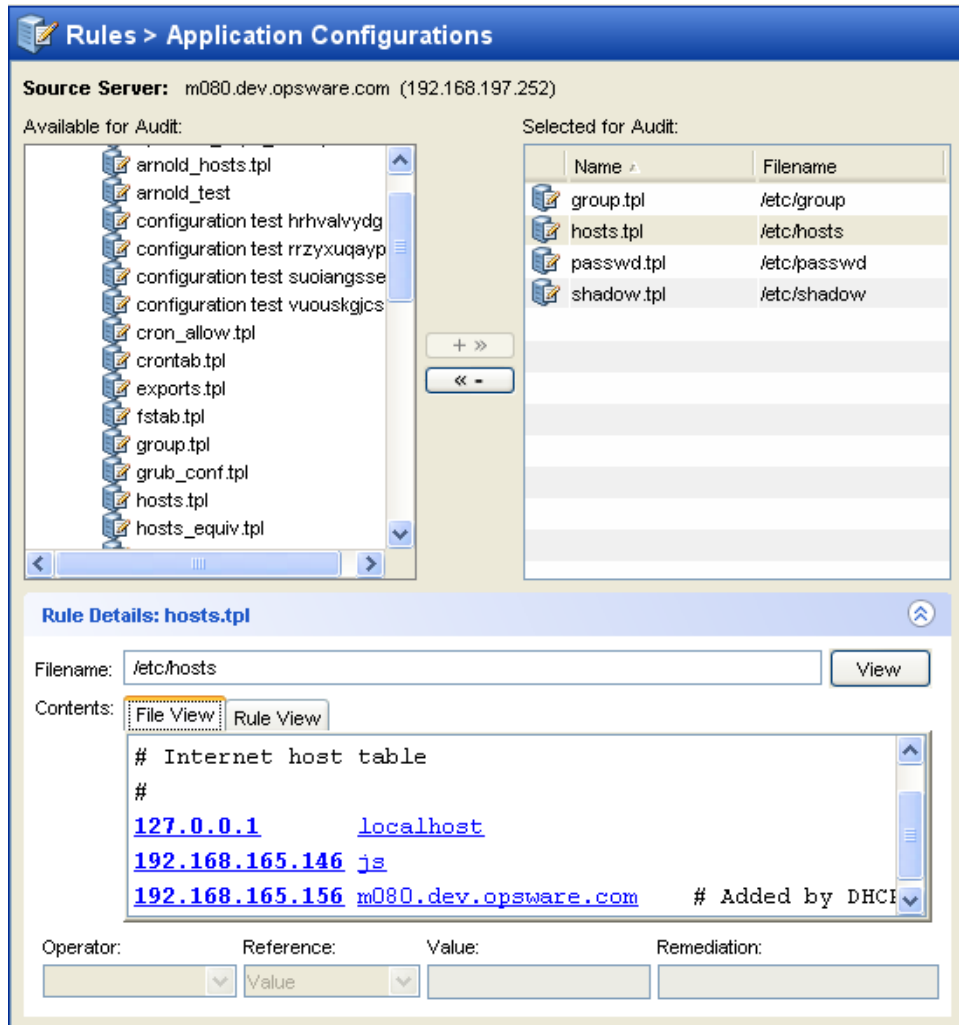
hosts file configuration, so you would choose that server as the source for your audit. This means that you will use the hosts file on that server as the basis of your application configuration audit rule.

To create an application configuration rule, perform the following tasks:

- 1** Create the new audit from any one of the methods for creating an audit listed at “Creating an Audit” on page 113.
- 2** Select a source for the audit. The source selected for the audit will determine what types of rules, if any, you can create for an application configuration. You must choose a source, or you will not be able to configure the application configuration rule.
- 3** In the Audit window, from the View pane, select Rules ► Application Configurations.
- 4** In the content pane of the audit or snapshot specification window, expand the top level node in the Available section and select an application configuration.
- 5** Click the right arrow button to move the application configuration into the Selected for Audit section.
- 6** In the Selected for Audit or Snapshot Specification section, select the application configuration.
- 7** Click **View**. (If you cannot view the contents of the configuration file, you might need to enter the correct path in the Filename section.) You see the contents of the configuration file in the File View tab.

For example, if you view a UNIX hosts file, you would see something similar to that shown in Figure 2-3:

Figure 2-7: Application Configuration Audit Rule for hosts File



You can see the contents – the IP address/host name pairs – from the source hosts file, highlighted in blue text.

- 8** In order to create an audit rule for this configuration file, you need to choose a key-value pair from the hosts file on the source server (the server you choose as the source for the audit).

- 9** To create this rule, first select an IP addresses in the File View tab content area. In the example in Figure 2-3, you would select an IP address such as 127.0.0.1. After you select the IP address, the element becomes highlighted in dark blue. This means that the element is selected but has not yet had a rule created from it.

(For more information on the color scheme used when configuring an application configuration audit rule, see Table 2-2 on page 136.)

To create a rule that will look for this IP address (127.0.0.1) in the hosts file on the target server, select the IP address in the contents area. Notice that the value in the Operator field below is set to blank. This means that the value has not yet been added to the rule. To add the value to the rule, you can either double-click it, or enter the following parameters in the rule expression area below the contents:

- **Operator:** Choose = (equals). When you change the operator to =, then the value immediately becomes added to the rule. If you change the operator back to no selection, then the value is immediately removed from the rule.
- **Reference:** Choose Value.
- **Value:** Enter 127.0.0.1.
- **Remediation:** Enter 127.0.0.1.

This expresses that you want to look for an IP address with the value of 127.0.0.1. If this is not found, then the remediation should be 127.0.0.1, so you can add this to any host files on the target servers that do not contain this IP address.

- 10** Next, select the host name in the File View tab area, and in the Rule section, set the following parameters:

- **Operator:** Choose = (equals).
- **Reference:** Choose Value. (If you choose a custom attribute here for the rule definition, this custom attribute must also exist on the target servers or the audit for this rule will fail.)
- **Value:** Choose host.
- **Remediation:** Choose host. This adds the final part of the rule that will check the target server for the key-value pair of IP address 127.0.0.1 matched with host.

11 Now, select the Rules View tab. The rule will be expressed as:

“Check that there is an entry where IP address is equal to value 127.0.0.1 and Hostnames contains an entry equal to value host.”

This rule is what will be used to audit the hosts file on the target server or snapshot specification.

12 To configure more application configuration rules, select more application configurations from the Available for Audit section.

13 To finish configuring the audit, define other rules and set the target servers, schedule, and notification for the audit.

14 Save the audit.

15 To run the audit, from the **Actions** menu, select **Run audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Application Configuration Audit Rule Color Scheme

When you first view an application configuration, all elements that can be used to build an audit rule will appear in blue underlined text. After you start selecting and building rules, then the colors will change. Table 2-2 describes the color scheme used for configuring application configuration audit rules.

Table 2-2: Application Configuration Audit Rule Color Scheme

TEXT COLOR	DESCRIPTION
Blue underlined	This shows all elements in the source configuration file that can be used in a rule.
Dark blue	This shows an element is selected but has no rule has been associated with it.
Light blue	This shows all that you add an element to a rule.
Medium blue	This shows all that an element is both selected and has a rule associated with it.

Table 2-2: Application Configuration Audit Rule Color Scheme (continued)

TEXT COLOR	DESCRIPTION
Green	<p>This shows all that the element is a primary key and is related to the current selected element. This means that the element will be used in the same rule that the current selected element will be used in.</p> <p>If the currently selected element is given a comparison value (=, contains, matches...) then the other elements with the green text will automatically be given a comparison value of “=”.</p> <p>An example of this would be:</p> <pre>127.0.0.1 localhost</pre> <p>If localhost is selected, then 127.0.0.1 would be green. If localhost is given a comparison value, then 127.0.0.1 will also be given an automatic comparison value, giving you a rule such as:</p> <p>There is an entry where ip is equal to 127.0.0.1 AND hostname is equal to localhost.</p>
Bold	This represents a primary key.
Italicized	This shows a custom attribute or Opsware attribute.

Configuring COM+ Rules

To configure a Windows COM+ object rule, select the COM+ objects that you want to audit or snapshot on a target server. You can also choose and audit COM+ categories.

The rules categorizes the COM+ objects based on an attribute of the object, where the COM object specifies zero or more categories. The Opware SAS displays all COM+ objects in one node in the Rules section of the COM+ object tree.

If you would like to be able to remediate COM+ rules in your audit or snapshot results, make sure you select the Archive all associated files option when you select the COM+ object or category.

To configure a COM+ rule, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit listed in “Creating an Audit” on page 113.
- 2** Select an Audit Source: Server, Snapshot, or None. (Some audit rules, such as application configuration, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► COM+.
- 4** In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a COM+ object or object category.
- 5** Click the right arrow button to move the COM+ object or object category into the Selected for Audit section. All COM+ object or object categories you select will be audited on the target servers or snapshot specification.



If you want to be able to remediate COM+ rules in your audit or snapshot results, select the Archive all associated files option when you select the COM+ object or category.

- 6** To finish configuring the audit, define any other COM+ object or object category rules you want and set the target servers, schedule, and notification for the audit.
- 7** Save the audit.
- 8** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Configuring Custom Scripts Rules

The custom script rule allows you to define your own script (batch, Python 1.5.2, or Visual Basic) to get and compare values used in an audit, audit policy, or snapshot specification. You can also write your own remediation scripts.

When you configure a custom script rule, you specify the target value, which is the expected values you want the script to return. The audit can gather this information in two ways:

- **Comparison-Based Audit:** Execute the script on the source server. The return values from the script (exit code or standard output) are compared with the output of the script after it has run on the target server or servers. This option is named: Source.
- **Value-Based Audit:** Specify your own value. This is compared with the output of the script after it has run on the target server. You can enter this value manually, if you know what the expected results of the script should be, or, you can execute the script on the source server and use those return values. When the audit is run, this value is compared with the returned results from the script after it has executed on the target server or servers. The option is named “Value.”

For an audit, you can also configure a remediation script, which can be used if differences are found between the rule and the value returned after the script has run on the target server.

For a snapshot, the script results will be generated by running the script (as defined in the rule detail) on target servers, and then captured in the snapshot. When you set up a snapshot specification, you can also add a remediation script. This type of script can be used to force remediation on target servers. You can execute the snapshot's remediation script on target servers on an individual server basis from the Snapshot window.

To configure a custom script rule, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit in “Creating an Audit” on page 113.
- 2** Select an Audit Source: Server, Snapshot, or None.
- 3** To build a script and define the audit rule, you can choose the following options:

Source

- **Rules:** Click **Add Rule** to add a new custom script rule.

Rule Details

- **Name:** Enter a name for the script.
- **Type of Script:** Choose from Batch, Python 1.5.2, or Visual Basic (VBS).
- **Script:** Type or copy and paste the script contents here. Or, click **Import Script** to import a script from your computer.


Success Criteria

- **Output:** Either Exit Code or Standard Output.
- **Operator:** Choose an Operator, such as equals (=), not equals (<>), less than (<), greater than (>), and so on.
- **Reference:** Choose the source of the script output.
 - **Source:** Select this option if you want the rule to execute the script on the source when an audit is run, and gets the value that the script requests. It will then compare that value with the value retrieved from the script that was run on the target server.

If you choose this option for a snapshot specification, then the script will run on the target, and the results of the script execution will be captured in the snapshot (results).

If the source of the audit is a snapshot, then the custom script rule will use the custom script definition configured in the snapshot specification.

- **Value:** Enter your own value. This option uses the value you enter and compares it with the value returned from the script after it is run on the target server. Using this option means that the script does not run on the source server at audit runtime. However, you can get the output from the script immediately from the

source server, if you click the eyedropper  icon. The returned value is displayed in the text box, which you can accept as is or edit to your liking.

If the source of the audit is a snapshot, then the custom script rule will use the Custom Script definition configured in the snapshot specification.

- **Server Attribute:** Select this option to compare a server attribute found on the source server with the output from the script that is run on the target server.
- **Custom Attribute:** Select this option to compare a custom attribute found on the target server with the output from the script that is run on the target server. Custom attributes for this option derive from the selected source server for the audit.

If you choose a custom attribute here for the rule definition, this custom attribute must also exist on the target servers or the audit for this rule will fail.

If you do not choose a source for the audit, then this list will be empty.

Remediation

- **Type of Script:** Choose from Batch, Python 1.5.2, or Visual Basic (VBS).
- **Script:** Type or copy and paste the script contents here. Or, click **Import Script** to import a script from your computer.

- 4** (Optional) You can add a remediation script to run if the audit comparison fails. The remediation will not be applied automatically; you can only run the remediation script from the audit results after the audit has run.

For a snapshot, the remediation script you define here can be executed on target servers on an individual server basis.

- 5** To finish configuring the audit, set the target servers, schedule, and notification for the audit.
- 6** Save the audit.
- 7** To run the audit, from the **Actions** menu, select **Run audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Configuring Event Logging Rules

Event logging rules allow you to gather specific application, system, and security event logging information. This version of the SAS Client allows you to configure the following event logging audit rules:

- Crash on Audit Failure Security Log Status
- Maximum Application Event Log Size
- Maximum Security Event Log Size
- Maximum System Event Log Size
- Security Log Near Capacity Warning

To configure event logging audit rules, perform the following steps:

- 1** Create the new audit using one of the methods listed at “Creating an Audit” on page 113.

- 2** Select an Audit Source: Server, Snapshot, or None. (Some audit rules, such as application configuration, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Event Logging.
- 4** In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select the Event Logging that you want to create a rule from.
- 5** Click the right arrow button to move the Event Logging rule object into the Selected for Audit section. All Event Logging rule objects you select will be audited on the target servers or snapshot specification.
- 6** Options for each event logging rule define different rule parameters:
 - Crash on Audit Failure Security Log Status
 - **Target Value:** Choose equals (=) or not equals (<>), a Reference from the Source server, or your own value. If you choose Value for the reference, then choose for the value either Disabled or Abled.
 - **Remediation Value:** Enabled or Disabled for crash on Full Log.
 - Maximum Application Even Log Size, Maximum Security Event Log Size, Maximum System Event Log Size.
 - **Target Value:** Choose equals (=) or not equals (<>), a Reference from the Source server or your own value. If you choose your own value for the reference, enter a value for the megabyte size for the log file.
 - **Remediation Value:** From Reference, choose Value and then enter a value that you want to remediate the log size to (in megabytes).
 - Security Log Near Capacity Warning
 - **Target Value:** Choose to set as equals (=) or not equals (<>), a Reference from the Source server or your own value. If you choose your own value for the Reference, enter a value for the megabyte size for the log file.
 - **Remediation Value:** From Reference, choose Value and then enter a value you want to remediate the log size to (in megabytes)
- 7** To finish configuring the audit, set the target servers, schedule, and notification for the audit.
- 8** Save the audit.
- 9** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Configuring File System Rules

The file system rule allows you to audit a server's file system and directory structure for the following evaluations:

- **Full contents of a file:** Audits the full contents of the selected file.
- **Checksum for each file:** Performs a Checksum on the contents of the selected file or files in a directory. You can choose to audit the entire contents of the file, or just the first 1MB of the file.
- **User and group access:** Audits the user and group access related to the file and directories.
- **Windows NT ACLs (Access Control List):** Audits the Windows Access Control List for files and directories.
- **File modification date:** Adds the file modification date to the audit.
- **Compare contents of subdirectories:** Includes contents of all subdirectories for a selected file system folder to the audit.
- **File/Wildcard directory:** Allows you to specify directories and files in the file system you want included in and excluded from the audit. For more information on how this option works, see "File System Inclusion and Exclusion Rules" on page 155.

There are two categories of file system rules that appear in the Available for Audit section of the Audit window. You can define the following specifications in an audit or snapshot:

- **File System:** These are comparison-based rules, which enable you to select a file system file or directory from the source of the audit or snapshot specification and compare these with the target servers. The purpose of this rule is to determine that the file or directory exists and its properties. You cannot set a target or remediation value in the rule.
- **Specific File System Rules:** These are value-based file system rules prebuilt into the SAS Client. They allow you to configure expected (target) and remediation values.

To configure file system rules, perform the following steps:

- 1** Create the new audit using one of the methods in "Creating an Audit" on page 113.
- 2** Select an Audit Source: Server, Snapshot, or None. (Some audit rules, such as application configuration, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► File System.

- 4** In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a folder or file to create a rule for.
- 5** Click the right arrow button to move the folder or file into the Selected for Audit section. All folders or files that you select will be used to audit or snapshot the target server.
- 6** In the Selected for Audit section, select a folder or file to apply a rule to.
- 7** In the Directory Options section, select file system rule options to apply to the selected folder or file. If you would like to reset the original settings of the source file system, select the Reset options to match those of the File System option.
- 8** (Optional) For folders, you can select a File/directory Wildcard option to specify files and directories that you want to include or exclude from the audit.

Click the **plus (+)** button to add a new rule, or click the **minus (-)** button to remove a rule. For more information on how to enter files and directories and how this affects the audit, see "File System Inclusion and Exclusion Rules" on page 155.
- 9** To finish configuring the audit, set the target servers, schedule, and notification for the audit.
- 10** Save the audit.
- 11** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see "Creating an Audit Policy" on page 166.

Configuring Hardware Rules

Configuring a hardware rule allows you to audit the following information about a server's hardware:

- **CPU:** Compare CPU type and specification of target server.
- **Memory:** Compare memory of the target server.
- **Storage:** Compare storage capacity on the target server.
- **Interfaces:** Compare all network interfaces attached to the device.

To configure hardware rules, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit listed in "Creating an Audit" on page 113.

- 2 Select an Audit Source: Server, Snapshot, or None. (Some audit rules, such as application configuration, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Hardware.
- 4 In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a hardware category to create a rule for.
- 5 Click the right arrow button to move the hardware item into the Selected for Audit section. All items that you select will be used to audit or snapshot the target server.
- 6 To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7 Save the audit.
- 8 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Configuring IIS Metabase Rules

The IIS Metabase audit rule allows you to select IIS Metabase objects and objects folders to compare in your audit. The audit will capture IIS Metabase object property information such as ID, name, path, attributes, and so on.

To configure IIS Metabase rules, perform the following steps:

- 1 Create the new audit using one of the methods for creating an audit listed at “Creating an Audit” on page 113.
- 2 Select an Audit Source: Server, Snapshot, or None. (Some audit rules, such as application configuration, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► IIS Metabase.
- 4 In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select an IIS Metabase folder or object to create a rule for.
- 5 Click the right arrow button to move the IIS Metabase folder or object into the Selected for Audit section. All items you select will be used to audit or snapshot the target server.
- 6 To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7 Save the audit.

- 8 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Configuring Operating System Rules

Operating system rules allow you to check dozens of different operating system properties and settings. You can check controller settings (for example, LDAP Server Signing Requirements), network configurations (for example, IP Source Routing Protection Level), auditing process tracking, alerts, clearing virtual memory page file, and so on.

While each rule is slightly different and requires its own configuration values, the basic parameters for each rule require that you define the Target Value – the expected value you want to find on the server – and an optional Remediation Value.

To configure operating system rules, perform the following steps:

- 1 Create the new audit using one of the methods in “Creating an Audit” on page 113.
- 2 Select an Audit Source: Server, Snapshot, or None.
- 3 In the Audit window, from the View pane, select Rules ► Operating System.
- 4 In the content pane of the Audit window, expand the top level node in the Available for Audit section and select an operating system rule that you want to create a rule from.
- 5 Click the right arrow button to move the operating system rule object into the Selected for Audit section. All operating system rule categories that you select will be audited on the target servers or snapshot specification.
- 6 For each operating system rule, you can define or set the following parameters:

Input Value


Some of the operating system rule checks require an input value as part of the configuration of the target value. For those rules, you will need to specify a success or failure which you can set to true or false. The Description section of the audit rule explains the CIS recommended values.

Target Value

Here you can specify the value that you expect to be on the target server or servers of the audit, or the value you want to capture in a snapshot. You can change the following parameters:

- **Operator:** If you want to build an expression from the output of the script, choose

an Operator, such as equals (=), not equals (<>), less than (<), greater than (>), and so on.

- **Reference:** Choose the source of the script output.
 - **Source:** This will use the value from the source server and compare that value to with the value found on the target server or servers.
 - **Value:** Enter your own value. This option uses the value you enter and compares it with the value returned on the target server. You can also choose to get the value from the source server if you click the eyedropper  icon. The returned value is displayed in the text box, which you can accept as is or edit to your liking.
 - **Server Attribute:** Select to compare a server attribute located on the source server.
 - **Custom Attribute:** Select to compare a custom attribute found on the target server.

Remediation Value

Each remediation value setting will be different depending on the type of rule, so choose accordingly.

- 7** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 8** Save the audit.
- 9** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Configuring Software Rules

Configuring a software rule allows you to audit the following information about a server's software:

- **Installed Packages:** Audit all packages installed on the target servers.
- **Installed Patches:** Audit all patches installed on the target servers.

To configure software rules, perform the following steps:

- 1** Create the new audit using one of the methods in "Creating an Audit" on page 113.
- 2** Select an Audit Source: Server, Snapshot, or None. (Some audit rules, such as application configuration, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Software.
- 4** In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select Installed Packages or Installed Patches.
- 5** Click the right arrow button to move the items into the Selected for Audit section. All items that you select will be used to audit or snapshot the target server.
- 6** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7** Save the audit.
- 8** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see "Creating an Audit Policy" on page 166.

Configuring Users and Groups Rules

The users and groups rule allows you to audit the following checks on the target of your audit or snapshot specification:

- **Disable CTRL-ALT-DEL Login:** Determines if the behavior requiring a user to select CTRL-ALT-DEL to log on is enabled or disabled.
- **Display Last Username at Login:** Determines if the display of the last user to log on is enabled or disabled.
- **Message Text for Users Attempting to Log On:** Specifies the expected text shown to users who attempt to log on to the server.

- **Message Title for Users Attempting to Log On:** Specifies the expected title for the text message shown to users who attempt to log on to the server.
- **Sharing and Security Model for Local Accounts:** Determines the sharing and security model for local user accounts.


To configure users and groups audit rules, perform the following steps:

- 1** Create the new audit using one of the methods in “Creating an Audit” on page 113.
- 2** Select an Audit Source: Server, Snapshot, or None. (Some audit rules, such as application configuration, must have a source.)
- 3** In the Audit window, from the View pane, select rules ► Users and Groups.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Users and Groups folder or key to create a rule for.
- 5** Click the right arrow button to move the Users and Groups rule into the Selected for Audit section. All items that you select will be used to audit or snapshot on the target server.
- 6** For each rule, specify the following:

Input Values

- Some rules require that you specify an extra input parameter value that you expect to find on the target server.

Target Values

- **Operator:** Select an operator for the target value rule, such as equals (=), not equals (<>), less than (<), greater than (>), and so on.
- **Reference:**
 - **Value:** Enter your own value. This option uses the value you enter and compares it with the value returned target server. You can also choose to get the value from the source server, if you click the eyedropper  icon. The returned value is displayed in the text box, which you can accept as is or edit to your liking.
 - **Source:** Use the value from the source server and compare that value to with the value found on the target server or servers. If your audit is using a snapshot as the source, then you will only be able to select rules from the snapshot specification and its results.
 - **Server Attribute:** Compares a server attribute located on the source server.

- **Custom Attribute:** Compares a custom attribute found on the target server.
- **Value:** Select a value for the rule.

Remediation Value

Specify a Remediation value for each rule by choosing an option or entering your own value.

- 7** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 8** Save the audit.
- 9** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Configuring Windows Registry Rules

The Windows Registry rule allows you to select Windows Registry folders and keys to compare in your audit. The audit compares the selected registry folders and keys and determines if these keys and folders exist on the target servers.

There are two categories of Windows Registry rules that you can define in an audit or snapshot specification. The following categories appear in the Available for Audit section of the Audit window:

- **Windows Registry:** These are comparison-based rules, which enable you to select a Windows Registry key or folder from the source of the audit or snapshot specification and compare these with the target servers. The purpose for this kind of rule is to determine if the Windows Registry key or folder exists and its properties. You cannot set a target or remediation value in the rule.

The Windows Registry object allows you to capture registry keys, values, and subkeys. A registry key is a directory that contains registry values, where registry values are similar to files within a directory. A subkey is similar to a subdirectory. The content area in this window excludes subkeys. The Opware Audit and Remediation feature supports the following Windows Registry keys: HKEY_CLASSES_ROOT, HKEY_CURRENT_CONFIG, HKEY_LOCAL_MACHINE, and HKEY_USERS.

Valid control characters audited and captured for the contents of the key entry (Data) include: #x9, #xA, [#xD, #x20-#xD7FF], [#xE000-#xFFFFD], and [#x10000-#x10FFFF]. Invalid control characters cannot be stored by the SAS Client and will be converted to XML entities and will display as &#;. For example, if the data value is 00 00 (in bytes), � will display in the audit or snapshot specification results.

- **Specific Windows Registry Rules:** These are value-based Windows Registry rules prebuilt into the Opsware SAS and they allow you to configure expected (target) and remediation values.

To configure Windows Registry audit rules, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit listed in “Creating an Audit” on page 113.
- 2** Select an Audit Source: Server, Snapshot, or None. (Some audit rules, such as application configuration, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Windows Registry.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Windows Registry folder or key to create a rule for.
- 5** Click the right arrow button to move the Windows Registry folder or key into the Selected for Audit section. All items that you select will be used to audit or snapshot the target server.
- 6** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7** Save the audit.
- 8** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Configuring Windows Services Rules

The windows service rule allows you to select windows services to compare in your audit or snapshot specification. The audit or snapshot specification compares the selected services with services on the target servers to determine if the services exist and if the services are started, stopped or disabled.

There are two categories of windows services rules that you can define in an audit or snapshot specification. The following rules appear in the Available for Audit section of the Audit window:

- **Windows Services:** These comparison-based rules enable you to select a service from the source of the audit or snapshot specification and compare them with the target servers. The purpose of windows services rule is to determine if the service exists and its settings. You cannot set a target or remediation value with this type of rule.

- **Other Windows Services Rules:** These value-based windows services rules prebuilt into the SAS Client allow you to configure expected (target) and remediation values.

To configure windows services rules, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit listed in “Creating an Audit” on page 113.
- 2** Select an Audit Source: Server, Snapshot, or None. (Some audit rules, such as application configuration, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Windows Services.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a windows services to create a rule for.
- 5** Click the right arrow button to move the selected windows services into the Selected for Audit section. All items that you select will be used to audit or snapshot on the target server.
- 6** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7** Save the audit.
- 8** To run the audit, from the **Actions** menu select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

Configuring The Opsware Network (TON) Rules

The Opsware Network (TON) rules (known by content developers as “pluggable checks”) are part of the TON Subscription Service and give you access to many different types of customized audit rules based on industry compliance standards.


The kinds of rules you have access to depend on your subscription, but can include such rules as the latest patch supplements for Microsoft Windows, current regulatory compliance policies (for example, FISMA, Sarbanes-Oxley), user-created rules from the TON developer community, daily updated vulnerability content, and so on.



If you are not subscribed to TON, you will not see any TON rules in your audits, audit policies, or snapshots. If you would like more information on how to subscribe to TON, contact your Opsware sales representative.

While each TON rule is slightly different and requires its own configuration values, the basic parameters for each TON rule require that you define the Target Value – the expected value you want to find on the server – and an optional Remediation Value.

To configure TON rules, perform the following steps:

- 1** Create an audit using one of the methods described in “Creating an Audit” on page 113.
- 2** Select an Audit Source: Server, Snapshot, or None.
- 3** In the Audit window, from the View pane, select Rules.
- 4** Depending on the kind of TON subscription you have, select one of the TON rule categories, indicated by the pluggable checks icon . For example, if you are subscribed to it, you would select the TON rule category named TON Compliance Content.
- 5** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a TON rule that you want to create a rule from. For example, if you are subscribed to TON Compliance Content, you could select the following: SOX ► Windows 2003 ► .NET Framework Support set to disabled.
- 6** Click the right arrow button to move the TON rule object into the Selected for Audit section. All TON rules that you select will be audited on the target servers or snapshot specification when you run the audit or snapshot specification.
- 7** For each rule, define or set the following parameters:


Input Value

Some TON rules require an input value as part of the configuration of the target value. For those rules, you will need to specify a success or failure which you can set to true or false. The Description section of the audit rule explains the recommended values.

Target Value

Here you can specify the value that you expect to be on the target server or servers of the audit, or the value you want to capture in a snapshot. You can change the following parameters:

- **Operator:** If you want to build an expression from the output of the script, choose an Operator, such as equals (=), not equals (≠), less than (<), greater than (>), and so on.

- **Reference:** Choose the source of the script output.
 - **Source:** This will use the value from the source server and compare that value to with the value found on the target server or servers.
 - **Value:** Enter your own value. This option uses the value you enter and compares it with the value returned on the target server. You can get the value from the source server if you click the eyedropper  icon. The returned value is displayed in the text box, which you can accept as is or edit to your liking.
 - **Server Attribute:** Select to compare a server attribute located on the source server.
 - **Custom Attribute:** Select to compare a custom attribute found on the target server.

Remediation Value

Each remediation value setting will be different depending on the type of rule, so choose accordingly.

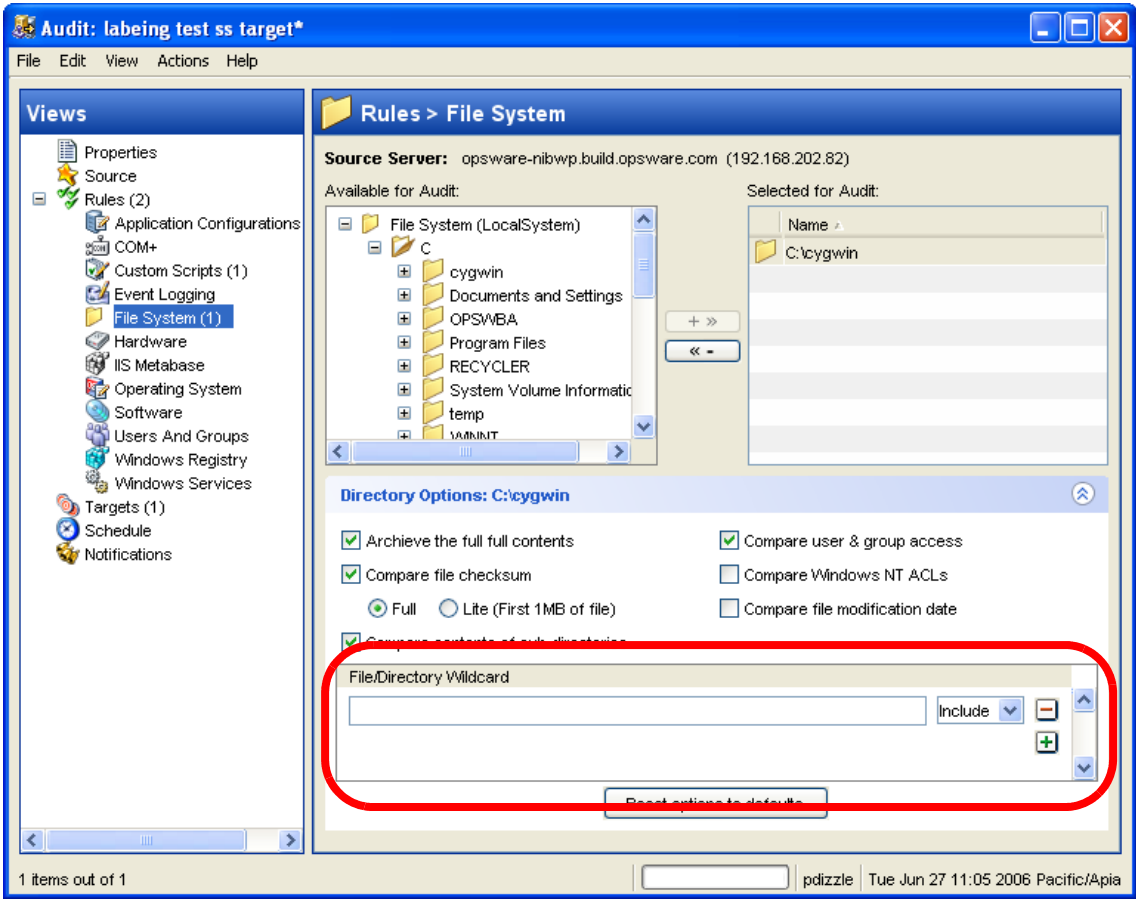
- 8** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 9** Save the audit.
- 10** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 166.

File System Inclusion and Exclusion Rules

When configuring a file system rule inside an audit, audit policy, or snapshot specification, you can specify the directories and files that you want included in and excluded from an audit or a snapshot. This section explains what the inclusion and exclusion rules are and how these rules are applied to the relative subset of the absolute path of the file.

Inclusions and exclusion rules inside of an audit's file system rule are found at the bottom of the audit or snapshot specification window, as shown in Figure 2-8.

Figure 2-8: File System File/Directory Wildcard Inclusion and Exclusion Rules



When you configure the file system rule in an audit or snapshot specification, you can enter inclusion/exclusion rules in the File/Directory Wildcard field. After you enter a rule, you can choose either Include or Exclude from the drop-down list. To add a new inclusion or exclusion rule, click the plus (+) button.

For information on how to create and configure file system rules for an audit or snapshot specification, see “Configuring File System Rules” on page 143.

Inclusion and Exclusion Rule Types

Opware Audit and Remediation provides the following types of inclusion and exclusion rules configuring a file system rule:

- A file-type rule applies to the file name path and contains neither a “/” or a “\”.
- A relative-type rule applies to the relative path and can contain a “/” for Unix and a “\” for Windows, and is not fully qualified.
- An absolute-type rule applies to the absolute path. In Unix, an absolute path begins with a “/”. In Windows, an absolute path begins with a volume letter that is followed by “:\” and is fully qualified, such as “C:\”, “d:\”, “f:\”, and so on. If you use a “/” (forward slash) for Windows paths, Opware Audit and Remediation will convert it to a “\” (backslash) to use it as a valid path.

Opware Audit and Remediation processes all exclusion rules first. After all exclusion rules are applied, then the inclusion rules are applied. The default for include is to include all objects in the file system. In many cases, inclusion rules might not even be processed because, combined with the exclusion rules (which occur first), they might become a moot point.

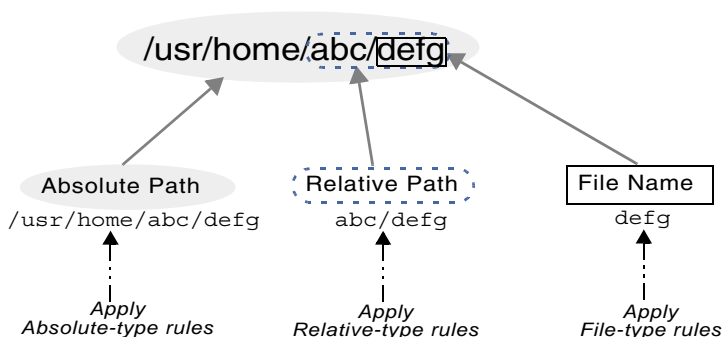
You can also use the asterisk (*) and the question mark (?) as valid wildcards in inclusion and exclusion rules. The wildcard character is a placeholder for matching a path, or one or more characters.

Depending on the type of inclusion and exclusion rule, the rule is applied only to the relevant subset of the absolute path of the file. In Opware Audit and Remediation, there is one top level for each snapshot or audit. Each file that you compare against the inclusion and exclusion rules has an absolute path. In Figure 2-9, the absolute path is `/usr/home/abc/defg`. A snapshot or an audit looks down the `/usr/home/abc/defg` absolute path and sees `abc/defg` as the relative path and `defg` as the file name. In this example, the inclusion and exclusion rules apply in the following manner:

- A file-type rule applies to the file name path `defg`.
- A relative-type rule applies to the relative path `abc/defg`.
- An absolute-type rule applies to the absolute path `/usr/home/abc/defg`.

See Figure 2-9 for an illustration of how Opsware Audit and Remediation applies the inclusion and exclusion rules to a relative subset of the path of the file.

Figure 2-9: How Inclusion and Exclusion Rules Apply



To best explain how these rules are applied, the following examples are provided.

A sample file system structure used in “Example: Including all .txt Files in a Snapshot or Audit” on page 157 and “Example: Including last temp.txt file and exclude all else” on page 159 is as follows:

```

/dir1/dir2/a
/dir1/dir2/b
/dir1/dir2/names.txt
/dir1/dir2/temp.txt
/dir1/dir2/version1.exe
/dir1/dir2/subdir/version2.exe
  
```

Example: Including all .txt Files in a Snapshot or Audit

If you want to include all files with the .txt extension in your snapshot or audit, your inclusion and exclusion rules would be:

- /dir1/dir2
- include *.txt (This is a file-type rule.)
- exclude * (This is a file-type rule.)

The following steps explain how Opsware Audit and Remediation iterates through the file structure and applies any corresponding inclusion and inclusion rules:

1. The * causes /dir1/dir2/a to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/a (a) and there is no match. The file is not included.

2. The * causes /dir1/dir2/b to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/b (b) and there is no match. The file is not included.
3. The * matches names.txt, but *.txt matches names.txt as well, which causes the file to be excluded.
4. Same as step 3.
5. Compare a to *, which is a match; compare a to a, which is a match. The file is included.
6. Compare b to *, which is a match; compare b to a which is not a match. The file is excluded.

These step numbers correspond to the paths in the sample file structure, with the numbering starting with the top-level path.

Example: Including Only File a in a Snapshot or Audit

If you want to include only the file in your snapshot or audit, your inclusion and exclusion rules would be:

- /dir1/dir2
- exclude * (This is a file-type rule.)
- include a (This is a file-type rule.)

The following steps explain how Opware Audit and Remediation iterates through the file structure and applies any corresponding inclusion and inclusion rules:

1. The * causes /dir1/dir2/a to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/a (a) and there is no match. The file is not included.
2. The * causes /dir1/dir2/b to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/b (b) and there is no match. The file is not included.
3. The * matches names.txt, but *.txt matches names.txt as well, which causes the file to be included.
4. Same as step 3.
5. Compare a to *, which is a match; compare a to a, which is a match. The file is included.
6. Compare b to *, which is a match; compare b to a which is not a match. The file is excluded.

These step numbers correspond to the paths in the sample file structure, with the numbering starting with the top-level path.

Example: Including last temp.txt file and exclude all else

If you want to include the last temp.txt file and exclude everything else in your snapshot or audit, your inclusion and exclusion rules would be:

- /dir1/dir2
- exclude * (This is a file-type rule.)
- include dir3/temp.txt (This is a relative-type rule.)

The following steps explain how Opsware Audit and Remediation iterates through the file structure and applies any corresponding inclusion and inclusion rules:

1. The * causes /dir1/dir2/a to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/a (a) and there is no match. The file is not included.
2. The * causes /dir1/dir2/b to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/b (b) and there is no match. The file is not included.
3. The * matches names.txt, but *.txt matches names.txt as well, which causes the file to be included.
4. Same as step 3.
5. dir3/temp.txt is dir3/temp.txt is compared against the relative portion of /dir1/dir2/dir3/temp.txt and there is a match.
6. Compare a to *, which is a match; compare a to subdir/version2.exe, which is not a match. The file is excluded.

These step numbers correspond to the paths in the sample file structure, with the numbering starting with the top-level path.

File System Rule Overlap

When you include a parent directory (with options) in a rule and a child directory (with different options) as additional parameters, the parent directory snapshot and the child directory snapshot will overlap each other as one snapshot. This logic also applies to Windows NT ACL collection and content collection options, and Windows Registry content collection options. The following examples explain how audit rules for a parent and child directory overlap.

Consider the following file system, where an ending forward slash (/) represents a directory:

```
/cust/app/bin/  
/cust/app/bin/file1  
/cust/app/bin/conf/  
/cust/app/bin/conf/conf1  
/cust/app/bin/conf/conf2  
/cust/app/bin/conf/dev/  
/cust/app/bin/conf/dev/conf3
```

Example A

If you create a snapshot using the following two rules:

Directory /cust/app/bin (recursive, no checksum)

Directory /cust/app/bin/conf (not recursive, checksum)

The snapshot will record the following file system information:

```
/cust/app/bin/ (directory)  
/cust/app/bin/file1 (no checksum)  
/cust/app/bin/conf/ (directory)  
/cust/app/bin/conf/conf1 (*checksum*)  
/cust/app/bin/conf/conf2 (*checksum*)  
/cust/app/bin/conf/dev/ (directory)  
/cust/app/bin/conf/dev/conf3 (no checksum)
```

As you can see, even though /cust/app/bin was recursive and had no checksum, the /cust/app/bin/conf directory overrode it and all files in that directory have checksums recorded for them.

Example B

If you create a snapshot using the following two audit rules (by switching the options used in Example A):

```
Directory /cust/app/bin (recursive, checksum)  
Directory /cust/app/bin/conf (not recursive, no checksum)
```

The snapshot will record the following file system information:

```
/cust/app/bin/ (directory)  
/cust/app/bin/file1 (checksum)  
/cust/app/bin/conf/ (directory)  
/cust/app/bin/conf/conf1 (*no checksum*)  
/cust/app/bin/conf/conf2 (*no checksum*)  
/cust/app/bin/conf/dev/ (directory)  
/cust/app/bin/conf/dev/conf3 (checksum)
```


Example C

If you create a snapshot using the following three audit rules (by adding a file option):

Directory /cust/app/bin (recursive, checksum)

Directory /cust/app/bin/conf (not recursive, no checksum)

File /cust/app/bin/conf/conf1 (checksum)

The snapshot will record the following file system information:

```
/cust/app/bin/ (directory)
/cust/app/bin/file1 (checksum)
/cust/app/bin/conf/ (directory)
/cust/app/bin/conf/conf1 (*checksum*)
/cust/app/bin/conf/conf2 (no checksum)
/cust/app/bin/conf/dev/ (directory)
/cust/app/bin/conf/dev/conf3 (checksum)
```

In this example, the very detailed audit rules for `conf1` override the `/cust/app/bin/conf` audit rule.

Audit Rule Exceptions

For most audit rules, you can create temporary or permanent rule exceptions on selected target servers (or groups of servers) in the audit. This means you can exclude specific rules on selected targets of the audit when the audit runs.

For example, in an audit that is auditing several servers, you might want to suspend one or more of the rules for a subset of the servers targeted by the audit. You might have a collection of Windows servers that are regularly audited to make sure that the IIS service is disabled, for example, to meet company security standards. Your audit is configured to check each of those servers to make sure IIS is disabled. If IIS is enabled on any of the servers, the audit will fail.

However, for a short period of time you might want to run a business application that requires the IIS service to be enabled in order to run on a few of the servers targeted in the audit. You can create a rule exception for the rule governing the IIS service and associate the exception with the servers that need to run the application. This ensures that the audit can still run and not fail when it encounters the servers that do have the IIS service enabled.

You can set an expiration date for the rule exceptions to make sure that when the rule exception is no longer needed or permitted, the rule will be applied to all servers in the audit. You can also write a reason for the exception and associate a ticket ID with it. Exceptions you create in one audit do not affect rules in any other audits.

Rules That Cannot Have Exceptions

Most audit rules can have exceptions created for them. However, rule categories that include ALL of a set of rules cannot have exceptions, such as:

- All Windows COM+ objects
- All IIS Metabase objects
- Hardware
- Installed Packages
- Installed Patches
- All Windows Services

Considerations When Applying Exceptions to Device Groups

When you set an audit rule exception for a device group, the exception will be applied to all servers in the group. It is possible that one of the servers in the group with the exception also belongs to another device group, which also happens to be the target of an audit that has no exceptions applied to it.

In this situation, the rule exception always applies to the server, even though the server also belongs to a device group with no exceptions. As a rule of thumb, keep in mind any servers in a device group that has a rule exception applied to it will have the audit rule excepted, whether or not the server belongs to another device group that is targeted by an audit and has the same rule applied without an exception.


Adding a Rule Exception to an Audit

To create an audit rule exception, select any of the rules configured in your audit and using the Add Rule Exception window, associate them with a target server in the audit. When you run the audit, the selected rule and the target servers or snapshots associated with the rule will not be applied.

You can also apply rule exceptions to device groups. You can set the rule exception to run indefinitely, or to expire at some future point in time. You can add a comment to explain why you are creating the exception, and also associate a ticket ID with the exception.

Some audit rules and audit rule collections cannot be excepted. For more information, see “Rules That Cannot Have Exceptions” on page 162.

To add a rule exception to an audit, perform the following steps:

- 1** First, create an audit. For information see “Creating an Audit” on page 113.
- 2** Configure audit rules for the audit. For information on configuring audit rules, see “Configuring Opsware Audit and Remediation Rules” on page 124.
- 3** From the audit view pane on the left, select the Exception  object.
- 4** Next, from the content pane, click **Add**.

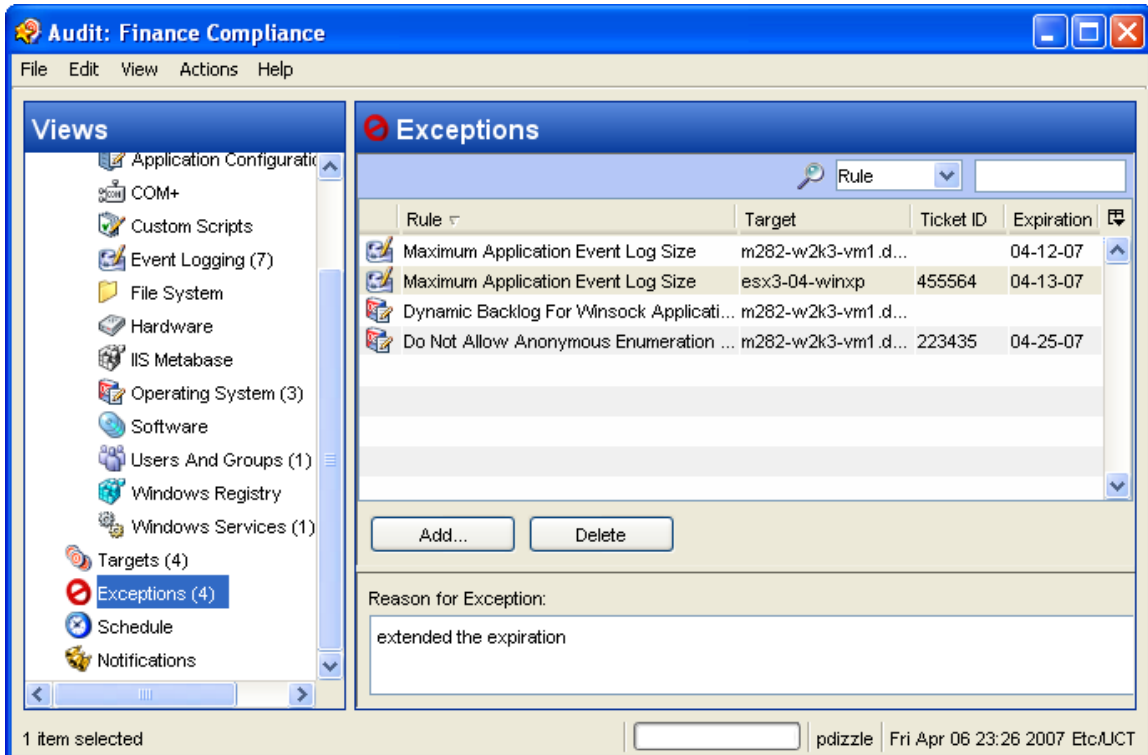


You can also select any rule in the Audit window, right-click, and select **Add Exception**. However, if the audit is referencing a linked audit policy, right-clicking a rule to add an exception will not work.

- 5** In the Add Exception window, from the Select Target Server section, select a server, multiple servers, or device groups to which you want to apply the rule exception.
- 6** Next, from the Select Rule section, select one or more rules you want to associated with the servers you selected in the previous step.
- 7** (Optional) In the Reason for Exception section, add an explanation.
- 8** (Optional) In the Ticket ID section, add the ticket ID associated with this exception.
- 9** In the Expires section, either enter a date to indicate when the exception expires, or select a date from the drop down list.
- 10** When you are finished configuring the exception, click **Add**.

- 11** You now see a list of rule exceptions that will be applied when you run the audit, as shown in Figure 2-10.

Figure 2-10: Audit Rule Exceptions Added to an Audit




Editing or Deleting a Rule Exception

You can edit an exception in one of two ways:


- Double-click the exception to modify the reason for the exception, the ticket ID, and the exception expiration date
- Click the **Add** to edit a rule (overwrite the existing rule)

To edit an exception, perform the following steps:

- 1** Open an audit window.
- 2** From the audit's View pane on the left, select the Exception  icon.
- 3** From Contents pane, double-click an exception.
- 4** In the Edit Exception window, you can edit any of the exceptions and servers or device groups they are assigned to. When you have edited the exception, click **Add**.

- 5 If you want to completely change the rule, click the **Add** button and then in the Add Exception window, change the rule by selecting target server and one or more rules. When you are finished, click **Add** to change the exception.

To delete an exception, perform the following steps:

- 1 Open an audit window.
- 2 From the audit view pane on the left, select the Exception  object.
- 3 From the Contents pane, select the exception you want to select, and then click **Delete**.

Audit Policies

Audit and Remediation allows you to create audit policies, which are a collection of rules that define a desired state of a server's configuration. An audit policy can be used inside an audit or snapshot specification, either through linking or importing. An audit policy is very similar – in fact, nearly identical – to an audit, but differs from an audit in that it does not contain any information about target servers or scheduling or notification.

An audit policy is like a reusable template that represents an ideal state of server configuration and defines specific compliance standard for servers in your facility. An audit policy is useful because it allows a policy setter to define server configuration compliance values, which can then be used by others in the context of an audit or snapshot specification.

You can create an audit policy from scratch, or you can save an existing audit as an audit policy, which extracts only the rules defined in an audit so it can be reused in other audits or snapshots. An audit policy can *link* into an audit or snapshot specification so whenever a change is made the audit using the policy will have the latest changes. An audit policy can also be *imported* into an audit or snapshot specification, without keeping the link to the source audit policy. When you import an audit policy into an audit, you can choose to replace any current values in the audit or merge rules from the audit policy with those in the audit or snapshot specification.

If you subscribe to The Opsware Network (TON), some of the latest industry compliance standards are defined as rules inside each new audit policy. For example, subscribing to TON Essential Content gives you access to regularly updated audit policies containing security best practices, such as CIS, NSA, and so on, and the Opsware patch supplement

for Microsoft Windows. Subscribing to the TON Subscription Service, you will be able to access the most current regulatory compliance audit policies (FISMA, Sarbanes-Oxley, etc.) and daily vulnerability alerts.

For information on subscribing to TON, contact your Opsware sales representative.

For information on creating rules for an audit policy, see “Configuring Opsware Audit and Remediation Rules” on page 124.

Creating an Audit Policy

When creating an audit policy, you have the option of creating the rule using either a live server or a snapshot. This allows you to use the rule from a known good server, or a snapshot of a known good server.

To create an audit policy, perform the following steps:

- 1** From the Navigation pane, select Library and then select Audit and Remediation.
- 2** In the Navigation pane, select audit Policies, and then select Windows or Unix.
- 3** Right-click inside the Content pane and from the **Actions** menu, select **New**.
- 4** In the Content pane, for the audit policy's properties, enter a name and description.
- 5** From the Views pane on the left side of the Audit Policy Window, select Source if you would like use a source server or snapshot to the audit policy on.
- 6** From the Content pane, select a source for the audit policy, and then click **Select**.
- 7** In the Select a Source window, select either a server or a snapshot, and then click **OK**.
- 8** From the Views pane, select rules to configure. For more information on how to configure specific rules, see “Configuring Specific Rules” on page 131.

Linking and Importing Audit Policies

You can import or save an audit policy into either an audit or snapshot specification in the three following ways (and create an audit policy):

- Linking an Audit Policy
- Importing an Audit Policy (replace or merge)
- Saving as Audit Policy

Linking an Audit Policy

Linking an audit policy into an audit or snapshot specification creates a link that populates an audit's or snapshot specification's rules with those of the audit policy. This is useful if a policy setter wants to define a server configuration policy in an audit and have others users link to the audit policy. If the policy setter makes any changes to the source audit or snapshot specification, then the changes will be reflected in the policy.

When an audit policy is linked into an audit or snapshot specification, the rules cannot be modified in the audit or snapshot specification.

If the audit or snapshot specification you are linking to already has some rules defined, then linking an audit policy will overwrite those existing rules.

To link an audit policy in an audit, perform the following steps:

- 1** Open an existing audit from the Library using one of the following methods:
 - From the Navigation pane, select Library ► Audit and Remediation ► Snapshot Specification, and then open the audit.
 - From the Navigation pane, open an existing snapshot specification from Library ► Audit and Remediation ► Snapshot Specification.
- 2** From the **Actions** menu, select **Link to Policy**.
- 3** If you are linking an audit policy into an audit or snapshot specification that already has had some rules defined, a message snapshot specification will display, explaining that you will overwrite any existing rule definitions. Click **Yes** to import the audit policy.
- 4** To save the audit or snapshot specification, from the **File** menu, select **Save**.

Importing an Audit Policy

Importing an audit policy into an audit or snapshot specification allows you to import (and optionally merge) an audit policy's rules into an audit or a snapshot specification, without keeping a link to the audit policy.

After you import an audit policy, there is no more connection to that audit policy, and any changes made to the source audit policy are not reflected where the audit policy was imported into.

To import an audit policy into an audit, perform the following steps:

- 1** Open an existing audit from the Library using one of the following methods:

- From the Navigation pane, select Library ► Audit and Remediation ► Audits, and then open the audit.
- From the Navigation pane, open an existing snapshot specification from Library ► Audit and Remediation ► Snapshot Specification.

2 From the **Actions** menu, select **Link to Policy**.

- 3** If the audit or snapshot specification already has rules defined, choose to either to overwrite the existing rules, or merge the audit policy rules with the existing rules:
- If you click **Yes**, then the audit policy will overwrite any existing rules in the audit or snapshot specification.
 - If you click **No**, then the audit policy will merge the audit policy rules with any existing rules. If any conflicts are found, then the audit policy rules will overwrite any existing rules.

4 To save the audit or snapshot specification, from the **File** menu, select **Save**.

Saving as Audit Policy

You can save an audit or a snapshot specification's rules as an audit policy, which can be then used by others in an audit or snapshot specification.

To save an audit or snapshot specification as an audit policy, perform the following steps:

- 1** Open an existing audit from the Library using one of the following methods:
- From the Navigation pane, select Library ► Audit and Remediation ► Audits, and then open the audit.
 - From the Navigation pane, open an existing snapshot specification from Library ► Audit and Remediation ► Snapshot Specification.
- 2** After you have configured the audit's or the snapshot specification's rules, from the **File** menu, select **Save As**.
- 3** In the Save As window, enter a name and description.
- 4** From the Type list, select Audit Policy.
- 5** Click **OK**. The audit policy is saved and can be accessed at Library ► Audit and Remediation ► Audit Policies.

Running an Audit

Running an audit will execute the selected audit on the target server, servers, or snapshot of the audit, and it will evaluate the targets according to the rules defined in the audit. You can run an audit from the following locations in the SAS Client:



- Running an Audit from the Library
- Running an Audit on a Server from All Managed Servers
- Re-running an Audit from Audit Results

Running an Audit from the Library

The Library contains all available audits that you can run, organized by operating system, either Windows or UNIX. The list of audits in the Library can be sorted by any of the columns (Name, Last Modified Date, and so on). The search tool (upper right of the window) can also be used to search the audit list by entering a name, ID, person who created the audit, and so on.

To run an audit from the Library, perform the following steps:

- 1** From the Navigation pane, select Library ► Audit and Remediation.
- 2** Select Audits, and either Windows or Unix.
- 3** Select the audit you want to run, right-click, and select **Run Audit**.
- 4** In the Run Audit window, step one shows you the name of the audit, the source server or snapshot being used in the Audit, the total number of rules defined in the audit, and all targets of the audit (servers and snapshot). Click **View Rule Details** to view the rule definitions.
- 5** Click **Next**.
- 6** In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 7** Click **Next**.
- 8** In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.

- 9** (Optional) You can specify if you want the email to be sent upon success of the audit job () or failure of the audit job ().
- 10** (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opsware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 11** Click **Next**.
- 12** In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

Running an Ad-Hoc Audit

An ad-hoc or single run audit is an audit that has been created, configured, modified and run, but that has not been saved.

For example, you can create, configure, and run an audit without saving it. This type of audit will audit the target servers and yield results. Such an audit is considered “ad hoc” until you save it. In another case, you could edit an existing audit and then run it without saving it. The audit will be run, and you will get audit results, but the original audit will remain intact (unless you save it).

Additionally, rule exceptions that are applied to ad-hoc audits are ignored.



For more information on audit rule exceptions, see “Audit Rule Exceptions” on page 161.

Running an Audit on a Server from All Managed Servers

You can run an audit from this location, if the server is being used as a target for an audit.

To run an audit from the All Managed Servers list, perform the following steps:

- 1** From the Navigation pane, select Devices and then select All Managed Servers.
- 2** Select a server. From the View drop-down list, select Audit and Remediation. The Details pane area will display below the Content pane.
- 3** From the Details pane Show drop-down list, select Audit - Server is Target.
- 4** Select an audit from the list, right-click, and select **Run Audit**.
- 5** In the Run Audit window, step one shows you the name of the audit, the source server or snapshot being used in the Audit, the total number of rules defined in the audit, and all targets of the audit (servers and snapshot). Click **View Rule Details** to view the rule definitions.

- 6** Click **Next**.
- 7** In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 8** Click **Next**.
- 9** In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 10** (Optional) You can specify if you want the email to be sent upon success of the audit job () or failure of the audit job ().
- 11** (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 12** Click **Next**.
- 13** In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.



Re-running an Audit from Audit Results

You can rerun an audit from an audit results if you would like to run the same audit another time.

Note that when you are viewing the results of an Audit or a Snapshot and re-run the audit from those results, the rules in the original audit may have changed after the results have been captured. Thus it is possible that you will be running the updated audit, and not necessarily the exact audit from which produced these results.

To rerun an audit, perform the following steps:

- 1** From the Navigation pane, select Library and then select Audit and Remediation.
- 2** In the Navigation pane, select Audit Results.
- 3** In the Content pane, select audit results and then select **Actions** ► **Re-Run audit**.
- 4** In the Run Audit window, step one shows you the name of the audit, the source server or snapshot being used in the Audit, the total number of rules defined in the audit, and all targets of the audit (servers and snapshot). Click **View Rule Details** to view the rule definitions.

- 5** Click **Next**.
- 6** In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 7** Click **Next**.
- 8** In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 9** (Optional) You can specify if you want the email to be sent upon success of the audit job () or failure of the audit job ().
- 10** (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 11** Click **Next**.
- 12** In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

Scheduling an Audit

Scheduling an audit requires specifying when you want an audit to be run (either once or as a recurring job) and who you want to receive email notification about the status of the job. You can also view, edit, and delete or cancel existing scheduled audits. When you delete a scheduled audit, all schedules that you have created associated with that audit will also be deleted.



You must have permissions to create, view, edit, and delete audit schedules. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Policy Setter's Guide* for more information.

Scheduling a Recurring Audit

After you have created, configured, and saved an audit, you can set up a schedule that specifies when you want the audit to run on a recurring basis. After the schedule is set, you can edit the schedule according to your needs.

To schedule a recurring audit, perform the following steps:

- 1** From the Navigation pane, select Library and select the By Type tab.
- 2** Select Audit and Remediation, and then select Audits.
- 3** Select an OS (Windows or UNIX) and then open an audit.
- 4** In the Views pane of the Audit window, select the schedule object.
- 5** In the Schedule section, choose to run the audit once, daily, weekly, monthly, or on a custom schedule. Parameters include:
 - **None:** No schedule will be set. To run the audit, select the audit, right-click, and select **Run Audit**.
 - **Daily:** Choose this option to run the audit on a daily basis.
 - **Weekly:** Choose the day or days of the week to run the audit.
 - **Monthly:** Choose the months to run the audit run, and the days of the month.
 - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the audit at 1:00 a.m. every weekday:


```
0 1 * * 1-5
```

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.
- 6** In the Time and Duration section, for each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose a date to end the audit schedule, select End and then choose a date. The Time Zone is set according to the time zone set in your user profile.
- 7** (Optional) Deselect the End option, if you want the audit schedule to run indefinitely.
- 8** To save the audit schedule, from the **File** menu, select **Save**. The audit will now run according to the defined schedule.

Editing an Audit Schedule

You can edit an audit schedule after you have created (or edited) and saved it.

To edit a scheduled audit, perform the following steps:

- 1** From the Navigation pane, select Jobs and Sessions.
- 2** Select Recurring Jobs.
- 3** From the drop-down list at the top of the Content pane, select Run Audit Task.
- 4** Open the scheduled audit job. The Audit Window displays.
- 5** Select the Schedule object in the Views pane to view the audit schedule.
- 6** To edit the audit Schedule, modify the following parameters:
 - **None:** No schedule will be set. To run the audit, select the audit, right-click, and select **Run Audit**.
 - **Daily:** Choose this option to run the audit on a daily basis.
 - **Weekly:** Choose the day or days of the week to run the audit.
 - **Monthly:** Choose the months to run the audit run, and the days of the month.
 - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the audit at 1:00 a.m. every weekday:

```
0 1 * * 1-5
```

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.
- 7** In the Time and Duration section, for each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose a date to end the audit schedule, select End and then choose a date. The Time Zone is set according to the time zone set in your user profile.
- 8** (Optional) Deselect the End option, if you want the audit schedule to run indefinitely.
- 9** To save the audit schedule, from the **File** menu, select **Save**. The audit will now run according to the defined schedule.

Viewing a Completed Audit Job

To view information on a completed audit job, perform the following steps:

- 1** From the Navigation pane, select Jobs and Sessions.
- 2** Select Job Logs.
- 3** The Content pane displays all jobs run in this Opsware core. To display only audit jobs, from the drop-down list at the top of the Content pane, select Run Audit Task. If you want to see only your scheduled audits, enter your user ID in the User ID field at the top of the Content pane.
- 4** Open an audit job to view the audit results, and then click **View Results**.

Viewing and Remediating Audit Results

An audit defines the server object configurations that you want to check on a server (according to the audit's defined rules); audit results are the end product of running an audit and show any differences between the audit rules and the actual server configuration values for each target server or snapshot.

The type of audit result and remediation that you can perform depends on the types of audit rule. You can view and remediate to types of audit results: server comparison or value-based.



If you have Audit Results with differences from Audits that were created in SAS 5.1, and you have upgraded to SAS 6.x, when you view those Audit Results in the upgraded version of the SAS Client, the Differences column in the Audit Results list will incorrectly display the value of -1 differences. To view the actual number of results, simply open the Audit Results window (double-click it) and you will see all the actual differences in the results.

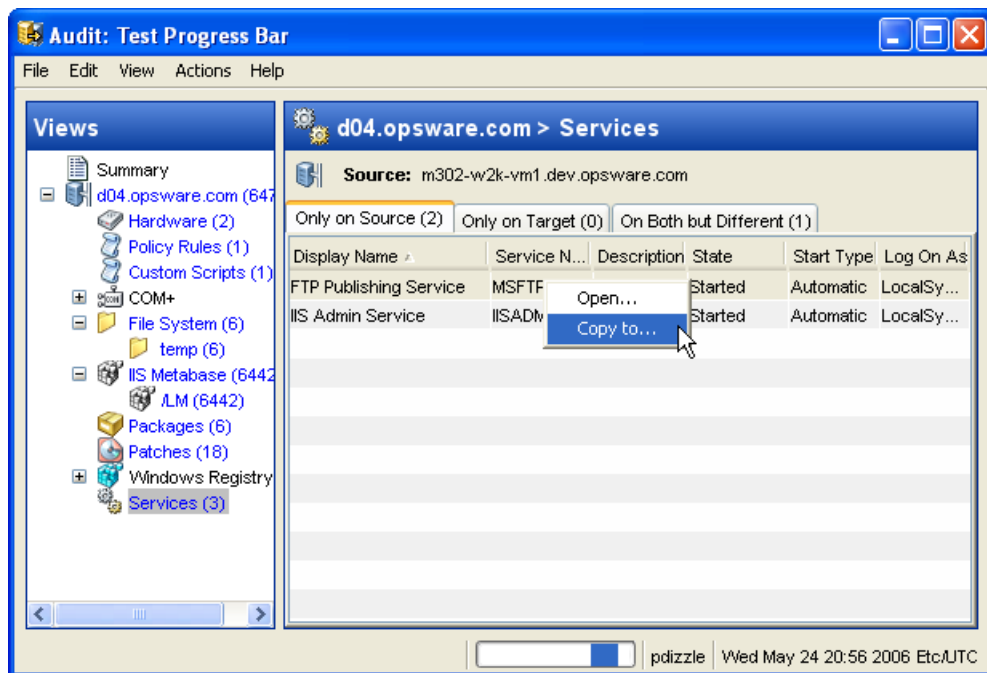
Viewing Comparison-Based Audit Results: “Copy To” Remediation

Audit results based on a server comparison audit allow you to view differences between the source server (or snapshot) and target servers or snapshot. If the audit results fails – that is, finds differences between source and target – you can remediate the differences. You can copy the rule values of the source objects in the audit and overwrite the values on the target (or add values that exist on the source, but do not exist on the target.)

For example, Figure 2-11 shows audit results for a windows services rule, the settings for a specific service (FTP Publishing) on the source do not exist on the target server, located under the Only On Source tab of the Audit Results Window.

From the Audit Results Window, you can select the Service rule, right-click, and choose Copy To, and the values from the rule will be remediated on the target server.

Figure 2-11: Audit Results For a Comparison-Based Audit Rule



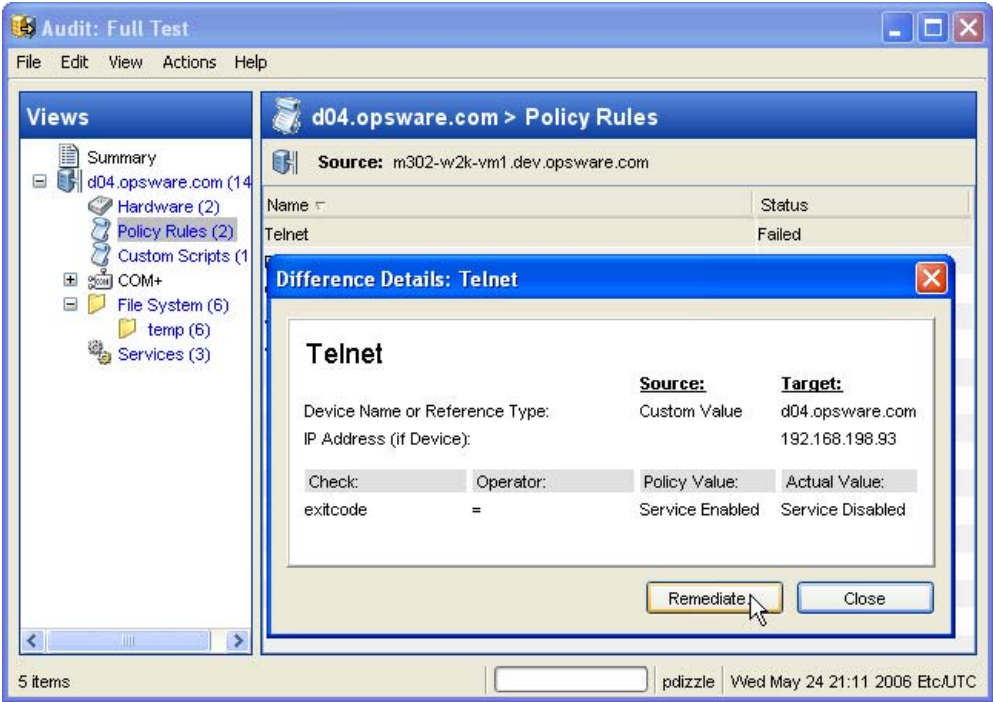
The Audit Results window shows all the objects defined in the audit in the Views pane. It also shows the audit results that failed, the differences found between the audit and the target servers are highlighted in light blue font.

Viewing Value-Based Audit Results - Audit Rule Remediation

Value-based audit results present the values returned from the target server. You can view the differences between what was defined as the expected value in the rule and the actual value found on the target server. When applicable, you can choose to remediate the value found on the target server. If you remediate then the values specified in the remediation area will be applied to the server object on the target server.

Figure 2-12 shows a value-based audit rule. In it, a windows services rule checks to see if the Telnet service is enabled on the target server. The audit results show a status of failed for the audit rule, which means that the rule checked to see if the service was running on the target, and it was found to be disabled. You can choose to Remediate the service. To do so, click the **Remediate** button. In this example, this will enable the service.

Figure 2-12: Audit Results for a Value-Based Audit rule



The Difference Details window displays the rule value compared with the actual value found on the target server, and allows you to remediate the differences.

Viewing and Remediating Differences of Audit Results Objects

For some server objects in an audit result, if the object exists on both the target and the source, and there are differences between them, then you can view those differences. You can also learn more what is different about them – and remediate them, if necessary. For some objects, you can view general differences, such as a service's status, the release number for a patch, a registry key's value, and so on. For other server objects, such as files, you can view the differences of the file's contents.

Viewing and Remediating File Audit Results Differences

For file system rules that were audited, you can view file content differences side by side and line by line. You can see the lines in a file that were added, deleted, or modified. If you want to remediate the results, you can choose to copy any files from the source server to the target.

To view and remediate contents of two files that differ in an audit, perform the following steps:

- 1** From the Navigation pane, open an Audit Results Window that has file system objects by selecting Library and then by selecting the By Type tab.
- 2** Select ► Audit and Remediation ► Audit Results.
- 3** In the Views pane, expand one of the target servers and select a result.
- 4** In the Content pane, expand a target server and select one of the results.
- 5** Next, in the Content pane, select the On Both but Different tab.
- 6** Select a file, right-click, and select View Differences.
- 7** In the Comparison window, select an item from the Encoding drop-down list to specify the character encoding of the data displayed.



If one of the files you are comparing exceeds 2MB in file size, Opware Audit and Remediation cannot display the file differences.

- 8** Click the arrows to find the first, next, previous, or last lines that were added, deleted, or modified. Differences are highlighted according to the following color scheme:
 - **Green:** This content was added.
 - **Blue:** This content was modified.

- **Red:** This content was deleted.
- **Black:** No changes were made to this content.

9 Click **Close** to close this window.

10 To remediate file differences, from inside the Audit Results window, select either the the Only On Source tab or On Both But Different tab, select a file, right-click and select Copy To.

11 In the Select Server window, select a server you want to copy the file from the source to, and then click **OK**.

Viewing and Remediating Server Object Audit Results Differences

For many server objects, when there are differences between the source object and the target object, you can view differences in object properties side by side. Each server object will show different windows, depending on the object and if the audit rule set was comparison-based (comparison between source and target) or value-based (comparison between user-defined audit rule and target).

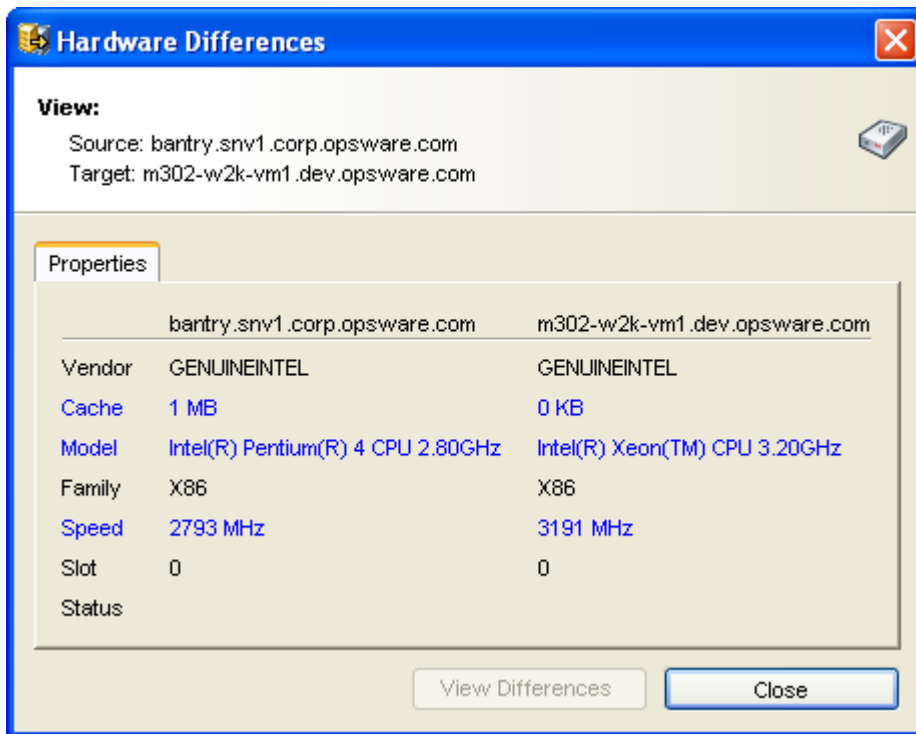
For some value-based audit rules, you can remediate the values on the target server.

To view the contents of two objects that differ, perform the following steps:

- 1** From the Navigation pane, open an Audit Results Window that has file system objects by selecting Library and then select the By Type tab.
- 2** Select ► Audit and Remediation ► Audit Results.
- 3** In the Views pane, expand one of the target servers and select a result.
- 4** In the Views pane, select an object.
- 5** In the Content pane, select the On Both but Different tab.
- 6** In the Content pane, select an object, right-click, and select **Open**. You will see a window that shows the differences between the object as defined the audit and the object on the target server.

The example in Figure 2-13 displays the audit Result differences for a CPU hardware audit rule (comparison-based rule) where the hardware on the target is shown to be different from the hardware on the source of the audit.


Figure 2-13: Comparison-Based Audit Results Difference: CPU



The window shows the source server information on the left, and the target server information on the right, with all the differences between the two CPUs listed in blue font. This type of audit rule cannot be remediated. For a value-based rule, the difference window will be slightly different and will also include a Remediate option, if

remediation is possible. This difference window displays the audit rule, including the policy value and the actual value found on the target server. The example in Figure 2-14 shows the differences for a value-based Windows Registry rule.

Figure 2-14: Rule-Based Audit Results Difference: Windows Registry



Since there is a difference between the two, click **Remediate** to make the target's actual value match the policy value.


- 7** To remediate the difference, click **Remediate**.
- 8** In the Remediate window, enter an optional Job ID, and then click **Remediate**.

Viewing Audit Results with Exceptions

If an audit contains rule exceptions, then the excepted rules are not applied to the targets of the audits. However, your audit results will show which of the rules in the audits are exceptions, including details about the rule exceptions.

The manner in which rule exceptions are displayed in audit results depends on the type of rule that has been excepted:

- Custom script and custom or pluggable check rule exceptions (such as those created by developers or provided by a TON Content Subscription) appear in the Contents pane of the Audit Results window. You can double-click the rule exception for details on the exception.
- All other rule exceptions, such as file system, registry settings, services, IIS Metabase, and COM+ rules, the Audit Results window will display an Exceptions icon in the Views

pane. You can select the Exceptions  icon and see the details of the exception in the Contents pane.

Searching for Audits

You can use the SAS Client Search tool to find audits in your facility. You can search for audits by name, by the operating system, and many other criteria.

To search for audits, perform the following steps:

- 1** From inside the SAS Client, ensure that the search pane is activated by selecting View ► Search pane.
- 2** From the top drop-down list, select Audit.
- 3** Click the green arrow button or ENTER to execute the search.
- 4** The results appear in the Content pane.
- 5** If you want to extend your search criteria, add new criteria in the search parameters section at the top of the Content pane. You can also save the search by clicking **Save**, or export the Search results to .html or .csv.

Understanding Snapshots

A snapshot captures the configuration of a managed server at a particular point in time, and provides a means of capturing the current state of a known working (or, not working) server. A snapshot is useful for capturing a server configuration that you know represents a desired state of configuration. You can compare the snapshot with other servers in your facility by using the snapshot in an audit.

A snapshot is also a useful way to back up a managed server, especially if you plan to make changes to the server and want to keep a record of it before you change anything.

In addition to recording information about objects on managed servers, a snapshot can contain the content of some objects. A server snapshot also identifies attributes of other objects on specific types of operating systems, such as the Windows Registry and Windows Services, application configurations, COM+ objects, hardware information, installed patches. You can even create custom scripts that gather data from the target managed servers.

Snapshot Specification and Snapshot

Snapshots are configured in similar way as you configure an audit. First you create a *snapshot specification*, which is like a template that defines exactly what you want to capture of a server's configuration. Then, you configure the snapshot specification's rules, and then run it. The results are a snapshot – a picture of a server's configuration. The main difference between a snapshot and an audit is that a snapshot takes a picture of a server's configuration, whereas an audit compares a server configuration with the rule values that you define.

You can schedule when you want a snapshot to be created (either once or as a recurring job) and who you want to receive email notification about the status of the job.

Snapshot Used in an Audit

You can use a snapshot in an audit to compare managed servers, groups of servers, and snapshots. By using a snapshot in an audit, you can compare a problematic server (target of the audit) with a known working server (snapshot as source for the audit). To further extend the audit definition, you can also define rules for server objects.

When a snapshot is used as the source for an audit, all server configuration values captured in the snapshot results are available to use as rules for the audit. For more information about using a snapshot in an audit, see “Configuring an Audit” on page 118.

Audit Policies and Snapshot Specification

An audit policy is collection of rules that defines a desired state of a server's configuration. An audit policy can be used inside a snapshot specification, either through linking or importing. An audit policy is useful because it allows a policy setter to define server configuration compliance values. These can then be used by others in the context of a snapshot specification.

An audit policy can be linked to an audit or snapshot specification, so whenever a change is made to the policy, the audit or snapshot specification using the policy will also reflect the latest changes. Or, an audit policy can be imported into a snapshot specification, without keeping the link to the source audit policy. When you import an audit policy into a snapshot specification, you can choose to replace any current values in the audit or merge values from the audit policy with those in the snapshot specification.

For more information on importing or linking an audit policy to a snapshot specification, see “Linking and Importing Audit Policies” on page 166.

Snapshot Specification Elements

An snapshot specification consists of the following elements:

- **Properties:** The name and description of the snapshot specification.
- **Targets:** The servers that you want to take a snapshot of – that is, capture the specific server configuration as defined in the snapshot specification's rules. You can choose as many servers and groups of servers as you want.
- **Source:** The source of a snapshot specification. If you choose a server then you can select server objects from that server as the basis of your snapshot. The source of a snapshot specification can be a server, or no source at all. (Some rules require a source server. Other rules can be defined by your own custom values without a source.)

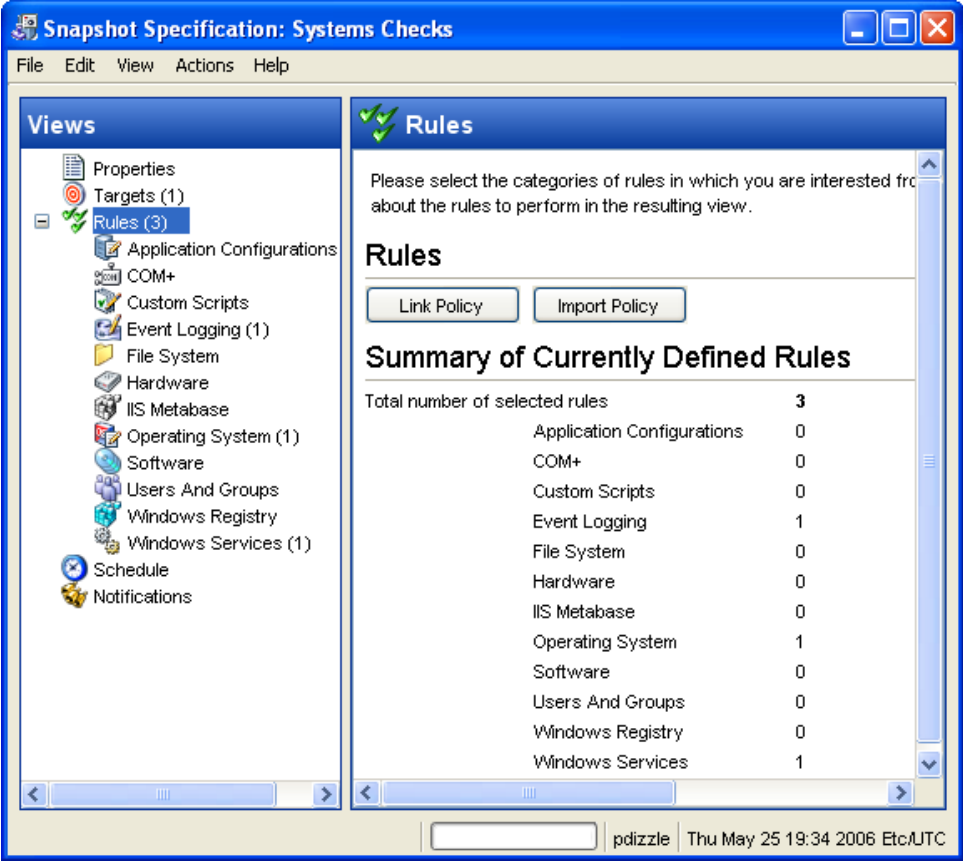
Note that the value of a source parameter is not used when taking a snapshot. It only has meaning when defining a snapshot specification.

- **Rules:** A check on a particular server object with a desired value and an optional remediation value. For example, you might check if a server contains a specific Windows Service, and if found, determine if the service is turned off. For a description of server objects that you can define rules for in a snapshot specification, see "Configuring Opsware Audit and Remediation Rules" on page 124.
- **Schedule:** The time the snapshot will run. You can run the snapshot specification as a job on a onetime basis, or on a recurring schedule.
- **Notifications:** The email notification send after the snapshot has run. You can base the notification on success, failure, or simply the completion of the snapshot specification job.

When you set up a snapshot specification, you select the objects to check for on the target server. You can also apply rules to these objects that define their desired configuration state. For some rules, you can define remediation values, in the event that the resulting snapshot is used as the source for an audit.

Figure 2-15 shows a snapshot specification that has three rules that will capture configuration information about the target server for event logging, operating system, and windows services.

Figure 2-15: Snapshot Specification Elements



The Snapshot Process

Taking a snapshot of a server configuration requires the two following basic steps:

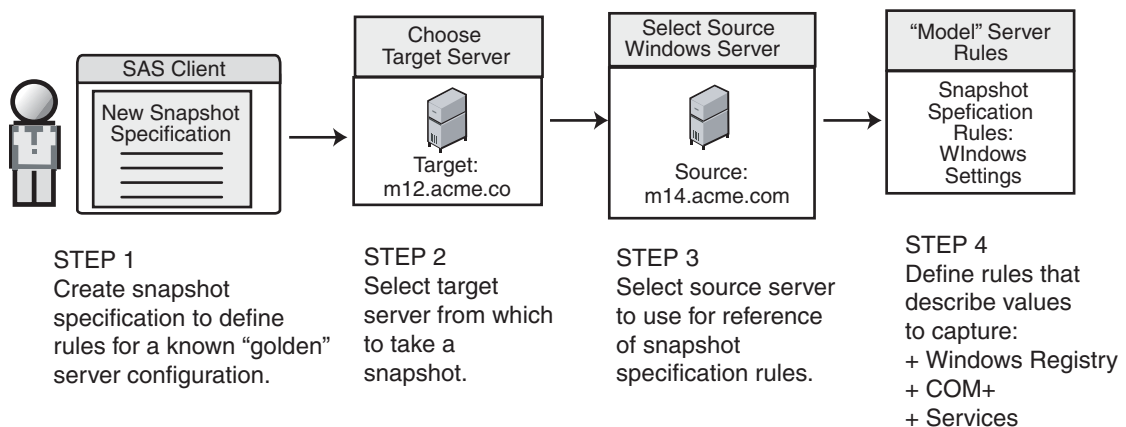
- Creating a snapshot specification, which is a template that defines the configuration parameters captured on a target server.
- Running the snapshot specification job that results in a snapshot.

Figure 2-16 illustrates an example of the snapshot process.

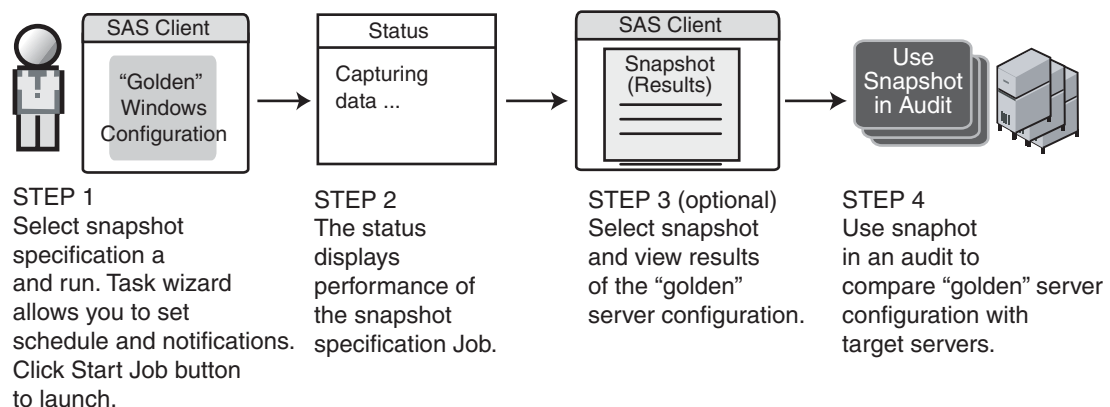
Figure 2-16: Snapshot Process

SNAPSHOT PROCESS - Windows Server Snapshot

Part A: Create Snapshot Specification to define “Golden” Server Configuration



Part B: Run Snapshot Specification Job and View Results in the Snapshot



Creating a Snapshot Specification

You can create a snapshot specification from two different locations inside the SAS Client, depending on your purpose. You can create a snapshot specification from the following locations inside the SAS Client:

- Creating a Snapshot Specification from a Server
- Creating a Snapshot Specification from the Library



You must have a set of permissions to create and modify snapshot specifications. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Policy Setter's Guide* for more information.

Creating a Snapshot Specification from a Server

When you create a new snapshot specification from a managed server, the snapshot specification will use the selected server as its source. You can choose several different server sources for the snapshot specification as you define the rules, or choose no source at all and define your own custom rules. Some rules, however, require a source.



To take a snapshot of a managed server, the server must be reachable and you must have access to the server.

To create a snapshot specification from a server, perform the following steps:

- 1** From the Navigation pane, select Devices and then select All Managed Servers.
- 2** Select a server, then select **Actions** ► **Create Snapshot Specification**.

Creating a Snapshot Specification from the Library

If you want to create a new snapshot specification and set all your own rules, create the audit from the SAS Client Library by performing the following steps:

- 1** From the Navigation pane, select Library and then select Audit and Remediation.
- 2** In the Navigation pane, select snapshot specifications, then Windows or Unix.

Configuring a Snapshot Specification

To configure a snapshot specification, performing the following tasks:

- Name and describe the snapshot specification.
- Choose target servers you want to take a snapshot of. You can choose to snapshot multiple servers or groups of servers.
- Configure your own custom rules, or choose settings from a source server to serve as the basis for the snapshot specification rules.
- Schedule the snapshot specification job to run once or on a recurring schedule.
- Set up email notifications to notify users when the snapshot specification job finishes successfully, if the job fails, or on both conditions.
- Save the snapshot specification.

Configuring a Snapshot Specification

To configure a snapshot specification, perform the following steps:

- 1** Create the new snapshot specification from one of the methods listed in “Creating a Snapshot Specification” on page 187.
- 2** In the snapshot specification Window, define the parameters by entering the following information:
 - **Properties:** Enter a name and description for the snapshot specification.
 - **Source:** Select a source for the snapshot specification. By default, the source server for the snapshot specification will be the managed server that you chose as the source for the snapshot specification. Browse the source server for values to populate the snapshot specification’s rules. You can also choose a different source server as the basis of the snapshot specification for each rule category, or no source at all. If you choose no source, you must define your own rules, or choose to link to an audit policy in the rules section.
 - **Rules:** Choose a rule category from the list to begin configuring your snapshot specification’s rules. Since each rule is unique and requires its own instructions, to configure specific rules, see “Configuring Opsware Audit and Remediation Rules” on page 124.

If you want to use an audit policy to define the rules of your snapshot specification, click either **Link Policy** or **Import Policy**. When you link an audit policy, the snapshot specification maintains a direct connection with the audit policy, so if any

changes are made to the policy, the snapshot specification will update it with the new changes. If you import an audit policy, the snapshot specification will use all the rules defined in the policy but will not maintain a link to the audit policy. For information on how to import or link to a snapshot specification, see “Linking and Importing Audit Policies” on page 166.

- **Targets:** Choose the Targets of the snapshot specification. These are servers or groups of servers that you want the configured snapshot specification rules to capture. To add a server or group of servers, click **Add**. To choose a source server to use to create the snapshot specification rules, click **Select**.
- **Schedule:** Choose to run the snapshot specification immediately, or on a recurring schedule. Choose whether you want to run it once, daily, weekly, monthly, or on a custom schedule. Parameters include:
 - **None:** No schedule will be set. To run the snapshot specification, select the snapshot specification, right-click, and select **Run snapshot specification**.
 - **Daily:** Choose this option to run the snapshot specification on a daily basis.
 - **Weekly:** Choose a day of the week to run the snapshot specification.
 - **Monthly:** Choose the months to run the snapshot specification.
 - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

```
0 0 * * 1-5
```

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.
 - **Time and Duration:** For each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the snapshot specification will keep running indefinitely. To choose an end date to end the snapshot specification schedule, select End, and from the calendar selector, choose a date. The Time Zone is set according to the time zone set in your user profile.

- **Notifications:** Enter the email addresses (separated by a comma or a space) of those you want to receive an email when the snapshot specification Job finishes running. You can choose to send the email notification on both success and the failure of the snapshot specification job (not the success of the audit rules). To add an email address, click **Add Notification Rule**.

- 3** When you have finished configuring the snapshot specification, from the **File** menu, select **Save**.



To prevent runaway processes, the snapshot process will time-out if it exceeds 60 minutes or if the data that is collected from a managed server exceeds one gigabyte (GB). If you specify that you want to collect the full contents of files in the selection criteria, the data collected might exceed the maximum size that can be successfully recorded in a snapshot.

Configuring Snapshot Specification Rules

For information on how to configure specific snapshot specification rules, see “Configuring Opware Audit and Remediation Rules” on page 124.

Saving a Snapshot as an Audit Policy

You can save selection criteria and use it in other snapshot specifications.

To save your snapshot selection criteria, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select Audit and Remediation.
- 2** From the Content pane, select the snapshot specifications tab.
- 3** Select the template that you want to make a copy of, and then open it.
- 4** In the Snapshot Specification window, select **Actions** ► **Save Selection Criteria As**.
- 5** In the Save Selection Criteria As window, enter a unique name.
- 6** (Optional) Enter a description of the Selection Criteria.
- 7** Click **Save** to save your Selection Criteria or click **Cancel** to close this window without saving your changes.

Deleting a Snapshot Specification

To conserve disk space, you should delete snapshot specifications that you no longer need from the Model Repository.

To delete an snapshot specification, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select Audit and Remediation.
- 2** From the Content pane, select the snapshot specifications tab.
- 3** Select one or more templates and then select **Actions** ► **Delete**.
- 4** In the Confirmation Dialog, click **Yes** to delete this snapshot specification, or click **No** if you do not want to delete it.



When you delete a snapshot specification, you do not delete any of the snapshots that were created from it. However, when you delete a snapshot specification, all schedules associated with that snapshot specification, will be deleted. See “Scheduling Snapshot Jobs” on page 192 in this chapter for more information.

Running a Snapshot Specification

When you run a snapshot specification, it captures from the target servers all configuration parameters configured in the rules. After you run a snapshot specification, the results of the snapshot job become a snapshot and can be viewed inside the snapshot.



To run a snapshot specification, perform the following steps:

- 1** From the Navigation pane, select Library and then select Audit and Remediation.
- 2** In the Navigation pane, select snapshot specifications, then Windows or Unix.
- 3** Select a snapshot specification, right-click, and select **Run Snapshot Specification**.
- 4** In the Run Snapshot Specification window, step one shows you the name of the snapshot, the total number of rules defined, and all targets). Click **View Rule Details** to view the rule definitions.
- 5** Click **Next**.

6 In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.

7 Click **Next**.

8 In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.

9 (Optional) You can specify if you want the email to be sent on success of the audit job () or failure of the audit job ().

10 (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.

11 Click **Next**.

12 In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

Scheduling Snapshot Jobs

A snapshot specification job enables you to specify when you want the SAS Client to create a snapshot (either once or on a recurring basis) and who you want to receive email notification about the status of the job. You can also view, edit, and delete existing snapshot specification schedules. When you delete a snapshot specification, all schedules associated with that snapshot specification will be deleted.

This section discusses the following topics:

- Scheduling a Recurring Snapshot Job
- Editing an Snapshot Job Schedule
- Viewing a Snapshot Job Schedule
- Deleting a Snapshot Job Schedule

Scheduling a Recurring Snapshot Job

After you have created, configured, and saved an snapshot specification, you can schedule snapshot specification a recurring snapshot job. After the schedule is set, you can edit the schedule according to your needs.

To schedule a recurring snapshot specification, perform the following steps:

- 1** From the Navigation pane, select Library and then Audit and Remediation.
- 2** Select a snapshot specification, select an OS (Windows or Unix), and then open it.
- 3** In the Snapshot Specification window, in the Views pane, select the schedule object.
- 4** In the Schedule section, choose to run the snapshot job immediately or on a recurring schedule. Choose to run it once, daily, weekly, monthly, or on a custom schedule. Parameters include:
 - **None:** No schedule will be set. To run the snapshot job, select the snapshot specification, right-click, and select **Run Audit**.
 - **Daily:** Choose to run the snapshot job on a daily basis.
 - **Weekly:** Choose a day of the week to run the snapshot specification job.
 - **Monthly:** Choose the months to run the snapshot specification job.
 - **Custom:** In the Custom Crontab string field, enter a string the indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

```
0 0 * * 1-5
```

An asterisk (*) in any of these fields represent all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.
- In the Time and Duration section, for each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the snapshot specification job will keep running indefinitely. To choose an end date to end the audit schedule, select End, and then choose an end date. The Time Zone is set according to the time zone set in your user profile.
- (Optional) Deselect the End option if you want the snapshot specification job to run indefinitely.

- 5 To save the snapshot specification job schedule, from the **File** menu select **Save**. The snapshot specification will now run according to the defined schedule.

Editing an Snapshot Job Schedule

You can edit a snapshot specification schedule after you have created (or edited) and saved it.

To edit a scheduled snapshot specification, perform the following steps:

- 1 From the Navigation pane, select Jobs and Sessions.
- 2 Select Recurring Jobs.
- 3 From the drop-down list at the top of the Content pane, select Run Snapshot Task.
- 4 Open the snapshot specification job. The Snapshot Specification window will open.
- 5 Select the Schedule object in the Views pane to view the snapshot specification job schedule.
- 6 To edit the snapshot specification job schedule, modify the following parameters:
 - **Schedule:** Choose to run the snapshot specification immediately, or on a recurring schedule. Choose to run it once, daily, weekly, monthly, or on a custom schedule. Parameters include:
 - **None:** No schedule will be set. To run the snapshot specification, select the snapshot specification, right-click, and select **Run snapshot specification**.
 - **Daily:** Choose to run the snapshot job on a daily basis.
 - **Weekly:** Choose the day of the week you want the snapshot job to run.
 - **Monthly:** Choose the months to run snapshot specification job.
 - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

```
0 0 * * 1-5
```

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.

- **Time and Duration:** For each type of schedule, specify the hour and minute, the day of the week (and month) you want the daily schedule to start. Unless you specify an end time, the snapshot specification job will keep running indefinitely. To choose a date to end the snapshot specification job schedule, select End and then choose a date. The Time Zone is set according to the time zone set in your user profile.
- (Optional) Deselect the End option if you want the snapshot specification schedule to run indefinitely.

- 7** To save the snapshot specification schedule, from the **File** menu select **Save**. The snapshot job will now run according to the defined schedule.

Viewing a Snapshot Job Schedule

To view information on a completed snapshot job, perform the following steps:

- 1** From the Navigation pane, select Jobs and Sessions.
- 2** Select Job Logs.
- 3** The Content pane displays all jobs that have been run on this Opsware core. To display only snapshot specification jobs, from the drop-down list at the top of the Content pane, select Run Snapshot Task. If you want to see only those snapshot specifications that you have scheduled or run, enter your user ID in the User ID field at the top of the Content pane.
- 4** Open a completed snapshot job. If you want to view the snapshot job schedule, select it, right-click, and select **Open**.

Deleting a Snapshot Job Schedule

To delete a snapshot job schedule, perform the following steps:

- 1** From the Navigation pane, select Jobs and Sessions.
- 2** Select Job Logs.
- 3** The Content pane displays all jobs that have been run on this Opsware core. To display only snapshot specification jobs, from the drop-down list at the top of the Content pane, select Run Snapshot Task. If you want to see only those snapshot specifications that you have scheduled or run, enter your user ID in the User ID field at the top of the Content pane.
- 4** To delete the schedule, select it, right-click, and select **Delete Schedule**.

Locating Snapshots

After you have created a snapshot, you can find it in several locations inside the SAS Client.

In the Library:

- 1** From the Navigation pane, select Library, then select the By Type tab.
- 2** Select Audit and Remediation ► Snapshots.
- 3** Select a snapshot and then open it.

In Jobs and Sessions:

- 1** From the Navigation pane, select Jobs and Sessions and then select Job Logs.
- 2** In the Content pane, select Run Snapshot Task from the Job Types drop-down list.
- 3** Select a snapshot task job in the list and then open it.
- 4** Wait until the job loads, and then select a server.
- 5** Click **View Results** to view the snapshot.

In the Server Explorer:

- 1** From the Navigation pane, select Devices and then All Managed Servers.
- 2** Select a server from the Content pane.
- 3** Select a server and then open it.
- 4** In the Server Explorer window, from the View pane, select Opsware Audit and Remediation.
- 5** In the Content pane, from the Show drop-down list, select snapshots. This shows a list of all snapshots taken on this server.
- 6** To view a snapshot, open it.

Searching for Snapshots

You can use the SAS Client Search tool to find snapshots in your facility. You can search for snapshots by name, by the operating system, and many other criteria.

To search for snapshots, perform the following steps:

- 1** From inside the SAS Client, ensure that the search pane is activated by selecting View ► Search pane.

- 2 From the top drop down list, select Snapshot.
- 3 Click the green arrow button or ENTER to execute the search. The results appear in the Content pane. If you want to extend your search criteria, you can add new criteria in the search parameters section at the top of the Content pane. You can also save the search by clicking **Save**, or export the Search results to .html or .csv.

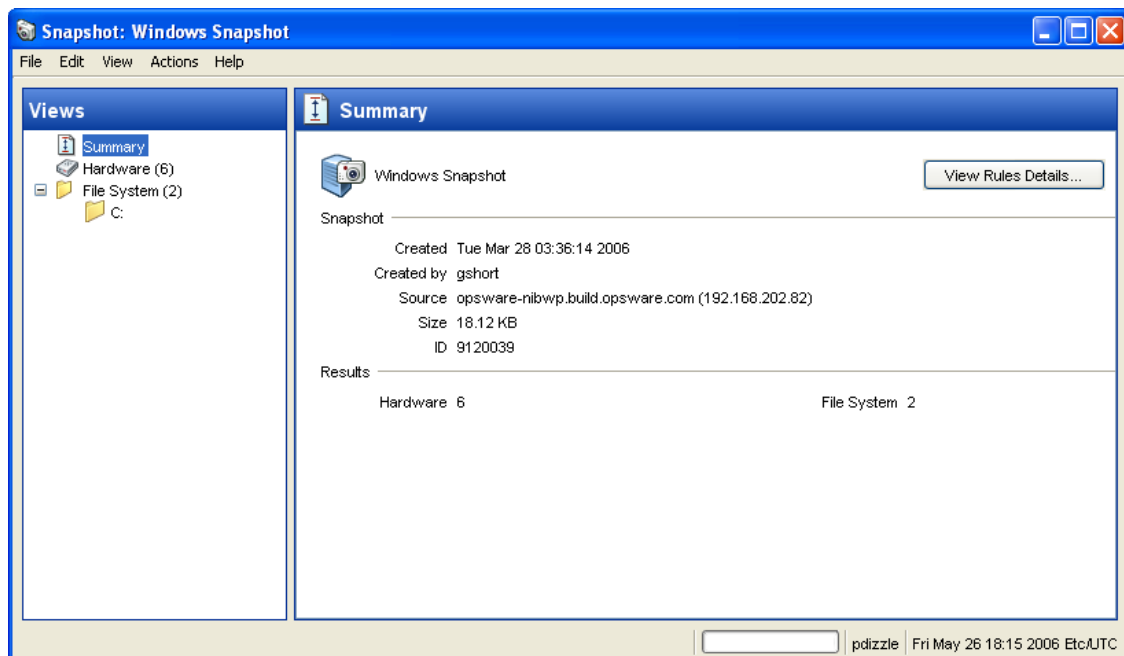
Viewing Snapshot Contents

You can view the contents of a snapshot and view detailed information about the server objects that were recorded.

To view the contents of a snapshot, perform the following steps:

- 1 From one of the starting points described in “Locating Snapshots” on page 196, open a snapshot.

Figure 2-17: Sample Snapshot Browser of a Windows Server



- 2 In the snapshot window, you can select:
 - **Summary:** Displays general information about a snapshot, such as the date and time the snapshot was created and by whom, the snapshot source (name of the managed server), the size of the snapshot file, and a snapshot ID number.

You can also click **View Rules Details** to see the snapshot specification, which this snapshot is based on.

- **Installed Hardware:** Information about the type of CPU processor and speed, cache size, memory size for SWAP and RAM, and storage devices that were recorded in the snapshot.
- **Installed Patches:** Displays information about the installed patches that were recorded in the snapshot, such as the patch type.
- **Installed Packages:** Displays information about the installed packages that were recorded in the snapshot, such as package type, package version, and release number.
- **Event Logging:** Displays security, application, and system log files recorded in the snapshot.
- **File System:** Displays the directories, file properties, attributes, and contents of the files recorded in the snapshot.



If a file in the snapshot exceeds 2MB in file size, Opware Audit and Remediation cannot display the file contents.

- **Windows Services:** Displays information about the running services recorded in a snapshot, such as the name, description, startup state, startup type, and log on account.
- **Windows Registry:** Displays information about Windows Registry entries in the snapshot, such as the registry key, registry value, and subkey. A registry key is a directory that contains registry values, where registry values are similar to files within a directory. A subkey is similar to a subdirectory. The content area in this window excludes subkeys. Opware Audit and Remediation supports the following Windows Registry keys: HKEY_CLASSES_ROOT, HKEY_CURRENT_CONFIG, HKEY_LOCAL_MACHINE, and HKEY_USERS.
- **COM+:** Displays information about Windows COM (Component Object Model) objects in the snapshot, such as the name and GUID (Globally Unique Identifier) of the object, and the path to the in-process server DLL.

Opware SAS provides warning messages that explain how Windows COM folders were processed. The following scenarios apply:

- When you create a snapshot and select a Windows COM folder that does not contain any objects, the snapshot window displays a summary. Opsware SAS displays a warning that the GUID (Globally Unique Identifier) for that folder is invalid, which means that the Windows COM folder does not contain any objects.
- When you create a snapshot specification and select a Windows COM+ object that does not exist on a target, Opsware SAS displays a warning that the folder is invalid.
- When you create a snapshot and select a Windows COM+ folder that does not contain any objects, Opsware SAS displays a warning that the folder is empty.
- **Metabase:** Displays information about IIS Metabase objects in the snapshot, such as the ID, name, path, attributes, and data of the object.
- **Custom Scripts:** Displays information about the custom script rule recorded in the snapshot.
- **Users and Groups:** Displays information about users and groups on servers, such as user name for last login, whether or not CTRL + ALT + DELETE is enabled, and so on.

3 Click **Close** to close the object browser.

Detaching a Snapshot From a Server

A snapshot is typically associated with the server (source) that it was generated from. If a server is going to be decommissioned, then all snapshots associated with the server will also be deleted. If you need to keep the snapshot but decommission its associated server, you can detach the snapshot from the server before you decommission the server.

To detach a snapshot from a server, perform the following steps:

- 1** From one of the starting points described in “Locating Snapshots” on page 196, select a snapshot.
- 2** Right-click and select **Actions** ► **Detach Snapshot**.
- 3** Click **OK** to save the snapshot in the Software Repository. After you save the snapshot, a general snapshot icon replaces the server snapshot icon.



When you decommission a managed server, all snapshots associated with that server will be deleted from the Software Repository.

Deleting a Snapshot

As a best practice, you should delete snapshots that you no longer need from the Software Repository to conserve disk space.



You must have read permissions for the snapshot to be able to delete it. To obtain these permissions, contact your Opware administrator. See the Opware® SAS Configuration Guide for more information.

To delete snapshots, perform the following steps:

- 1** Select a snapshot or select multiple snapshots and then select **Actions ► Delete**.
- 2** In the Confirmation Dialog, click **Yes** to delete this snapshot or click **No** if you do not want to delete it.



When you delete a snapshot, you do not delete the snapshot specification that was used to create it. See “Deleting a Snapshot Specification” on page 191 in this chapter for more information.

Copying Objects from a Snapshot to a Server

After viewing snapshot contents, you can copy certain objects to a target server. Opware Audit and Remediation allows you to copy directories, files, windows services (state only), IIS Metabase objects, COM+ objects and categories, and Windows Registry keys to a managed server.



In order to copy COM+ rule snapshot results from a snapshot to a server, you must have selected the Archive all associated files option when you configured the COM+ rule. Also the COM+ object being copied must not be in use by any application in order for the copy to remediation to work. For more information, see “Configuring COM+ Rules” on page 138.

Before you copy these objects over to a managed server, it is important to understand what actually gets copied to or created on the destination server:

- When you select a directory, only the directory will be copied to the destination server, excluding any files in that directory. For example, if dir1 contains file1 and file2, and you select dir1, Opsware Audit and Remediation copies only dir1 (not file1 and file2) to the destination server.
- When you select a file and its parent directory does not exist on the destination server, Opsware Audit and Remediation will create the directory on and copy the files to the destination server. For example, if you select file1 and dir1 does not exist on the destination server, Opsware Audit and Remediation will create dir1 on and copy file1 to the destination server.
- When you copy a Windows Services object, you copy the state of the service, such as started, stopped, paused, and so on. You can select one or more Windows Services objects for a single copy process.
- When you copy a Windows Registry object, you can select one or more registry keys and subkeys for a single copy process.
- ACLs are not copied along with COM+ objects or Microsoft IIS objects to the target server.



You must have write permission on the destination server to be able to copy an object to it. To obtain these permissions, contact your Opsware administrator. See the *Opsware[®] SAS Policy Setter's Guide* for more information.

Copying Objects to a Server from a Snapshot

To copy an object from a snapshot to a managed server, perform the following tasks:

- 1** From one of the starting points described in “Locating Snapshots” on page 196, open a snapshot.
- 2** In the Views pane, select a file system, Windows Services, or Windows Registry object.
- 3** In the Content pane, select one or more objects that you want to copy.
- 4** Select **Actions** ► **Copy To**.
- 5** In the Select Server window, select a destination server.



Use the search tool to dynamically filter this list by entering a server name, IP address, or operating system.

- 6 Click **Select** to copy the object to that managed server or click **Cancel** to close this window without saving your changes.



For other types of server objects, such as packages and patches, you can also create installable packages to update a destination server. See “Visual Packager” on page 245 in Chapter 5 for more information.

Chapter 3: Compliance Dashboard

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Compliance Dashboard
- Compliance Dashboard Terms and Concepts
- Compliance Dashboard Remediation
- Software Compliance
- Application Configuration Compliance
- Patch Compliance
- Audit Compliance
- Duplex Compliance
- Filtering and Sorting Compliance Dashboard Information
- Exporting the Compliance Dashboard

Overview of Compliance Dashboard

The Compliance Dashboard allows you to view at a glance the overall compliance levels for all the devices in your facility and helps you to remediate compliance problems. The Compliance Dashboard also displays compliance tests for software policies, application configurations, audits, patches, and duplex status – it enables you to create your own individual audit tests as well. Each of these compliance tests is based upon an Opware Server Automation System (SAS) “policy” (user or system defined) which defines a unique set up server or device configuration settings or values that help ensure that your IT environment is configured as it should be.

Generally speaking, a server or device is “compliant” if its actual configuration matches the configuration defined by a policy setter (or by the system) in the Opware SAS Web Client. For example, a policy setter can create a software policy that defines specific patches and packages that should be installed and how specific applications should be

configured on a server. The Compliance Dashboard shows you if the server's actual installed software and configuration settings match the configuration defined in the software policy. If everything defined in the policy matches what is configured on the server, then the Compliance Dashboard will show a green icon in the software policy column indicating full compliance. If the server configuration mismatches the policy, then the Compliance Dashboard will show a yellow or red icon, which means the server is out of compliance with the compliance test. From the Compliance Dashboard, you can find out where specifically the server is out of compliance and remediate the problem.

Most compliance policies are created and defined by a user, usually the policy setter of an organization (though sometimes an ad hoc policy might be created by a systems administrator). The policy setter creates audit or software policies and then a server's configurations are checked to ensure they are compliant with the policy attached to the server. One exception to this is the duplex compliance category, which is a "system defined" policy test that determines if the duplex settings of all a server's active interfaces match its corresponding switch ports.

Compliance Dashboard Usage: Proactive and Reactive

You can proactively use the Compliance Dashboard by viewing it on a regular basis to assess your servers' and devices' compliance levels, and to take the necessary action to fix any problems.

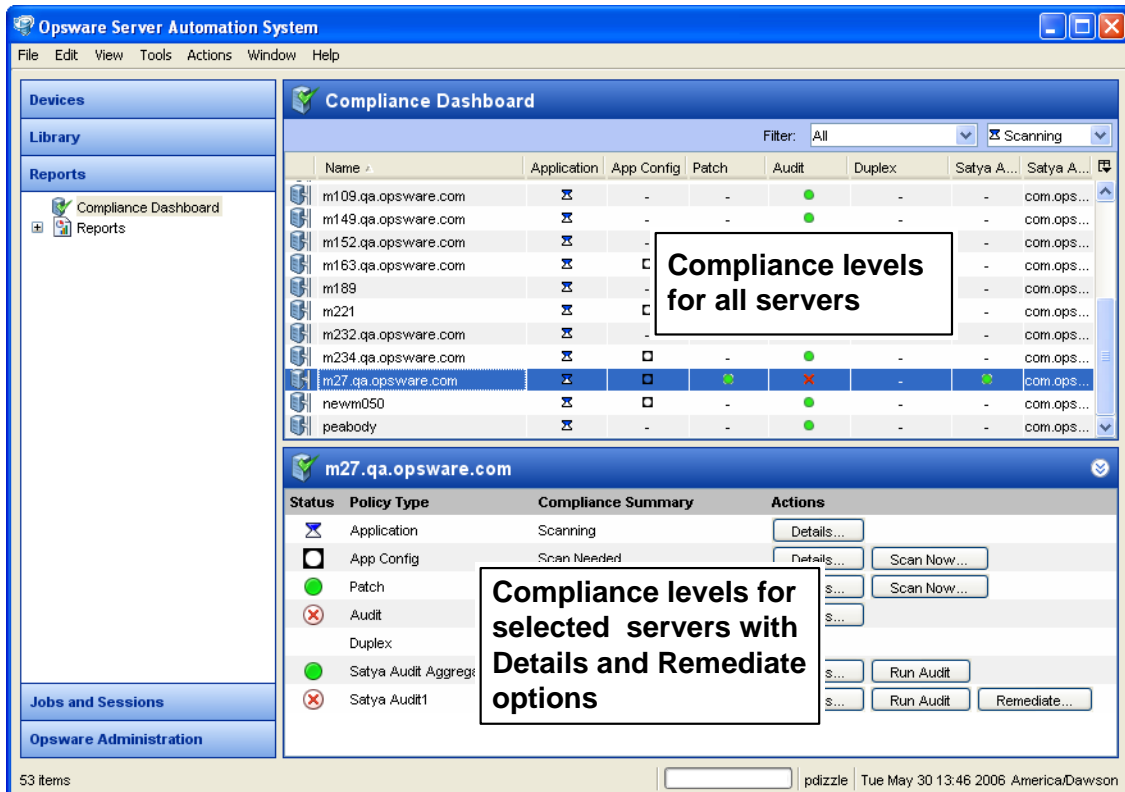
For example, you might use the Compliance Dashboard to view the status of an individually scheduled audit that makes sure a web application's configuration (such as Apache's http.conf file) meets the standards set by your group. In other words, you want to ensure that no one has changed the application's configuration. To verify that no unwanted changes have been made, you should regularly check the Compliance Dashboard for this scheduled audit and see if its compliance status has changed to red (non-compliant), and if so, view the audit results and remediate the problem.

In other situations, you can reactively use the Compliance Dashboard to answer a specific question or diagnose a specific problem. For example, you can create a scheduled audit that defines security standards for a set of servers in your facility. The audit will ensure that all Windows 2003 servers contain a specific patch. When Microsoft releases a new security patch, you want to identify the Windows 2003 servers that contain the new patch and those that do not. You can update the audit to contain the new patch and then browse the Windows 2003 servers in the Compliance Dashboard. When you rerun the audit, you can discover the servers that need the patch and remediate them by installing the new patch.

Viewing the Compliance Dashboard

To view the Compliance Dashboard, from the Navigation pane, select **Reports** ► **Compliance Dashboard**. Figure 3-1 shows what is displayed in the Compliance Dashboard.

Figure 3-1: The Compliance Dashboard



You can access Compliance Dashboard information from the following locations in the SAS Client:

- Reports section of the Navigation pane
- Server Explorer for an individual server
- Compliance view on all devices

General Compliance Dashboard Categories







The Compliance Dashboard displays compliance statuses for the following feature categories:

- **Software:** Software compliance is determined by whether the software policy definition matches what is installed on the server. A software policy defines patches, packages, and application configurations on a server, and may contain other software policies.
- **Application Configuration:** An application configuration's compliance is determined by whether or not the application configuration definition matches the application configurations on the server that the application is attached to. An application configuration defines the configuration settings and values for application configuration files.
- **Patch:** A server's patch compliance is determined by whether or not the patch policy definition matches what patches are installed on the server. Patch policies define patches that should be installed on a server.
- **Audit:** Audit compliance represents an aggregate of all audits that run on a recurring schedule, and the Compliance Dashboard indicates whether or not the rules defined in the audit match what is installed and configured on the target server or servers. In addition, each audit with a recurring schedule can appear as a column in the dashboard, if the user chooses to show an individual test. These can also be hidden.
- **Duplex:** Duplex compliance determines whether or not the duplex settings of all a server's active network interfaces match their corresponding switch ports.

Compliance Dashboard Statuses

Table 3-1 lists and defines all possible compliance statuses shown in the Compliance Dashboard.

Table 3-1: Compliance Dashboard Compliance Status Statuses

ICON	COMPLIANCE STATUS DESCRIPTION
	<p>Compliant: Compliance scan ran successfully and the actual server configuration matches the compliance criteria defined in the policy.</p> <p>It is possible, however, that actual server configurations or policy information might have changed from the last time you viewed the Compliance Dashboard. To get the latest compliance data from the core, from the View menu, select Refresh. (Or, press F5.)</p>
	<p>Partial: Compliance scan ran successfully, but server configuration did not fully pass the compliance criteria defined in the policy.</p> <p>You will see this status for patch policies, if the patch policy has exceptions defined in it.</p>
	<p>Noncompliant: Compliance scan ran and the actual server configuration did not match the criteria defined in the policy.</p>
	<p>Scan Failure: Compliance scan was unable to run.</p>
	<p>Scan Needed: Results unavailable, perhaps because a compliance scan was never run (for example, on a new installation), or the configuration on the server changed since the last time information was reported to the Compliance Dashboard.</p>
	<p>Scanning: Compliance scan currently being run.</p>
–	<p>No Tests Defined: No policies of this type are attached to the server. Or, in the case of duplex, the compliance scan is unable to determine duplex status on the device.</p>

Refreshing to Get the Latest Compliance Information

When you first select the Compliance Dashboard, the information displayed shows the latest information reported on the Compliance Dashboard from the Opsware core for each compliance category. It is possible, however, that actual server configurations might have changed since you selected to view the Compliance Dashboard. Or, it is possible that a policy has changed since you last viewed the Compliance Dashboard.

If this is the case, then the compliance tests need to be re-run which will generate new data for the Compliance Dashboard to display.

As a best practice, it's useful to refresh the Compliance Dashboard to ensure that you are looking at the latest compliance information in your core. To get the latest compliance data from the core, from the **View** menu, select **Refresh**, click **Refresh**, or press F5.

Compliance Dashboard Terms and Concepts

- **Compliance:** The degree to which an server's actual configuration conforms to the configuration as defined in a compliance policy.
- **Compliance Dashboard:** Displays all managed servers in your facility and their compliance statuses. The Compliance Dashboard allows you to view, at a glance, the overall state of server configuration compliance in your facility.
- **Compliance Statuses:** Indicates the compliance status for a feature category – in other words, reports the difference between what should be (compliance policy) and what actually is (server configuration). For example, the software compliance test displays compliant if all configurations defined in the policy match the server configuration. Compliance statuses include: Compliant, Partial, Noncompliant, Scan Failure, Scan Needed, Scanning, Not Applicable. For more information on these statuses, see "Compliance Dashboard Statuses" on page 209.
- **Compliance Scan Results:** The results of a compliance scan. These results report the compliance status, details, and can also include a remediate option.
- **Compliance Policy:** The user-defined (or system-defined) configuration that expresses the desired state for a server or device configuration or setting. For example, a patch policy defines the specific patches that should be installed on a computer. An audit policy might define that a certain Windows service should be disabled at all times. Or, a system defined compliance policy would be the Duplex compliance category, which

reports whether or not the duplex settings of all a server's active interfaces match its corresponding switch ports.

- **Compliance Scan:** The mechanism that runs a scan and returns information to populate the compliance center dashboard for a device or group of devices. A compliance scan could simply check to see what patches are installed on a computer and return the results.

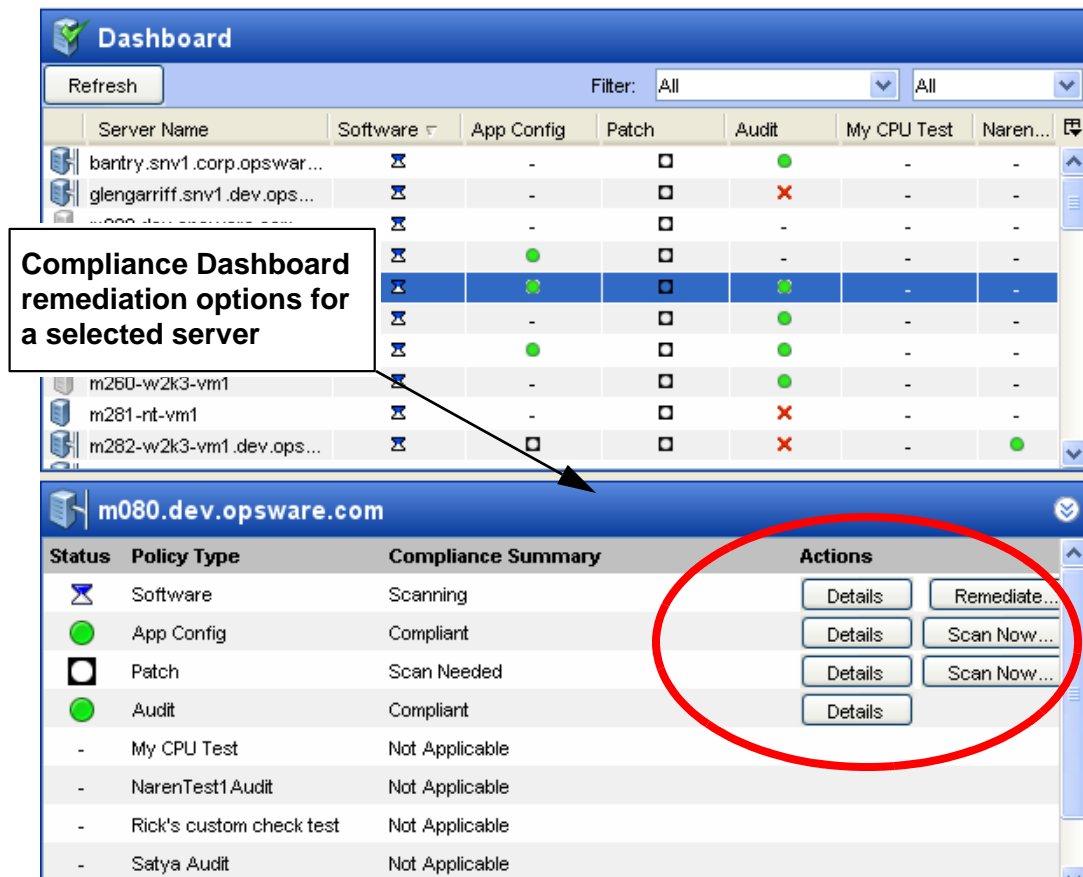
Compliance Dashboard Remediation

The main purpose of using the Compliance Dashboard is to determine if your servers are in compliance with the various policies set for them, and to remediate those server configurations that are not in compliance with your organization's standards.

Generally speaking, the act of "remediating" a server or device means finding how and where a server or device is out of compliance, and fixing the server's or device's configuration – making sure that the actual configuration conforms to the compliance policy.

Using the Compliance Dashboard for each compliance test, you can perform a set of remediate actions, which appear as buttons in the Details pane when you select a server in the Compliance Dashboard. Figure 3-2 shows where a server's compliance remediation options are located in the SAS Client.

Figure 3-2: Compliance Dashboard Showing Remediation Options in the Details Pane



When you select a server in the Compliance Dashboard, the Details pane allows you to perform some of the following general actions (depending upon the compliance test):

- **Details:** Launches the Server Explorer for the server and displays the server's actual configuration for this test. For example, if you click the **Details** button for patch compliance, the Server Explorer will appear showing the patch policies attached to the server.
- **Scan Now:** Initiates a compliance scan, which will compare the compliance policy definitions with what is installed or configured on the server. You will not be able to use this option for the audit compliance test, since you cannot run all scheduled audits at the same time. For individually scheduled audits, however, you can click **Run Audit** to run the specific audit job.
- **Remediate:** Remediates the compliance policy with the actual server configuration. In most cases, this will launch a wizard that allows you to ensure that the server conforms to the compliance policy. This might include installing software, reconfiguring application configurations, and uninstalling software or patches. In the case of the duplex compliance category, this button will launch a shell to the managed server.
- **Run Audit:** Runs the individually scheduled audits on the selected server.



Software Compliance

Software compliance indicates whether or not all software policies attached to the selected server are compliant with the actual server configuration. A software policy includes installed packages and patches, application configurations, and other software policies. If the actual server configuration does not match the software policy definitions, then the server's software policies are considered out of compliance.

For more information on creating and using software policies, see "Software Management" on page 365

Understanding Software Compliance Status

A software policy is either compliant or non-compliant. This means that if any of the patches, packages, or application configurations on the server that the policy is attached to do not match the software policy, then the server is considered out of compliance with the policy. Specifically, software policy compliance is defined as:

- **Compliant:** If a server is compliant with all its software policies (the server configuration matches the software policy), then the policy is considered compliant and the Dashboard will display the compliant icon .
- **Non-compliant:** If any one of the definitions in the software policy does not match with what is installed on the server, then the server is considered non-compliant and the Compliance Dashboard will display the non compliant icon .

Software Compliance Remediate Options

For software policies in the Compliance Dashboard, you can perform the following remediation actions:

- **Details:** Opens the Device Explorer and shows the software policies attached to the selected server. To remediate the policy with the actual server configuration, from the **Actions** menu in the Device Explorer, select **Remediate**.
- **Scan Now:** Starts a compliance scan to determine if the server configuration is out of compliance with the software policy.

- **Remediate:** Opens the Remediate wizard for the server, and lists available software policies. For more information on how to run the Remediate wizard, see “Remediating Software Policies” on page 374.



Application Configuration Compliance

An Application Configuration manages application configuration files on a managed server. Application configuration compliance indicates whether or not all of the Application Configurations attached to a server are compliant with the actual application configuration files on the server. If the actual server configuration does not match the Application Configuration definitions, then the server’s Application Configurations are considered out of compliance.

For more information on creating and using Application Configurations, see “Application Configuration Management” on page 399.

Understanding Application Configuration Compliance Status

An Application Configuration is either compliant or non-compliant. If any of the configuration files on the server do not match the Application Configuration definitions, then the server is considered out of compliance with its attached Application Configurations. Specifically, Application Configuration compliance is defined as:

- **Compliant:** If a server is compliant to with all its Application Configurations (the configuration files on the server matches the Application Configuration definitions), then the server is considered compliant and the Dashboard will display the compliant icon  .
- **Non-compliant:** If any of the Application Configuration definitions do not match the configuration files on the server, then the server’s Application Configurations are considered non-compliant and the Compliance Dashboard will display the non-compliant icon  .

Application Configuration Compliance Remediate Options

For Application Configurations in the Compliance Dashboard, you can perform the following remediate actions:

- **Details:** Opens the Device Explorer for the selected server to the Installed Applications view. This view shows you all Application Configurations that have been attached to server. To browse specific instances of the applications, from the View pane in the Device Explorer, expand the Configured Applications folder.
- **Scan Now:** Starts an Application Configuration audit job. After the scan has finished, you can see what the compliance status is.




Patch Compliance

Patch Compliance determines whether all patches in a patch policy and a patch policy exception were installed successfully on a managed server. To test patch compliance, servers are scanned to determine whether they conform to their attached policies and exceptions, based on compliance statuses and rules. If any of the patches defined in the patch policy do not match what is actually installed on the server, then the server's patch policies are considered out of compliance.

For more information on creating and using patches and patch policies, see "Patch Management for Windows" on page 239 or "Patch Management for Unix" on page 329

Understanding Patch Compliance Status

Specifically, patch policy compliance is defined as:

- **Compliant:** If a server is compliant with all its patch policies (the patches installed on the server match the patch policy definitions), then the server is considered compliant for patch and the Compliance Dashboard will display the compliant icon .
- **Non-compliant:** If any of the patch policy definitions do not match the actual patch installed on the server, then the server's patch policies are considered non-compliant and the Compliance Dashboard will display the non-compliant icon .
- **Partial:** You will see this status for patch policies, if the patch policy has exceptions defined in it indicated by this icon .

Patch Remediate Options

For patch policies in the Compliance Dashboard, you can perform the following remediation options:

- **Details:** Opens the Server Explorer showing the Patch Policies view. This shows the patch policies attached to the selected server. If you want to remediate the policy with the actual server configuration, from the **Actions** menu in the Server Explorer, select **Remediate**.
- **Scan Now:** Starts a compliance scan to determine if the server configuration is out of compliance with the patch policy.
- **Remediate:** Opens the Patch Remediate wizard for the server, with available patch policies selected. From this wizard, you can remediate the server to ensure that it has all the Patches defined in the policy installed.

Audit Compliance

The Compliance Dashboard displays the following two types of compliance for audits:

- **All Scheduled Audits:** Displays a single column of all scheduled audits in the Compliance Dashboard by default. This status enables you to view, at a glance, the total compliance status of all audits that you have scheduled to run on the selected server. A server can be the target for several audits, and this compliance test provides a roll up of compliance status for all audits being run against the selected server.



You will only see the compliance status for those audits that have been scheduled on servers that your user has access to. Any servers you do not have access to will not be represented in the Compliance Dashboard in the audit roll up.

- **Individual Audits:** Displays individual audits scheduled to run on the selected server on a per-audit basis. These audits will not appear by default and must be activated to display in the Compliance Dashboard. You must have access to view the server where the audit is running in order to see it displayed in the Compliance Dashboard.



For more information on creating and using audits, see "Audit and Remediation" on page 103.

Audit Compliance Status: All Scheduled and Individual

Audit compliance for all scheduled audits is represented in the Compliance Dashboard by the two following statuses:

- **Compliant:** If all scheduled audits run against the selected servers are successful – that means, the configurations of all servers being audited match the rule values defined in the audit – then the audit column in Compliance Dashboard will be shown with a green icon. If there are no scheduled audits being run against the selected server, then the Compliance Dashboard will display the compliant icon .
- **Failed (Non Compliant):** All of audit rules for all scheduled audits being run against the selected server do not match the actual server configuration values.
Failed compliance is represented with the non-compliant icon .

Audit compliance for individual audits is represented in the Compliance Dashboard by the two following statuses:

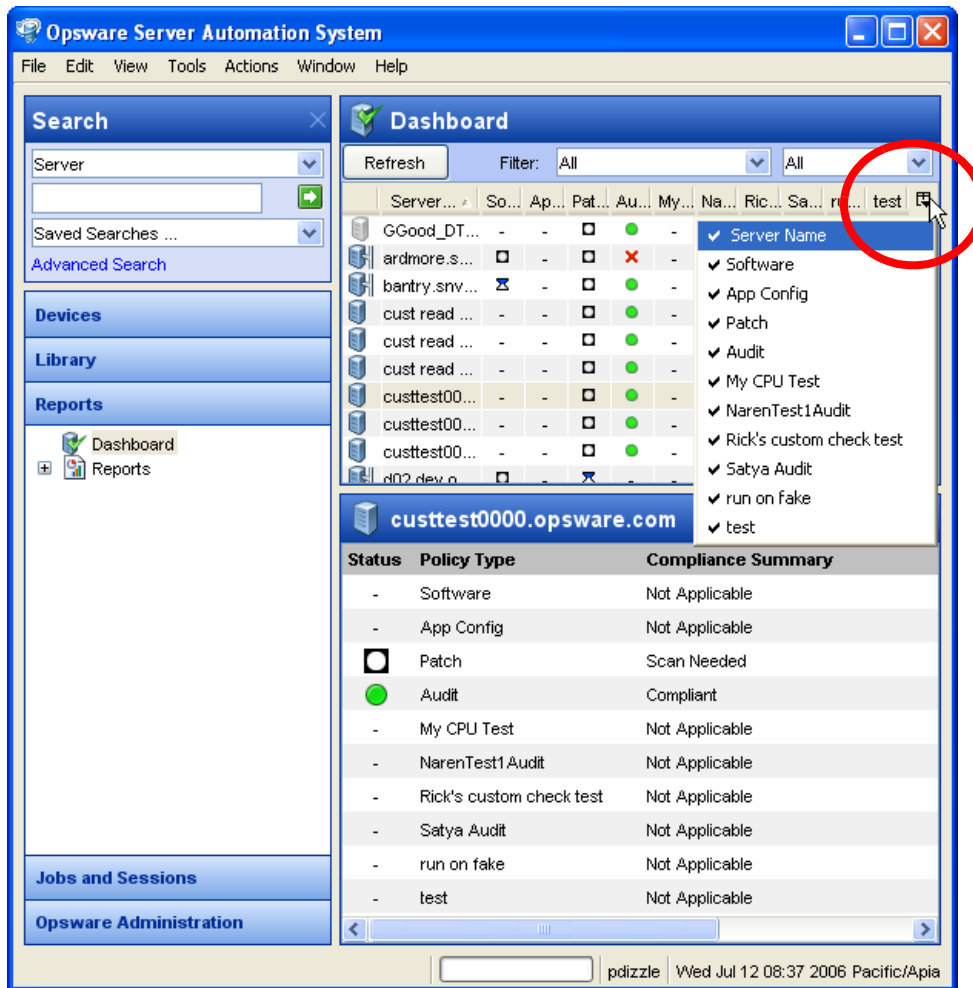
- **Compliant:** If an audit is successful – the selected server's configuration matches the audit's rules' definitions – then the Audit column in the Compliance Dashboard will display the compliant icon .
- **Non Compliant:** If all of audit rules for a single audit did not match the actual server configuration values, then the Compliance Dashboard will display the non-compliance icon .

Showing Individual Audits

By default, the Compliance Dashboard shows compliance status only for all scheduled audits. If you would like the Compliance Dashboard to display compliance for an individual scheduled audit, perform the following steps:

- 1 To view the Compliance Dashboard, from the Navigation pane, select **Reports** ► **Compliance Dashboard**.
- 2 In the Contents pane, in the far upper right side, select the column selected drop-down list to display all Compliance Categories. See Figure 3-3.

Figure 3-3: Compliance Dashboard Category Selector



- 3 Select an individual audit by the name that you want to display in the Compliance Dashboard.

Audit Remediate Options

The types of remediation you can perform with audits in the Compliance Dashboard vary according to the type of audit.

For the roll up of all scheduled audits (listed as audit in the Compliance Dashboard) run against a selected server, you can perform the following remediation action:

- **Details:** Opens the Server Explorer window, shows the Audit & Remediation view, and lists all audits that use the selected server as its source.

For all each individually scheduled audit, you can perform the following remediate actions:

- **Details:** Opens the Audit window and displays the audit configuration.
- **Run Audit:** Opens the Run Task window, which allows you to run the audit again.
- **Remediate:** Opens the Audit Results window, which shows you the results of the audit. From this location, you can see the specific audit results related to the selected server and if available, perform audit remediation. For more information on how to remediate audit results, see "Viewing and Remediating Audit Results" on page 175.


Duplex Compliance


Duplex compliance determines whether or not the duplex of a server's active interfaces match their corresponding switch ports. For example, if one of the server's network interfaces is set to full duplex and its switch port is set to half duplex, then the server is out of compliance.

Duplex compliance is determined by the following criteria:

- **Compliant:** A server is compliant if the duplex of all of its active network interfaces matches its corresponding switch ports.

In this case, the Compliance Dashboard will display the compliant icon .

- **Non-Compliant:** If one of a server's network interfaces does not match one of its corresponding network switch ports, then the duplex compliance state will be non-compliant, and the Compliance Dashboard will display the non-compliant icon .

- **Scan Needed:** When you first install the SAS Client, the duplex compliance scan will not have been run. To initiate the first scan, click **Scan Now** or wait until the first scan runs. (Scans run every 24 hours.)
- **Unknown:** If the duplex of either the server interface or the network switch port cannot be determined, then the duplex compliance level will be unknown, and the Compliance Dashboard will display the Scan Needed icon .

Duplex Remediate Options

For patch policies in the Compliance Dashboard, you can perform the following options:

- **Details:** Opens the Server Explorer showing the Hardware view.
- **Scan Now:** Initiates a compliance scan to determine duplex compliance.
- **Remediate:** Opens a shell terminal to server where the duplex mismatch occurs so you can troubleshoot problem.

For more information, see “Duplex Mismatch” on page 368.

Filtering and Sorting Compliance Dashboard Information

To better view the information displayed in the Compliance Dashboard, you have several options for filtering the compliance results you can show only the compliance tests you want to see, or you can show or hide any of the compliance tests, including individually scheduled audits.

For example, you might be interested in specific compliance tests and their statuses, such as, all non-compliant patch policies. You can select the patch compliance test and then select the non-compliant compliance status, and filter the Compliance Dashboard to show only that information. The results will show only those server whose patch policies are out of compliance.



Each audit with a recurring schedule can be displayed as a column in the Compliance Dashboard, if an administrator with sufficient authority adds it.

In the All Managed Servers list, you can show the following Compliance Dashboard compliance tests: Software, App Config (Application Configuration), Patch, and Audit.

Filtering Compliance Dashboard Compliance Tests

To filter the Compliance Dashboard to display specific compliance tests, perform the following steps:

- 1** To view the Compliance Dashboard, from the Navigation pane, select **Reports** ► **Compliance Dashboard**.
- 2** In the Contents pane, from the Filter drop-down list, select a compliance test, such as Software, App Config, and so on.
- 3** Next, from the second drop-down list, select a compliance status, such as non-compliant. The Compliance Dashboard displays only those compliance tests with the selected status.

Showing/Hiding Specific Compliance Tests


By default, the Compliance Dashboard only displays the main compliance tests: Software, App Config (Application Configuration), Patch, Audit, and Duplex (if your core is NAS-enabled). You can choose to show or hide any of these tests. You can also choose to show individually scheduled audits.

To show individually scheduled audits, perform the following steps:

- 1** To view the Compliance Dashboard, from the Navigation pane, select **Reports** ► **Compliance Dashboard**.
- 2** Use the column selector (upper right corner of the table) and select a compliance test. When a test has a check mark next to it, it will display in the Compliance Dashboard. Select it again to hide it.

Showing Compliance in All Managed Servers

To show or hide compliance tests in the All Managed Servers list, perform the following steps:

- 1** From the Navigation pane, select **Devices** ► **All Managed Servers**.
- 2** In the Contents pane, in the far upper right side, select the column selector  drop-down list to display all Compliance Categories.

- 3 At the bottom of the column list, You can choose to show any of the compliance categories, such as Software, App Config, and so on.

Exporting the Compliance Dashboard

If you want to view all the information displayed in the Compliance Dashboard in a file, you can export all compliance results to either .html or .csv.

To export the Compliance Dashboard to a file, perform the following steps:

- 1 To view the Compliance Dashboard, from the Navigation pane, select **Reports** ► **Compliance Dashboard**.
- 2 Right-click inside the Contents pane of the Compliance Dashboard and select **Export**.
- 3 In the Export Dashboard window, enter a name for the file, and choose if you want to export to .html or .csv. You can also change the encoding if you want the saved file to use a specific encoding scheme.
- 4 Click **Export**.

Chapter 4: SAS Client Reports

IN THIS CHAPTER

This section discusses the following topics:

- Overview of SAS Client Reports
- Reports Features
- Opsware SAS Client Reports
- User Permissions
- Launching the Reports Feature
- Reports Display
- Running a Report
- Report Results

Overview of SAS Client Reports

The Opsware SAS Client Reports feature provides comprehensive, real-time information about managed servers, network devices, software, patches, customers, facilities, operating systems, compliance policies, and users and security in your environment. These parameterized reports are presented in graphical and tabular format, and are actionable—which means that you can perform appropriate actions on objects, such as a policy or an audit, within the report. These reports are also exportable to your local file system (as .html and .xls files) to facilitate use within your organization.

This section contains information about the types of SAS Client reports, how to modify report parameters, how to run the reports, and how to perform actions in the report results.

Reports Features

SAS Client Reports enable you to perform enterprise health assessments by providing the following features:

- Actionable reports that enable you to take the appropriate action on objects within the reports. For example, in the list view of a compliance report, you can select a server and open a Remote Terminal or Server Explorer to browse it, perform an audit, create a snapshot, create a package, and so on.
- A single entry point in the SAS Client Dashboard for all reports.
- Reports that are data-secured—controlled by the user's permissions. You can view all objects that you have read permissions for. You can perform actions on objects that you have write permissions for.
- Reports that are exportable to .html and .xls formats. You can export reports to your local file system for use within your organization.

Opware SAS Client Reports

Table 4-1 lists the SAS Client Reports by report folders.

Table 4-1: SAS Client Reports

REPORT FOLDER	REPORT TITLE
Server Reports	Servers by Customer
	Servers by Facility
	Servers by Manufacturer
	Servers by Model
	Servers by Operating System
	Servers by Use

Table 4-1: SAS Client Reports (continued)

REPORT FOLDER	REPORT TITLE
Virtualization Reports	Virtualization by Virtual Technology All Virtual Servers Solaris 10 Virtual Servers by Hypervisors (zones only) Resource Allocation by Hypervisors (zones only) VMware ESX 3 Virtual Servers by Hypervisors (VMs only) Resource Allocation by Hypervisors (VMS only)
Compliance Reports	Audit Policy Compliance (All Servers) Audit Policy Compliance by Customer Audit Policy Compliance by Facility Patch Policy Compliance (All Servers) Patch Policy Compliance by Customer Patch Policy Compliance by Facility Software Policy Compliance (All Servers) Software Policy Compliance by Customer Software Policy Compliance by Facility Audit Compliance - Devices by Audit Device Audit Compliance Detail

Table 4-1: SAS Client Reports (continued)

REPORT FOLDER	REPORT TITLE
SOX	SOX Compliance Summary
	Audit Results With Failures
	Audit Results Without Failures
	Current Users
	Defined Patch Policies
	Defined Server Audit Policies
	Defined Server Audits
	Defined Software Policies
	Deleted Users
	Jobs With an Associated Ticket ID
	Jobs Without an Associated Ticket ID
	Servers Audited Without Failures
	Servers in Compliance With Their Patch Policies
	Servers in Compliance With Their Software Policies
	Servers Not Audited or Audited With Failures
	Servers Not in Compliance With Their Patch Policies
	Servers Not in Compliance With Their Software Policies
	Servers With Associated Audits
	Servers With Attached Patch Policies
	Servers With Attached Software Policies
Servers Without Associated Audits	
Servers Without Attached Patch Policies	
Servers Without Software Patch Policies	
	Users Created in the Last 30 Days
	User Groups Membership

Table 4-1: SAS Client Reports (continued)

REPORT FOLDER	REPORT TITLE
User and Security Reports	Client and Feature Permissions Customer/Facility Permissions and Device Group Permission Overrides User Groups Memberships
Network Reports	Connections by Network Device Connections by Server Duplex Compliance (All Servers) Duplex Compliance by Customer Duplex Compliance by Facility

See the following documentation for more information about the SAS Client features that support information in these reports:

- “Software Management” on page 365
- “Audit and Remediation” on page 103
- “Exploring Servers and Groups in SAS Client” on page 141
- “Patch Management for Windows” on page 239
- “Patch Management for Unix” on page 329
- “NAS Integration” in the *Opware[®] SAS User’s Guide: Server Automation*
- “Server Management in SAS Web Client” in the *Opware[®] SAS User’s Guide: Server Automation*

User Permissions

Reports are controlled by the user's permissions. You can view all objects that you have read permissions for, and you can perform actions on objects that you have write permissions for.

To view or run a network report, NAS Integration must be installed. See "NAS Integration" in the *Opware® SAS User's Guide: Server Automation*.

To view or run a user and security report, system administrator permissions are required.

Launching the Reports Feature

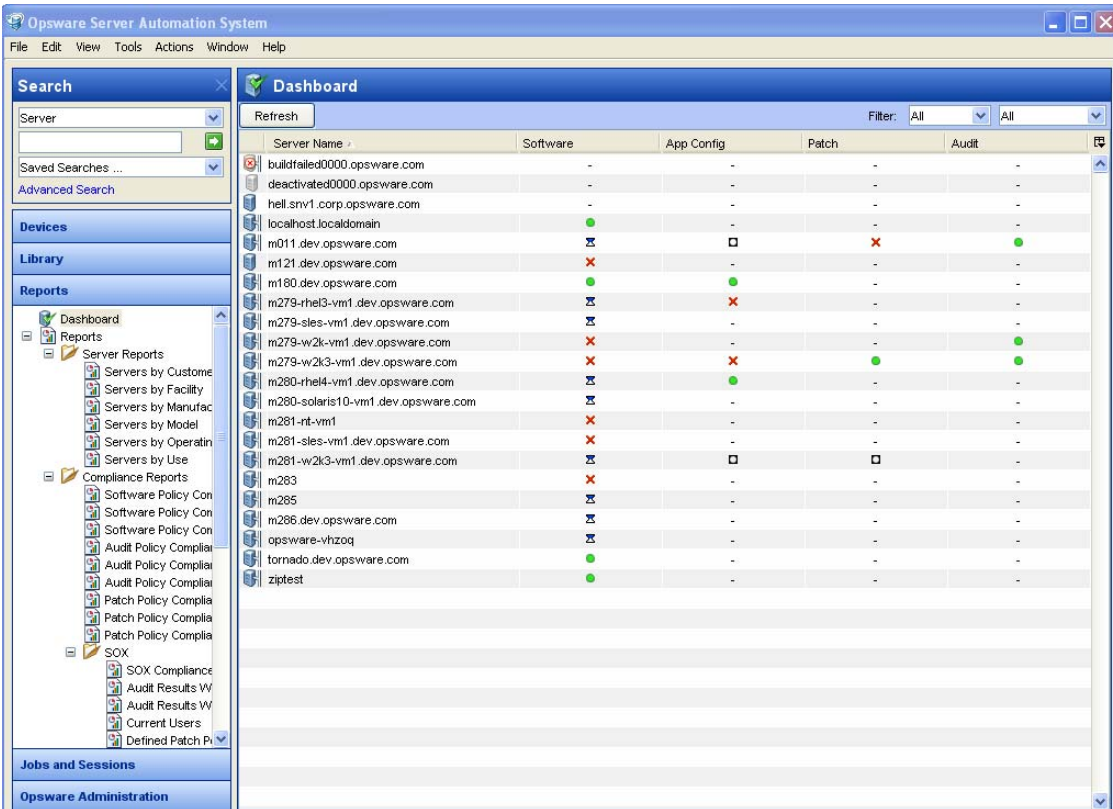
To launch the Reports feature, perform one of the following steps:

- From the **View** menu, select **Reports** ► **Dashboard**.
- From the **View** menu, select **Reports** ► **Reports**.
- From the Navigation pane, select Reports.

Reports Display

The Reports feature display consists of a Search pane, a Dashboard, report parameters, report folders, and report parameters.

Figure 4-1: The Reports Feature Display



Search Pane

In the Reports feature, you can use the SAS Client Search feature to find reports by defining specific filter criteria. See “SAS Client Search” in the *Opware® SAS User's Guide: Server Automation*.

Dashboard

The Dashboard is the default content pane that displays when you select Reports from the Navigation pane. See “Compliance Dashboard” on page 205.

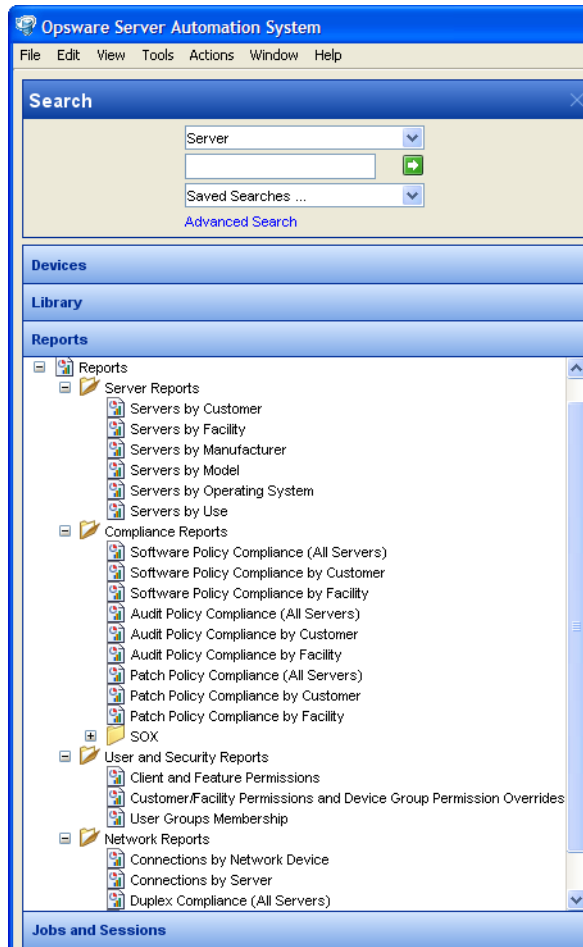
Report Folders

Reports are organized in the following folders according to regulatory or IT best practice standards:

- **Server Reports:** This folder contains reports about servers by customer, facility, manufacturer, model, operating system, and server usage.
- **Compliance Reports:** This folder contains reports about compliance for software policies, audit policies, and patch policies by servers, customer, and facility.
- **SOX Reports:** This folder contains reports about compliance standards based on Sarbanes-Oxley, including the COSO process model and the CobiT control model.
- **Network Reports:** This folder contains reports about connections and duplex compliance for network devices and servers. You must have Opware NAS installed to see this folder in the Navigation pane.
- **User and Security Reports:** This folder contains reports about client and feature permissions; customer, facility, and device group permissions; and user group memberships. You must have system administrator permissions to see this folder in the Navigation pane.
- **Custom Reports:** This folder contains any custom reports you have created. For more information on creating custom reports, check the Opware Knowledge Base.

- Figure 4-2 illustrates the Report folders in the Navigation pane, including the reports you will find in each folder.

Figure 4-2: Report Folders



Report Parameters

Many reports require input parameters in order to be run. For reports that require parameters, you can run the report with its default parameter values or modify the parameter values. If you want to run a report that includes or excludes certain servers, customers, or hardware models, you need to specify this criteria in the report parameters. See “Running a Report” on page 234.

Running a Report

To run a report, perform the following steps:

- 1** From the Navigation pane, select Reports.
- 2** Expand the Reports folder and then expand the Server Reports and Virtualization Reports.
- 3** Select one of the virtualization report listed in the folder.
- 4** If there are no report parameters in the Content pane, click **Run**.
- 5** If there are report parameters in the Content pane, you can either use the default parameters or change them:
 - To use the default report parameters, click **Run** to run the report.
 - To change the report parameters, see “Modifying Report Parameters” on page 234.

Modifying Report Parameters

To modify the default parameters and run a report that includes certain servers, customers, hardware models, and so on, perform the following steps:

- 1** In the drop-down list for (the Server, Customer, Model, and so on), select Contains, Equals, Begins With, or Ends With.
- 2** (Optional) Select the ellipsis button to open the Select Values window.
- 3** In the Select Values window, select a value in the Available or Selected pane and then use the directional buttons to include it in or exclude it from your search criteria.
- 4** Click **OK** to save your changes.
- 5** Click **Run** to run the report.



If data cannot be found to run the report, a “No records to display!” error displays.

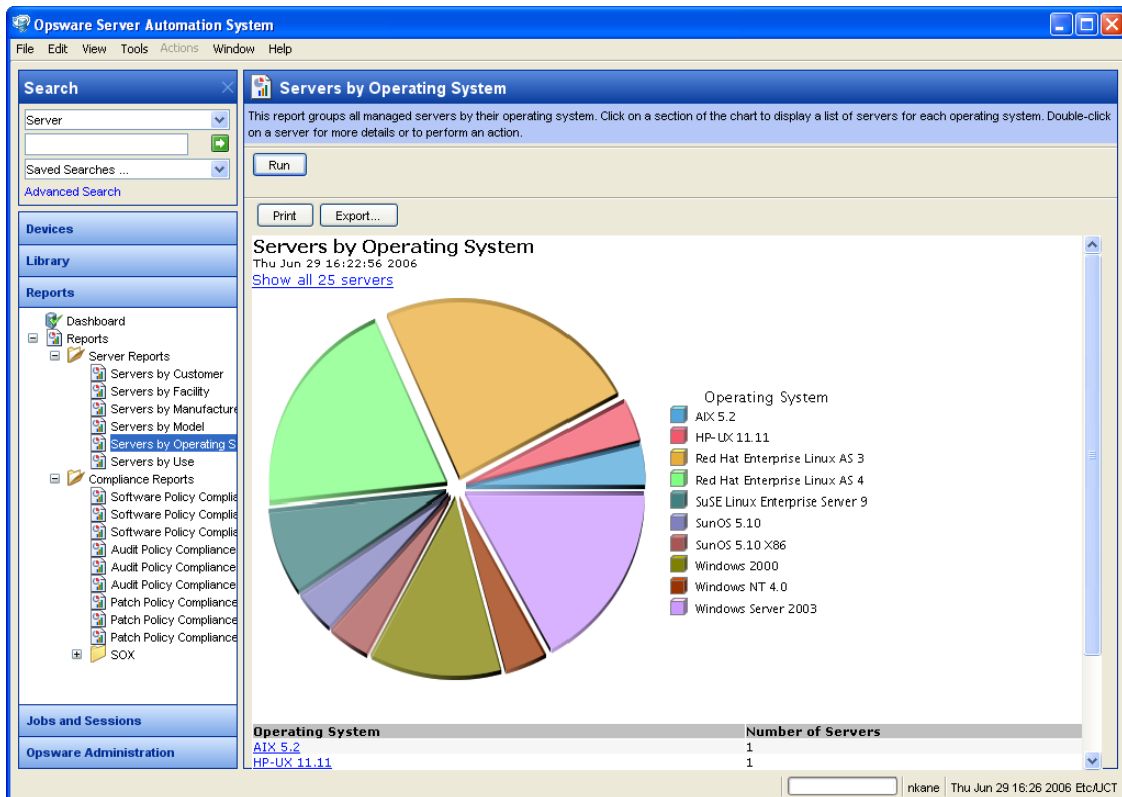
Report Results

Report results initially appear in a graphical or list view. The graphical report is an overview of available data for this report displayed in a pie chart or in a bar graph. You can drill down for more detail in the chart or graph by clicking on any of the sections or bars. For example, you can drill down to individual servers that appear in a report and get detailed information about them.

Graphical Report


A graphical report is a pie chart or a bar graph. Click on a section of the chart or graph to drill down for more details or to perform an action. You can also click on the “Show all <number> servers” link to display a list of servers. See Figure 4-3 and Figure 4-4 for examples.

Figure 4-3: Pie Chart



To display the corresponding list view of a bar graph, click on the front part of the bar. Do not click on the top or shaded part of the bar.

Figure 4-4: Bar Graph



List Report

A list report is a tabular display of information. Double-click on a row in the list, such as a server, audit, or policy, for more detail or to perform an action. See Figure 4-5 for an example.

Figure 4-5: List Report

The screenshot shows the 'Opware Server Automation System' window. The main area displays a report titled 'Software Policy Compliance by Customer'. Below the title, there is a search filter for 'Customer' set to 'Equals' and a 'Run' button. Below that are 'Print' and 'Export...' buttons. The report content is a table titled 'Server Listing' with the following data:

Hostname	IP Address	Operating System	Customer	Compliance Level	Date Scanned
m208.dev.opsware.com	192.168.165.130	Red Hat Enterprise Linux AS 3	Not Assigned	Not Applicable	
erendir.snv1.corp.opsware.com	192.168.9.224	Red Hat Linux 7.3	Not Assigned	Not Applicable	
BDay_LTXP.snv1.corp.opsware.com	192.168.8.186	Windows XP	Not Assigned	Not Applicable	
m282-rhel4-vm1.dev.opsware.com	192.168.202.68	Red Hat Enterprise Linux AS 4	Not Assigned	Not Applicable	
andy	192.168.165.132	Red Hat Enterprise Linux AS 3	Not Assigned	Not Applicable	
m279-rhel2-vm1.dev.opsware.com	192.168.202.36	Red Hat Enterprise Linux AS 2.1	Not Assigned	Not Applicable	
m282-nt-vm1	192.168.202.63	Windows NT 4.0	Not Assigned	Not Applicable	
loki.corp.opsware.com	192.168.9.128	Red Hat Linux 8.0	Not Assigned	Not Applicable	
m279-solaris10-vm1.dev.opsware.com	192.168.202.41	SunOS 5.10 X86	Not Assigned	Not Applicable	
fake_Satya0001.opsware.com	157.23.29.86	AIX 4.3	Not Assigned	Not Applicable	
m049.dev.opsware.com	192.168.194.58	AIX 5.2	Not Assigned	Not Applicable	
fake_Satya0001.opsware.com	158.190.233.168	Windows 2000	Not Assigned	Not Applicable	
fakeSatya0001.opsware.com	214.146.204.139	Windows 2000	Not Assigned	Not Applicable	
m047	192.168.194.56	HP-UX 11.11	Not Assigned	Not Applicable	

Exporting a Report

You can export a report for use in other applications in your environment and attach a report for email distribution. Depending on the report format, you can export a report to your local file system as an .html file or an .xls file. You can export a graphical report to .html only. You can export a list report to .html and .xls.

To export a report, perform the following steps:

- 1** From the report, click **Export** to open the Save window.
- 2** In the Save in field, enter a location that identifies where you want to save the file to, or select from the drop-down list.

- 3** Enter a file name.
- 4** Select the file type, such as .html or .xls.
- 5** Click **Save**.

Printing a Report

To print a report, perform the following steps:

- 1** From the report, click **Print** to open the Print window.
- 2** Use the default print options or modify them, and then click **OK**.

Chapter 5: Patch Management for Windows

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Patch Management for Windows
- Patch Management Roles for Windows
- Patch Management Process
- Patch Properties
- Policy Management
- Patch Compliance
- Patch Administration for Windows
- Locales for Windows Patching
- Patch Installation
- Patch Uninstallation

Overview of Patch Management for Windows

The Patch Management for Windows feature enables you to identify, install, and remove Microsoft® Windows patches and maintain a high level of security across managed servers in your organization. With the Opware SAS Client user interface, you can identify and install patches that protect against security vulnerabilities for the Windows 2000, Windows 2003, and Windows NT4.0 operating systems. These patches include Service Packs, Update Rollups, and hotfixes. This feature also supports patching on 64 bit for Windows 2003 operating systems and for 32 bit for Windows XP operating systems.

This section contains information about how to install Windows patches using patch policies and how to uninstall patches using a sequence of tasks. It also contains information about running patch compliance scans and generating patch policy compliance reports.

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

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

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

Patch Management is a fully integrated component of Opsware SAS. It leverages the Opsware SAS server automation features. Opsware SAS, for example, maintains a central database (called the Model Repository) that has detailed information about every server under management, the patches and software installed on the servers, and the patches and software available for installation. You can use this data to determine the severity of your exposure to a newly discovered threat, and to help you assess the benefits of rolling out a patch versus the costs in downtime and testing requirements.

By automating the patching procedure, the Patch Management feature can reduce the amount of downtime required for patching. Opsware SAS also allows you to schedule patch activity, so that patching occurs during off-peak hours.

Patch Management for Windows Features

Opsware SAS automates patch management by providing the following features:

- A central repository where patches are stored and organized in their formats
- A database that includes information on every patch that has been applied
- Customized scripts that can be run before and after a patch is installed
- Advanced search abilities that identify servers that require patching
- Auditing abilities for tracking the deployment of important patches

These features enable you to browse patches by a certain operating system, schedule patch downloads and installations, set up email notifications, preview a patch installation, use policies and remediation to install patches, and export patch information to a reusable file format.

Types of Patch Browsing

The SAS Client interface organizes Microsoft patches by operating systems and displays detailed vendor security information about each patch, such as Microsoft Security Bulletins. You can browse patches by the date Microsoft released the patch, by the severity level, by the Security Bulletin ID, QNumber, and so on. You can also browse all patches that are installed on a server, and view and edit patch metadata.

Scheduling and Notifications

In Patch Management, you can separately schedule when you want patches imported from Microsoft (either automatically or on demand) into Opware SAS and when you want these patches downloaded to managed servers. As a best practice, patch installations are typically scheduled for a time that causes minimal disruption to an organization's business operation. If you are installing one patch on one server, the installation operation will start only after the download operation has completed.

Patch Management also allows you to set up email notifications that alert you whether the download and installation operations completed, succeeded, or failed. When you schedule a patch installation, you can also specify reboot preferences to adopt, override, postpone, or suppress the vendor's reboot options.

Patch Policies and Exceptions

To provide flexibility in how you identify and distribute patches on managed servers or groups of servers, Patch Management allows you to create patch policies that define groups of patches that you need to install. By creating a patch policy and attaching it to a server or a group of servers, you can effectively manage which patches get installed where in your organization. If you want to include or exclude a patch from a patch installation, Patch Management allows you to deviate from a patch policy by specifying that individual patch in a patch policy exception. An additional patch is one that is not already specified in the patch policy and is one that you want to include in (add to) the patch installation. A patch that you want to exclude from a patch installation is one that is already specified in a patch policy and is identified in the patch policy exception as one you do not want installed. In cases where it is already known that a certain Windows patch may cause a server or application to malfunction, you should create a patch policy exception to exclude it from being installed on that server or on all servers that have that application.

Patch Installation Preview

While Patch Management allows you to react quickly to newly discovered security vulnerabilities, it also provides support for strict testing and standardization of patch installation. After you have identified patches to install, Patch Management allows you to simulate (preview) the installation before you actually install a patch. This preview process tells you whether the servers that you selected for the patch installation already have that patch installed. In some cases, a server could already have a patch installed if a system administrator had manually installed it. After this type of patch installation, if a compliance scan has not been run or the installed patch has not been registered, Opware SAS does not know about it. The preview process for an up-to-date report of the patch state of servers. The preview process also reports on patch dependency and supersedence information, such as patches that require certain Windows products, and patches that supersede other patches or are superseded by other patches.

Patch Policy Remediation

Patch Management also provides a solution for remediating servers that are not operating properly due to installed patches. If installed patches cause problems, even after being tested and approved, Patch Management allows you to uninstall the patches in a safe and standardized way. Patch Management allows you to specify uninstall options that control server reboots and the execution of uninstall commands, and pre-uninstall and post-uninstall scripts. Similar to previewing a patch installation, you can also preview a patch uninstallation.

Exporting Patch Data

To help you track the patch state of servers or groups of servers, Patch Management allows you to export this information. This information can be exported in a comma-separated value (.csv) file and includes details about when a patch was last detected as being installed, when a patch was installed by Opware SAS, the patch compliance level, what patch policy exceptions exist, and so on. You can then import this information into a spreadsheet or database to perform a variety of patch analysis tasks.

Library

The SAS Client Library provides flexibility in searching for and displaying Microsoft patches by operating system, severity level, release date, bulletin ID, and so on. See Figure 5-1. The number in parenthesis is the total number of patches (for that operating system version) that were uploaded from the Microsoft web site. In the Content pane, a

dimmed patch icon indicates that the patch has not yet been uploaded to the Library. Use the column selector to control the columns of patch metadata data that you want to display.

Since the Library is integrated with Microsoft patch metadata, you can review vendor information (in real-time) in the Preview pane.

Figure 5-1: Windows Patches in the Opware SAS Client Library

The screenshot displays the Opware Server Automation System interface. The main window is titled "Opware Server Automation System" and has a menu bar with "File", "Edit", "View", "Tools", "Actions", "Window", and "Help".

On the left side, there is a "Search" panel with a search box containing "Server", a search button, and a "Saved Searches ..." dropdown. Below it is a "Library" panel with tabs for "By Type" and "By Folder". The "By Type" view is selected, showing a tree structure of categories: Application Configuration, Software Policies, Audit and Remediation, and Patches. The "Patches" category is expanded, showing sub-categories: AIX, HP-UX, Solaris, and Windows. The "Windows" category is further expanded to show: Windows 2000 (664), Windows NT 4.0 (586), Windows Server 2003 (434), Windows Server 2003 x64 (32), and Windows XP (164). The "Windows Server 2003 (434)" item is highlighted with a red box.

On the right side, there is a "Windows Server 2003" panel with a "View: Properties" dropdown. Below it is a table of patches:

Name	Type	Severity
Q893066	Windows H...	Critical
Q890923	Windows H...	Critical

Below the table is a "Q890923" panel with a "General" section containing the following details:

- Name: Q890923
- Type: Windows Hotfix
- OS: Windows 2003
- Size: 3.85 MB
- Opware ID: 7800040
- Availability: Available
- Title: Cumulative Security Upd
- KB #: 890923
- Bulletin: MS05-020

Patch Management for Windows Prerequisites

The managed servers that will be patched have the following requirements:

- Either Microsoft Core XML Services (MSXML) 3.0 (or later) or Internet Explorer (IE) 6.0 (or later) must be installed on the managed servers. These versions of MSXML and IE support the Microsoft XML parser and related DLL files that are required for the native Microsoft Baseline Security Analyzer (MBSA) tool (mbsacli.exe). Opware SAS uses version 2.0 of the MBSA tool for patch management. (From Opware SAS 5.5 through 6.0, version 1.2.1 of MBSA was also used.) Vendor-recommended patches that are installed during the patch remediation process are based on MBSA 2.0.

- Windows Installer 3.1 must be installed on the managed servers. This installer is available at the following URL:

<http://support.microsoft.com/kb/893803/>

- On the managed servers, the Automatic Update service must be set to either Automatic or Manual. To set a Windows service, from the Windows Control Panel select Administration Tools ► Services. This service setting is required because Patch Management relies on the MBSA 2.0 scanning engine (mbsacli) to detect installed and recommended patches. If the Automatic Update service is disabled, mbsacli will not work properly and the patching process will not continue after reboot. In this situation, the Agent is unable to report a complete set of installed and recommended patches.
- For Windows 2000 managed servers, SP4 must be installed. Servers with earlier service packs are not supported by Patch Management.
- For Windows XP managed servers, SP2 must be installed.
- To use Patch Management on managed servers with Opware Agent versions earlier than 6.1, the language (locale) of the managed server must be either English, Japanese, or Korean. To set the language, on the managed server, open the Control Panel, open the Regional and Language Options window, select the Regional Options tab, and select an item from the drop-down list at the top.
- Specific versions of the Opware Agent are required to support the functions of Patch Management, as listed in Table 5-1.

Table 5-1: Opware Agent Requirements

PATCH FUNCTIONALITY	OPSWARE AGENT VERSION
Install Patch	4.5 or later

Table 5-1: Opsware Agent Requirements

PATCH FUNCTIONALITY	OPSWARE AGENT VERSION
Uninstall Patch	4.5 or later
Remediate	5.5 or later

Microsoft Patch Database

The Microsoft patch database contains information about released patches and how they should be applied. Patch Management compares all Windows servers to this database to enable the policy setter to determine the patches that must be applied.

Microsoft posts patches on its web site on the second Tuesday of each month, unless a special circumstance requires an immediate release. Windows patches released on *patch Tuesday* are available immediately to import into Opsware SAS. Before Patch Management can install a patch on a managed server, the patch must be downloaded from the Microsoft web site and imported into the Software Repository. You can download and import patches with either the Opsware SAS Client or with a script.

Once every 24 hours, the Opsware Agent on a Windows server compares the server's current state against the Microsoft patch database (based on the latest version of the MBSA) that has been imported into Opsware SAS by the patch administrator. The Opsware Agent reports the results of that comparison and then stores the data in the Model Repository. When a user requests a patch compliance scan of a Windows server, the data is retrieved from the Model Repository and displayed in the SAS Client. By storing the data in the Model Repository, rather than performing an actual comparison on the server itself when a user requests an analysis, the data can be quickly retrieved and displayed.

If you perform a patch analysis of a Windows server immediately after importing a new version of the Microsoft patch database, the analysis does not yet include the data from the new patch database. Instead, Opsware SAS reports the data from the last time that the Opsware Agent recorded the results of its comparison. For example, the Opsware 5.5 Agent on a Windows server uses Microsoft's latest detection engine (MBSA 2.0) to identify installed patches. If you used a previous version of the Opsware Agent to create a package of installed patches (from a server snapshot), a previous version of Microsoft's detection engine (MBSA 1.2.1) was used. Because different versions of MBSA were used to identify patches installed on a Windows server, you should expect to see a difference between the list of installed patches that the SAS Client displays and the installed patches in the package that was created from a snapshot.



While MBSA 2.0 can include programs that are not patches in the Microsoft patch database, such as Malicious Software Removal Tool entries, these programs are excluded from Patch Management.

Opsware SAS Integration

When a server is brought under management by Opsware SAS, the Opsware Agent installed on the server registers the server's configuration, including installed patches, with Opsware SAS. (The Opsware Agent repeats this registration every 24 hours.) This information, which includes data about the exact operating system version, hardware type, installed software and patches, is immediately recorded in the Model Repository. Also, when you first provision a server with Opsware SAS, the same data is immediately recorded.

When a new patch is issued, you can use the SAS Client to immediately identify which servers require patching. Opsware SAS provides a Software Repository where you upload patches and other software. Users access this software from the SAS Client to install patches on the appropriate servers.

After a server is brought under management, you should install all Windows patches by using the Patch Management feature. If you install a patch manually, Opsware SAS does not have data about that patch until the next software registration. If you install a patch manually, it can take as long as 24 hours until the data about that server in the Model Repository is up-to-date. However, whenever you install patches with Opsware SAS, the Opsware Agent immediately updates the information about the server in the Model Repository.

You cannot use Opsware SAS to uninstall a patch that was not installed by using the Patch Management feature.

Support for Windows Patch Testing and Installation Standardization

Opsware SAS offers features to minimize the risk of rolling out patches. When a patch is initially imported into Opsware SAS, its status is marked as untested (Limited) and only administrators with the required permissions can install it.

The patch administrator then defines patch installation and uninstallation options and tests the patch. Only after the patch is tested and the patch administrator marks it as available for use (Available) can other administrators install the patch.

The Patch Management feature allows you to standardize the way that patches are installed and uninstalled, thereby preventing ad-hoc installation procedures. Patch administrators standardize patch installation by providing pre-install and post-install scripts, install and uninstall flags, reboot instructions, and how to handle error codes from the pre-install and post-install scripts

Supported Windows Patch Types

The following table lists the Windows patch types that Patch Management supports.

Table 5-2: Windows Patch Types

OS VERSIONS	PATCH TYPES
Windows NT 4.0	Windows Hotfix Windows OS Service Pack
Windows 2000	Windows Hotfix Windows OS Service Pack Update Rollup
Windows 2003	Windows Hotfix Windows OS Service Pack Update Rollup
Windows XP	Windows Hotfix Windows OS Service Pack Update Rollup

Supporting Technologies for Patch Management

Patch Management uses patching utilities and technologies for each supported Windows operating system. Opware SAS uses these tools behind the scenes. This allows you to perform patch management through a single interface, without having to worry about invoking a number of different patching utilities.

The following patch management and installation tools are used for the supported Windows operating systems:

- **mbsacli.exe**: Lists and verifies patches that are installed on a managed server. It also detects the application files that are already installed on a managed server and, subsequently, recommends the correct patch to install if multiple patches have the same QNumber.
- **msiexec.exe**: Installs and uninstalls MSI packages.
- **qchain.exe**: Enables a single reboot when you are installing more than one hotfix.
- **unzip.exe**: Extracts info-zip compatible zip archives.
- **Windows Update Agent**: Enables access to the Microsoft framework for patch updates.

See "Importing Windows Patch Utilities" on page 291.

Windows Hotfixes

After a Microsoft Windows hotfix is imported into Opware SAS, you can specify options to reboot the server when a hotfix is installed or uninstalled. A Windows hotfix typically requires a reboot if it updates system files. This reboot enables Opware SAS to use the newly updated system files.

When a hotfix is installed along with other hotfixes, this process is called hotfix chaining. If one or more hotfixes require that the server is rebooted, the reboot can sometimes be postponed until all hotfixes have been installed. The user performing the installation must first run qchain.exe before performing the reboot. This ensures that the Pending File Rename Queue is correctly ordered.

Postponing reboots is not always possible, due to a defect in qchain.exe that was resolved in December 2002. All Windows hotfixes created after May 2001 included the Pending File Rename Queue manipulation logic in qchain.exe. Therefore, all hotfixes created between May 2001 and December 2002 are vulnerable to the same qchain.exe defect. See the Microsoft Article for Q815062.

If a Windows Service Pack or Security Rollup Package is being installed in the same hotfix chaining process, a reboot is required. This reboot cannot be postponed. Before the reboot that is associated with this package occurs, qchain.exe must be run.

When multiple hotfixes are chained by Opsware SAS, the setting that specifies that a reboot on install is required for each hotfix is honored. Opsware SAS analyzes the set of hotfixes being installed to determine whether one or more reboots can be postponed until the end of the chaining operation.



If you are installing a Windows hotfix that does not support the `-z` flag, remember to use the `/-z` option to prevent the Patch Management feature from passing in the `-z` flag.

Opsware SAS examines the date each hotfix was created to determine whether any associated reboot can be safely postponed until the end of the chained installation.

Opsware SAS will *not* change the installation order of the chained hotfixes (as an attempt to further reduce the number of reboots), whether or not Service Pack or Security Rollup Packages are being installed in the chained operation.

When Opsware SAS installs a hotfix in isolation (not as part of a chained installation operation), Opsware SAS honors the value of the reboot on the installation operation.

Opsware SAS runs qchain.exe on the managed server after the installation of each Windows hotfix and before any associated reboot. This guards against problems associated with an incorrectly ordered Pending File Rename Queue. This problem could occur if another hotfix was installed on the managed server outside of Opsware SAS.

Searching for Patches and Policies

In the SAS Client, you can search for information about your operational environment by using the SAS Client Search feature. The Search feature enables you to search for patches, patch policies, servers, and so on. See “SAS Client Search” in the *Opsware® SAS User’s Guide: Server Automation*.

Patch Management Roles for Windows

Opsware SAS provides support for rigorous change management by assigning the functions of patch management to several types of users in an organization. These users include a policy setter, a patch administrator, and a system administrator.

Policy Setter

The policy setter is a member of a security standards group that reviews patch releases and identifies the vendor patches that will be included in the organization's patch policies. A policy setter is responsible for reviewing the latest security threats and the patches that vendors have released to address these problems. A policy setter is generally known as an expert in the operating systems and applications that they manage, and is able to assess the necessity of applying patches issued by vendors. A policy setter is also able to diagnose common problems that arise after patches are installed, allowing for a thorough test of the patch application process.

Patch Administrator

The patch administrator has the authority to import, test, and edit patch options. The patch administrator is often referred to as the security administrator in an organization. A patch administrator is granted specific permissions to import patches into Opsware SAS to test the patches and then mark them as available for use. Basic users can import patches, but they cannot install them or mark them as available. Patch administrators are also able to edit patch options (such as installation scripts) through patch management. Other types of users are not allowed to import or edit patches. Typically, a patch administrator imports the Microsoft patch database and tests patches on non-production reference hardware. After testing the patches and determining that the patches are safe to apply to production systems, a patch administrator marks the patches available in the Library and then advises the system administrators that they must apply the approved patches.

System Administrator

The system administrator installs patches (that have been approved for use) uniformly and automatically, according to the options that the patch administrator specifies. The system administrator is an Opsware user who is responsible for the day-to-day maintenance of the servers in a deployment. These users are not required to have the same level of expertise in low-level system details as the policy setter and patch administrator. Because the patch administrator has set up the patch installation, the

system administrators can attach policies to servers, set an exception for a patch, and install patches on a large number of managed servers. They are responsible for searching for servers that require the approved patch, installing the patches, and verifying that the patches were successfully installed. The system administrator can import patches but cannot install a patch until the patch administrator has marked it as available. The system administrator can also uninstall patches.



These responsibilities are enforced by assigning permissions for managing patches in Opsware SAS. To obtain these permissions, contact your Opsware administrator. See the *Opsware[®] SAS Administration Guide*.

Patch Management Process


The Windows patching process consists of several key phases: setup, policy management, patch compliance, and deployment. Setup steps include getting the Microsoft database (patches and metadata) into Opsware SAS, identifying products you want to track patches for, and configuring patch compliance. Policy management steps include investigating released patches, creating and updating patch policies or exceptions, marking patches available to use, and attaching policies or exceptions to servers or groups of servers. Patch compliance steps include running compliance scans to determine whether a server is out of compliance, remediating policies, setting up installation options, and installing applicable patches. To deploy patches on demand, you

can import the required patches, test them, update policies, create new policies, mark them as available to use, specify install options, and install the required patches. Figure 5-2 and Figure 5-3 illustrate these phases and steps.

Figure 5-2: Windows Patching Process: Part A and Part B

WINDOWS PATCHING PROCESS

Part A: Set Up Patch Management



Part B: Create and Attach Patch Policies to Servers

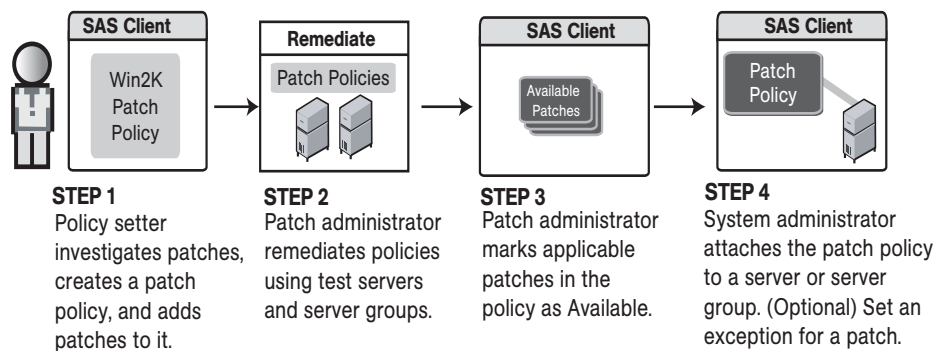
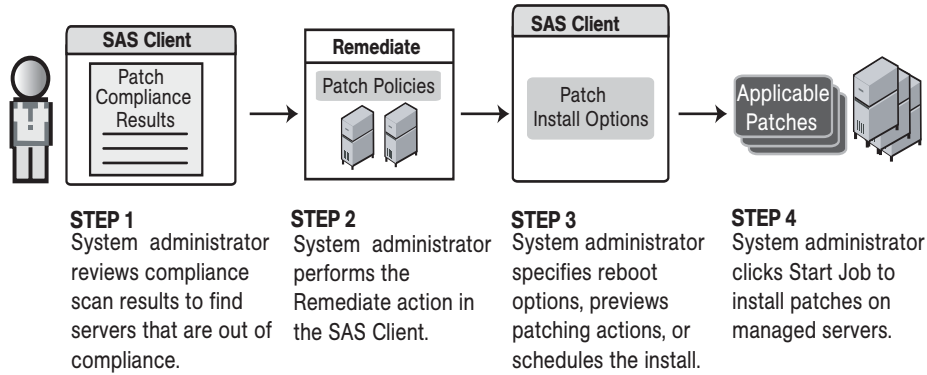


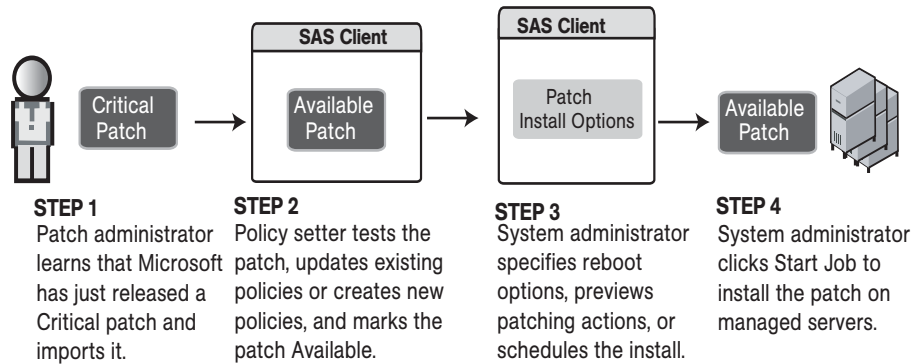
Figure 5-3: Windows Patching Process: Part C and Part D

WINDOWS PATCHING PROCESS

Part C: Install Patches By Remediating Policies



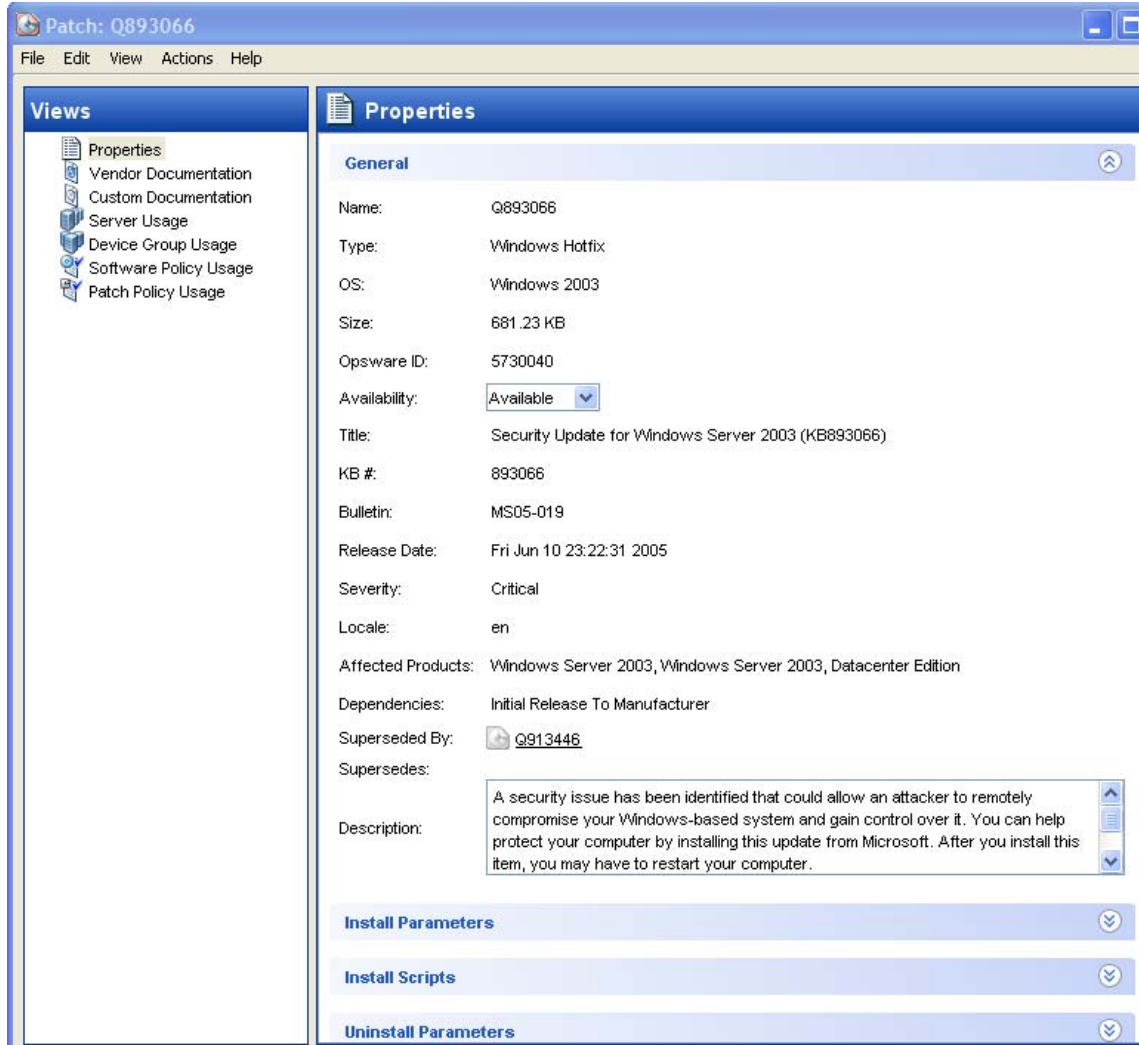
Part D: Install Patches on Demand



Patch Properties

Patch Management displays detailed information (properties) about a patch.

Figure 5-4: Windows Patch Properties



Patch properties include the following information:

- **Name:** The Microsoft name of the patch, such as QNumber, Windows 2000 Service Pack 4, and so on.
- **Type:** The type of patch, such as Windows Hotfix or Windows Update Rollup.
- **OS:** The Windows operating systems that are known to be affected by this patch.

- **Size:** The size of the patch file, in kilobytes (KB) or in megabytes (MB).
- **Opware ID:** The Opware SAS unique ID for the patch.
- **Availability:** The status of a patch within Opware SAS, which can be one of the following:
 - **Not Imported:** The patch is listed in the Microsoft Patch Database, but has not been imported (uploaded) into Opware SAS.
 - **Limited:** The patch has been imported into Opware SAS but cannot be installed. This is the default patch availability.
 - **Available:** The patch has been imported into Opware SAS, tested, and has been marked available to be installed on managed servers.
 - **Deprecated:** The patch cannot be added to patch policies or set as a patch policy exception but can still be installed.
- **Title:** The title of the Microsoft Knowledge Base article for this patch.
- **KB #:** The Microsoft Knowledge Base article ID number for this patch.
- **Bulletin** (Optional): The Microsoft Security Bulletin ID number for this patch.
- **File Name:** The name of the .exe for this patch.
- **Release Date:** The date that Microsoft released this patch.
- **Severity** (Optional): The Microsoft severity rating for this patch, which can be one of the following:
 - **Critical:** A patch that if exploited could allow the propagation of an internet worm, without user action.
 - **Important:** A patch that if exploited could result in a compromise of the confidentiality, integrity, or availability of user data, or of the integrity or availability of processing resources.
 - **Moderate:** A patch that if exploited could result in minimal impact. Exploitability is mitigated to a significant degree by certain factors, such as default configuration, auditing, or difficulty of exploitation.
 - **Low:** A patch that is difficult to exploit or if exploited, could result in minimal impact.
- **Locale:** The locale this patch applies to.

- **Affected Products:** Information from MBSA that identifies other Microsoft software that is known to be affected by this patch.
- **Dependencies:** Microsoft products that this patch requires. The patch cannot be installed if these products do not already exist on the server.
- **Superseded By** (Optional): A list of patches that this patch is superseded by. This relationship does not apply to MBSA 1.2.1 patches.
- **Supersedes** (Optional): A list of patches that this patch supersedes. This relationship does not apply to MBSA 1.2.1 patches.

Patch Dependencies and Supersedence

Patch metadata identifies all known dependency and supersedence relationships between patches and Windows products, and between patches and other patches. Dependency relationships identify Windows products that must already exist on a server before you can install a certain patch. Supersedence relationships identify patches that supersede or are superseded by other patches. In Patch Management, *supersedes* means that one patch replaces another and *superseded by* means that the patch you are installing is replaced by another patch.

For all MBSA 2.0 patches, Patch Management analyzes this information to determine the viability of a patch installation. For example, if you are remediating patches and a superseding patch is already installed, the patch will not be installed. If you try to install a superseded patch and the superseding patch is available and included in a patch policy, the superseded patch will not be installed. Patch Management does not analyze this information for MBSA 1.2.1 patches.



Patch Management does not detect whether two patches are mutually exclusive, which is when either one can be installed but not both. Subsequently, Patch Management does not prevent you from installing both patches on a server. This means that you may be able to install both a superseded patch and a superseding patch on a server.

Viewing Windows Patches

The SAS Client displays information about Microsoft Windows patches that have been imported into Opware SAS.

To view information about a patch, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand Patches and select a specific Windows operating system.
The Content pane will display all of the patches listed in the Microsoft Patch Database for the Windows operating system that you selected.
- 3** (Optional) Use the column selector to sort the patches according to Name, Type, Severity, Availability, Release Date, and Bulletin Number.
- 4** In the Content pane, open a patch to view its properties in the Patch window.

Editing Windows Patch Properties

You can edit a patch's Description, Availability, Install Parameters, and Uninstall parameters. Due to the nature of the type of patch, some properties are not editable. For example, you cannot turn the reboot-on-install option of a Windows Service Pack off.

The Availability property indicates the status of the patch in Opware SAS. If the Availability is Not Imported, you cannot change this property.

You can set the install and uninstall parameters on either the patch properties page or in the Patch Actions only when you are installing or uninstalling one patch at a time. The parameters on the properties page are saved in the Model Repository, but the parameters in Patch Actions are used only for that action. The parameters in Patch Actions override those on the patch properties page.

To edit the patch properties, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3** In the Content pane, open a patch to view its properties in the Patch window.
- 4** Edit any of the following fields: Description, Availability, and the Install and Uninstall parameters.

- 5 From the **File** menu, select **Save** to save your changes.

Importing Custom Documentation for a Patch

The Custom Documentation view of a patch displays text files that have been imported from the local file system. Non-plain text file types, such as .html or .doc, are not supported.

To import your own documentation for a patch, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 In the Content pane, open a patch to view its properties in the Patch window.
- 4 From the Views pane, select Custom Documentation.
- 5 From the **Actions** menu, select **Import Custom Documentation** or click **Import**.
- 6 In the Import Custom Documentation window, locate a text file and specify encoding.
- 7 Click **Import**.

Deleting Custom Documentation for a Patch

The Custom Documentation view of a patch displays text files that have been imported from the local file system. Non-plain text file types, such as .html or .doc, are not supported.

To delete custom documentation for a patch, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 In the Content pane, open a patch to view its properties in the Patch window.
- 4 From the Views pane, select Custom Documentation.
- 5 From the **Actions** menu, select **Delete Custom Documentation**.
- 6 In the Delete Custom Documentation window, click **Delete**.

Finding Vendor-Recommended Windows Patches

To find the patches that Microsoft recommends for a particular server (based on MBSA 2.0), perform the following steps:

- 1 From the Navigation pane, select Devices ► Servers ► All Managed Servers.
- 2 From the View drop-down list, select Patches.
- 3 From the Content pane, select a server that is running Opware Agent 5.5 and a Windows 2000 with Service Pack 3 (or higher) operating system or a Windows 2003 operating system.
- 4 From the Preview pane, select Patches Recommended By Vendor from the drop-down list. This displays the types of patches for the selected server.

Finding Servers That Have a Windows Patch Installed

To find the servers that have a particular patch installed, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list in the Content pane, select Server Usage.
- 5 From the Show drop-down list for the selected patch, select Servers with Patch Installed.

You can browse a server in this list to view a list of all installed patches. Please note that this list may display a more complete list of installed patches than the list you will find in the Windows Add or Remove Programs utility.

Finding Servers That Do Not Have a Windows Patch Installed

To find the servers that do not have a particular patch installed, perform the following steps:

- 1 From the Navigation pane, select Library ► Patches.

- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list, select Server Usage.
- 5 From the Show drop-down list, select Servers without Patch Installed.

Importing a Patch

Windows patches are downloaded from the Microsoft web site and then imported (uploaded) into Opware SAS. To see if a patch has been imported, view the patch's Availability property. The Availability of an imported patch is either Limited, Available, or Deprecated. A patch can be imported with the SAS Client or with a script. For information about the script, see "Automatically Importing Windows Patches" on page 260.

To import a patch with the SAS Client, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Package Repository.
- 2 Expand the Package Repository and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.
- 4 To import a patch directly from the Microsoft web site, from the **Actions** menu, select **Import ► Import from Vendor**.

The Import from Vendor window displays the URL of the patch's location on the Microsoft web site. You can override this URL, as needed.

Or

To import a patch that has already been downloaded to your local file system, from the **Actions** menu, select **Import ► Import from File**.

In the file browser window, locate the patch.

- 5 Click **Import**.

Automatically Importing Windows Patches

The `populate-opware-update-library` shell script downloads the Microsoft Patch Database and patches from the Microsoft site. The script also imports the database and patches into Opware SAS. (To be imported, a patch must be in the Microsoft Patch

Database that has been imported into the Software Repository.) Optionally, the script sets the initial status (Available or Limited) of newly imported patches. The script can also filter the patches imported according to operating system (such as Windows 2003). The functionality of the script is also available in the SAS Client, as described in “Importing the Microsoft Patch Database” on page 289.

To run the `populate-opsware-update-library` script, you need to log onto the Software Repository server as `root`. Typically, you schedule the script to run periodically as a `cron` job on the Software Repository server. To end users of the SAS Client, the patches imported with the script appear to have been automatically imported. Do not run concurrent instances of the script.

The `populate-opsware-update-library` script is in the following directory:

```
/opt/opsware/mm_wordbot/util/
```

Table 5-3 describes the script's options.

Table 5-3: Options of `populate-opsware-update-library`

OPTION	DESCRIPTION
<code>--spin hostname-or-IP</code>	Hostname or IP address of Data Access Engine (spin) host. Default value: spin
<code>--theword hostname-or-IP</code>	Hostname or IP address of Software Repository (theword) host. Default value: theword
<code>--cert_path file-path</code>	File specification of cert file to be used for Spin connection. Default value: <code>/var/opt/opsware/crypto/wordbot/wordbot.srv</code>
<code>--ca_path file-path</code>	File specification of CA file to be used for Spin connection. Default value: <code>/var/opt/opsware/crypto/wordbot/opsware-ca.crt</code>
<code>--verbose</code>	Display copious output, including patches skipped during the upload.

Table 5-3: Options of `populate-opsware-update-library` (continued)

OPTION	DESCRIPTION
<code>--no_nt4</code>	Do not process NT4 patches.
<code>--no_w2k</code>	Do not process W2K patches.
<code>--no_w2k3</code>	Do not process W2K3 patches.
<code>--no_w2k3x64</code>	Do not process Windows 2003 (64 bit) patches.
<code>--no_xp</code>	Do not process Windows XP (32 bit) patches.
<code>--use_proxy_url url</code>	When downloading binaries, connect via this proxy URL.
<code>--proxy_userid userid</code>	Basic-auth userid to provide to proxy server.
<code>--proxy_passwd passwd</code>	Basic-auth passwd to provide to proxy server.
<code>--set_available</code>	Set availability status to Available when uploading patches. The <code>--set_available</code> and <code>--set_limited</code> options cannot be specified at the same time.
<code>--set_limited</code>	Set availability status to Limited when uploading patches.
<code>--no_hotfixes</code>	Do not upload hotfixes.
<code>--no_servicepacks</code>	Do not upload servicepacks.
<code>--no_updaterollups</code>	Do not upload updaterollups.
<code>--no_wsusscan_upload</code>	Do not upload the MBSA 2.0 patch database.
<code>--wsusscan_url_override url</code>	Download the MBSA 2.0 patch database from this URL.
<code>--update_all</code>	Refresh the patches already uploaded into Opsware SAS.

Table 5-3: Options of populate-opsware-update-library (continued)

OPTION	DESCRIPTION
<code>--download_only path</code>	Download files from the vendor's web site to the specified path (directory), but do not upload them into Opsware SAS. The files are downloaded into the <code>platform_ver/locale</code> subdirectory beneath the specified path.
<code>--upload_from_update_root path</code>	Upload files from the specified path (directory), not from the vendor's web site. The script looks for patches in the <code>platform_ver/locale</code> subdirectory beneath the specified path. If it cannot find the patch in the that subdirectory, the script looks for the patch in the specified path. If a patch is not found, the script skips the patch and does not upload it. This option is ignored if <code>--download_only</code> is also specified.
<code>--help</code>	Display the syntax of this script.

Exporting a Windows Patch

To export a patch from Opsware SAS to the local file system, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3** From the Content pane, select a patch.
- 4** From the **Actions** menu, select **Export**.
- 5** In the Export Patch window, enter the folder name that will contain the patch file in the File Name field.
- 6** Click **Export**.

Exporting Windows Patch Information

You can export information about patches installed on a server and patches recommended by the vendor. You can also export information from patches recommended by the vendor along with model information on the selected server (such as patch policies or patch policy exceptions). The following information is exported into a .csv file:

- **Server Name:** The name of the managed server.
- **OS:** The operating system of the server.
- **Service Pack:** The service pack level of the server being reported, such as Service Pack 0, Service Pack 1, and so on.
- **KB#:** The Microsoft Knowledge Base Article number for the patch.
- **Bulletin:** The MSYY-XXX ID associated with a hotfix, such as MS05-012, MS06-012, and so on. If the MSYY-XXX ID is unknown, this column will be blank.
- **Description:** A brief description of the purpose of the patch.
- **Time Queried:** The last software registration by the Agent.
- **Time Installed:** The time that the patch was installed.
- **Type:** The patch type.
- **Compliance Level:** An integer that represents the compliance level.
- **Compliance:** Text that displays when you place your cursor over the Compliance column in the Patch Preview pane.
- **Exception Type:** The type of exception, such as Always Install or Never Install.
- **Exception Reason:** A description that explains the purpose of the exception.



Patch Management will display all of the text, including commas, from the Description field displayed in the Patch Properties window in the Description column in the .csv file. To preserve commas in the Description column and keep all text together in that column, double quotes will be converted to single quotes. This does not distort the semantics of the patch description.

To ensure that all of the text about a patch displays in the Description field in the .csv file, Patch Management surrounds the entire description (that you see in the Patch Properties window) with double quotes.

To export the patch information to a .csv file, perform the following steps:

- 1** From the Navigation pane, select **Devices** ► **All Managed Servers**.
- 2** From the Content pane, select one or more managed servers.
- 3** From the Show drop-down list, select an option.
- 4** From the **Actions** menu, select **Export Patch Info to CSV**.
- 5** In the Export to CSV window, navigate to a folder and enter the file name.
- 6** Verify that the file type is Comma Separated Value Files (.csv). If you did not include the .csv extension in the file name field, Patch Management will append it only if you have the .csv file type selected.
- 7** Click **Export** to save the patch information in a .csv file or click **Cancel** if you do not want to export the patch information.

Deleting a Patch

When you delete a patch, it is removed from Opsware SAS, but it is not uninstalled from managed servers. A patch cannot be deleted if it is attached to a policy or if an exception has been set for it.



Do not delete all of the patches from Opsware SAS. If you do so accidentally, contact your Opsware, Inc. support representative for assistance in importing the patches back into Opsware SAS.

To delete a patch, perform the following steps:

- 1** From the Navigation pane, select **Library** ► **By Type** ► **Patches**.
- 2** Expand **Patches** and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3** From the Content pane, select a patch.
- 4** From the **Actions** menu, select **Delete Patch**.

- 5 In the Delete Patches windows, click **Delete**.

Policy Management

In Patch Management, patch policies and patch policy exceptions enable you to customize patch distribution in your environment. Policies and exceptions define the Windows patches that should be installed or not installed on certain managed servers.

You can choose to have patching in your server environment comply to the model that these policies and exceptions define or you can choose to deviate from this model. If you choose to deviate from the patch policies and exceptions and perform ad hoc patch installs, then you need to remediate. The remediation process ensures that the applicable patches get installed on servers.

Patch Policy

A patch policy is a group of patches that you want to install on Opware SAS managed servers. All patches in a patch policy must apply to the same Windows operating system.

A patch policy provides broad flexibility for distributing patches. For example, you can create a patch policy that contains security patches that you want to distribute only to servers used by your sales force. You can also create a patch policy that contains security patches that are applicable to specific software that is already installed on a server, such as Exchange Server, Internet Information Services (IIS), SQL Server, and so on. Or, you can create a patch policy that includes all patches ranked critical (by Microsoft) and installs them on all servers that are used by everyone in your organization.



If you do not want to create a patch policy, you can use the vendor-recommended set of patches (by operating system) as a default patch policy, such as the patches provided by MBSA.

You can attach as many patch policies as you want to servers or groups of servers. If several policies are attached to one server, the installation logic is cumulative—all patches listed in all attached policies will be installed on the server. The Remediate window allows you to select an individual patch policy to remediate. You do not have to remediate all policies attached to a server. You cannot nest patch policies.

If a description of the patch policy is defined, it is recorded in the server's patched state (in the Model Repository). This information enables Patch Management to report on patch policies for patch compliance purposes. The patch compliance process compares patch policies with corresponding patch policy exceptions.

Patch Management supports the following types of patch policies:

- **User-defined patch policy:** This allows an Opware SAS user to specify the patches that are included in a policy. User-defined patch policies can be edited or deleted by a user who has permissions.

A user-defined patch policy allows a policy setter to opt out of patches. The policy setter can create a (user-defined) patch policy that is a subset of all available patches (that are in a vendor-recommended patch policy). This enables the policy setter to apply only those patches that their environment needs.

- **Vendor-recommended patch policy:** Membership of patches is defined by MBSA recommendations on a server-by-server basis. Vendor-recommended patch policies are system defined and cannot be edited or deleted by a user.



You can only export user-defined patch policies. You cannot export vendor-recommended patch policies.

Patch policies have the following characteristics:

- All patches in a patch policy must apply to the same operating system, such as Windows.
- A patch policy is associated with an operating system version, such as Windows 2003.
- A patch policy has a name and can (optionally) include a description that explains its purpose.
- A patch policy can be either user-defined or vendor-defined.
- A patch policy does not have sub-policies. There is no inheritance.
- A patch policy is Customer Independent, which means that patches in the policy can be installed on any managed server, no matter what customer is associated with it. See the *Opware[®] SAS User's Guide: Server Automation*.
- A patch policy is always public.

- A patch policy can be attached to zero or more servers or public device groups.
- More than one patch policy can be attached to a server or public device group.
- Only user-defined patch policies can be created, edited, and deleted by a user who has permissions.

Patch Policy Exception

A patch policy exception identifies a single patch that you want to explicitly include or exclude from a specific managed server, along with an optional reason for why the exception exists. The patch in a patch policy exception must apply to the same Windows operating system that the established patch policy is attached to.

A patch policy exception allows you to deviate from an established patch policy (one that is already attached to a server or a group of servers). You can do this by deselecting or adding individual patches to a server. Since patch policy exceptions override all patch policies attached to a server, you can use them to intentionally deviate from a patch policy on a server-by-server basis.

If a reason for a patch policy exception is defined, the description is recorded in the server's patched state (in the Model Repository). This information enables Patch Management to report on patch policy exceptions for patch compliance purposes. The patch compliance results explain how patch policy exceptions compare with corresponding established patch policies. All users who have access to the managed server can view attached patch policy exceptions.

Patch Management supports the following types of patch policy exceptions:

- **Always Installed:** The patch should be installed on the server, even if the patch is not in the policy.
- **Never Installed:** The patch should not be installed on the server, even if the patch is in the policy.



If you ever need to override a patch policy exception, you can manually install a patch.

The following information summarizes characteristics of a patch policy exception:

- A patch policy exception can (optionally) include a description that explains its purpose.
- A patch policy exception can have a rule value of Never Installed or Always Installed.
- A patch policy exception can be set for one patch and one server of the same operating system version. If a patch policy exception is set for a public device group and a server in that group does *not* match the operating system version specified in the patch policy exception, the patch policy exception is *not* applied.
- A patch policy exception can be set, copied, and removed by users who have permissions.

Precedence Rules for Applying Policies

By creating multiple patch policies and patch policy exceptions (that are either directly attached to a server or attached to a group of servers), you control the patches that should be installed or not installed on a server. A precedence hierarchy in Patch Management delineates how a patch policy or a patch policy exception is applied to a patch installation. This hierarchy is based on whether the patch policy or patch policy exception is attached at the server or device group level.

The following precedence rules apply to policies and exceptions:

- Patch policy exceptions that are directly attached to a server always take precedence over patch policies that are directly attached to a server.
- Patch policies that are directly attached to a server take precedence over patch policies and patch policy exceptions that are attached to a public device group.
- Patch policy exceptions that are attached to a public device group take precedence over patch policies that are attached to a public device group.
- If a server is in multiple public device groups, a Never Installed patch policy exception type always take precedence over an Always Installed patch policyexception type for the same patch.

Remediation Process

To ensure patch compliance, Patch Management identifies vulnerable managed servers and simultaneously deploys patches to many servers when a remediation process is performed. The remediation process examines and applies an entire patch policy (including multiple policies) to the managed servers that it is attached to. A policy must be attached to a server or a group of servers before you can remediate the policy with that server or group.



The remediation process requires that the selected managed server is running Opware Agent 5.5 and a Windows 2000 Service Pack 3 (or higher) operating system or a Windows 2003 operating system. You cannot use the remediation process if the selected managed server is running a Windows NT4.0 operating system, a Windows 2000 RTM (no service pack), Service Pack 1, or Service Pack 2 operating system, or if the server is not running Opware Agent 5.5. Use the Install Patch window to install patches on servers that are running these operating systems or Opware Agents 4.5 or earlier.

As a best practice, each time you review the latest Microsoft patch releases and subsequently update a patch policy (by adding new patches to a policy), you should perform remediation. In these situations, a remediation process provides demand forecasting information. This allows you to determine how patch policy changes will impact servers that this policy is attached to.

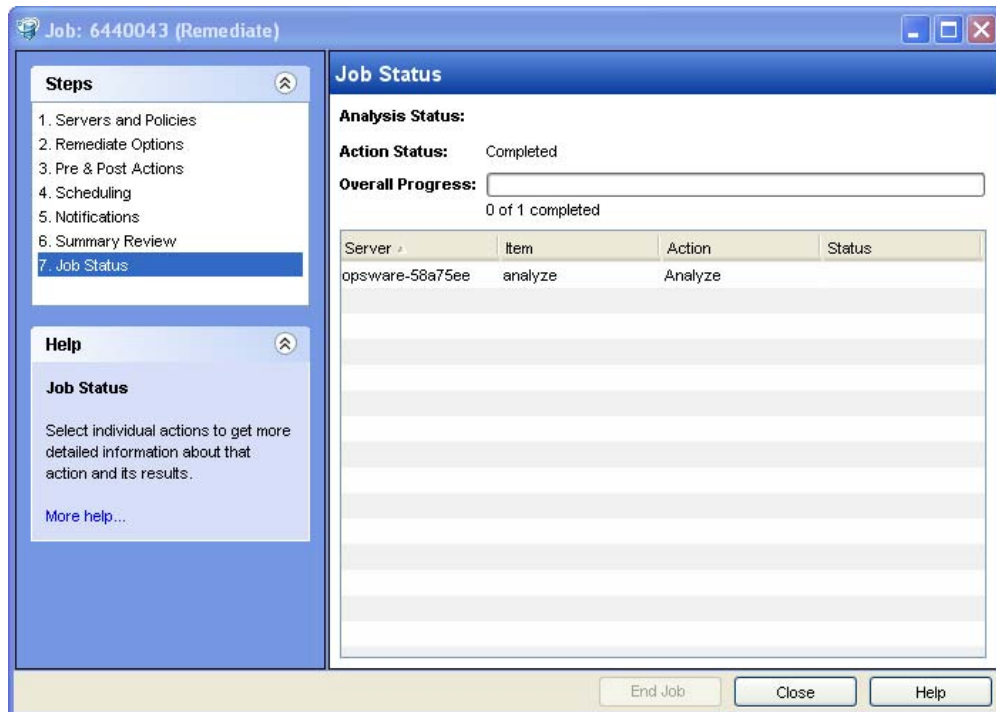
If the remediation process discovers any (applicable) missing patches, these patches will be installed on the servers.

If a patch was installed as part of a patch policy, the remediation process will not uninstall it. However, if a patch was installed as part of a software policy and it is no longer in the software policy, the remediation process will uninstall it.

After Opware SAS determines the packages that need to be installed to complete the remediation process, remediation uses a set of standard system utilities to complete the operation. See “Supporting Technologies for Patch Management” on page 248.

To help you optimally manage the remediation conditions, Patch Management allows you to specify remediate options and pre and post actions, and set up ticket IDs and email notifications that alert you about the status of the remediate process. The Remediate window guides you through setting up these conditions.

Figure 5-5: Remediate Window



Remediating Patch Policies

This action installs the patches in a policy that has been attached to managed servers. (This action does not uninstall patches.) A patch policy can be overridden by an exception, which indicates that a patch is either always or never installed on a particular server.

When you invoke the remediation process for a group of servers, patches will only be remediated if any server in the group is running Opsware Agent 5.5 and a Windows 2000 with Service Pack 3 (or higher) operating system or a Windows 2003 operating system. The Remediate option is not available in the Actions menu if the selected server is not running Opsware Agent 5.5 and a Windows 2000 with Service Pack 3 (or higher) operating system or a Windows 2003 operating system.

To remediate a patch policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patch Policies
- 2** Expand Patch Policies and select a specific Windows operating system. The Content pane will display all patch policies associated with that operating system.
- 3** From the Content pane, open a patch policy.
- 4** From the View drop-down list, select Server Usage.
- 5** From the Show drop-down list in the Content pane, select Servers with Policy Attached.
- 6** From the Preview pane, select one or more servers.
- 7** From the **Actions** menu, select **Remediate**. The first step of the Remediate window appears: Servers and Device Groups.

For instructions on each step, see the following sections:

- Setting Remediate Options
- Setting Reboot Options for Remediation
- Specifying Pre and Post Install Scripts for Remediation
- Scheduling a Patch Installation for Remediation
- Setting Up Email Notifications for Remediation
- Previewing a Remediation

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8** Click **Start Job** to launch the remediation job.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If you leave the Remediate window open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select **Refresh** from the **View** menu to update information in the Patch Preview pane.

Setting Remediate Options

You can specify the following remediate policy option:

“Do not interrupt the remediate process even when an error occurs with one of the policies.”

To set this option, perform the following steps:

- 1** From the Remediate window, click **Next** to advance to the Remediate Options step.
- 2** Select one of the following Staged Install Options:
 - Continuous:** Run all phases as an uninterrupted operation.
 - Staged:** Allow download and installation to be scheduled separately.
- 3** Select the Error Options check box if you want the remediation process to continue even when an error occurs with any of the patches or scripts. As a default, this check box is not selected.
- 4** Click **Next** to go to the next step or click **Cancel** to close the Remediate window.

Setting Reboot Options for Remediation

To minimize the downtime that server reboots can cause, you can control when servers reboot during a patch installation.

You can specify the reboot options in the following two places in the SAS Client:

- Install Parameters tab of the patch properties window
- Pre & Post Actions step of the Remediate window



When you are selecting reboot options in the Remediate window, Opware, Inc. recommends that you use Microsoft’s reboot recommendations, which is the “Reboot servers as specified by patch properties” option. If it is not possible to use the Microsoft reboot setting, select the single reboot option, which is the “Do not reboot servers until all patches are installed” option. Failure to do this can result in the MBSA incorrectly reporting which patches are installed on the server until the next reboot occurs (outside of Opware control).

The following options in the Remediate window determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Remediate window; they do not change the Reboot Required option, which is on the Install Parameters tab of the Patch Properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Hold all server reboots until after all packages are installed and/or uninstalled:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1** From the Remediate window, click **Next** to advance to the Pre & Post Actions step.
- 2** Select one of the Reboot Options.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Remediate window.

Specifying Pre and Post Install Scripts for Remediation

For each patch remediation, you can specify a command or script to run before or after remediation. A pre-install script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-install script fails, the patches would not be installed. A pre-install script could also be used to shut down a service or application before it is patched. A post-install script could be used to perform a certain cleanup process on the managed server.

You can specify the following types of scripts to run on the managed server before or after a remediation process:

- **Pre-Download:** A script that runs before patches are downloaded from Opware SAS to the managed server. This is available only if you select Staged in the Remediate Options step.

- **Post-Download:** A script that runs after patches are downloaded from Opsware SAS to the managed server and before the patch is installed. This is available only if you select Staged in the Remediate Options step.
- **Pre-Install:** A script that runs before patches are installed on the managed server.
- **Post-Install:** A script that runs after patches are installed on the managed server.

To specify a pre-install script, perform the following steps:

1 From the Remediate window, click **Next** to advance to the Pre & Post Actions step.

2 Select the Pre-Install tab.

You may specify different scripts and options on each of the tabs.

3 Select the Enable Script check box. This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.

4 Select either Saved Script or Ad-Hoc Script from the drop-down list.

A Saved Script has been previously stored in Opsware SAS with the Opsware SAS Web Client. To specify the script, click **Select**.

An Ad-Hoc script runs only for this operation and is not saved in Opsware SAS. Select the Type, such as .bat. In the Script box, enter the contents of the script, including the drive letter of where the script is located, such as `echo dir>> C:\temp\preinstall1.log`. If you do not enter a drive letter, the default is %SYSTEMDRIVE%, which is where the system folder of Windows is installed.

5 If the script requires command-line flags, enter the flags in the Command text box.

6 In the User section, if the system is not Local System, select Name.

7 Enter the system name, your password, and the Domain name.

8 To stop the installation if the script returns an error, select the Error check box.

9 Click **Next** to go to the next step or click **Cancel** to close the Remediate window

Scheduling a Patch Installation for Remediation

You can schedule when you want patches installed and when you want patches downloaded.

To schedule a patch installation, perform the following steps:

- 1 From the Remediate window, select the Scheduling step. To reach this step, you must have completed the Pre & Post Actions step.



By default, the Scheduling step displays only the scheduling options for the installation phase. If you selected Staged in the Remediate Options step, the scheduling options for the download phase will also be displayed.

- 2 Select one of the following Install Phase options:
 - **Run Task Immediately:** This enables you to perform the download or installation immediately.
 - **Run Task At:** This enables you to specify a date and time that you want the download or installation performed.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Remediate window.

Setting Up Email Notifications for Remediation

You can set up email notifications to alert users when the download and installation operations complete successfully or with errors.

To set up email notifications, perform the following steps:

- 1 From the Remediate window, click **Next** to advance to the Notifications step.
- 2 To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3 To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase. If you selected Staged in the Remediate Options step, the notification status for the download phase is also displayed.
- 4 Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5 Click **Next** to go to the next step or click **Cancel** to close the Remediate window.



If you previously selected Staged in the Remediate Options step, the Notifications pane displays notification options for both the download and installation phases.

Previewing a Remediation

The remediate preview process provides an up-to-date report about the patch state of servers. The remediate preview is an optional step that lets you see the patches that will be installed on managed servers. This preview process verifies whether the servers you selected for the patch installation already have that patch installed (based MBSA 2.0). In some cases, a server could already have the patch installed if a system administrator had manually installed it, which means that Patch Management does not know about it.

In the Preview, the servers, device groups, and patches that are listed in the Summary Step window will be submitted to remediation when you click **Start Job**. Patches that are not recommended by the vendor will be excluded from this list. If there are other patches in the policy with the same QNumber, only the vendor-recommended patch is displayed.

This list shows patches and their associated servers (regardless of any patch policy and server group membership changes that may have occurred). If you preview a remediation, this same list of servers, device groups, and patches will be used, even if changes have occurred to the patch policy or server group memberships.

If you modify parameters in the Remediate window after you have already clicked **Preview**, the preview process will produce an invalid summary of simulated patching actions. For example, if you have already clicked **Preview** and you add patches, patch policies, servers, or device groups, you must click **Preview** again for results that include your changes.



The remediation preview does not report on the behavior of the server as though the patches have been applied.

To preview a remediation, perform the following steps:

- 1** From the Remediate window, click **Next** to advance to the Summary Review step.
- 2** Verify the information displayed for the Servers, Device Groups, and Patches at the top of the window.
- 3** (Optional) Click **Preview** to see the separate actions that will be performed when the patch is installed. To view the details of a previewed action, select a row in the table.
- 4** To launch the installation job, click **Start Job**.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected a specific time, the job will run then.

- 5** The Job Progress displays in the Remediate window.

The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:

- **Analyze:** Opsware SAS examines the patches needed for the installation, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.
- **Download:** The patch is downloaded from Opsware SAS to the managed server.
- **Install:** After it is downloaded, the patch is installed.
- **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
- **Run Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the installation.
- **Install & Reboot:** When a patch will be installed is also when the server will be rebooted.
- **Verify:** Installed patches will be included in the software registration.

- 6** To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See the *Opsware® SAS User's Guide: Server Automation* for more information on browsing job logs.

- 7** Click **Stop Job** to prevent the job from running or click **Close** to close the Remediate window. You can stop a job only if it is scheduled.

Verifying Patch Policy Compliance

To determine whether a managed server complies with patch policies and exceptions, perform the following steps:

- 1** From the Navigation pane, select Devices ► All Managed Servers.
- 2** From the Content pane, select Patches from the View drop-down list, select a server, and examine the Compliance column.

The Compliance column shows one of the following statuses:

- **Non-Compliant:** This status indicates one of the following conditions:
 - The patch is installed on the server and the exception is Never Install.
 - The patch is not installed on the server but is in the policy.
- **Partial:** The policy and exception do not agree, and the exception reason field is empty.
- **Compliant:** This status indicates one of the following conditions:
 - The patch is installed on the server and is in the policy.
 - The patch is installed on the server and the exception is Always Install.
 - The patch is not installed on the server and the exception is Never Install.
 - The patch is not installed on the server, but another patch with the same QNumber is in the policy.
 - The patch is not installed on the server, it is in the policy, but is not recommended by the vendor. (If a patch is not recommended by the vendor, it is treated as if it has a Never Install exception.)
- **No Indicator:** This status indicates one of the following conditions:
 - The patch is installed on the server, is not in a patch policy, and there is no exception.
 - The patch is not installed on the server, is not in a patch policy, and there is no exception.

The preceding list does not list all of the possible permutations of patch installations, policies, and exceptions. To view a table listing all permutations, see “Setting the Patch Policy Compliance Level” on page 291.

Creating a Patch Policy

A patch policy is a set of patches that should be installed on a managed server. When it is first created, a patch policy contains no patches and is not attached to servers.

To create a patch policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patch Policies.
- 2** Select a specific Windows operating system.
- 3** From the **Actions** menu, select **Create Patch Policy**.

The name of the policy you just created is New Patch Policy n, where n is a number based on the number of New Patch Policies already in existence.

- 4** From the Content pane, open the New Patch Policy.
- 5** (Optional) In the Name field of the Properties, enter a name that describes the purpose or contents of the policy.

Deleting a Patch Policy

This action removes a patch policy from Opsware SAS but does not remove or uninstall patches from managed servers. You cannot delete a patch policy if it is attached to servers or groups of servers. You must first detach the policy from the servers or groups of servers before removing it from Opsware SAS.

To delete a patch policy from Opsware SAS, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patch Policies.
- 2** Select a specific Windows operating system.
- 3** From the Content pane of the main window, select a policy.
- 4** From the **Actions** menu, select **Delete Patch Policy**.

Adding a Patch to a Patch Policy

This action adds a patch to a patch policy, but does not install the patch on a managed server. The patch will be installed when the policy is remediated.

To add a patch to a patch policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patch Policies.
- 2** Select a specific Windows operating system and view the list of Windows patches.
- 3** From the Content pane, select the patch.
- 4** From the View drop-down list, select Patch Policies.
- 5** From the Show drop-down list, select Policies without Patch Added.
- 6** Select a policy. From the **Actions** menu, select **Add to Patch Policy**.
- 7** In the Add to Patch Policy window, click **Add**.

Removing a Patch from a Patch Policy

This action only removes a patch from a patch policy. This action does not uninstall the patch from a managed server and does not remove the patch from Opsware SAS.

To remove a patch from a patch policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Select a specific Windows operating system and view the list of Windows patches.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list, select Patch Policies.
- 5** From the Show drop-down list, select Policies with Patch Added.
- 6** Select a patch. From the **Actions** menu, select **Remove from Patch Policy**.
- 7** In the Remove Patch from Policy window, select the policy and click **Remove**.

Attaching a Patch Policy to a Server

This action associates a patch policy with a server (or group of servers). You must perform this action before you remediate a policy with a server (or group of servers).

To attach the policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patch Policies.
- 2** Select a specific Windows operating system and view the list of Windows patch policies.
- 3** From the Content pane, select a patch policy.
- 4** From the View drop-down list, select Server Usage (or Device Group Usage).
- 5** From the Show drop-down list, select Servers with Policy Not Attached (or Server Groups with Policy Not Attached).
- 6** From the Preview pane, select one or more servers.
- 7** From the **Actions** menu, select **Attach Server**.
- 8** Click **Attach**.

Detaching a Patch Policy from a Server

This action does not delete the patch policy and does not uninstall patches from a managed server.

To detach the policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patch Policies.
- 2** Select a specific Windows operating system and view the list of Windows patch policies.
- 3** From the Content pane, select a patch policy.
- 4** From the View drop-down list, select Server Usage (or Device Group Usage).
- 5** From the Show drop-down list, select Servers with Policy Attached (or Server Groups with Policy Attached).
- 6** From the Preview pane, select one or more servers.
- 7** From the **Actions** menu, select **Detach Server**.
- 8** Click **Detach**.

Setting a Patch Policy Exception

A patch policy exception indicates whether the patch is installed during the remediation process. (The Install Patch and Uninstall Patch actions ignore patch policy exceptions.) A patch policy exception overrides the policy. You can specify an exception for a particular patch and server (or group of servers), but not for a patch policy.

To set a patch policy exception, perform the following steps:





- 1** From the Navigation pane, select Devices ► All Managed Servers.
- 2** Select a server.
- 3** From the Content pane, select a server.
- 4** From the View drop-down list, select Patches.
- 5** From the Preview pane, select a patch.
- 6** From the **Actions** menu, select **Set Exception**.
- 7** In the Set Policy Exception window, select the Exception Type:

- **Never Install:** The patch should not be installed on the server, even if the patch is in the policy.
 - **Always Install:** The patch should be installed on the server even if the patch is not in the policy.
- 8** (Optional) In the Reason field, enter an explanation. This explanation is displayed when you move the cursor over the Exception column in the Preview pane. The Patches with Exceptions option must be selected.
 - 9** Click **OK**.

Finding an Existing Patch Policy Exception

You can search for managed servers that already have patch policy exceptions attached to them, and you can search for patches that have exceptions.

To find an existing patch policy exception, perform the following steps:

- 1** From the Navigation pane, select Devices ► All Managed Servers.
- 2** From the View drop-down list, select Patches.
- 3** From the Content pane, select a server.
- 4** From the Show drop-down list, select Patches with Policies or Exceptions or Patches with Exceptions.
- 5** In the Exception column, move the cursor over the icon to display the reason for this exception. The following icons indicate the type of patch policy exception:
 -  An always install exception on a patch/server association.
 -  An always install exception inherited to a server from a group of servers/patch association.
 -  A never install exception on a patch/server association.
 -  A never install exception inherited to a server from a group of servers/patch association.

Copying a Patch Policy Exception

To copy an exception between servers or groups of servers, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand the Patches and select a specific Windows operating system.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list, select Server Usage (or Device Group Usage).
- 5** From the Show drop-down list, select Servers with Exception (or Server Groups with Exception).
- 6** From the Preview pane, select a server. This server is the source of the copied exception.
- 7** From the **Actions** menu, select **Copy Exception**.
- 8** In the Copy Policy Exception window, select the target servers or device groups.

These servers are the destinations of the copied exception. If this operation would result in replacing an existing exception, a message displays asking you to confirm whether this is the preferred action.
- 9** Click **Copy**.

Removing a Patch Policy Exception

To remove a patch policy exception, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand the Patches and select a specific Windows operating system.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list, select Servers.
- 5** From the Show drop-down list, select Servers with Exception.
- 6** From the Preview pane, select a server.
- 7** From the **Actions** menu, select **Remove Exception**.

Patch Compliance

Patch Management performs conformance tests (compliance checks) against managed servers and public device groups to determine whether all patches in a policy and a policy exception were installed successfully. To enforce patch compliance, servers are scanned to determine whether they conform to their attached policies and exceptions. To optimize patch compliance information for your organization, you can set the patch compliance levels and edit the rules of the customized patch compliance level.

Patch Compliance Scans

A patch compliance scan compares patches that are installed on a server with patch policies and patch policy exceptions that are attached to that server. The results of this scan show you the servers that are in compliance (have all required patches installed) and the servers that are out of compliance (do not have all required patches installed).

A patch compliance scan occurs only when you request the scan or according to a schedule you have set up. There is no default schedule for running patch compliance scans in Opsware SAS. Opsware, Inc. recommends that you run or schedule patch compliance scans based on the dynamics of your patching environment. For example, if you updated a patch policy or installed a patch outside of (by not using) Opsware SAS, a compliance scan is required because the Opsware model has been changed and the compliance information must now be recalculated. Patch Management indicates these types of conditions by displaying Scan Needed. In this case, instead of waiting for the scan schedule to iterate, you can request the scan to update your patch compliance information.

To indicate whether a server is in compliance or out of compliance, Patch Management displays the following patch compliance icons:



The server is compliant for all patches. Patches in policies attached to the server are all installed on that server.



The server is partially compliant for patches. An exception has been set for these patches.



The server is not compliant for patches. Patches in policies attached to the server are not installed on that server.

Patch Compliance Levels

Patch compliance levels define your patch compliance rules. Results of a patch compliance scan can include only policies, both policies and exceptions, or your own customized level.

Patch Management supports the following compliance levels:

- **Policy Only:** Verifies whether the patches installed on a server comply with the patch policies.
- **Policy and Exception:** Verifies whether the patches installed on a server comply with the patch policies and any exceptions. The Partial (yellow) icon is displayed if the policy and exception do not agree and the exception does not have data in the Reason field.
- **Customized:** Verifies the rules that you edited for this compliance level.

Patch Compliance Rules

Patch compliance rules are the conditions that determine the compliance icons that are displayed in the Managed Server window.

Patch Management supports the following compliance rules:

- **Patch Added to Policy:** The patch has been added to the patch policy.
- **Patch Installed on Server:** The patch has been installed on the managed server.
- **Exception Type:** The Exception Type can have the following values:
 - **Always Installed:** The patch should be installed on the server, even if the patch is not in the policy.
 - **Never Installed:** The patch should not be installed on the server, even if the patch is in the policy.
 - **None:** An exception has not been specified for the patch and server.
- **Exception Reason:** A description entered in the Exception Reason of the Set Policy Exception window. In the Patch Compliance Rules window, the Exception Reason can have the following values.
 - **Yes:** The Exception Reason has data.
 - **No:** The Exception Reason is empty.
 - **N/A:** An exception has not been specified for the patch and server.

- **Compliance Result:** The icon that indicates the result of the patch compliance scan. These icons are displayed in the Managed Server window.

Patch Compliance Reports

To help troubleshoot problems that involve patches and patch compliance, you can run and examine several patch compliance reports that are based on Sarbanes-Oxley (SOX) standards. By using the Reports feature in the SAS Client, you can produce the following patch compliance reports. These reports identify whether all patches in a policy and a policy exception were installed successfully on managed servers in your environment.

Defined Patch Policies

This report lists patch policies by name, customer, and operating system, and includes the total number of patch policies.

Patch Policy Compliance (All Servers)

This report groups all managed servers by their patch policy compliance level to show compliant and non-compliant servers.

Patch Policy Compliance by Customer

This report lists all servers by the customer they belong to and then by the patch policy compliance level.

Patch Policy Compliance by Facility

This report groups all managed servers by the facility they belong to and then by the patch software policy compliance level.

Servers in Compliance With Their Patch Policies

This report lists all managed servers that are in compliance with all of their attached patch policies.

Servers Not in Compliance With Their Patch Policies

This report lists all managed servers that are not in compliance with their attached patch policies.

Servers With Attached Patch Policies

This report lists all managed servers that have one or more patch policies attached, and includes the total number of servers with attached patch policies.

Servers Without Attached Patch Policies

This report lists all managed servers that do not have any patch policies attached, and includes the total number of servers without any attached patch policies.



See the *Opsware® SAS User's Guide: Server Automation* for information about how to run, export, and print these reports.

Patch Administration for Windows

You can customize patch administration for Windows to best support your environment in the following manner:

- You can specify whether you want patches immediately available for installation by using a command-line script or the SAS Client.
- You can import the Microsoft patch database (on demand) by using a command-line script or the SAS Client.
- You can track (and import) only patches that apply to certain Microsoft products or particular locales.
- You can import and export Windows patch utilities.
- You can manually launch (on demand) or schedule periodic policy compliance scans to determine the patch state of your managed servers.
- You can customize the icon display of policy compliance scan results.

Setting the Patch Availability

You can set the default patch availability with either the SAS Client or a command-line script. The default used by the script overrides the default set by the SAS Client. For information about the script, see “Automatically Importing Windows Patches” on page 260.

To set the default value for the Availability of a newly imported patch, perform the following steps:

- 1 From the Navigation pane, select Opsware Administration.
- 2 Select Patch Settings.
- 3 For the Patch Availability for Imported Patches, select either Available or Limited. The default is Limited.

If the patch is Available, it can be installed on managed servers. If the patch is Limited, it has been imported into Opsware SAS and can be installed only by a patch administrator who has the required permissions. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide* for an explanation of these permissions.

Importing the Microsoft Patch Database

You can import the Microsoft Patch Database by using a command-line script or the SAS Client. For information about the script, see “Automatically Importing Windows Patches” on page 260.

To import the database with the SAS Client, perform the following steps:

- 1 From the Navigation pane, select Opsware Administration.
- 2 Select Patch Settings.
- 3 To import the database from the Microsoft web site, click **Import from Vendor**.

A window appears with the default URL for the location of the database on the Microsoft web site. Click **Import**. To re-import a new version of the Microsoft database that is released monthly, you must use the default URL.

- 4 To import the database from the local file system, click **Import from File**.

A file browser window appears. Go to the folder containing the `wsusscan.cab` (MSBA 2.0) file and click **Import**. This file must have been previously downloaded from the Microsoft web site and copied to the local file system.



To be imported, a patch must be in the Microsoft Patch database that has already been imported into the Software Repository.

Selecting Windows Products to Track for Patching

This operation limits the patches tracked by Opsware SAS to specific Windows products. After performing this operation, the next time the Microsoft Patch Database is imported, any new patches listed by Opsware SAS are limited to the products that you select. Patches that were previously listed by Opsware SAS are still tracked. You can also track patches for all MBSA 2.0 products.

To limit the patches tracked to specific Windows operating systems, run the command-line script that automatically imports patches. For more information about the script, see “Automatically Importing Windows Patches” on page 260.

To select the Windows products to track for patching, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Settings.
- 3** Select the Windows MBSA tab.
- 4** Click **Edit**.
- 5** In the Edit Patch Properties window, use the include and exclude arrows to select the products whose patches you want to track and then click **Select**.

Scheduling a Patch Policy Compliance Scan

To schedule or modify a patch compliance scan, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Compliance Settings.
- 3** In the Patch Policy Compliance Scan Schedule section, click **Edit**.
- 4** In the Schedule Compliance Scan window, select Enable Compliance Scan.
- 5** In the Schedule drop-down list, select the frequency of the scans.

If you select Custom, specify the crontab string with the following values:

- Minute (0-59)
- Hour (0-23)
- Day of the month (1-31)
- Month of the year (1-12)

- Day of the week (0-6 with 0=Sunday)
- Any of these fields can contain an asterisk to indicate all possible values. For example, the following crontab string runs the job at midnight every weekday:

```
0 0 * * 1-5
```

The crontab string can also handle serial (1,2,3,4) as well as range (1-5) values. For more information, consult the crontab man pages on a Unix computer.

- 6** In the Start Time field, specify the time you want the job to begin.
- 7** In the Time Zone drop-down list, select a default time zone for the job execution time or accept the default time zone. The default time shown converts the scheduled time to the time zone set in your user preferences. If you do not set a preferred time zone, the time zone is derived from the Opsware SAS core server, which is typically UTC.
- 8** In the Day(s) to Run field, select one or more days of the week that you want the scan to run.
- 9** Click **OK**.

Setting the Patch Policy Compliance Level

The patch policy compliance level defines your patch compliance rules. To view these rules or to set the patch policy compliance level, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Compliance Settings.
- 3** Select one of the following compliance levels: Policy and Exception, Policy Only, or Customized.

Importing Windows Patch Utilities

You can import the following Windows utilities from your local file system into Opsware SAS:

- mbsacl.exe
- parsembsacl20.exe
- qchain.exe

- WindowsUpdateAgent20-x86.exe
- WindowsUpdateAgent20-x64.exe
- wusscan.dll

Initially, these files are imported into Opsware SAS during the installation of the core. To import a Windows patch utility, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Settings.
- 3** In the Patch Utilities section, select a utility and then click **Import Utility Update**.

Exporting Windows Utility Files

You can export the following Windows patch utilities from Opsware SAS to your local file system:

- mbsacli.exe
- parsembsacli20.exe
- qchain.exe
- WindowsUpdateAgent20-x86.exe
- WindowsUpdateAgent20-x64.exe
- wusscan.dll

To export a Windows patch utility, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Settings.
- 3** In the Patch Utilities section, select one or more utilities and then click **Export Utility**.

Editing the Customized Patch Policy Compliance Level

Of the three compliance levels, only the Customized level can be edited. To edit this level, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Compliance Settings.
- 3** From the Compliance Level, select Customized.

- 4** In the Patch Policy Compliance Setting section, click **Edit**.
- 5** Select the Compliance Level icons that you want to change in the Compliance Result column: Non-Compliant, Compliant, No Indicator, or Partial.
- 6** Click **Apply** and then click **Close**.

Locales for Windows Patching

The locale of a patch identifies the language of the Windows servers that should receive the patch. A patch with the same name might be available for different locales. For example, a patch named Q123456 might be available for servers running the English and Japanese versions of Windows. Although they have the same name, the patches installed on the English and Japanese servers are different binaries.

Patch Management supports multiple locales in the same Opsware multimaster mesh. To install a patch on Windows servers with different locales, you specify the patch by name. During the installation (or policy remediation), Opsware SAS matches the locale of the patch with the locale of each managed server. You do not need to repeat the installation for each locale.

Supported Locales

Patch Management supports Windows patches of the following locales:

- English (en)
- Japanese (ja)
- Korean (ko)

Overview of Locale Configuration Tasks

By default, Patch Management supports only the English locale. To set up Patch Management for non-English locales, step through the instructions in the following sections:

- “Configuring the Opsware Core for Non-English Locales” on page 294
- “Selecting the Locales of Patches to Import” on page 294

- “End User Requirements for Non-English Locales” on page 295

Configuring the Opsware Core for Non-English Locales

This task requires `root` access to core servers and a restart the OCC core component. To configure the core for non-English locales, perform the following steps on each core server running the OCC component:

- 1** Log onto the server as `root`.
- 2** With a text editor, in `/etc/opt/opsware/occ/psrvr.properties`, change the line for `pref.user.locales` to the following:
`pref.user.localesAllowed=en;ja;ko`
- 3** Restart the OCC component of the core:
`/etc/init.d/opsware-sas restart occ.server`
- 4** In a text editor, open the following file:
`/opt/opsware/occclient/jnlp.tmp1`
- 5** For the Japanese language, In the `<resources>` section of the `jnlp.tmp1` file, add the following XML element:

```
<property name="com.opsware.ngui.font.japanese" value="Arial Unicode MS"/>
```

- 6** For the Korean language, In the `<resources>` section of the `jnlp.tmp1` file, add the following XML element:

```
<property name="com.opsware.ngui.font.korean" value="Arial Unicode MS"/>
```

- 7** In the `/opt/opsware/occclient` directory, if the following files exist, delete them:
`$HOST_ja.jnlp`
`$IP_ja.jnlp`
`$HOST_ko.jnlp`
`$IP_ko.jnlp`
- 8** Follow the steps in “Selecting the Locales of Patches to Import” on page 294.

Selecting the Locales of Patches to Import

Follow the instructions in “Configuring the Opsware Core for Non-English Locales” on page 294 before performing the steps in this section.

This operation selects the locales of the Windows patches to import into Opsware SAS. The selections take effect the next time patches are imported into Opsware SAS. After the patches have been imported, they can be installed on managed servers. If you remove locales from the list with this operation, patches with those locales that have already been imported are not removed from Opsware SAS.

To select the locales of the Windows patches to import into Opsware SAS, perform the following steps:

- 1** In the SAS Client, from the Navigation pane, select Opsware Administration.
- 2** Select Patch Settings.
- 3** On the Windows MBSA tab, select Patch Locales.
- 4** Click **Edit**.
- 5** In the Edit Patch Locales window, use the include and exclude arrows to select the locales whose patches you want to import. If you want to select a locale that is not listed in “Supported Locales” on page 293, contact Opsware Inc. Support.
- 6** Click **Select**.
- 7** Follow the instructions in “End User Requirements for Non-English Locales” on page 295.

End User Requirements for Non-English Locales

To view non-English fonts in the SAS Client, end users must perform the following steps:

- 1** The end user verifies that the Windows desktop running the SAS Client uses the Arial Unicode MS font.
- 2** After the Opsware Administrator performs the steps in “Configuring the Opsware Core for Non-English Locales” on page 294, the end user logs onto the SAS Web Client and goes to the My Profile page,
- 3** On the My Profile page, the end user updates the Locale field on the User Identification tab. For example, if the Opsware Administrator configured the core for Japanese, then the end user sets the Locale field to Japanese.

Patch Installation

Patch Management provides the following two phases in the patch installation process:

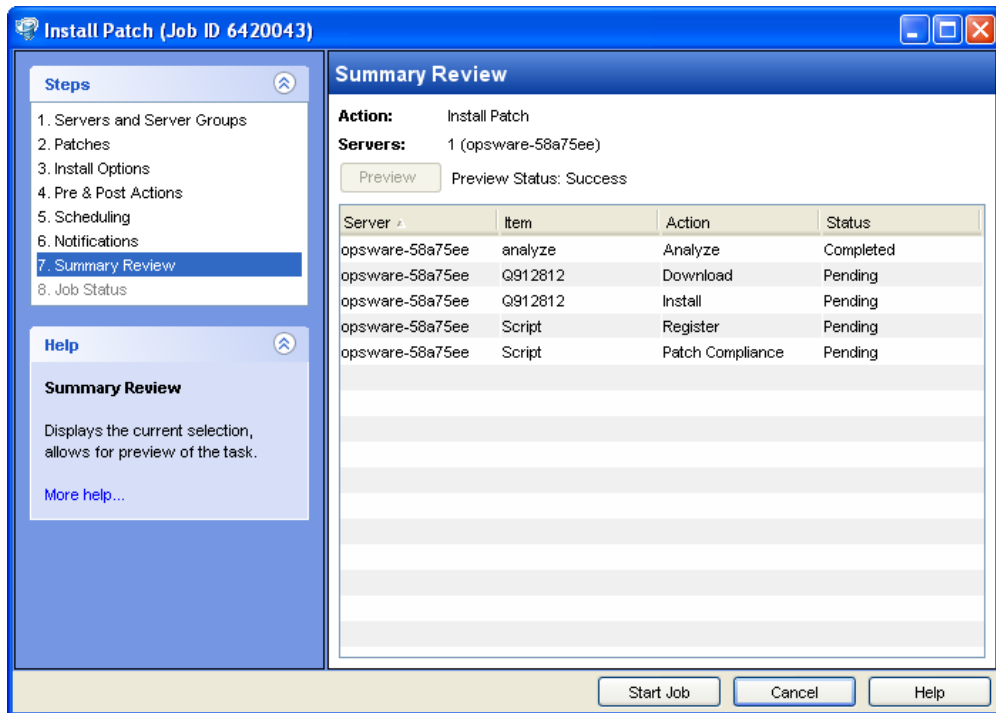
- **Download Phase:** This is when the patch is downloaded from Opsware SAS to the managed server. This phase is commonly referred to as the staging phase.
- **Installation Phase:** This is when the patch is installed on the managed server. This phase is commonly referred to as the deployment phase.

You can specify whether you want the installation to occur immediately after the patch is downloaded (staged) or you can schedule the installation to occur at a later date and time. Patch Management also supports the need for best-effort installations of multiple patches by allowing you to specify that the patch installation process will continue even when an error occurs with one of the patches.

Patch Management displays the name of the command (.exe file and any predefined command-line arguments) that the Opsware Agent runs on the managed server to install the patch. You can override these default command-line arguments.

To help you optimally manage Windows patch installation, Patch Management allows you to manage server reboot options, specify pre and post installation scripts, simulate (preview) a patch installation, and set up email notifications to alert you about the status of the installation process. The Install Patch window guides you through setting up these conditions.

Figure 5-6: Install Patch Window



Installation Flags

You can specify installation flags that are applied whenever a Windows patch is installed. However, Opsware SAS also uses default installation flags and requires that patches are installed with these flags. You must therefore be certain that you do not specify any installation flags that override or contradict the default flags passed by Opsware SAS. See “Setting Windows Install Options” on page 300 for information about how to specify commands and flags.



Some Windows hotfixes do not support the `-z` flag, some do not support the `-q` flag, and some do not support either. In such cases, you must use a special expression: `/-z` or `/-q` or `/-z -q` respectively. This prevents the Patch Management feature from passing in the `-z` or `-q` or `-z -q` flag. By default, Opware SAS adds `/z /q` to the command line arguments when installing patches. To override this, specify `/-z /-q`. For example, if you prefer to not suppress the reboot, specify `/-z`.

The following table lists the default installation flags that Opware SAS uses.

Table 5-4: Default Installation Flags

WINDOWS PATCH TYPE	FLAGS
Windows Hotfix	<code>-q -z</code>
Windows Security Rollup Package (treated identically to a Hotfix by the Patch Management feature)	<code>-q -z</code>
Windows OS Service Pack	<code>-u -n -o -q -z</code>

Application Patches

The Patch Management feature does not allow you to apply a patch to an operating system for which the patch is not intended. When you are installing an application patch, Patch Management does not automatically filter out servers that do not have the corresponding application installed. Although Patch Management does not prevent you from doing so, you should not attempt to apply application patches to servers that do not have the necessary applications installed. If a patch is for an application that is not installed on the server, the patch will not be applied and an error message will display, such as "There was an error with package <name of the package>".

If an application patch is intended for an application that is running on more than one version of the same operating system, you cannot apply the patch to all of the servers at the same time. An application patch is associated with only one operating system version. You must first select the patch for one operating system, select the servers where the application is installed, and apply the patch. You must repeat this process for each version of the operating system where the application is installed.

Similarly, when uninstalling application patches that are installed on multiple versions of the same operating system, you cannot uninstall all of the patches at the same time. You must repeat the uninstallation process for each version of the operating system where the patch is installed.

Service Packs, Update Rollups, and Hotfixes

When you try to install a Service Pack, Update Rollup, or a Hotfix, there is a known delay when a confirmation dialog displays. Since the Opsware Agent is installing or uninstalling the patch, it cannot respond to the confirmation dialog. The Agent will time out an installation or uninstallation process if you do not click **OK** in the confirmation dialog. For Hotfixes, the Agent will time out if five minutes have lapsed and you have not clicked **OK** in the confirmation dialog. For Service Packs and Update Rollups, the Agent will time out if 60 minutes have lapsed and you have not clicked **OK** in the confirmation dialog.

To prevent this from happening, patch install and uninstall commands should have arguments that invoke silent mode installs and uninstalls. By default, the -q flag is set.

Installing a Windows Patch

Before a patch can be installed on a managed server, it must be imported into Opsware SAS and its status must be Available. Administrators who have the required permissions can install patches that are marked Limited.



You must have a set of permissions to manage patches. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide*.

You can perform the installation by explicitly selecting patches and servers, and you can install a patch even if the patch policy exception is Never Install.

To install a patch on a managed server, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand the Patches and select a specific Windows operating system.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list, select Servers (or Device Groups).

- 5 From the Show drop-down list, select Servers without Patch Installed (or Device Groups without Patch Installed).
- 6 From the Preview pane, select one or more servers.
- 7 From the **Actions** menu, select **Install Patch**.

The first step of the Install Patch window appears: Servers and Device Groups. For instructions on each step, see the following sections:

- Setting Windows Install Options
- Setting Reboot Options for a Windows Patch Installation
- Specifying Install Scripts for a Windows Patch Installation
- Scheduling a Windows Patch Installation
- Setting Up Email Notifications for a Windows Patch Installation
- Previewing a Windows Patch Installation
- Viewing Job Progress of a Windows Patch Installation

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8 When you are ready to launch the installation job, click **Start Job**.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If the Install Patch window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select **Refresh** from the **View** menu to update information in the Patch Preview pane.

See “Remediating Patch Policies” on page 271 for another method of installing a patch.

Setting Windows Install Options

You can specify the following types of patch installation options:

- Perform the patch installation immediately after the patch is downloaded or at a later date and time.
- Do not interrupt the patch installation process even when an error occurs with one of the patches.

- Use different command-line options to perform the installation.

To set these options, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Install Options step.
- 2** Select one of the following Staged Install Options:
 - **Continuous:** This allows you to run all phases as an uninterrupted operation.
 - **Staged:** This allows you to schedule the download and installation to run separately.
- 3** Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.
- 4** In the Install Command text box, enter command-line arguments for the command (.exe file) that is displayed. By default, Opware SAS adds /z /q. If you want to override these install flags, enter /-z /-q in the text box.
- 5** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Setting Reboot Options for a Windows Patch Installation

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is installed on it, completely suppress all server reboots, or postpone reboots until all patches have been installed.



When you are selecting reboot options in the Install Patch window, Opware, Inc. recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option. If it is not possible to use the Microsoft reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option. Failure to do this can result in MBSA incorrectly reporting the patches that are installed on the server until the next reboot occurs (outside of Opware control).

The following options determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Install Patch window; they do not change the Reboot Required option, which is on the Install Parameters tab of the patch properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install:** Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Hold all server reboots until after all packages are installed and/or uninstalled:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2** Select one of the Rebooting Options.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Specifying Install Scripts for a Windows Patch Installation

For each patch, you can specify a command or script to run before installation or after installation. A pre-install script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-install script fails, the patch would not be installed. A pre-install script could also be used to shut down a service or application before it is patched. A post-install script could be used to perform a certain cleanup process on the managed server.

You can also specify the following types of scripts to run on the managed server before or after an installation or download phase:

- **Pre-Download:** A script that runs before patches are downloaded from Opsware SAS to the managed server. This is available only if you select Staged in the Install Options step.
- **Post-Download:** A script that runs after patches are downloaded from Opsware SAS to the managed server and before the patch is installed. This is available only if you select Staged in the Install Options step.
- **Pre-Install:** A script that runs before patches are installed on the managed server.
- **Post-Install:** A script that runs after patches are installed on the managed server.

To specify a pre-install script, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2** Select the Pre-Install tab. You may specify different scripts and options on each of the tabs.
- 3** Select Enable Script. This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
- 4** Select either Saved Script or Ad-Hoc Script.

A Saved Script has been previously stored in Opsware SAS with the Opsware SAS Web Client. To specify the script, click **Select**.

An Ad-Hoc script runs only for this operation and is not saved in Opsware SAS. Select the Type, such as .bat. In the Script box, enter the contents of the script, including the drive letter of where the script is located, such as `echo dir>> C:\temp\preinstall1.log`. If you do not enter a drive letter, the default is `%SYSTEMDRIVE%`, which is where the system folder of Windows is installed.

- 5** If the script requires command-line flags, enter the flags in the Command text box.
- 6** Specify the information in the User section. If you choose a system other than Local System, enter the User Name, Password, and Domain. The script will be run by this user on the managed server.
- 7** To stop the installation if the script returns an error, select the Error check box.
- 8** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Scheduling a Windows Patch Installation

Since the two phases of patching can be decoupled, you can schedule that you want patches installed independently of when patches are downloaded.

To schedule a patch installation, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Scheduling step.

By default, the Scheduling step displays only the scheduling options for the installation phase. If you selected Staged in the Install Options step, the scheduling options for the download phase will also be displayed.
- 2** Select one of the following Install Phase options:
 - **Run Task Immediately:** This enables the system to perform a preview analysis in the Summary Review step. The scheduling option for the download phase is **Run Immediately Following Download**.
 - **Run Task At:** This enables you to specify a later date and time that you want the installation or download performed.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.





A scheduled patch installation can be cancelled (prior to its execution), even if the patch download has already completed.

Setting Up Email Notifications for a Windows Patch Installation

You can set up email notifications to alert users when the download and installation operations complete successfully or with errors.

To set up email notifications, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Notifications step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.

- 3** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase.
- 4** Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.



If you previously selected Staged in the Install Options step, the Notifications pane displays notification options for both the download and installation phases.

Previewing a Windows Patch Installation

The installation preview process provides an up-to-date report about the patch state of servers. The installation preview is an optional step that lets you see the patches that will be installed on managed servers and the type of server reboots that are required. This preview process verifies whether the servers that you selected for the patch installation already have that patch installed (based on the MBSA). In some cases, a server could already have the patch installed if a system administrator had manually installed it, which means that Patch Management does not know about it.

The preview process also reports on dependency and supersedence information, such as patches that require certain Windows products, and patches that supersede other patches or are superseded by other patches. If a dependency is not met, Patch Management will display an error message indicating this condition. For example, if a managed server is running Windows 2000 Service Pack 3 (or higher) or Windows 2003, and an Opware SAS 5.5 Agent, Patch Management will report that a dependency has not been fulfilled. If you try to install a patch for Service Pack 4 and your server is using Service Pack 3, the remediate preview will display a “Will Not Install” error message to indicate this discrepancy. The Install Patch window allows superseded patches to be installed.



The installation preview does not report on the behavior of the server as though the patches have been applied.

To preview a patch installation, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Summary Review step.
- 2** Verify the information displayed for the Servers, Device Groups, and Patches at the top of the window.
- 3** (Optional) Click **Preview** to see the separate actions that will be performed when the patch is installed. To view the details of a previewed action, select a row in the table.
- 4** Click **Start Job** to launch the installation job or click **Cancel** to close the Install Patch window without launching the installation.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Windows Patch Installation

You can review progress information about a patch installation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Job Progress step. This will start the installation job.

The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:

- **Analyze:** Opware SAS examines the patches needed for the installation, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.
- **Download:** The patch is downloaded from Opware SAS to the managed server.
- **Install:** After it is downloaded, the patch is installed.
- **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
- **Pre/Post Install/Download Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstallation.
- **Install & Reboot:** When a patch is installed, the server is also rebooted.
- **Verify:** Installed patches will be included in the software registration.

- 2** To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See the *Opware® SAS User's Guide: Server Automation* for more information about browsing job logs.
- 3** Click **Stop Job** to prevent the job from running or click **Close** to close the Install Patch window.

Patch Uninstallation

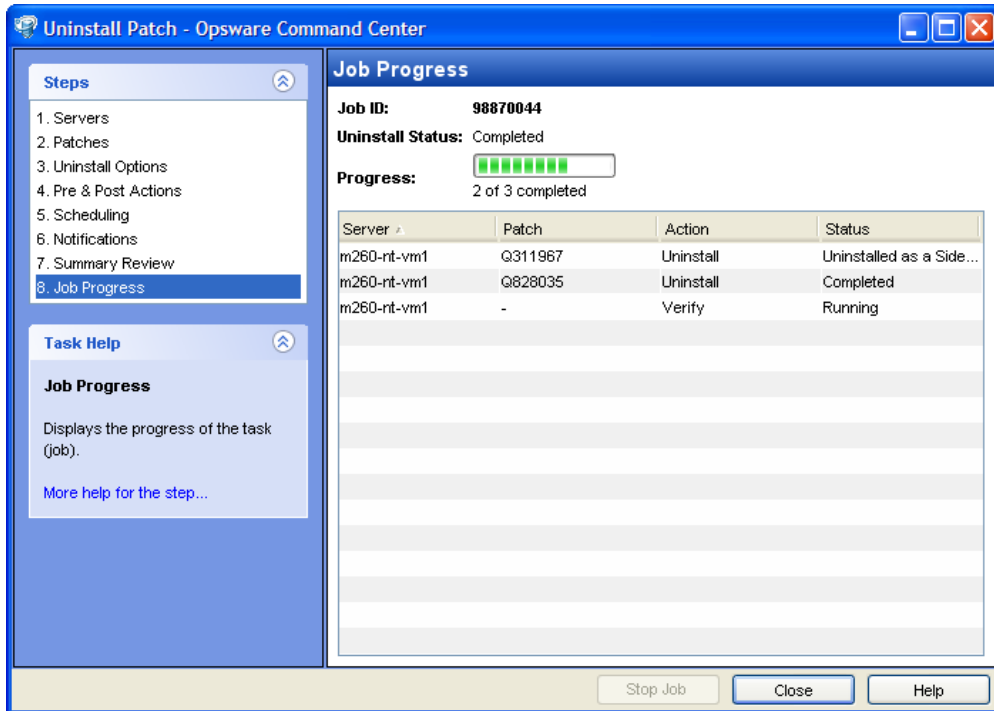
Patch Management provides granular control over how and under what conditions Windows patches are uninstalled (removed) from managed servers. To minimize problems, you can only uninstall one patch at a time. You cannot use Opware SAS to uninstall a patch that was not installed by using the Patch Management feature.

To help you optimally manage these conditions, Patch Management allows you to do the following:

- Manage server reboot options, and pre and post installation scripts.
- Simulate (preview) a patch uninstallation.
- Set up email notifications to alert you about the status of the uninstallation process.

The Uninstall Patch window guides you through setting up these conditions.

Figure 5-7: Uninstall Patch Window



Uninstallation Flags

You can specify uninstallation flags that are applied whenever a Windows patch is uninstalled. However, Opware SAS also uses default uninstallation flags and requires that patches are uninstalled with these flags. You must therefore be certain that you do not specify any uninstallation flags that override or contradict the default flags passed by Opware SAS.



Some Windows hotfixes do not support the `-z` flag, some do not support the `-q` flag, and some do not support either. In such cases, you must use a special expression: `/-z` or `/-q` or `/-z -q` respectively, to prevent the Patch Management feature from passing in the `-z` or `-q` or `-z -q` flag. By default, Opware SAS adds `/z /q` to the command line arguments when uninstalling patches. To override this, specify `/-z /-q`. For example, if you prefer to not suppress the reboot, specify `/-z`.

- The following table lists the default uninstallation flags that Opsware SAS uses.

Table 5-5: Default Uninstallation Flags

WINDOWS PATCH TYPES	FLAGS
Windows Hotfix	-q -z
Security Rollup Package	-q -z
Windows OS Service Pack	Not uninstallable

Uninstalling a Windows Patch

To remove a patch from a managed server, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand the Patches and select a specific Windows operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list, select Servers.
- 5 From the Show drop-down list, select Servers with Patch Installed.
- 6 From the Preview pane, select one or more servers.
- 7 From the **Actions** menu, select **Uninstall Patch**. The first step (Servers) in the Uninstall Patch window appears.

For instructions on each step, see the following sections:

- Setting Uninstall Options
- Setting Uninstall OptionsSetting Reboot Options for a Windows Patch Uninstallation
- Specifying Install Scripts for a Windows Patch Uninstallation
- Scheduling a Windows Patch Uninstallation
- Setting Up Email Notifications for a Windows Patch Uninstallation
- Viewing Job Progress of a Patch Uninstallation

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8 When you are ready to launch the uninstallation job, click **Start Job**.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If the Uninstall Patch window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select Refresh from the View menu to update information in the Patch Preview pane.

Setting Uninstall Options

You can specify the following types of patch uninstallation options:

- Do not interrupt the patch uninstallation process even when an error occurs with one of the patches.
- Use different command-line options to perform the uninstallation.

To set these options, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Uninstall Options step.
- 2 Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.
- 3 In the Uninstall Command text box, enter command-line arguments for the command (.exe file) that is displayed. By default, Opware SAS adds /z /q. If you want to override these uninstall flags, enter /-z /-q in the text box.
- 4 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Setting Reboot Options for a Windows Patch Uninstallation

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is removed from it, completely suppress all server reboots, or postpone reboots until all patches have been uninstalled.



When you are selecting reboot options in the Uninstall Patch window, Opsware, Inc. recommends that you use Microsoft's reboot recommendation. This is the "Reboot servers as specified by patch properties" option. If it is not possible to use the Microsoft reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option. Failure to do this can result in MBSA incorrectly reporting which patches are installed on the server until the next reboot occurs (outside of Opsware control).

The following options determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Uninstall Patch window; they do not change the Reboot Required option, which is on the Uninstall Parameters tab of the patch Properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install:** Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Hold all server reboots until after all packages are installed and/or uninstalled:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2** Select one of the Rebooting Options.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Specifying Install Scripts for a Windows Patch Uninstallation

For each patch, you can specify a command or script to run before uninstallation or after uninstallation. A pre-uninstall script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-uninstall script fails, the patch would not be removed from a server. A pre-uninstall script could also be used to shut down a service or application before it is removed from a server. A post-uninstall script could be used to perform a certain cleanup process on the managed server.

You can specify the following types of scripts to run on the managed server before or after a patch uninstallation:

- **Pre-Uninstall:** A script that runs before the patch is removed from a managed server.
- **Post-Uninstall:** A script that runs after the patch is removed from a managed server.

To specify a script, perform the following steps:

1 From the Uninstall Patch window, click **Next** to advance to the Pre & Post Actions step.

2 Select the Pre-Uninstall or Post-Uninstall tab.

You may specify different scripts and options on each of the tabs.

3 Select Enable Script.

This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.

4 Select either Saved Script or Ad-Hoc Script.

A Saved Script has been previously stored in Opware SAS with the SAS Web Client. To specify the script, click **Select**.

An Ad-Hoc script runs only for this operation and is not saved in Opware SAS. Select the Type, such as .bat. In the Script box, enter the contents of the script, including the drive letter of where the script is located, such as `echo dir>> C:\temp\preinstall1.log`. If you do not enter a drive letter, the default is %SYSTEMDRIVE%, which is where the system folder of Windows is installed.

5 If the script requires command-line flags, enter the flags in Commands.

6 Specify the information in the User section. The script will be run by this user on the managed server.

7 To stop the uninstallation if the script returns an error, select Error

Scheduling a Windows Patch Uninstallation

You can remove a patch from a server immediately, or at a later date and time.



To schedule a patch uninstallation, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Scheduling step.
- 2** Select one of the following Install Phase options:
 - **Run Task Immediately:** This enables you to perform the uninstallation in the Summary Review step.
 - **Run Task At:** This enables you to specify a later date and time that you want the uninstallation performed.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Setting Up Email Notifications for a Windows Patch Uninstallation

You can set up email notifications to alert users when the patch uninstallation operation completes successfully or with errors.

To set up email notifications, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Notifications step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the uninstallation phase.
- 4** Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Previewing a Windows Patch Uninstallation

The uninstallation preview process provides an up-to-date report about the patch state of servers. The uninstallation preview is an optional step that lets you see the patches that will be removed from managed servers. This preview process verifies whether the servers you selected for the patch uninstallation have that patch installed (based on the MBSA).



The uninstallation preview process does not report or simulate the behavior of a system with patches removed from the server.

To preview a patch uninstallation, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Summary Review step.
- 2** Verify the information displayed for the Servers, Device Groups, and Patches at the top of the window.
- 3** (Optional) Click **Preview** to see the separate actions that will be performed when the patch is uninstalled. To view the details of a previewed action, select a row in the table.
- 4** Click **Start Job** to launch the job or click **Cancel** to close the Uninstall Patch window without launching the uninstallation.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Patch Uninstallation

You can review progress information about a patch uninstallation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Job Progress step. The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:
 - **Analyze:** Opware SAS examines the patches needed for the uninstallation, checks the managed servers for the most recent patches installed, and

determines other actions it must perform.

- **Uninstall:** The patch is uninstalled.
 - **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
 - **Pre/Post Uninstall Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstallation.
 - **Uninstall & Reboot:** When a patch is installed, the server is also rebooted.
 - **Verify:** Installed patches will be included in the software registration.
- 2** To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See the *Opware[®] SAS User's Guide: Server Automation* for more information on browsing job logs.
 - 3** Click **Stop Job** to prevent the job from running or click **Close** to close the Uninstall Patch window.

Chapter 6: Patch Management for Unix

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Patch Management for Unix
- Patch Management Roles for Unix
- Patch Management for Specific Unix Operating Systems
- Patch Properties
- Software Policies
- Patch Administration for Unix
- Patch Installation
- Patch Uninstallation

Overview of Patch Management for Unix

The Patch Management for Unix feature enables you to identify, install, and remove patches, and maintain a high level of security across managed servers in your organization. With the SAS Client user interface, you can identify and install patches that protect against security vulnerabilities for the AIX, HP-UX, and Solaris operating systems.

This section contains information about how to install and uninstall Unix patches using software policies. It also contains information about generating patch policy compliance reports.

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

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

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

Patch management is a fully integrated component of Opware SAS. It leverages the Opware SAS server automation features. Opware SAS, for example, maintains a central database (called the Model Repository) that has detailed information about every server under management, the patches and software installed on the servers, and the patches and software available for installation. You can use this data to determine the severity of your exposure to a newly discovered threat, and to help assess the benefits of rolling out a patch versus the costs in downtime and testing requirements.

By automating the patching procedure, the Patch Management feature can reduce the amount of downtime required for patching. Opware SAS also allows you to schedule patch activity, so that patching occurs during off-peak hours.

Patch Management for Unix Features

Opware SAS automates patch management by providing the following features:

- A central repository where patches are stored and organized in their formats
- A database that includes information on every patch that has been applied
- Customized scripts that can be run before and after a patch is installed
- Advanced search abilities that identify servers that require patching
- Auditing abilities that enable security personnel to track the deployment of important patches

These features enable you to browse patches by a certain operating system, schedule patch downloads and installations, set up email notifications, preview a patch installation, use software policies and remediation to install and uninstall patches, and export patch information to a reusable file format.

Types of Patch Browsing

The Opware SAS Client interface organizes Unix patches by operating systems and displays detailed vendor security information about each patch. You can browse patches by patch type, availability, platform version, and so on. You can also browse all patches that are installed on a server, and view and edit patch metadata.

Scheduling and Notifications

In Patch Management, you can separately schedule when you want patches uploaded into Opware SAS and when you want these patches downloaded to managed servers. As a best practice, patch installations are typically scheduled for a time that causes minimal disruption to an organization's business operation. If you are installing one patch on one server, the installation operation will start only after the download operation has completed.

Patch Management also allows you to set up email notifications that alert you whether the download and installation operations completed, succeeded, or failed. When you schedule a patch installation, you can also specify reboot preferences to adopt, override, postpone, or suppress the vendor's reboot options.

Software Policies

Software policies enable you to customize patch distribution in your environment. They define the Unix patches that should be installed or not installed on certain managed servers. See "Software Management" on page 365 for more information about creating software policies to install Unix patches.

Patch Installation Preview

While Patch Management allows you to react quickly to newly discovered security vulnerabilities, it also provides support for strict testing and standardization of patch installation. After you have identified patches to install, Patch Management allows you to simulate (preview) the installation before you actually install a patch. This preview process tells you whether the servers that you selected for the patch installation already have that patch installed. In some cases, a server could already have a patch installed if a system administrator had manually installed it. The preview process provides an up-to-date report of the patch state of servers.

Software Policy Remediation

Patch Management also provides a solution for remediating servers that are not operating properly due to installed patches. If installed patches cause problems, even after being tested and approved, Patch Management allows you to uninstall the patches in a safe and standardized way. Patch Management allows you to specify uninstall options that control server reboots and the execution of uninstall commands, and pre-uninstall and post-uninstall scripts. Similar to previewing a patch installation, you can also preview a patch uninstall. See "Software Management" on page 365 for more information about remediating software policies.

Exporting Patch Data

To help you track the patch state of servers or groups of servers, Patch Management allows you to export this information. This information can be exported in a comma-separated value (.csv) file and includes details about when a patch was last detected as being installed, when a patch was installed by Opsware SAS, the patch compliance level, what patch policy exceptions exist, and so on. You can then import this information into a spreadsheet or database to perform a variety of patch analysis tasks.

Opsware SAS Integration

When a server is brought under management by Opsware SAS, the Opsware Agent installed on the server registers the server's hardware and software configuration with Opsware SAS. (The Opsware Agent repeats this registration every 24 hours.) This information, which includes data about the exact OS version, hardware type, installed software and patches, is immediately recorded in the Model Repository. Also, when a server is initially provisioned with Opsware SAS, the same data is immediately recorded.

When a new patch is issued, you can use Opsware SAS to immediately identify the servers that require patching. The Opsware SAS Client provides a software repository where you upload patches and other software. Users access this software from the Opsware SAS Client to install patches on the appropriate servers.

After a server is brought under management, you should install all patches by using the Patch Management feature. If you install a patch manually, Opsware SAS does not have data about that patch until the next software registration. If you install a patch manually, it can take as long as 24 hours until the data about that server in the Model Repository is up-to-date.

Whenever you install or uninstall software or patches with Opsware SAS, however, the Opsware Agent immediately updates the information about the server in the Model Repository.

Support for Unix Patch Testing and Installation Standardization

Opsware SAS offers features to minimize the risk of rolling out patches. First, when a patch is uploaded into Opsware SAS, its status is marked as untested and only administrators with special privileges can install it.

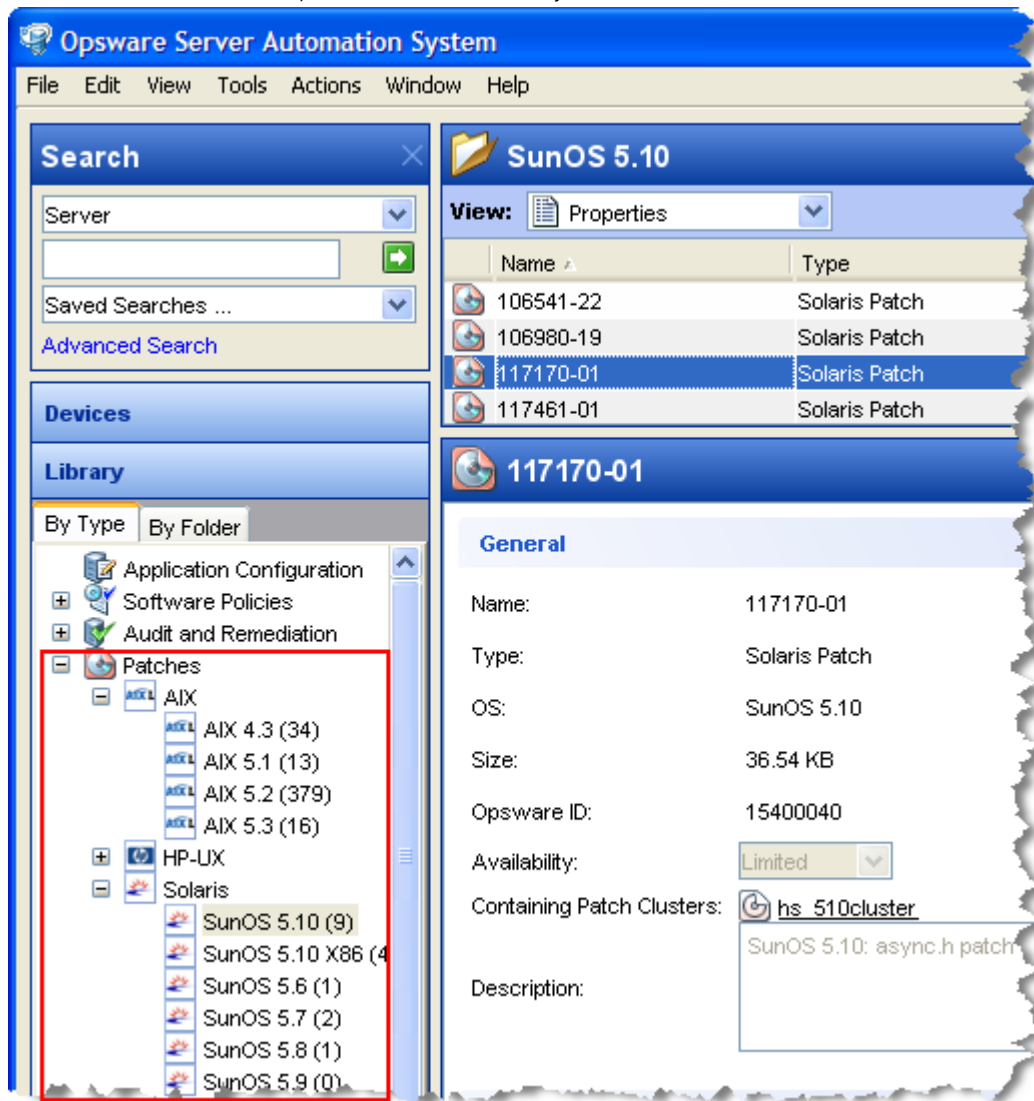
The patch administrator then defines patch installation and uninstallation options and tests the patch. Only after the patch is tested and the patch administrator marks it as available for use can other administrators install the patch.

The Patch Management feature allows you to standardize the way that patches are installed and uninstalled, thereby preventing ad-hoc installation procedures. Patch administrators standardize patch installation by providing pre-install and post-install scripts, install and uninstall flags, reboot instructions, and how to handle error codes from the pre-install and post-install scripts.

Library

The SAS Client Library provides flexibility in searching for and displaying Unix patches by name, type of patch, operating system, relationship to other packages, and so on. See Figure 6-1. The number in parenthesis is the total number of patches (for that operating system version) that were uploaded from the Unix web site. Use the column selector to control which columns of patch metadata data to display.

Figure 6-1: Unix Patches in the Opware SAS Client Library



Search Feature

In the SAS Client, you can search for any information about your operational environment that is available in Opware SAS using the SAS Client Search feature. The Search feature enables you to search for patches, software policies, servers, and so on. See “SAS Client Search” in the *Opware® SAS User's Guide: Server Automation*.

Patch Management Roles for Unix

Opware SAS provides support for rigorous change management by assigning the functions of patch management to the patch administrator and the system administrator:

- The patch administrator (often referred to as the security administrator) has the authority to upload, test, and edit patch options.
- The system administrator applies the patches (that have been approved for use) uniformly, according to the options that the patch administrator specifies.



Only the patch administrator should have the Patches permission, which gives access to advanced features. To obtain these permissions, contact your Opware administrator. See the Permissions Reference appendix in the *Opware® SAS Administration Guide*.

Patch Administrator

In most organizations, patch administrators are responsible for reviewing the latest security threats and the patches that vendors have released to address these problems. The patch administrators are generally experts in the operating systems and applications that they manage, and are able to assess the necessity of applying patches issued by vendors. They are able to diagnose common problems that arise after patches are installed, allowing them to thoroughly test the patch application process.

In Opware SAS, patch administrators are granted specific permissions that allow them to upload patches into Opware SAS to test the patches and then mark them as available for use. Basic users can upload patches, but they cannot install them or mark them as available. Patch administrators are also able to edit patch options (such as installation scripts) through patch management. Other types of users are not allowed to upload or edit patches.

Typically, the patch administrator uploads patches and then tests them on non-production reference hardware. After testing the patches and determining that the patches are safe to apply to production systems, they mark the patches as available in the Opsware SAS Client, and then advise the system administrators that they must apply the approved patches.

System Administrator

System administrators are responsible for the day-to-day maintenance of the servers in a deployment. These users are not required to have the same level of expertise in low-level system details as the patch administrator.

Because the patch administrator has set up the patch installation, the system administrators can apply the patches to a large number of servers with a few mouse clicks. They are responsible for searching for the servers that require the approved patch, installing the patch, and verifying that the patches are installed successfully.

Patch Management for Specific Unix Operating Systems

The types of patches and their underlying technologies can vary according to the vendor of the operating system. This section discusses the vendor-specific details for Unix patch management in Opsware SAS.

Supported Unix Versions and Patch Types

The Patch Management feature supports all of the operating system versions that Opsware SAS supports, except for Linux.

Linux does not support patches in the ordinary sense. The packages are not patchable. Instead, new versions of the RPM are delivered. Linux systems that Opsware SAS manages are therefore not viewable through the Patch Management feature interfaces. New Linux packages and updates should be managed and applied through the software page.

The following table shows the Unix versions and the patch types that the Patch Management feature supports.

Table 6-1: Supported Unix Versions and Patch Types

UNIX VERSIONS	PATCH TYPES
AIX 4.3	AIX Update Fileset APARs
AIX 5.1	AIX Update Fileset APARs
AIX 5.2	AIX Update Fileset APARs
AIX 5.3	AIX Update Fileset APARs
HP-UX 11.00	HP-UX Patch Fileset HP-UX Patch Product
HP-UX 11.11	HP-UX Patch Fileset HP-UX Patch Product
HP-UX 11.23	HP-UX Patch Fileset HP-UX Patch Product
Solaris 6	Solaris Patch Solaris Patch Cluster
Solaris 7	Solaris Patch Solaris Patch Cluster
Solaris 8	Solaris Patch Solaris Patch Cluster
Solaris 9	Solaris Patch Solaris Patch Cluster

Table 6-1: Supported Unix Versions and Patch Types (continued)

UNIX VERSIONS	PATCH TYPES
Solaris 10	Solaris Patch Solaris Patch Cluster

Underlying Technologies for Patch Management on Unix

Although the utilities vary, Opware SAS enables you to perform patching tasks by using a single interface. Opware SAS models the way it treats patches by the way the underlying utility treats patches. For example, if the Solaris patchadd utility is not able to install one patch contained in a patch cluster, the Solaris utility continues to install the remaining patches in the patch cluster. Opware SAS respects this behavior and allows that patch installation operation to continue. Any patches that are not installed are reported at the end of the installation operation.

The following table shows the patch management and installation tools that are used for each of the supported Unix systems.

Table 6-2: Supporting Technologies for Patch Management on Unix

SOLARIS	AIX	HP-UX
Patchadd installs Solaris patches	Installp installs and uninstalls filesets	Swlist lists patch products, files, products, and filesets
Patchrm uninstalls Solaris patches	Lslpp lists installed LPPs	Swinstall installs a depot
Showrev lists installed Solaris patches	Instfix lists installed APARS	Swremove removes a depot
Pkgadd installs Solaris packages		

Table 6-2: Supporting Technologies for Patch Management on Unix (continued)

SOLARIS	AIX	HP-UX
Pkginfo lists installed Solaris packages		

AIX Patches

AIX periodically releases Authorized Program Analysis Reports (APARs), which specify what update filesets (contained in LPPs) are necessary to fix an identified problem. An APAR only specifies the minimum version of an update fileset required to fix a problem; an APAR can therefore be satisfied with later versions of the same filesets. To maintain compatibility, however, Opware SAS always adopts the fileset with the lowest version number that meets the minimum version that APAR specifies. If a later version of the update fileset is uploaded, Opware SAS still associates the earlier version of the fileset with the APAR.

When uploading an LPP, Opware SAS recognizes which APARs the filesets contained in the LPP belong to. An entry is created for the APAR in the Patch Management feature when the first fileset associated with an APAR is uploaded. (In some cases, a fileset is associated with more than one APAR. An entry is created for each APAR the fileset is associated with, if the entry does not already exist.)

If you want to install all LPPs that APAR specifies, you must make certain to upload all of the specified LPPs into the Patch Management feature.

If you do not upload all of the LPPs that APAR specifies, it is still possible for the system administrator to browse for an APAR and install the partial set of LPPs that are uploaded. In such cases, the administrator receives a warning that the filesets for the APAR are not all installed.



The Patch Administrator must first upload and test an LPP before it is generally available in Opware SAS. The new fileset is integrated into the APAR only after the LPP is tested and approved. Even though the APAR is updated automatically, you still maintain control over the exact filesets that are allowed to be installed on your managed servers.



APAR update filesets cannot be installed on a server if the server does not already have the base filesets for which the update filesets are intended.

If, however, a server has a partial set of the base filesets, the APAR can be applied and only the applicable filesets for the base filesets are installed. For example, if an APAR specifies four update filesets to update four base filesets, and you attempt to apply the APAR to a server that has only three of the base filesets, three of the four update filesets from the APAR are installed.

When installing an AIX update fileset, the Patch Management feature normally applies the fileset, which allows it to be rejected (uninstalled.) If you want to commit the fileset instead (so that it cannot be removed), use the `-c` option here.



Since update filesets can be included in folders, global read permissions are required to view and edit AIX update filesets. See "Software Management Setup" in the *Opware® SAS Policy Setter's Guide* for information about how to use folders.

Solaris Patches

A Solaris patch cluster contains a set of selected patches for a specific Solaris release level. Ordinarily, after a patch cluster is installed, it is not possible to search for a particular patch cluster. The patches do not contain any metadata that relate them to the patch cluster in which they were originally bundled. You can only search for the individual patches.

If you install a Solaris patch cluster by using the Patch Management feature, however, Opware SAS keeps track of the patch cluster in the Model Repository. You can therefore search for a patch cluster to determine if a full patch cluster is installed. If you installed the patch cluster with the Patch Management feature, you can uninstall individual patches in the cluster. You cannot uninstall a patch cluster.

HP-UX Patches

HP-UX patches are delivered exclusively as depots, which are patch products that contain patch filesets. The depot is uploaded directly into Opware SAS by using the Patch Management feature.

If a depot is already uploaded and attached to a node, it cannot be uploaded by using the Patch Management feature. If you want to upload the depot by using the Patch Management feature, you must detach a depot from any nodes that it is attached to, and then delete it from the Software Repository.

Patch Uploads for Unix

Before a Unix patch can be installed on a managed server with Opsware SAS, the patch must be uploaded into the SAS Client Library. Uploading patches is the responsibility of the patch administrator. See the *Opsware® SAS Administration Guide* and the *Opsware® SAS Policy Setter's Guide* for information about how to upload Unix patches and the importing software process.

Patch Uploads for Specific Unix Versions

When a patch is uploaded, you associate the patch with a specific version of an operating system. When you upload a Solaris patch, for example, you must select the version of the Solaris operating system that this patch applies to, such as Solaris 5.6 or 5.9. You can only install this patch on servers that are running that version of the operating system.

If, for any reason, you need to install a given patch across servers running different versions of the same operating system, you need to upload the patch multiple times and associate the patch with each of the operating system versions that the patch applies to.

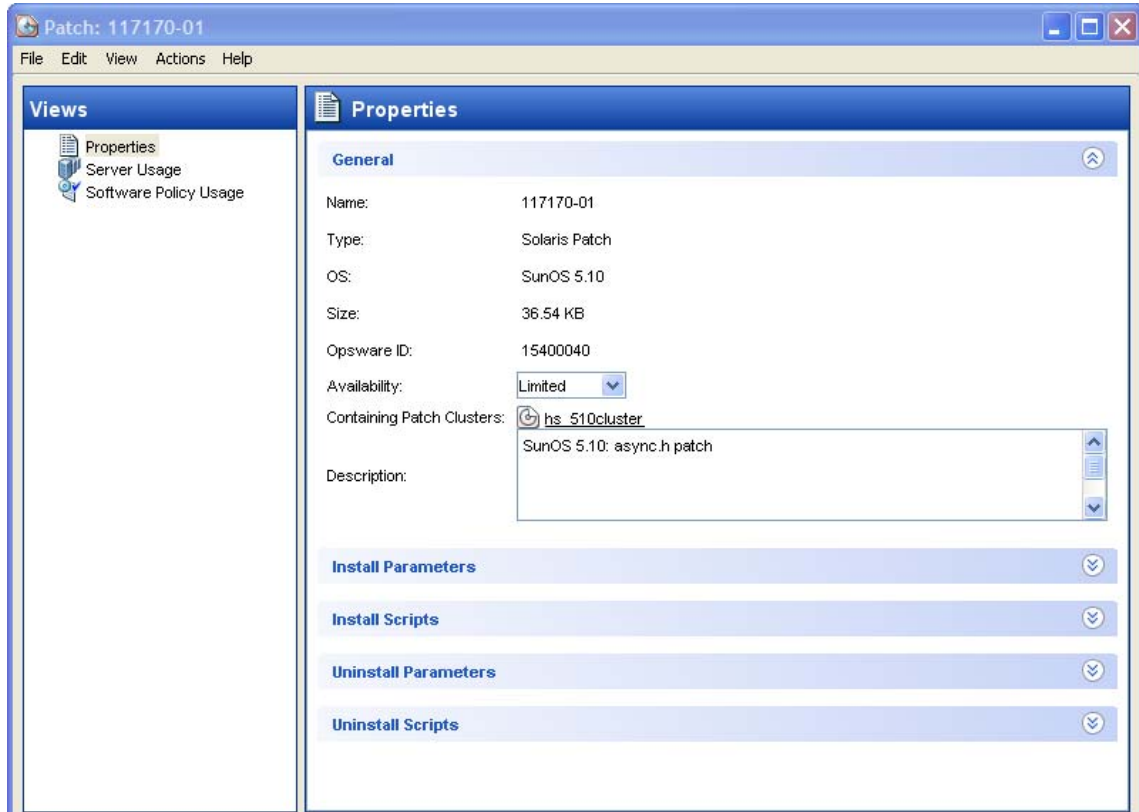
For example, if the same Solaris patch needs to be installed on servers running Solaris 2.7 and 2.8, you must upload the patch two times. The first time that you upload the patch, you associate it with the Solaris 2.7. You then repeat the procedure and associate the patch with Solaris 2.8. (This procedure also allows you to specify different installation options. The different versions of the same operating system can sometimes require different installation scripts, installation flags, and so on.)

In the case of application patches, it is even more common that you need to upload a patch multiple times. A Solaris patch for Oracle, for example, often needs to be applied to instances of Oracle running on slightly different versions of the Solaris operating system.

Patch Properties

Patch Management displays detailed information (properties) about a patch.

Figure 6-2: Unix Patch Properties



Patch properties include the following information:

- **Name:** The Unix name for the patch.
- **Type:** The type of Unix patch. Table 6-1 identifies these patch types.
- **OS:** The Unix operating systems that are known to be affected by this patch.
- **Size:** The size of the patch file, in kilobytes (KB) or in megabytes (MB). Size is not shown for AIX APARs.
- **Opware ID:** The Opware SAS unique ID for the patch.
- **Availability:** The status of a patch within Opware SAS, which can be one of the following:

- **Limited:** The patch has been imported into Opware SAS but cannot be installed. This is the default patch availability.
- **Available:** The patch has been imported into Opware SAS, tested, and has been marked available to be installed on managed servers.
- **Deprecated:** The patch cannot be added to patch policies or set as a patch policy exception but can still be installed.
- **Containing** (Optional): Depending on the selected patch type, this is the relationship to other packages. For example, for AIX update filesets, this field displays Containing LPPS/APARS.
- **Description:** A brief description of the Solaris patch cluster.

Viewing Unix Patches

The SAS Client displays information about Unix patches that have been imported into Opware SAS.

To view information about a patch, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand Patches and select a specific Unix operating system.
- 3** (Optional) Use the column selector to sort the patches according to Name, Type, Availability, and Description.
- 4** In the Content pane, open a patch to view its properties in the Patch window.

Editing Unix Patch Properties

You can edit a patch's Description, Availability, Install Parameters, and Uninstall parameters. Due to the nature of the type of patch, some properties are not editable.

The Availability property indicates the status of the patch in Opware SAS.

You can set the install and uninstall parameters on either the patch properties page or in the Patch Actions only when you are installing or uninstalling one patch at a time. The parameters on the properties page are saved in the Model Repository, but the parameters in Patch Actions are used only for that action. The parameters in Patch Actions override those on the patch properties page.

To edit the patch properties, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.

- 2** Expand Patches and select a specific Unix operating system. The Content pane will display all patches associated with that operating system.
- 3** In the Content pane, open a patch to view its properties in the Patch Window.
- 4** Edit any of the following fields: Description, Availability, and the Install and Uninstall parameters.
- 5** From the **File** menu, select **Save** to save your changes.

Finding Servers That Have a Unix Patch Installed

To find out which servers have a particular patch installed, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand Patches and select a specific Unix operating system. The Content pane will display all patches associated with that operating system.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list in the Content pane, select Server Usage.
- 5** From the Show drop-down list for the selected patch, select Servers with Patch Installed.

Finding Servers That Do Not Have a Unix Patch Installed

To find out which servers do not have a particular patch installed, perform the following steps:

- 1** From the Navigation pane, select Library and then select Patches.
- 2** Expand Patches and select a specific Unix operating system. The Content pane will display all patches associated with that operating system.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list, select Server Usage.
- 5** From the Show drop-down list, select Servers without Patch Installed.

Exporting a Patch

To export a patch from Opware SAS to the local file system, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.

- 2 Expand Patches and select a specific Unix operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.
- 4 From the **Actions** menu, select **Export**.
- 5 In the Export Patch window, enter the *folder* name that will contain the patch file in the File Name field.
- 6 Click **Export**.

Deleting a Patch

This action removes a patch from Opsware SAS, but does not uninstall the patch from managed servers. A patch cannot be deleted if it is attached to a policy.



Do not delete all of the patches from Opsware SAS. If you do so accidentally, contact your Opsware, Inc. support representative for assistance in uploading all of the patches back into Opsware SAS.

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Unix operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.
- 4 From the **Actions** menu, select **Delete Patch**.
- 5 In the Delete Patches windows, click **Delete**.

Software Policies

In Patch Management for Unix, software policies enable you to customize patch distribution in your environment. Software policies define which Unix patches should be installed or not installed on certain managed servers.

If you use software policies and you also perform ad hoc patch installs, you must run the remediate process to install all applicable patches on servers. See “Software Management” on page 365 for more information about creating and remediating software policies to install Unix patches.

Patch Compliance Reports

To troubleshoot and resolve patch compliance problems, you can run and examine several patch compliance reports by using the Reports feature in the SAS Client. The following patch compliance reports identify whether all patches in a software policy were installed successfully on managed servers in your environment.

Patch Policy Compliance (All Servers)

This report groups all managed servers by their patch policy compliance level to show compliant and non-compliant servers.

Patch Policy Compliance by Customer

This report lists all servers by the customer they belong to and then by the patch policy compliance level.

Patch Policy Compliance by Facility

This report groups all managed servers by the facility they belong to and then by the patch software policy compliance level.



See the *Opware® SAS User's Guide: Server Automation* for information about how to run, export, and print these reports.

Patch Administration for Unix

You can customize patch administration for Unix to best support your environment by setting the availability flag.

Setting the Default Patch Availability

You can set the default patch availability with the SAS Client. The default used by the script overrides the default set by the SAS Client. See the *Opware® SAS Administration Guide* for information about the script.

To set the default value for the Availability of a newly imported patch, perform the following steps:

- 1** From the Navigation pane, select Opware Administration.
- 2** Select Patch Configuration.

- 3** For the Default Availability for Imported Patches, select either Available or Limited. The default is Limited.

If the patch is Available, it can be installed on managed servers. If the patch is Limited, it has been imported into Opware SAS and can be installed only by a patch administrator who has the required permissions. To obtain these permissions, contact your Opware administrator. See the *Opware[®] SAS Administration Guide*.

Patch Installation

The patch installation process consists of the following two phases:

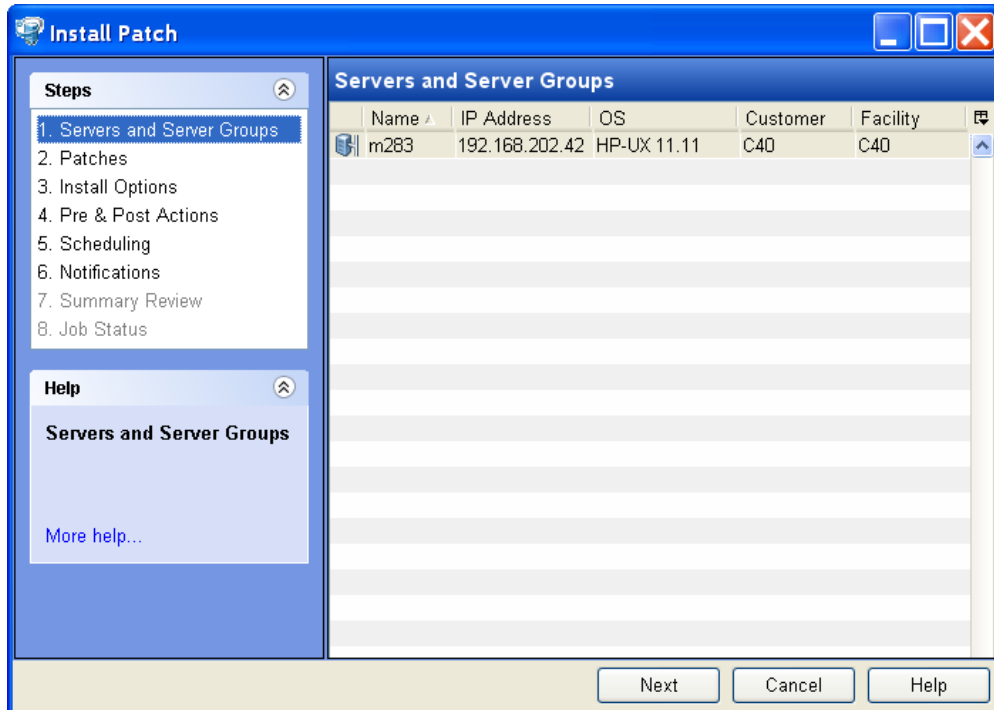
- **Download Phase:** This is when the patch is downloaded from Opware SAS to the managed server. This phase is commonly referred to as the staging phase.
- **Installation Phase:** This is when the patch is installed on the managed server. This phase is commonly referred to as the deployment phase.

You can specify whether you want the installation to occur immediately after the patch is downloaded (staged) or you can schedule the installation to occur at a later date and time. Patch Management also supports the need for best-effort installations of multiple patches by allowing you to specify that the patch installation process will continue even when an error occurs with one of the patches.

Patch Management displays the name of the command that installs the patch. The Opware Agent runs this command on the managed server. You can override the default command-line arguments that you want to perform the installation.

To optimally manage Unix patch installations, Patch Management allows you to manage server reboot options, and pre and post installation scripts, simulate (preview) a patch installation, and set up email notifications to alert you about the status of the installation process. The Install Patch window guides you through setting up these conditions.

Figure 6-3: Install Patch Window



Installation Flags

You can specify installation flags that are applied whenever a Unix patch is installed. However, Opware SAS also uses default installation flags and requires that patches are installed with these flags. You must therefore be certain that you do not specify any installation flags that override or contradict the default flags passed in by Opware SAS. See "Setting Unix Install Options" on page 351 for information about how to specify commands.

The following table lists the default installation flags that Opware SAS uses.

Table 6-3: Default Installation Flags

UNIX PATCH TYPE	FLAGS
AIX	-a -Q -g -X -w
HP-UX	None

Application Patches

The Patch Management feature does not allow you to apply a patch to an operating system for which the patch is not intended. When you are installing an application patch, Patch Management does not automatically filter out servers that do not have the corresponding application installed. Although Patch Management does not prevent you from doing so, you should not attempt to apply application patches to servers that do not have the necessary applications installed. If a patch is for an application that is not installed on the server, the patch will not be applied and an error message will display, such as “There was an error with package <name of the package>”.

If an application patch is intended for an application that is running on more than one version of the same operating system, you cannot apply the patch to all of the servers at the same time. An application patch is associated with only one operating system version. You must first select the patch for one operating system, select the servers where the application is installed, and apply the patch. You must repeat this process for each version of the operating system where the application is installed.

Similarly, when uninstalling application patches that are installed on multiple versions of the same operating system, you cannot uninstall all of the patches at the same time. You must repeat the uninstallation process for each version of the operating system where the patch is installed.

Installing a Unix Patch

Before a patch can be installed on a managed server, it must be imported into Opware SAS and its status must be Available. Administrators who have the required permissions can install patches that are marked Limited.



You must have a set of permissions to manage patches. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide*.

You can perform the installation by explicitly selecting patches and servers.

To install a patch on a managed server, perform the following steps:

- 1** From the Navigation pane, select Library and then select Patches.
- 2** Expand the Patches and select a specific Unix operating system.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list, select Servers (or Server Groups).
- 5** From the Show drop-down list, select Servers without Patch Installed (or Server Groups without Patch Installed).
- 6** From the Preview pane, select one or more servers.
- 7** From the **Actions** menu, select **Install Patch**.

The first step of the Install Patch window appears: Servers and Server Groups.

For instructions on each step, see the following sections:

- Setting Unix Install Options
- Setting Reboot Options for a Unix Patch Installation
- Specifying Install Scripts for a Unix Patch Installation
- Scheduling a Unix Patch Installation
- Setting Up Email Notifications for a Unix Patch Installation
- Previewing a Unix Patch Installation
- Viewing Job Progress of a Unix Patch Installation

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8** When you are ready to launch the installation job, click **Start Job**.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If the Install Patch window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select **Refresh** from the **View** menu to update information in the Patch Preview pane.

Setting Unix Install Options

You can specify the following types of patch installation options:

- Perform the patch installation immediately after the patch is downloaded or at a later date and time.
- Do not interrupt the patch installation process even when an error occurs with one of the patches.
- Use different command-line options to perform the installation.

To set these options, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Install Options step.
- 2** Select one of the following Staged Install Options:
 - Continuous:** This allows you to run all phases as an uninterrupted operation.
 - Staged:** This allows you to schedule the download and installation to run separately.
- 3** Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.
- 4** In the Install Command text box, enter command-line arguments for the command that is displayed.
- 5** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Setting Reboot Options for a Unix Patch Installation

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is installed on it, completely suppress all server reboots, or postpone reboots until all patches have been installed.



When you are selecting reboot options in the Install Patch window, Opware, Inc. recommends that you use the Unix reboot recommendations, which is the “Reboot servers as specified by patch properties” option. If you cannot use the Unix reboot setting, select the single reboot option, which is the “Do not reboot servers until all patches are installed” option.

The following options determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Install Patch window; they do not change the Reboot Required option, which is on the Install Parameters tab of the patch properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install:** Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Do not reboot servers until all patches are installed:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2** Select one of the Rebooting Options.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Specifying Install Scripts for a Unix Patch Installation

For each patch, you can specify a command or script to run before installation or after installation. A pre-install script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-install script fails, the patch would

not be installed. A pre-install script could also be used to shut down a service or application before it is patched. A post-install script could be used to perform a certain cleanup process on the managed server.

You can also specify the following types of scripts to run on the managed server before or after an installation or download phase:

- **Pre-Download:** A script that runs before patches are downloaded from Opsware SAS to the managed server. This is available only if you select Staged in the Install Options step.
- **Post-Download:** A script that runs after patches are downloaded from Opsware SAS to the managed server and before the patch is installed. This is available only if you select Staged in the Install Options step.
- **Pre-Install:** A script that runs before patches are installed on the managed server.
- **Post-Install:** A script that runs after patches are installed on the managed server.

To specify a pre-install script, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2** Select the Pre-Install tab. You may specify different scripts and options on each of the tabs.
- 3** Select Enable Script. This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
- 4** Select either Saved Script or Ad-Hoc Script.

A Saved Script has been previously stored in Opsware SAS with the Opsware SAS Web Client. To specify the script, click **Select**.
- 5** If the script requires command-line flags, enter the flags in the Command text box.
- 6** Specify the information in the User section. If you choose a system other than Local, enter the User Name, Password, and Domain. The script will be run by this user on the managed server.
- 7** To stop the installation if the script returns an error, select the Error check box.
- 8** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Scheduling a Unix Patch Installation

Since the two phases of patching can be decoupled, you can schedule when you want patches installed (deployed) to occur independently of when patches are downloaded (staged).

To schedule a patch installation, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Scheduling step.

By default, the Scheduling step displays only the scheduling options for the install phase. If you selected Staged in the Install Options step, the scheduling options for the download phase will also be displayed.
- 2** Select one of the following Install Phase options:
 - **Run Task Immediately:** This enables the system to perform a preview analysis in the Summary Review step. The scheduling option for the download phase is **Run Immediately Following Download**.
 - **Run Task At:** This enables you to specify a later date and time that you want the installation or download performed.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.





A scheduled patch installation can be cancelled (prior to its execution), even if the patch download has already completed.

Setting Up Email Notifications for a Unix Patch Installation

You can set up email notifications to alert users when the download and installation operations complete successfully or with errors.

To set up email notifications, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Notifications step.
- 2** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase.
- 3** Enter a Ticket ID to be associated with a Job in the Ticket ID field.

- 4 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.



If you previously selected Staged in the Install Options step, the Notifications pane displays notification options for both the download and installation phases.

Previewing a Unix Patch Installation

The installation preview process provides an up-to-date report about the patch state of servers. The installation preview is an optional step that lets you see what patches will be installed on managed servers and what type of server reboots are required. This preview process verifies whether the servers you selected for the patch installation already have that patch installed. In some cases, a server could already have the patch installed if a system administrator had manually installed it, which means that Patch Management does not know about it.

The preview process also reports on dependency information, such as patches that require certain Unix products, and patches that obsolete other patches or are obsoleted by other patches. If a dependency is not met, Patch Management will display an error message indicating this condition.



The installation preview does not report on the behavior of the server as though the patches have been applied.

To preview a patch installation, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Summary Review step.
- 2 Verify the information displayed for the Servers, Server Groups, and Patches at the top of the window.
- 3 (Optional) Click **Preview** to see the separate actions that will be performed when the patch is installed. To view the details of a previewed action, select a row in the table.
- 4 Click **Start Job** to launch the installation job or click **Cancel** to close the Install Patch window without launching the installation.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Unix Patch Installation

You can review progress information about a patch installation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Job Progress step. This will start the installation job.

The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:

- **Analyze:** Opware SAS examines the patches needed for the installation, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.
 - **Download:** The patch is downloaded from Opware SAS to the managed server.
 - **Install:** After it is downloaded, the patch is installed.
 - **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
 - **Pre/Post Install/Download Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstallation.
 - **Install & Reboot:** When a patch will be installed is also when the server will be rebooted.
 - **Verify:** Installed patches will be included in the software registration.
- 2** To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See the *Opware® SAS User's Guide: Server Automation* for more information about browsing job logs.
 - 3** Click **Stop Job** to prevent the job from running or click **Close** to close the Install Patch window.

Patch Uninstallation

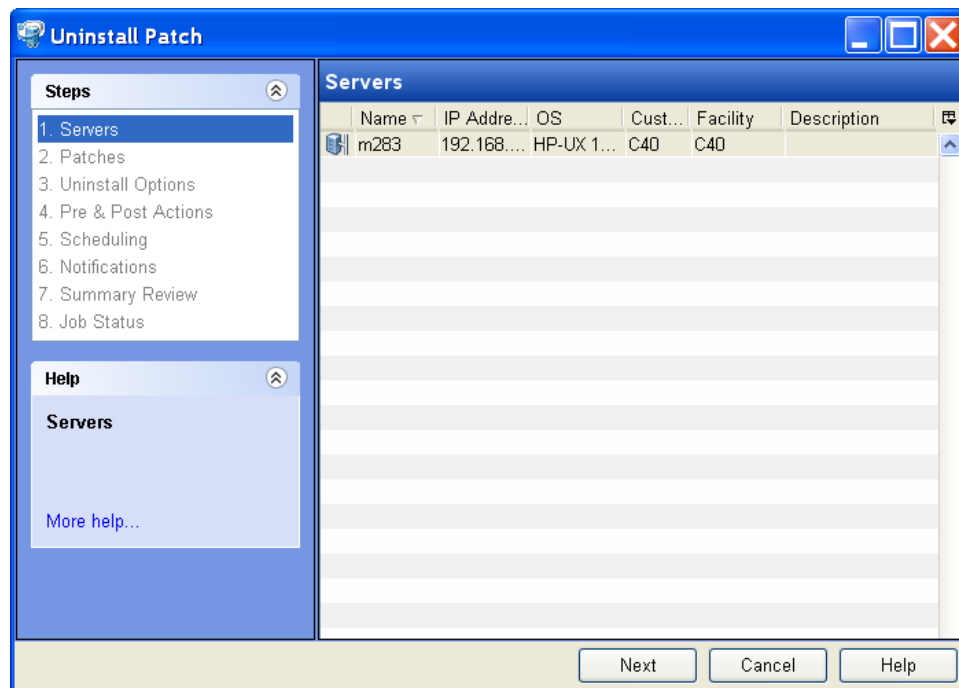
Patch Management provides granular control over how and under what conditions Unix patches are uninstalled (removed) from managed servers. To minimize problems, you can only uninstall one patch at a time. You cannot use Opware SAS to uninstall a patch that was not installed by using the Patch Management feature.

To help you optimally manage these conditions, Patch Management allows you to do the following:

- Manage server reboot options, and pre and post installation scripts.
- Simulate (preview) a patch uninstallation.
- Set up email notifications to alert you about the status of the uninstallation process.

The Uninstall Patch window guides you through setting up these conditions.

Figure 6-4: Uninstall Patch Window



Uninstallation Flags

You can specify uninstallation flags that are applied whenever a Unix patch is uninstalled. However, Opware SAS also uses default uninstallation flags and requires that patches are uninstalled with these flags. You must therefore be certain that you do not specify any uninstallation flags that override or contradict the default flags passed by Opware SAS.

The following table lists the default uninstallation flags that Opware SAS uses.

Table 6-4: Default Uninstallation Flags

OPERATING SYSTEM/PATCH TYPES	FLAGS
AIX	-u -g -X
AIX Reject Options	-r -g -X
HP-UX	None

Uninstalling a Unix Patch

To remove a patch from a managed server, perform the following steps:

- 1** From the Navigation pane, select Library and then select Patches.
- 2** Expand the Patches and select a specific Unix operating system.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list, select Servers.
- 5** From the Show drop-down list, select Servers with Patch Installed.
- 6** From the Preview pane, select one or more servers.
- 7** From the **Actions** menu, select **Uninstall Patch**.

The first step of the Uninstall Patch window appears: Servers.

For instructions on each step, see the following sections:

- Setting Reboot Options for a Unix Patch Uninstallation
- Specifying Pre and Post Install Scripts for a Unix Patch Uninstallation
- Scheduling a Unix Patch Uninstallation
- Setting Up Email Notifications for a Unix Patch Uninstallation
- Viewing Job Progress of a Patch Uninstallation

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8** When you are ready to launch the uninstallation job, click **Start Job**.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

- If the Uninstall Patch window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select Refresh from the View menu to update information in the Patch Preview pane.

Setting Uninstall Options

You can specify the following types of patch uninstallation options:

- Do not interrupt the patch uninstallation process even when an error occurs with one of the patches.
- Use different command-line options to perform the uninstallation.

To set these options, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Uninstall Options step.
- 2** Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.
- 3** In the Uninstall Command text box, enter command-line arguments for the command that is displayed.
- 4** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Setting Reboot Options for a Unix Patch Uninstallation

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is removed from it, completely suppress all server reboots, or postpone reboots until all patches have been uninstalled.



When you are selecting reboot options in the Uninstall Patch window, Opsware, Inc. recommends that you use the Unix reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the window. If it is not possible to use the Unix reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option in the window.

The following options determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Uninstall Patch window; they do not change the Reboot Required option, which is on the Uninstall Parameters tab of the patch properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install:** Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Do not reboot servers until all patches are installed:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2** Select one of the Rebooting Options.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Specifying Pre and Post Install Scripts for a Unix Patch Uninstallation

For each patch, you can specify a command or script to run before uninstallation or after uninstallation. A pre-uninstall script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-uninstall script fails, the patch would not be removed from a server. A pre-uninstall script could also be used to shut down a service or application before it is removed from a server. A post-uninstall script could be used to perform a certain cleanup process on the managed server.

You can specify the following types of scripts to run on the managed server before or after a patch uninstallation:

- **Pre-Uninstall:** A script that runs before the patch is removed from a managed server.

- **Post-Uninstall:** A script that runs after the patch is removed from a managed server.

To specify a script, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Pre & Post Actions step.

- 2** Select the Pre-Uninstall or Post-Uninstall tab.

You may specify different scripts and options on each of the tabs.

- 3** Select Enable Script.

This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.

- 4** Select either Saved Script or Ad-Hoc Script.

A Saved Script has been previously stored in Opsware SAS with the SAS Web Client. To specify the script, click **Select**.

- 5** If the script requires command-line flags, enter the flags in Commands.

- 6** Specify the information in the User section. The script will be run by this user on the managed server.

- 7** To stop the uninstallation if the script returns an error, select Error.

Scheduling a Unix Patch Uninstallation

You can schedule that a patch will be removed from a server immediately, or at a later date and time.

To schedule a patch uninstallation, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Scheduling step.

- 2** Select one of the following Install Phase options:



- **Run Task Immediately:** This enables you to perform the uninstallation in the Summary Review step.
- **Run Task At:** This enables you to specify a later date and time that you want the uninstallation performed.

- 3** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Setting Up Email Notifications for a Unix Patch Uninstallation

You can set up email notifications to alert users when the patch uninstallation operation completes successfully or with errors.

To set up email notifications, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Notifications step.
- 2** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the uninstallation phase.
- 3** Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 4** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Previewing a Unix Patch Uninstallation

The uninstallation preview process provides an up-to-date report about the patch state of servers. The uninstallation preview is an optional step that lets you see what patches will be removed from managed servers. This preview process verifies whether the servers you selected for the patch uninstallation have that patch installed.



The uninstallation preview process does not report or simulate the behavior of a system with patches removed from the server.

To preview a patch uninstallation, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Summary Review step.
- 2** Verify the information displayed for the Servers, Server Groups, and Patches at the top of the window.
- 3** (Optional) Click **Preview** to see the separate actions that will be performed when the patch is uninstalled. To view the details of a previewed action, select a row in the table.
- 4** Click **Start Job** to launch the job or click **Cancel** to close the Uninstall Patch window without launching the uninstallation.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Patch Uninstallation

You can review progress information about a patch uninstallation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Job Progress step. The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:
 - **Analyze:** Opsware SAS examines the patches needed for the uninstallation, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.
 - **Uninstall:** The patch is uninstalled.
 - **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
 - **Pre/Post Uninstall Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstallation.
 - **Uninstall & Reboot:** When a patch will be installed is also when the server will be rebooted.
 - **Verify:** Installed patches will be included in the software registration.
- 2** To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See the *Opsware[®] SAS User's Guide: Server Automation* for more information on browsing job logs.
- 3** Click **Stop Job** to prevent the job from running or click **Close** to close the Uninstall Patch window.

Chapter 7: Software Management

IN THIS CHAPTER

This section contains the following topics:

- Overview of Software Installation
- Software Installation Process
- Ways to Install Software in Opsware SAS
- Installing Software Using a Software Policy
- Uninstalling Software Using a Software Policy
- Overview of Software Template
- Overview of Running ISM Controls
- Software Policy Compliance
- Software Policy Reports

Overview of Software Installation

Opsware SAS automates the time-consuming process of installing software on managed servers. In the SAS Client, using software policies, you can install software and configure applications across a large number of managed servers with a minimum amount of downtime. Opsware SAS allows you to specify in a software policy the packages and patches to be installed, and the configurations to be applied to the managed servers.

When you apply a software policy to a server, the packages and patches are installed and the application configurations are applied on the managed server in a single step. In a software policy, you can also set the installation order among the software resources in a software policy, and set custom attributes and ISM controls for servers. See the *Opsware[®] SAS Policy Setter's Guide* for information about creating software policies.

To install software in Opsware SAS, you must attach a software policy to servers or groups of servers. When you remediate a server or group of servers, the patches, packages, and application configurations specified in the attached policy are automatically installed and

applied respectively. During remediation, you can separate the download and installation stages of software deployment, specify the reboot operations, schedule the download and installation stages, set email notifications, and associate a ticket ID with the job. The remediation process allows you preview the installation of software before you actually install the software on servers. See "Overview of Software Policies Remediation" on page 372 in this chapter for more information.

You can uninstall any software that you installed by using the SAS Client. To uninstall a software, you must detach a software policy from a server and then remediate the server against that software policy. See "Detaching a Software Policy from a Server" on page 382 in this chapter for more information.

The Software Management feature also enables you to run software compliance scans to determine the compliance status of managed servers with respect to a software policy and then remediate non-compliant servers. See "Software Policy Compliance" on page 396 in this chapter for more information.

The Reporting feature in Opware SAS allows you to generate reports that provide summaries of the software policy compliance across servers. After you generate reports, you can print the reports, export the reports to .html and .xls, and perform actions on the results. See "Software Policy Reports" on page 397 in this chapter for more information.

This section contains information about how to install software using a software policy. It also contains information about running software compliance scans and generating software policy compliance reports. See *Opware® SAS Policy Setter's Guide* for information about uploading packages, and creating and managing software policies.

Software Installation Process

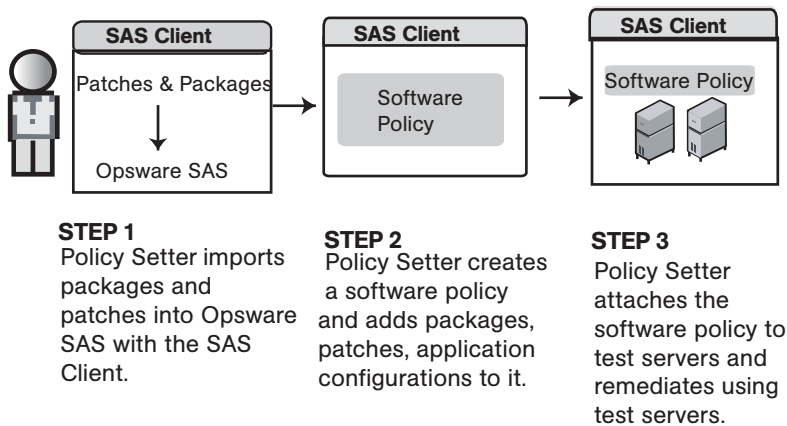
The software installation process, as shown in Figure 7-1, consists of installing software to managed servers by attaching software polices to managed servers and then remediating the servers against those software policies. This phase includes tasks such

as running software compliance scans to determine the compliance status of servers to remediate non-compliant servers, and generating software compliance reports across servers.

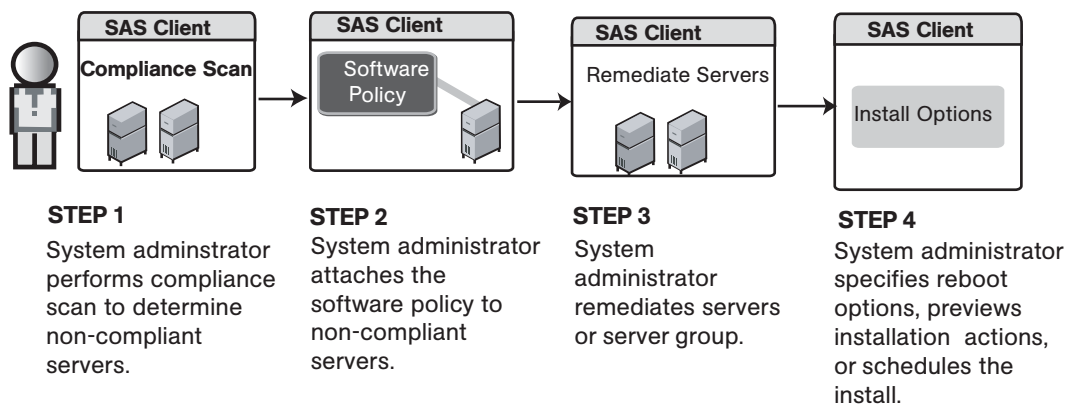
Figure 7-1: Software Management Process

SOFTWARE MANAGEMENT PROCESS

Part A: Set Up Software Policies



Part B: Attach Software Policies to Servers and Remediate



Ways to Install Software in Opware SAS

Opware SAS provides several ways to install software and configure applications. In the SAS Client, you can perform the following tasks:

- Use a software policy to install software and configure applications on a managed server. See “Installing Software Using a Software Policy” on page 368 in this chapter for more information.
- Install software and configure applications on a managed server. See “Installing Software” on page 380 in this chapter for more information.
- Select a single patch and install it directly on a managed server. See “Patch Management for Windows” on page 239 in this chapter for more information.
- Use Visual Packager to prepare installable software packages. See the *Opware® SAS Policy Setter's Guide* for information about Visual Packager.
- Use Application Configuration Management to configure applications on a managed server. See “Application Configuration Management” on page 399 in this chapter for more information.

Installing Software Using a Software Policy

Installing software by using a software policy includes the following steps:

- Attaching a software policy to a server
- Remediating a server against a software policy

Attaching a Software Policy to a Server

When you attach a software policy to a server or group of servers, the software policy is associated with that server or group of servers. This action does not install the software contained in the software policy. To install the software, you must remediate the server with the software policy. See “Remediating Software Policies” on page 374 in this chapter for more information.



You must have a set of permissions to attach a software policy to a server. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

Perform the following steps to attach a software policy to a server:

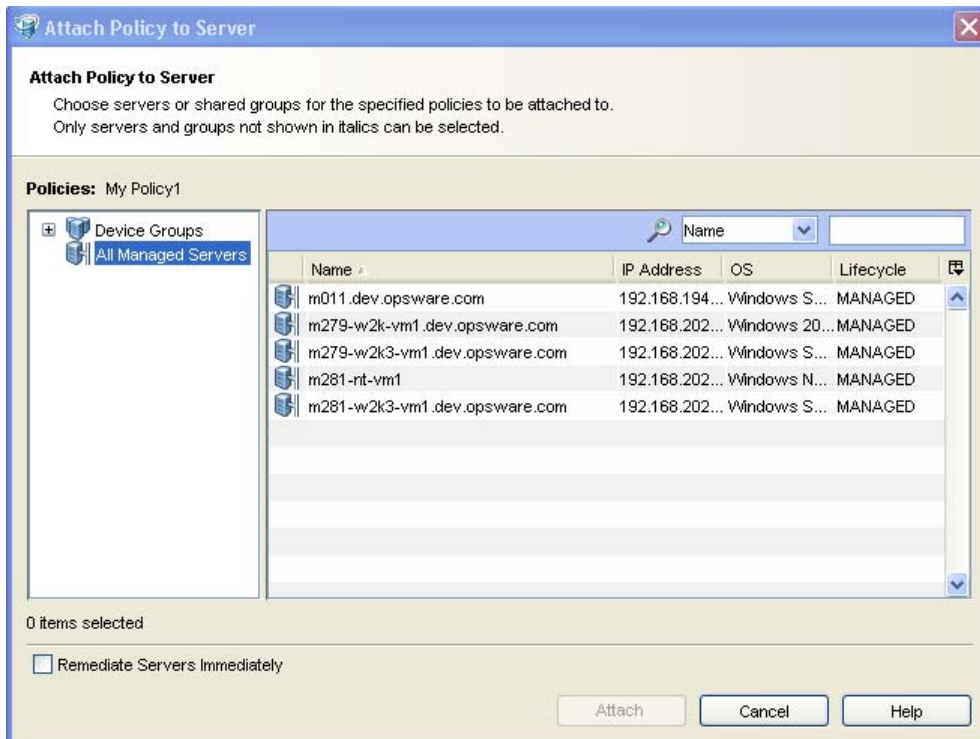
- 1** From the Navigation pane, select **Library ► By Type ► Software Policies**. The software policies appear in the Content pane.
- 2** From the Content pane, select the software policy.
 1. Open the software policy. The Software Policy Window appears.
 2. From the View pane, select Server Usage.
 3. From the View drop-down list, select Servers Attached to Policies.
 4. From the Content pane, select a server.

Or

1. From the View drop-down list in the Content pane, select Server Usage.
2. From the Show drop-down list in the Details pane, select Servers Attached to Policy.
3. Select a server.

- 3 From the **Actions** menu, select **Attach Policy to Server**. The Attach Policy to Server window appears as shown in Figure 7-2:

Figure 7-2: The Attach Policy to Server Window in the SAS Client



- 4 In the Attach Policy to Server window, select servers or device groups and then click **Attach**. You can only select servers that are not in italics. Servers in italics indicate that you do not have the permission to attach a software policy to the server.
- 5 (Optional) Select “Remediate Servers Immediately” to remediate the servers against the software policy. Selecting this option displays the Remediate window. See “Remediating Software Policies” on page 374 in this chapter for more information.

Attaching a Server to a Software Policy

When you attach a server or group of servers to a software policy, the software policy is associated with that server or group of servers. This action does not install the software contained in the software policy. To install the software, you must remediate the server with the software policy. See “Remediating Software Policies” on page 374 in this chapter for more information.



You must have a set of permissions to attach a server to a software policy. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide* for more information.

1 From the Navigation pane, select **Devices** ► **All Managed Servers**. The server list appears in the Content pane.

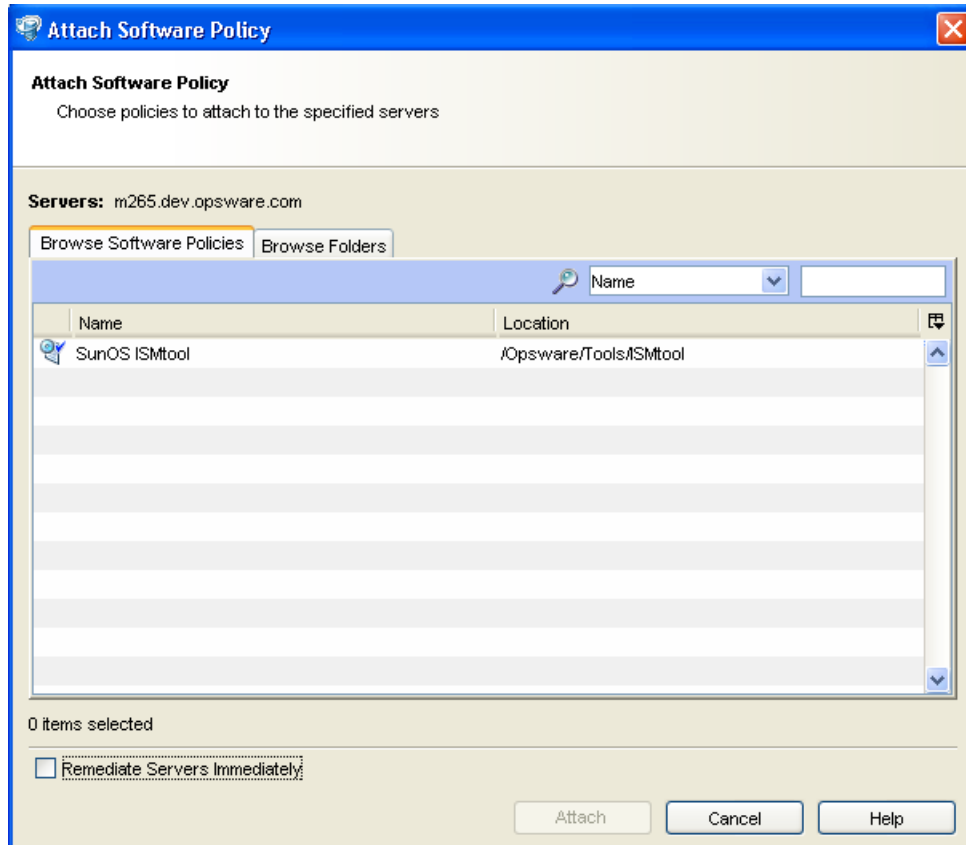
Or

From the Navigation pane, select **Devices** ► **Device Groups**. The device group list displays in the Content pane.

2 From the Content pane, select a server or a device group.

3 From the **Actions** menu, select **Attach** ► **Software Policy**. The Attach Server to Policy window appears as shown in Figure 7-3.

Figure 7-3: The Attach Policy Window in the SAS Client



- 4 Select Browse Software Policies and then select the software policies from the list.

Or

Select Browse Folders and then select the software policies from the folder hierarchy.

- 5 Click **Attach**.

- 6 (Optional) Select "Remediate Servers Immediately" to remediate the servers against the software policy. Selecting this option displays the Remediate window. See "Remediating Software Policies" on page 374 in this chapter for more information.

Overview of Software Policies Remediation

The remediation process installs the packages and patches, and applies the configurations specified in a software policy to a server. (A software policy must be attached to a server or a group of servers before you can remediate the software policy with that server or group of servers.) When you detach a software policy from a server and remediate, then the remediation process uninstalls the software in a software policy.

The remediation process allows you to specify remediation options and pre and post installation scripts required for the remediation process, schedule the download and the installation phase of the remediation process, set up email notifications to alert you about the status of the remediation process, and associate a Ticket ID with each remediation process.



You must have a set of permissions to remediate policies. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide* for more information.

The Remediate window allows you to remediate the servers against the software policies and define the conditions for remediation.

Ways to Open the Remediate Window

From the server list:

- 1 From the Navigation pane, select **Devices** ► **All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices ► Device Groups**. The device group list appears in the Content pane.

- 2** From the Content pane, select a server or device group.
- 3** From the **Actions** menu, select **Remediate**. The Remediate window appears.

From the software policies list:

- 1** From the Navigation pane, select **Library ► By Type ► Software Policies**. The software policy list appears in the Content pane.
- 2** From the Content pane, select a software policy.
 1. From the View drop-down list, select Server Usage.
 2. From the Show drop-down list in the Details pane, select Servers Attached to Policy.
 3. Select servers and then select **Remediate** from the **Actions** menu. The Remediate window appears.

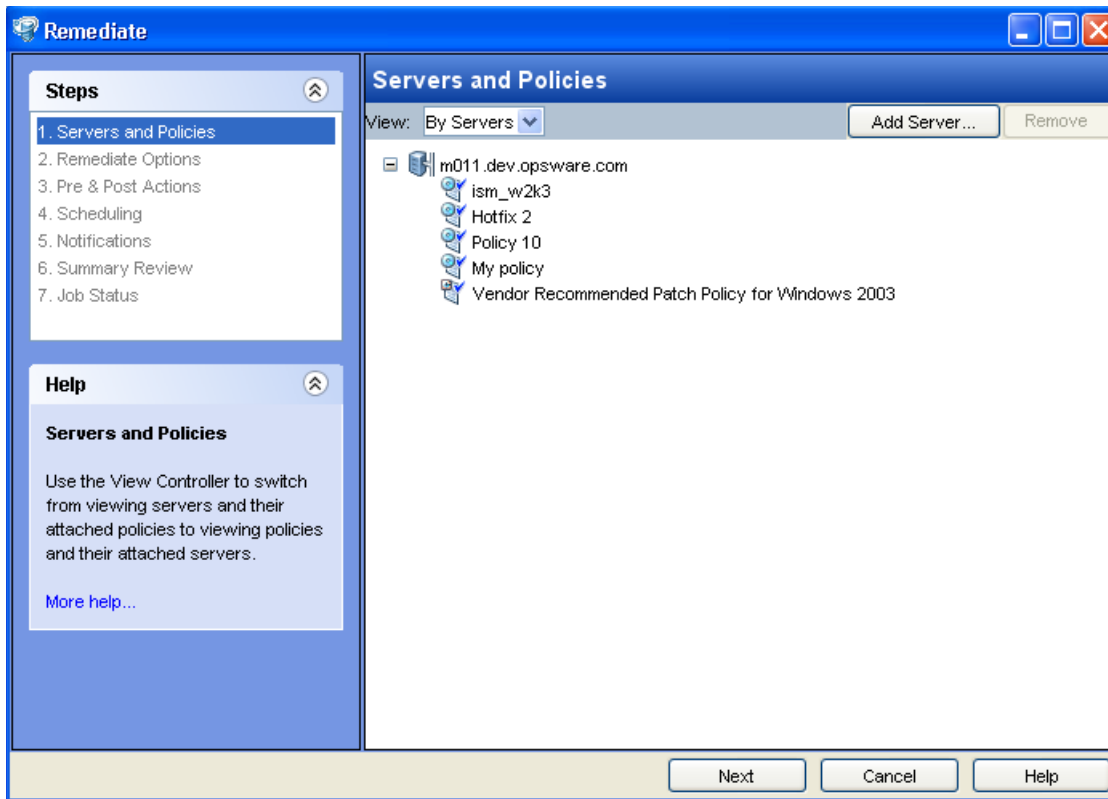
Or

1. From the Content pane, open a software policy. The Software Policy Window appears.
2. From the View pane, select Server Usage.
3. From the View drop-down list, select Servers Attached to Policies.
4. Select servers and then select **Remediate** from the **Actions** menu. The Remediate window appears.

Remediating Software Policies

The Remediate window as shown in Figure 7-4, allows to you remediate the servers against the software policies and consists of the following steps:

Figure 7-4: The Remediate Window in the SAS Client



- Selecting Servers and Policies for Remediation
- Setting Remediate Options
- Specifying Pre and Post Actions for Remediation
- Scheduling Software Policy Remediation
- Setting Email Notifications for Remediation
- Previewing Software Policy Remediation
- Viewing Job Status


Selecting Servers and Policies for Remediation

This step allows you to specify the servers (with software policies attached) for remediation. In this step, you can add and remove servers from the list, view all the application policies attached to a server, and remove software policies attached to servers.

Perform the following steps to select servers and policies for remediation:

- 1** Open the Remediate window from one of the methods described in “Ways to Open the Remediate Window” on page 372.
- 2** In the Remediate window, select the Servers and Policies step. The servers with attached software policies and patch policies appear.

A software policy is represented by the icon .

A patch policy is represented by the icon .

You can also view a list of policies with attached servers by selecting By Policies from the View drop-down list.

- 3** (Optional) Click **Include Server** to add servers to the list or select a server and click **Exclude** to remove servers from the list.
- 4** Select servers with attached software policies.
- 5** Click **Next** to proceed to the Remediate Options step.

Setting Remediate Options

In this step, you can separate the download and installation stage of the remediate policies process. You can choose to continue with the remediate process if an error occurs during the installation or uninstallation of any software contained in the software policy.

Perform the following steps to set the options for remediation:

- 1** From the Remediate window, click **Next** to advance to the Remediate Options step.
- 2** Select one of the following Installation Staging options:
 - **Continuous:** Run all phases as an uninterrupted operation.
This option allows you to run the download and installation step continuously.
 - **Staged:** Allow download and install to be scheduled separately.
This option allows you to separate the download and installation step.

- 3** Select "Attempt to continue running if an error occurs", if you want the remediate process to continue even when an error occurs with any of the package, patches or scripts. By default, this check box is not selected.
- 4** Click **Next** to proceed to the Pre and Post Actions step.

Specifying Pre and Post Actions for Remediation

In this step, you can specify the reboot actions required for the remediation process. You can control when to reboot servers during remediation to minimize the downtime caused by server reboots.

In this step, you can specify the scripts to run on a server before or after remediation. The scripts include:

- **Pre-Download:** A script that runs before packages or patches are downloaded from Opware SAS to the server. This option is available only if you selected Staged in the Remediate Options step.
- **Post-Download:** A script that runs after packages or patches are downloaded from Opware SAS to the server and before the package or patch is installed. This option is available only if you selected Staged in the Remediate Options step.
- **Pre-Install:** A script that runs before packages or patches are installed on the server.
- **Post-Install:** A script that runs after packages or patches are installed on the server.

Perform the following steps to specify pre and post actions for remediation:

- 1** From the Remediate window, click **Next** to advance to the Pre and Post Actions step.
- 2** Select one of the following Reboot options:
 - Reboot servers as dictated by package properties
This option allows you to reboot servers depending on the reboot option specified in the package properties.
 - Hold all server reboots until after all packages are installed and uninstalled
If the reboot option is selected in the package properties, this option allows you to reboot the servers after all the packages are installed and uninstalled. If the reboot option is not selected in the package properties, this option does not reboot the server after all the packages are installed and uninstalled.
 - Suppress all reboots

This option allows you to suppress the reboots even if the reboot option is selected in the package properties.



If a software policy contains multiple non RPM type packages with the option "reboot =yes" selected for every package in the Package Properties window, and the option "Reboot as dictated by package properties" selected in the Remediate window, then remediating a sever with the software policy will reboot the server every time a package is installed. If a software policy contains multiple RPM type packages with the option reboot =yes selected for every RPM package in the Package Properties window, and the option "Reboot as dictated by package properties" selected in the Remediate window, then remediating a sever with software policy will reboot the server only once after all the RPM packages are installed.

- 3** Select the Pre-Install tab or Post-Install tab. You may specify different scripts and options on each of the tabs. If you selected the Staged option in the Remediate Options step, the Pre-Download and Post-Download tabs are also displayed.
- 4** Select **Enable Script**. Selecting Enable Script enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
- 5** Select Saved Script or Ad-Hoc Script from the drop-down list. A Saved script is stored in Opware SAS after you upload the script to Opware SAS. An Ad-Hoc script is intended only for one operation and is not stored in Opware SAS.
- 6** If you selected Saved Script from the drop-down list, click **Select** to specify the script. The Select Script window appears. Select the scripts to run and click **Select**.
- 7** If you selected Ad-Hoc Script from the drop-down list, select the type from the Type drop-down list and then enter the contents of the script in the Script field.
- 8** Enter the command-line flags in the Command field, if required.
- 9** Enter a script time-out value in minutes in the Script Timeout field.
- 10** In the User section, select Root to execute the script as root. To execute the script as a specified user, select Name and enter the user name and then the password.
- 11** Select "Stop job if script returns an error" to stop the installation if the script returns an error.
- 12** Click **Next** to proceed to the Scheduling step.

Scheduling Software Policy Remediation

In this step, you can schedule the installation and download stage to be run immediately or at a specified date and time.

Perform the following steps to schedule the remediation process:

- 1** From the Remediate window, click **Next** to advance to the Scheduling step.
- 2** By default, the Scheduling step displays only the scheduling options for the installation phase. If you selected the Staged option in the Remediate Options step, the scheduling options for the download phase are also displayed.

Select one of the following Install Phase options:



- **Run Task Immediately:** This option allows you to download or install immediately.
- **Run Task At:** This option allows you to specify the date and time to download or install.

- 3** Click **Next** to proceed to the Notification step.

Setting Email Notifications for Remediation

In this step, you can set email notifications to alert users on the success or failure of the download and installation stage of the remediation process. You can associate a Ticket ID with the remediation process.

Perform the following steps to set email notifications:

- 1** From the Remediate window, click **Next** to advance to the Notification step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase. If you selected the Staged in the Remediate Options step, the notification status for the download phase is also displayed.
- 4** Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5** Click **Next** to go to the Summary Review step.

Previewing Software Policy Remediation

In this step, you can view a summary of the remediation process and have the option preview the remediation process.

The preview option allows you to view a detailed list of actions performed on a server as a result of installation or uninstallation of software. It displays information for each server that is selected for remediation. Preview shows the packages and patches that will be installed on or uninstalled from a server, the application configurations that will be applied to a server, the dependency information required for the packages or patches to be installed, the reboots required during the remediation process, and the scripts that will be executed during the remediation process.



When you remediate a sever with a software policy containing Unix patches, in the preview details the software policy compliance is displayed for the server. When you remediate a sever with a software policy containing Windows patches, in the preview details the software policy compliance and patch policy compliance are displayed for the server.

Perform the following steps to preview the remediation process:

- 1** From the Remediate window, click **Next** to advance to the Summary Review step.
- 2** Verify the summary information displayed for the remediation process at the top of the window.
- 3** (Optional) Click **Preview** to view the separate actions that will be performed during the remediate policy process. To view the details of each of the actions, select a row in the table. The details for each action appear.
- 4** Click **Start Job** to remediate the servers.

Viewing Job Status

In this step, you can view the summary information for the progress of a job and the individual status of each action required to be performed for the job to be completed.

Perform the following steps to view the job status:

- 1** From the Remediate window, click **Start Job** to advance to the Job Status step.

- 2** If you selected Run Task Immediately in the Scheduling step, the job begins immediately. If you scheduled the job for a later time, the job will run at the scheduled time. The job progress appears in the Remediate window.
- 3** To view the details of each action, select a row in the table. The details for each action appear.
- 4** Click **End Job** to stop the job from running or click **Close** to close the Remediate window.



You can also view all your jobs from the job logs in the SAS Client. See the *Opware® SAS User's Guide: Server Automation* for information about job logs.

Installing Software



You must have a set of permissions to attach a server to a software policy. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

Perform the following steps to install software on a server:

- 1** From the Navigation pane, select **Devices ► All Managed Servers**. The server list appears in the Content pane.

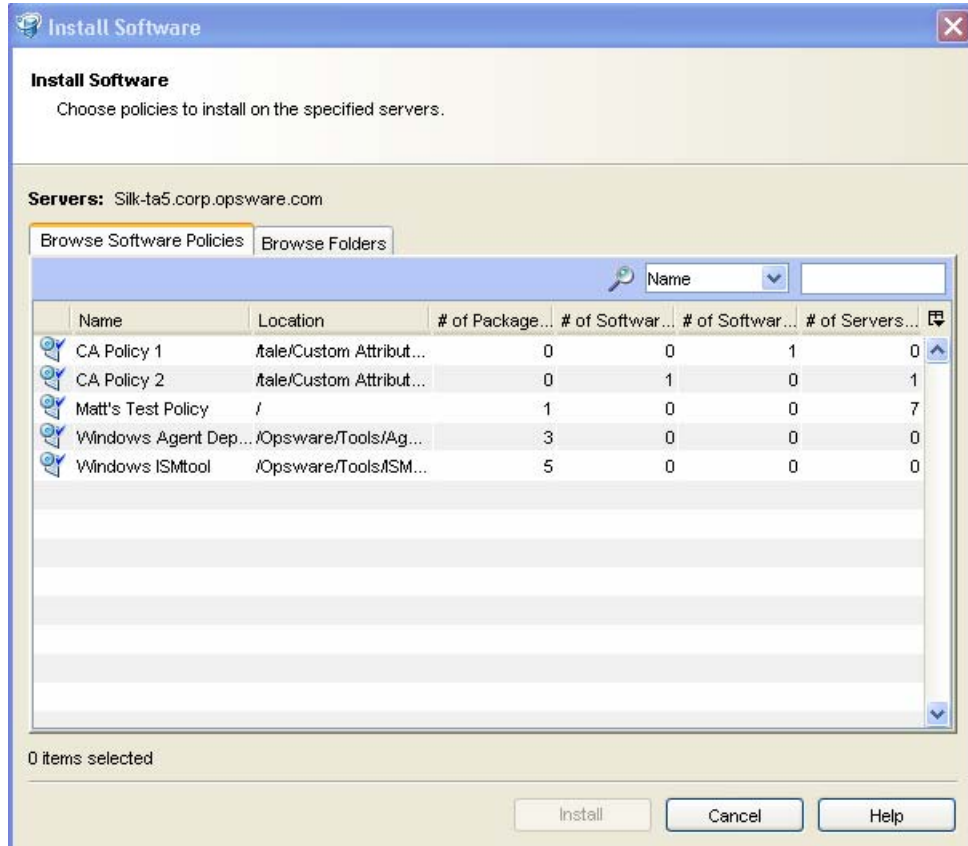
Or

From the Navigation pane, select **Devices ► Device Groups**. The device group list appears in the Content pane.

- 2** From the Content pane, select a server or device group.

- 3 From the **Actions** menu, select **Install ► Install Software**. The Install Software window appears.

Figure 7-5: Install Software Window in the SAS Client



- 4 Select **Browse Software Policies** and then select the software policies from the list.
Or
Select **Browse Folders** and then select the software policies from the folder hierarchy.
- 5 Click **Install**. The Remediate window appears.
- 6 Perform the steps listed in “Remediating Software Policies” on page 374 to remediate the server.

Uninstalling Software Using a Software Policy

Uninstalling software by using a software policy includes the following steps:

- Detaching a software policy to a server
- Remediating a server against a software policy

Detaching a Software Policy from a Server

Detaching a software policy from a server does not delete the policy or uninstall the software from a server. To uninstall the software, you must detach the software policy from the server and then remediate the server with the software policy. See “Remediating Software Policies” on page 374 in this chapter for more information.



You must have a set of permissions to detach a software policy from a server. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

Perform the following steps to detach a software policy from a server:


- 1** From the Navigation pane, select **Devices ► All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices ► Device Groups**. The device group list appears in the Content pane.
- 2** From the Content pane, select a server or a device group.
- 3** From the View drop-down list, select Software Policies.
- 4** From the Show drop-down list, select Policies Attached to Servers to display the software policies attached to the server.

- 5 From the **Actions** menu, select **Detach**. The Detach Software Policy window appears as shown in Figure 7-6.

Figure 7-6: The Detach Software Policy Window in the SAS Client



- 6 Click **Detach**.
- 7 (Optional) Select “Remediate Servers Immediately” to remediate the servers against the software policy. Selecting this option will display the Remediate window. See “Remediating Software Policies” on page 374 in this chapter for more information.

Overview of Software Template

Opsware SAS allows you to install software by using a software template. A software template can only contain other software policies. A software template is not persistently associated with a server or group of servers. When you install a software template to a server or group of servers, the software policies specified in the software template are installed. If you update a software template, servers that already had the software

template applied are not automatically modified to match the updated software template. You must install the software template again to reflect the changes made to the software template on the server.

A software template has the following features:

- A software template is not associated with a server or group of servers.
- A software template contains other software policies.
- A software template is associated with an operating system family.
- Software templates are located in folders.
- Custom attributes can be set on a software template.

Installing software on a server by using a software template consists of the following steps:

- Creating a software template

See the *Opsware® SAS Policy Setter's Guide* for information about creating a software template.

- Adding software policies to a software template

See the *Opsware® SAS Policy Setter's Guide* for information about adding software policies.

- Installing the software template

See "Installing Software Using a Software Template" on page 385 in this chapter for more information.

Ways to Open the Install Software Templates Window

From the server list

- 1** From the Navigation pane, select **Devices** ► **All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices** ► **Device Groups**. The device group list appears in the Content pane.

- 2** From the Content pane, select a server or device group.

- 3 From the **Actions** menu, select **Install ► Install Software Template**. The Install Software Templates window appears.

From the software policies list

- 1 From the Navigation pane, select **Library ► By Type ► Software Policies**. The software policy list appears in the Content pane.
- 2 From the Content pane, select a software template.
- 3 From the **Actions** menu, select **Install Software Template**. The Install Software Templates window appears.

Installing Software Using a Software Template

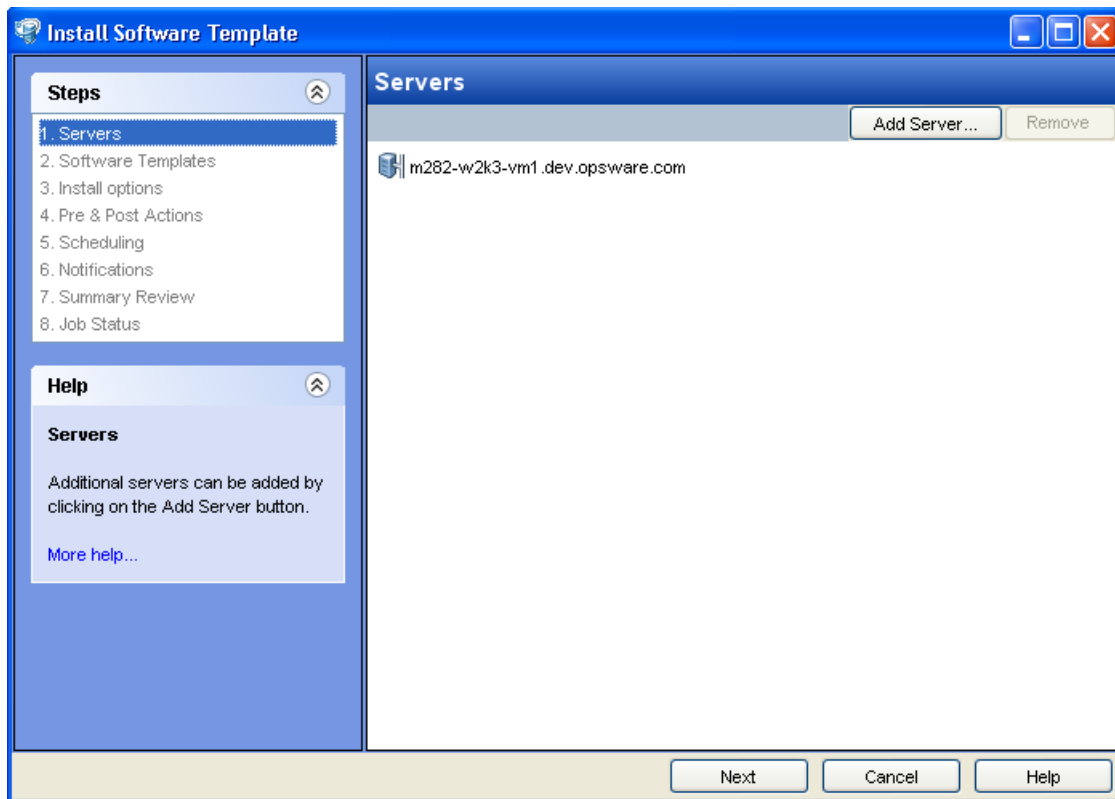


You must have a set of permissions to install to software template. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide* for more information.

The Install Software Templates window allows to you install a software template on a server.

The Install Software Template, as shown in Figure 7-7, allows you to install the software on a server and consists of the following steps:

Figure 7-7: The Install Software Template Window in the SAS Client





- Selecting Servers
- Selecting Software Templates
- Specifying Install Options
- Specifying Pre and Post Actions for Installation
- Scheduling Installation
- Setting Email Notifications
- Previewing Software Installation
- Viewing Job Status



If you access the Software Templates window from the server list, the first step in the window is Selecting Servers. If you access the Software Templates window from the software policies list, the first step in the window is Selecting Software Templates.



In the SAS Client, a software policy is represented by the icon . A software policy is represented by the icon .

Selecting Servers

In this step, you can specify the servers for installing the software template.

Perform the following steps to select servers:

- 1** In the Install Software Policy window, select the Servers step.
- 2** (Optional) Click **Include Server** to add additional servers to the list or click **Exclude** to remove servers from the list.
- 3** Select the servers.
- 4** Click **Next** to proceed to the Software Templates step.

Selecting Software Templates

In this step, you can specify the software templates to install on servers.

Perform the following steps to select software templates:

- 1** From the Software Templates window, click **Next** to advance to the Software Template step.
- 2** In the Software Template window, click **Add Template**. The Attach Software Policy window appears.
- 3** Select the software templates to be installed on the servers.
- 4** (Optional) Click **Remove** to remove any software templates.
- 5** Click **Next** to proceed to the Install Options Step.

Specifying Install Options

In this step, you can separate the download and installation stage of software installation. You can choose to continue with the software installation if an error occurs during the installation of any software contained in a software template.

Perform the following steps to set the installation options:

- 1** From the Software Template window, click **Next** to advance to the Install Options step.
- 2** Select one of the following Installation Staging options:
 - **Continuous:** Run all phases as an uninterrupted operation.
This option allows you to run the download and installation step continuously.
 - **Staged:** Allow download and installation to be scheduled separately.
This option allows you to separate and schedule the download and installation step.
- 3** Select "Attempt to continue running if an error occurs" if you want the installation to continue even when an error occurs with any of the package, patches, or scripts. By default, this check box is not selected.
- 4** Click **Next** to proceed to the Pre and Post Actions step.

Specifying Pre and Post Actions for Installation

In this step, you can specify the reboot actions required for installing software. You can control when to reboot servers during installation to minimize the downtime caused by server reboots. You can also specify the following types of scripts to run on a server before or after software installation:

- **Pre-Download:** A script that runs before packages or patches are downloaded from Opware SAS to the server. This option is available only if you selected Staged in the Install Options step.
- **Post-Download:** A script that runs after packages or patches are downloaded from Opware SAS to the server and before the package or patch is installed. This option is available only if you selected Staged in the Install Options step.
- **Pre-Install:** A script that runs before packages or patches are installed on the server.
- **Post-Install:** A script that runs after packages or patches are installed on the server.

Perform the following steps to specify the pre and post actions for installing software:

- 1** From the Software Template window, click **Next** to advance to the Pre and Post Actions step.
- 2** Select one of the following Reboot options:
 - Reboot servers as dictated by package properties

This option allows you to reboot servers depending on the reboot option specified in the package properties.

- Hold all server reboots until after all packages are installed and uninstalled

If the reboot option is selected in the package properties, this option allows you to reboot the servers after all the packages are installed and uninstalled. If the reboot option is not selected in the package properties, this option does not reboot the server after all the packages are installed and uninstalled.

- Suppress all reboots

This option allows you to suppress the reboots even if the reboot option is selected in the package properties.

- 3** Select the Pre-Install tab or Post-Install tab. You can specify different scripts and options on each of the tabs. If you selected the Staged option in the Install Options step, the Pre-Download and Post-Download tabs are also displayed.
- 4** Select **Enable Script**. Selecting Enable Script enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
- 5** Select Saved Script or Ad-Hoc Script from the drop-down list. A Saved script is stored in Opware SAS after you upload the script to Opware SAS. An Ad-Hoc script is intended only for one operation and is not stored in Opware SAS.
- 6** If you selected Saved Script from the drop down list, click **Select** to specify the script. The Select Script window appears. Select the scripts to run and click **Select**.
- 7** If you selected Ad-Hoc Script from the drop-down list, select the script type from the Type drop-down list and enter the contents of the script in the Script field.
- 8** Enter the command-line flags in the Command field, if required.
- 9** Enter a script time-out value in minutes in the Script Timeout field.
- 10** In the User section, select Root to execute the script as root. To execute the script as specified user, select Name and enter the user name and the password.
- 11** Select "Stop job if script returns an error" to stop the installation if the script returns an error.
- 12** Click **Next** to proceed to the Scheduling step.

Scheduling Installation

In this step, you can schedule the install and download phase to be run immediately or at a specified date and time.

Perform the following steps to schedule software installation:

- 1** From the Software Template window, click **Next** to advance to the Scheduling step.
- 2** By default, the Scheduling step displays only the scheduling options for the install phase. If you selected Staged in the Install Options step, the scheduling options for the download phase are also displayed.

Select one of the following Install Phase options:

- **Run Task Immediately:** This option allows you to download or install immediately.
- **Run Task At:** This option allows you to specify the date and time to download or install.

- 3** Click **Next** to proceed to the Notification step.

Setting Email Notifications

In this step, you can set email notifications to alert users on the success or failure of the download and installation phase of software installation. You can also associate a Ticket ID with the software installation job.

Perform the following steps to set up email notifications:

- 1** From the Software Template window, click **Next** to advance to the Notification step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.

- 3** To set the notification status on the success of the job, select the  icon.

To set the notification status on the failure of the job, select the  icon.

By default, the Notification step displays only the notification status for the installation phase. If you selected Staged in the Install Options step, the notification status for the download phase is also displayed.

- 4** Enter a Ticket ID to be associated with the installation job in the Ticket ID field.
- 5** Click **Next** to go to the Summary Review step.

Previewing Software Installation

In this step, you can view a summary of the software installation. This step provides you with an option to preview the software installation.

The preview option allows you to view a detailed list of actions performed on a server as a result of installation of software. Preview shows the packages and patches that will be installed on a server, the application configurations that will be applied to a server, the dependency information required for the packages or patches to be installed, the reboots required during the software installation process, and the scripts that will be executed during the software installation process.

Perform the following steps to preview the installation process:

- 1** From the Software Template window, click **Next** to advance to the Summary Review step.
- 2** Verify the summary information displayed for the installation process at the top of the window.
- 3** (Optional) Click **Preview** to view the separate actions that will be performed during the software installation. To view the details of each of the actions, select a row in the table. The details for each action appear.
- 4** Click **Start Job** to install the software template.

Viewing Job Status

In this step, you can view the summary information for the progress of a job and the individual status of each action required for the job to be completed.

Perform the following steps to view the job status:

- 1** From the Software Template window, click **Start Job** to advance to the Job Status step.
- 2** If you selected Run Task Immediately in the Scheduling step, the job begins immediately. If you scheduled the job for a later time, the job will run at the scheduled time. The job progress appears in the Software Template window.
- 3** To view the details of each action, select a row in the table. The details for each action appear.
- 4** Click **End Job** to stop the job from running or click **Close** to close the Software Template window.



You can also view all your jobs from the job logs in the SAS Client. See the *Opsware[®] SAS User's Guide: Server Automation* for information about browsing job logs.

Overview of Running ISM Controls

The Run ISM Control window in the SAS Client allows you to run the control scripts in an ISM (Intelligent Software Module).

To run the control scripts in an ISM, you must add the ISM package to a software policy first and then attach the software policy to a server.

See the *Opware® SAS Policy Setter's Guide* for information about adding an ISM package to a software policy. See "Attaching a Software Policy to a Server" on page 368 in this chapter for more information.



You must have a set of permissions to run an ISM Control. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

Ways to Open the Run ISM Control Window

From the server list:

1 From the Navigation pane, select **Devices** ► **All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices** ► **Device Groups**. The device group list appears in the Content pane.

2 From the Content pane, select a server or device group.

3 From the **Actions** menu, select **Run ISM Control**. The Run ISM Control window appears.

From the software policies list:

1 From the Navigation pane, select **Library** ► **By Type** ► **Software Policies**. The software policy list appears in the Content pane.

2 From the Content pane, select a software policy containing an ISM.

1. From the View drop-down list, select Server Usage.

2. From the Show drop-down list in the Details pane, select Servers Attached to Policy. Select servers and then select **Run ISM Control** from the **Actions** menu. The Run ISM Control window appears.


Or

1. From the Content pane, open a software policy containing ISM package. The Software Policy Window appears.
2. From the View pane, select Server Usage.
3. From the View drop-down list, select Servers Attached to Policies.
4. Select servers and then select **Run ISM Control** from the **Actions** menu. The Run ISM Control window appears.

Running ISM Controls

The Run ISM Control window, as shown in Figure 7-8, allows you to run an ISM Control on a server and consists of the following steps:

Figure 7-8: The Run ISM Control Window in the SAS Client



- Selecting Servers
- Selecting Control Parameters

- Scheduling ISM Control Script Execution
- Setting Email Notifications
- Viewing Job Status

Selecting Servers

In this step, you can specify the servers for running an ISM Control.

Perform the following steps to select servers:

- 1** In the Run ISM Control window, select the Servers.
- 2** (Optional) Click **Include Server** to add additional servers to the list or click **Exclude** to remove servers from the list.
- 3** Select the servers.
- 4** Click **Next** to proceed to the Control Parameters step.

Selecting Control Parameters

In this step, you can select a control script in an ISM package to be executed.

Perform the following steps to select the control parameters:

- 1** From the Run ISM Control window, click **Next** to advance to the Control Parameters step.
- 2** From the Software Policy drop-down list, select an ISM package.
- 3** From the Control script drop-down list, select a control script. The drop-down list contains only the control scripts assigned to the ISM package selected in the previous step.
- 4** In the Parameters section, the name of a parameter matches the name of its corresponding custom attribute name. The value of a custom attribute determines the value of the parameter.
- 5** Click **Next** to proceed to the Scheduling step.

Scheduling ISM Control Script Execution

In this step, you can schedule an ISM Control script to be run immediately or at a specified date and time.

Perform the following steps to schedule the ISM Control script execution:



- 1** From the Run ISM Control window, click **Next** to advance to the Scheduling step.

- 2** Select one of the following options:
 - **Run Task Immediately:** This option allows you to run the ISM control script immediately.
 - **Run Task At:** This option allows you to specify the date and time to run the ISM control script.
- 3** Click **Next** to proceed to the Notification step.

Setting Email Notifications

In this step, you can set email notifications to alert users on the success or failure of ISM control script execution. You can associate a Ticket ID with the ISM Control script execution job.

Perform the following steps to set email notifications:

- 1** From the Run ISM Control window, click **Next** to advance to the Notification step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3** To set the notification status on the success of a job, select the  icon.
To set the notification status on the failure of a job, select the  icon.
- 4** Enter a Ticket ID to be associated with the job in the Ticket ID field.
- 5** Click **Next** to go to the Summary Review step.

Viewing Job Status

In this step, you can view the summary information for the progress of a Job and the status of each action required for the Job to be completed.

Perform the following steps to view the job status:

- 1** From the Run ISM Control window, click **Start Job** to advance to the Job Status step.
- 2** If you selected Run Task Immediately in the Scheduling step, the job begins immediately. If you scheduled the job for a later time, the job will run at the scheduled time. The job progress appears in the Run ISM Control window.
- 3** To view the details of each action, select a row in the table. The details for each action will appear.

- 4 Click **End Job** to stop the Job from running or click **close** to close the Run ISM Control window.



You can also view all your jobs from the job logs in the SAS Client. See *Opware® SAS User's Guide: Server Automation* for information about job logs.





Software Policy Compliance

Software compliance indicates whether or not the software policies attached to the selected server are compliant with the actual server configuration. A software policy scan compares the actual configuration of the server with the software policies attached to that server. If the actual server configuration does not match the software policies attached to a server, then the server is said to be out of compliance with the software policies.

A server can be either compliant or non-compliant with respect to a software policy attached to it. If the server's configuration does not match the packages, patches, and application configurations defined in a software policy (attached to that server), then the server is said to be non-compliant with that software policy.

In the SAS Client, when you perform a software compliance scan, the scan indicates the server's overall compliance state as a result of all the software policies attached to the server. Even if only one software policy attached to the server is not compliant, the server is said to be non-compliant. You can then view the non-compliant server and remediate the server against that software policy.

The SAS Client displays the following compliance information for a software policy:

- **Compliant:** If all the software policies attached to a server are compliant, the server is compliant and is represented by the icon .
- **Non-compliant:** If one of the software policies attached to a server is not compliant, the server is non-compliant and is represented by the icon .
- **Scan Started:** The software compliance information is currently being calculated and is represented by the icon .
- **Scan Needed:** The software compliance information needs to be calculated or the compliance information might be inaccurate and is represented by the icon .

In the SAS Client, you can perform a software compliance scan from the Compliance Dashboard or from the server list. See “Compliance Dashboard” on page 205 in Chapter 3 for information about how to perform a software compliance scan.

See “Performing a Software Compliance Scan” on page 397 in this chapter for information about performing a compliance scan from the server list.

Performing a Software Compliance Scan



You must have a set of permissions to perform a software compliance scan. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

Perform the following steps to scan a server for software compliance:

- 1** From the Navigation pane, select **Devices** ► **All Managed Servers**. The server list appears in the Content pane.
- 2** From the Content pane, select the server.
- 3** (Optional) From the View drop-down list, select **Software Policies**.
- 4** From the Content pane, select the column selector drop-down list and select **Software**.
- 5** From the **Actions** menu, select **Scan** ► **Scan Software Compliance**. The compliance status of the server appears in the server list.

After you perform a software compliance scan, you can view the software policies that are not compliant and then remediate the server against that software policy. See “Remediating Software Policies” on page 374 in this chapter for more information.

Software Policy Reports

The Reporting feature in Opware SAS allows you to generate reports that provide a summary of the software policy compliance across servers. You can also generate reports that provide information about software policies on a given server. After you generate reports, you can print them, export the reports to .html and .xls, and perform actions on the results.

The Opware SAS allows you to run the following software policy reports:

- **Software Policy Compliance:** This report groups all managed servers by their software policy compliance level to show compliant and non-compliant servers.
- **Software Policy Compliance By Customer:** This report lists all servers by the customer they are associated with and then by the software policy compliance level.
- **Software Policy Compliance By Facility:** This report displays a chart of all servers by the facility they are associated with and then by the software policy compliance level.
- **Defined Software Policies:** This report lists all the software policies by name and their location in the folder hierarchy.
- **Servers With Attached Software Policies:** This report lists all servers that have one or more software policies attached.
- **Servers In Compliance With Their Software Policies:** This report lists all servers that are in compliance with all of their attached software policies.
- **Servers Not In Compliance with their Software Policies:** This report lists all servers that are not in compliance with all of their attached software policies.
- **Servers Without Attached Software Policies:** This report lists all servers that have no software policies attached.

See "SAS Client Reports" on page 225 in Chapter 4 for information about how to run and view reports in the SAS Client.

Chapter 8: Application Configuration Management

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Application Configuration Management (ACM)
- Application Configuration Creation and Use
- ACM Components
- Application Configuration Inheritance
- Sequence Merging and Inheritance
- Using ACM

Overview of Application Configuration Management (ACM)

Opware Application Configuration Management (ACM) enables you to create templates that help manage configuration files associated with applications. Using ACM, you can manage and update application configuration files from a central location. This ensures that applications in your facility are accurately and consistently configured.

For example, you can create an Application Configuration and set its values, and then push those values to all instances of that application in your facility, whether that application resides on a single server or on groups of servers. You can also check live servers against your Application Configuration and view any differences between the desired state of the application's configuration and the actual state of the application's configuration. If you would like to make a change, simply edit the values and push the changes.

In addition, ACM supports rollback. Because ACM creates a record of the application configuration before the configuration change is made, you can rollback to the original Application Configuration.

ACM also allows you to configure different instances of the same application in your facility. Because an Application Configuration can be attached to several application instances across multiple servers, you can modify default values by customer and facility. For example, you can create default application configuration values across your entire facility, and then make changes to specific instances of the application configuration contained in different facilities and for specific customers.

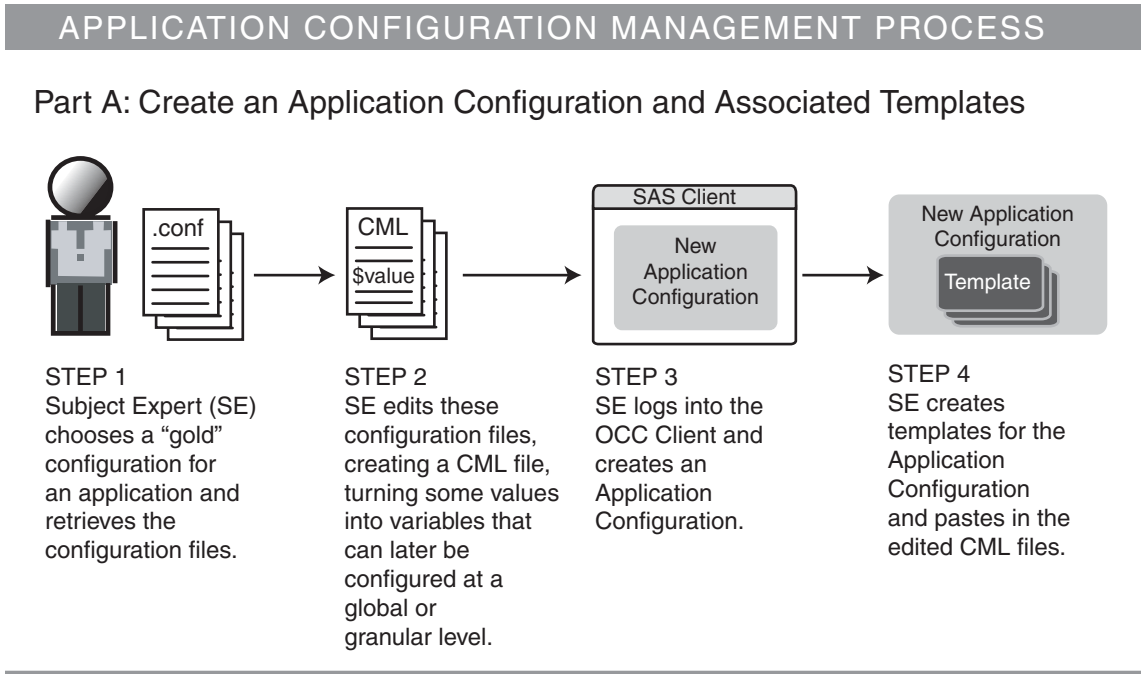
Application Configuration Creation and Use

Using an Application Configuration enables you to manage and modify configuration files for applications on your Opware-managed servers. The process of using ACM follows these general steps:

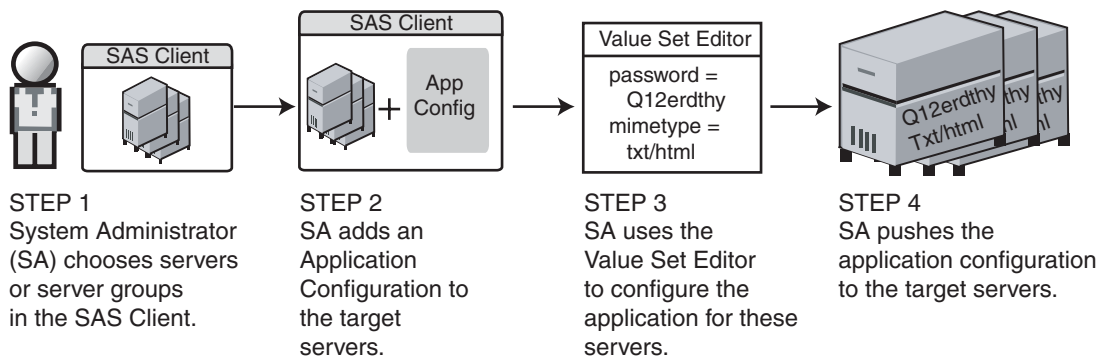
1. **Determine which applications you want to manage:** Your first step is to choose applications to manage. For example, for the iPlanet Web server, you might want to manage the following configuration files: password.conf, obj.conf, mimetypes, and magnus.conf. To manage these iPlanet configuration files with ACM, you need to make templates out of each configuration file.
2. **Create CML files:** For each application file, create a CML file based upon the actual configuration file you want to manage.
3. **Create configuration templates:** Once you have created all of your CML files from the configuration files, create a Application Configuration Template for each CML file inside the SAS Client.
4. **Create application configuration to hold templates:** Once all the configuration files associated with an application have Application Configuration Templates, add them to an Application Configuration. An Application Configuration is a container that houses multiple Application Configuration Templates.
5. **Set the default values:** Next, set the Application Configuration's default values at various levels in the Application Configuration hierarchy, such as at the customer or facility level, or individually at the application instance level on a server.
6. **Attach application configuration to a server (or group):** Once you have created and configured your Application Configuration, attach it to each server (or group of servers) where you are managing application files.

7. **Compare the actual configuration files with the configuration template:** You can easily compare a Application Configuration Template with the actual configuration file on the server and see if any changes have been made. This comparison shows manually changed configuration files or configuration values that have been changed, but not pushed.
8. **Push configuration changes:** No changes are made to the actual configuration files on the server until you push those changes to the server where the Application Configuration files are stored. Application configuration changes can be pushed to individual servers or groups of servers

Figure 8-1: Application Configuration Creation and Usage Process



Part B: Configure and Push Application Configurations to Servers



ACM Components

Application Configuration Management consists of the following main components:

- Configuration Template
- Application Configuration

- Value Set Editor
- Configuration Markup Language (CML)

Configuration Template

An Application Configuration Template is set of values that represent the configuration file of an application. Using the ACM tool in the SAS Client, you can edit the values in the configuration template and push those changes to the actual configuration file on the server.

Using Opsware's Configuration Markup Language (CML), an application expert modifies an application's configuration file and turns it into a Application Configuration Template. In this form, Opsware SAS can then make changes to the actual configuration file on the server. When you make changes to the Application Configuration Template and then push the changes to a server, ACM replaces a section of text in the configuration template with the desired value.

Application Configuration

In many cases, an application has multiple configuration files. For each application managed by ACM, create an Application Configuration that holds all Application Configuration Templates associated with the application. The application configuration aggregates all those templates in a single location.

In addition to configuration templates, an Application Configuration can be configured to contain and execute pre/post configuration scripts.

Value Set Editor

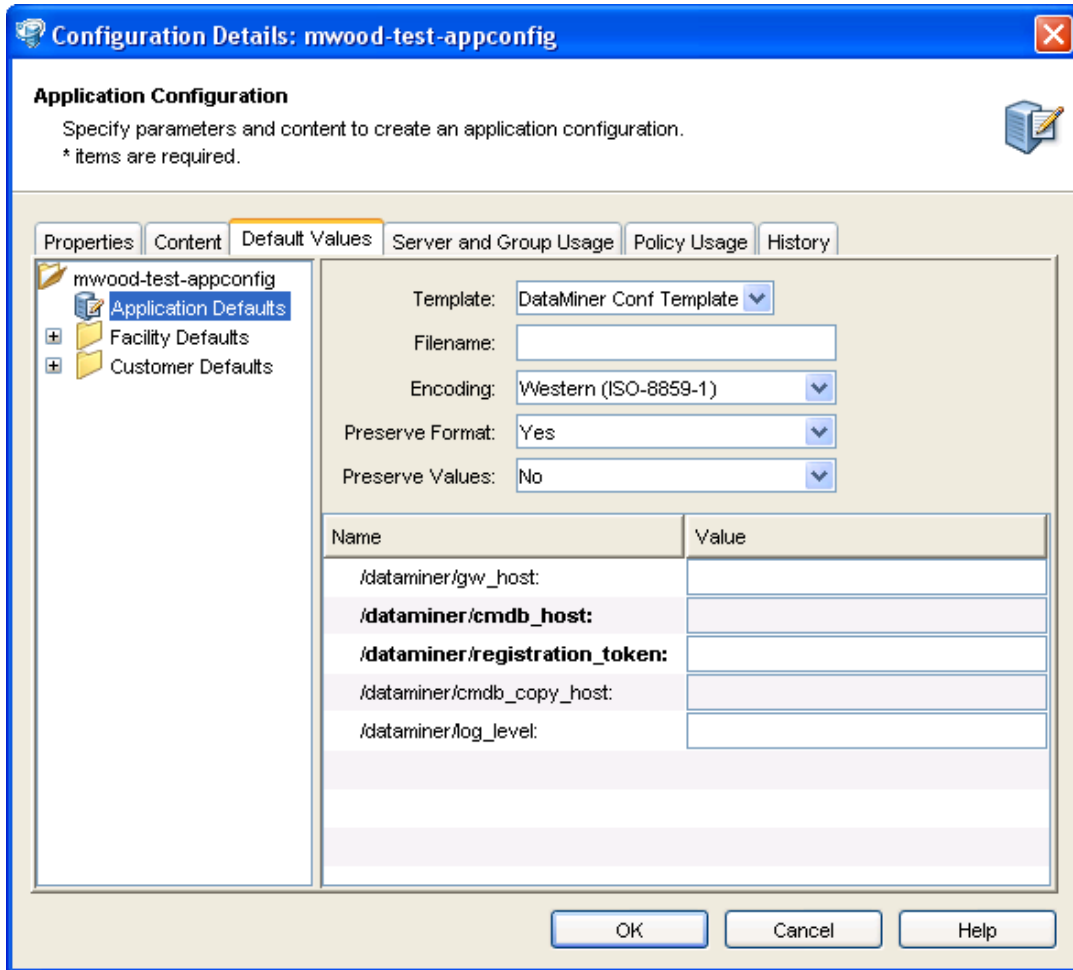
The Value Set Editor enables you to specify the values for each configuration file. Each entry inside a configuration file is represented inside the value set editor as an element, which consists of a name-value pair. The entire collection of elements in a configuration file is referred to as the configuration file's value set – that is, all the elements and their names and values in the file.

You can edit value set elements for an application configuration at two of the following levels:

- **Default Values Level:** The value set elements you edit at this level are applied across all instances of the application that the application configuration is attached to. (These can, however, be overridden by customer or facility.) You access the value

set editor at the default level by selecting the Application Configuration feature from inside the SAS Client and double-clicking an Application Configuration.


Figure 8-2: Application Configuration Default Values



Inside the Value Set Editor, elements that are required will appear in bold.

- **Server Explorer or Device Groups Browser:** Value set elements you edit at this level replaced the actual values in the configuration files on the server when changes are applied (pushed) to a server.

Figure 8-3: Value Set Editor in the Server Explorer



The left side of the Server Explorer's Configured Applications enables you to browse and select an Application Configuration to edit. If an application has more than one instance, then those instances are displayed as children of the main application. The values you edit at the parent application level apply to all instances of the application on the server. You can also edit the values of individual instances of the application.

Value Set Editor Fields

The Value Set Editor contains the following fields:

- **Template:** This enables you to choose the template you want to edit. (Some application configurations can contain multiple Application Configuration Templates.)
- **Filename:** The name of the configuration file on the managed server that is being managed by the Application Configuration Template. If no name is set, then the file name is inherited from its parent in the inheritance hierarchy. If no file name is set

anywhere in the application configuration hierarchy, then the file name listed in the Application Configuration Template is used. This field is set so the Application Configuration knows the configuration file it is supposed to manage. If you have multiple instances of an application on a server, then indicate the full pathname for each configuration file here.

- **Encoding:** Choose a character encoding for the source configuration file that the Application Configuration will be managing. The default encoding is used is the encoding used on the managed server. (Note that UTF-16 encoding is not supported in the SAS Client.)
- **Preserve Format:** Choose this option if you want to both keep comments and preserve as much of the original ordering and spacing of the actual configuration file on the target server. The Application Configuration feature will attempt to preserve as much of the target source file as possible, but may not be able to preserve all comments and formatting. This options is also required if your Application Configuration uses the `@!partial-template@` CML tag. For more information on how to use CML, see the CML tutorial located in the *Opware® SAS Policy Setter's Guide*.
- **Preserve Values:** Choose this option if you want to preserve the values contained in the actual configuration file on the server. With this option selected, the actual file's values will serve as default values for the template, and will be used unless overridden by values at some level of the inheritance hierarchy. In other words, if you would like to preserve a value of the configuration file on the server, then choose this option and leave the value blank in all scope levels. By default, this option is turned off.
- **Show Inherited Values:** Choose this option if you want to show what values are being inherited from higher levels of inheritance hierarchy. When turned off, you will only see the values set at the current level of the Application Configuration inheritance scope. When turned on, you will see both values set at the current level and those that are inherited. This view is read-only.

Value Set Editor Columns

- **Name:** This is the element name from the configuration file. A name can be a simple type, a list of simple types, or a multidimensional list. Multidimensional list names are displayed beneath their parent. Elements that are required appear in bold font. You can double-click to show or hide multidimensional lists. To add another entry to a list type value, right-click the parent and choose **Add Item**. Elements that are required will appear in bold.

- **Value:** Lists all values for each value set in the Application Configuration. You can either enter a literal value or choose an attribute from the Server's settings, such as customer name, customer ID, chassis ID, device ID, and so on. If you leave a setting blank, then the setting is inherited from its parent or ancestor (given that a parent or ancestor has settings configured). To use an Opsware SAS or custom attribute for the value, click the browse (...) button to access the Set Value dialog box.
- **Inherited From:** Indicates where the value is inherited from. The value is applied at the server instance level or inherited from its ancestors in ascending order. The order is server instance, server, group instance, group, customers facility, and application default. However, if Preserve Values option is set in the Value Set Editor, then the configuration file on the server becomes the outermost level of the inheritance hierarchy.

Configuration Markup Language (CML)

To create a Application Configuration Template, you need to transform an application's configuration so that all its value sets become variables. See your Opsware Administrator for more information about using CML.

Application Configuration Inheritance

There are two means of controlling how an application configuration's values are applied and inherited:

- **Default Values Level:** Changes to an Application Configuration's values at this level apply to all instances of the application on all servers. You can, however, override the application configuration by customer or facility.
- **Application Level on a Managed Server:** Changes to an Application Configuration's values at this level apply to applications on a specific server, either globally to all instances of the application, or individually to specific instances of applications on the server.

Application Configuration Default Values

From the Configuration Details dialog box, you can set configuration values at the root level, and further control the scope of the configuration at the customer and facility level.

You can access this level of configuration by opening an application configuration from the SAS Client. Changes made here affect only the application configuration and do not affect the actual configuration file on the server until you push the changes onto a server using the Server Explorer. Figure 8-2 shows the application configuration hierarchy.

Application Defaults apply to all instances of all applications everywhere in the managed server environment, on all managed servers and groups of servers. These defaults are subdivided into the two following groups:


- **Facility:** This applies to all applications (and all instances) existing on servers that belong to a specific facility. Facility settings inherit the Application Configuration default values unless otherwise specified.
- **Customer:** This applies to all applications existing on servers that belong to a specific Customer. Customer settings inherit the facility and then the Application Configuration default values unless otherwise specified.

Application Instance Values

From the Server Explorer (or Device Groups Browser), you can manage configuration values for all or individual instances of an application on a specific managed server. Application configurations at this level inherits default values from the Application Configuration, unless you override them.

You can access this level of configuration by selecting the application or application instance from the Server Explorer ► Configured Applications. These Application Configurations represent actual instances of the application and its configuration on the server. Changes made here can be applied directly to the server when you click **Push**. Figure 8-4 shows the Application Configuration hierarchy at the server level.

Figure 8-4: Application Configuration Inheritance Hierarchy at the Server Level




Application configuration inheritance on a managed server adheres to the following hierarchy:

- **Group of Servers:** This applies to all applications on all servers within the specific group of servers. Configuration values are inherited from the Application Configuration default values unless otherwise specified. (For example, if this group of servers belongs to a specific customer, it inherits the values of that customer.)
 - **Group of Servers Application Instance:** This applies to a specific instance of an application on all servers in the specific group of servers. This instance inherits configuration values from the application defaults and any other application configuration default values, unless otherwise specified.
- **Server:** This applies to all applications on the server. The instance inherits configuration values from application defaults on the managed server from the group of servers it belongs to (if it belongs to a group), and any Application Configuration default values.
 - **Server Instance:** This applies only to the specific instance of the application on the specific server. This instance inherits configuration values from application defaults, from defaults server settings, from the group of servers the server belongs to (if it belongs to a group), and any other Application Configuration default values.

Application Configuration Inheritance Visualized

Figure 8-5 illustrates how Application Configuration values are inherited.

Figure 8-5: Application Configuration Inheritance



Sequence Merging and Inheritance

Because Application Configuration values can be set across many different levels in the Application Configuration inheritance hierarchy (also referred to as the inheritance scope), it is important that you be able control the way multiple sequence values are merged together when you push an Application Configuration on to a server.

ACM allows you to control the way sequence values are merged across inheritance scopes. This means that you can, for example, add some values to a sequence in the Customer scope, Group scope, and the Server scope, and all the values will be merged together to form the final sequence.

The manner in which sequence values are merged is controlled by special tags in the CML template, using three different sequence merge modes:

- **Sequence Replace:** Sequence values from more specific scopes completely replace those from less specific scopes. This occurs for both sequences of sets and lists.
- **Sequence Append:** For lists, values at more general scopes are appended (placed after) to those at more specific scopes. Duplicates, if present, are not removed. For sets, the behavior is the same, except duplicates are merged. For lists, duplicates are identified according to child elements marked with the `primary-key` tag, and then merged. For scalars, this is done by simply removing duplicate values, leaving only the value from the most specific scope (the last occurrence is the merged sequence). This is the default mode, and will be used if nothing else is specified.
- **Sequence Prepend:** Works the same as append, but values at more general scopes are prepended (placed before) to those at more specific scopes.

For example, with these two sets:

- "a, b" – At a more specific (inner) level of the inheritance scope, for example, server instance level.
- "c, d" – At a more general (outer) of the inheritance scope, for example, the server group level.

When the application configuration template is pushed onto the server, the merging results would be:

- Sequence replace: "a, b"
- Sequence append: "a, b, c, d"

- Sequence prepend: “c, d, a, b”

Sequence aggregation occurs not only between scopes, but also within a scope itself. This is evident if there are duplicate values within a sequence of namespaces.

Sequence Replace

In the Replace merge mode (CML tag “`sequence-replace`”), the contents of a sequence defined at a particular scope replace those of less specific scopes, and no merging is performed on the individual elements of the sequence.

For example, if the `sequence-replace` tag has been set for a list in an Application Configuration Template CML source, then values set for that list at the server instance level will override, or replace, those set at the group level and at the Application Configuration default values level.

For example, if a list in an `etc/hosts` file was defined at the group level (outer) as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1 localhost
/system/dns/host/1/hostnames/2 mymachine
/system/dns/host/2/ip           10.10.10.10
/system/dns/host/2/hostnames/1 loghost
```

And the same list was defined at the device scope (inner), as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1 localhost
/system/dns/host/1/hostnames/2 mymachine.mydomain.net
/system/dns/host/2/ip           10.10.10.100
/system/dns/host/2/hostnames/1 mailserver
```

If template had defined the `/system/dns/host` element with the `sequence-replace` tag, the final results of the configuration file on the server after the push would be:

```
127.0.0.1 localhost mymachine.mydomain.net
10.10.10.100 mailserver
```

Sequence Append

When the append list merge mode (CML tag “`sequence-append`”) is used for sequences, the values at more general scopes are appended (placed after) those of more specific scopes. Sequence append mode is the default mode for merging list values. If nothing is specified in the CML of the template, the sequence append will be used.

If a list in an `etc/hosts` file was defined at the group level (outer) as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1 localhost
/system/dns/host/1/hostnames/2 mymachine
/system/dns/host/2/ip           10.10.10.10
/system/dns/host/2/hostnames/1 loghost
```

And the same list was defined at the device scope (inner), as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1 localhost
/system/dns/host/1/hostnames/2 mymachine.mydomain.net
/system/dns/host/2/ip           10.10.10.100
/system/dns/host/2/hostnames/1 mailserver
```

Using the value sets from the above example, if the `/system/dns/host` element was a list with the `sequence-append` tag set in the Application Configuration Template, the final results of the configuration file on the server after the push would be:

```
127.0.0.1 localhost mymachine.mydomain.net
10.10.10.100 mailserver
127.0.0.1 localhost mymachine
10.10.10.10 loghost
```

But since it is not allowable for a hosts file to contain duplicate entries, the `/system/dns/host` element will have to be flagged in the Application Configuration Template as a set rather than a list, because sets do not allow duplicates. To avoid duplication of the list values in the example, the Application Configuration Template author would use the Primary Key option.

Primary Key Option in Sequence Merging

When operating in append mode on sets, new values in more specific scopes are appended to those of less specific ones, and duplicate values are merged with the resulting value placed in the resulting sequence according to its position in the more specific scope.

How this affects merged sequence values depends on what kind of data is contained in the sequence:

- For elements in a sequence which are scalars, the value from the most specific scope is used. In other words, values at the server instance level would replace the values at the group level.

- For elements which are namespace sequences, the value is obtained by applying the merge mode specified for that element (in this example, append) based upon matching up the primary fields.

To avoid the duplication of the `/system/dns/host/.ip` value, the Application Configuration Template author would use the CML `primary-key` option. With this option set, ACM will treat entries with the same value for `/system/dns/host/.ip` as the same and merge their contents.

In the example above, the final results of the configuration file on the server after the push would be:

```
127.0.0.1 localhost mymachine.mydomain.net mymachine
10.10.10.100 mailserver
10.10.10.10 loghost
```



Since it is possible to have a set without primary keys, if there are scalars in the sequence, then an aggregation of all scalar values will be used as the primary key. If there are no scalars, then the aggregation of all values in the first sequence will be used as the primary key. Although this is an estimate, in most cases the values will be merged effectively. To ensure that the correct values are used as primary keys, we recommend that you always explicitly set the primary key in a sequence.

Sequence Prepend

When the append list merge mode (CML tag “`sequence-prepend`”) is used for sequences, the values at more general scopes are prepended (placed before) those those of more specific scopes.

For example, if a sequence in an `etc/hosts` file was defined at the group level (outer) as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1 localhost
/system/dns/host/1/hostnames/2 mymachine
/system/dns/host/2/ip           10.10.10.10
/system/dns/host/2/hostnames/1 loghost
```

And the same sequence was defined at the device scope (inner), as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1 localhost
/system/dns/host/1/hostnames/2 mymachine.mydomain.net
```

```
/system/dns/host/2/ip          10.10.10.100
/system/dns/host/2/hostnames/1 mailserver
```

If the `/system/dns/host` element was a set with the `sequence-prepend` tag set in the Application Configuration Template, the final results of the configuration file on the server after the push would be:

```
10.10.10.10 loghost
127.0.0.1 mymachine localhost mymachine.mydomain.net
10.10.10.100 mailserver
```

To find out how sequences are handled in an Application Configuration Template before you push, you need to look at the contents of the CML template source. For information on how to examine the CML contents of an Application Configuration Template, see “Viewing Application Configuration Template Sources” on page 421.

If you would like to see preview the results of sequence merging before you push, see “Comparing a Template Against an Actual Configuration File” on page 438.

Using ACM

This section contains the following tasks:

- Creating an Application Configuration
- Creating a Configuration Template
- Searching for Application Configurations
- Viewing Application Configuration Template Sources
- Adding or Removing Configuration Templates
- Deleting Application Configurations
- Loading a Template File
- Setting a Configuration Template to Run as a Script
- Specifying Template Order
- Editing Default Values for an Application Configuration
- Attaching an Application Configuration to a Server or Group
- Setting Application Configuration Values on a Server or Group
- Loading Existing Values into a Configuration Template

- Pushing Changes to a Server or Group
- Scheduling an Application Configuration Push
- Comparing Two Configuration Templates
- Comparing a Template Against an Actual Configuration File
- Scanning Configuration Compliance
- Scheduling a Configuration Compliance Scan
- Rolling Back to a Previous State

Creating an Application Configuration

An application configuration can contain one or more Application Configuration Templates (and scripts). Because an application is likely to have more than one configuration file and thus necessitate multiple Application Configuration Templates, you need to create an application configuration to organize and manage your templates from a single location.

If you only want to manage a single configuration file with a single Application Configuration Template, you still need to create an Application Configuration to deploy the template on a server.

To create an application configuration, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2** Select Application Configuration, and then select the Application Configurations tab.
- 3** From the **Action** menu, select **New**.
- 4** In the Properties tab of the Configuration Detail dialog box, specify the following properties:
 - **Name:** This field enables you to name the Application Configuration. (This is required.)
 - **Description:** This field enables you to describe the Application Configuration.
 - **Version:** This section enables you to give a version number to the Application Configuration. This value is set by the person who creates and modifies the Application Configuration. (This version number is not incremented automatically.)
 - **OS:** This allows you to limit the use of the Application Configuration to specific operating systems. The Available list indicates the operating systems you can

associate with the Application Configuration. The Selected list shows the operating systems currently associated with the Application Configuration. Click the arrow to add or remove an operating system to the Application Configuration. Once you add an operating system, then only servers using those operating systems will be able to use the Application Configuration. If you do not want this Application Configuration to be associated with an operating system, select OS Independent.

- **Customers:** This option enables you to limit the use of the application configuration to a specific customer. The Available list of platforms indicates the customers currently supported for the Application Configuration. The Selected list shows the customers associated with the Application Configuration. Click the arrow to add or remove customers from the Application Configuration. If you do not want this Application Configuration to be associated with a customer, select Customer Independent.
- **Notes:** This section allows you to add notes to the Application Configuration.
- **Created:** The date that the Application Configuration was created.
- **Created By:** The user who created the Application Configuration.
- **Last Modified:** The date that the Application Configuration was last modified.
- **Modified By:** The user who last modified the Application Configuration.
- **Tested:** This option allows you to indicate that the Application Configuration has successfully been pushed to a server and that it works.

- 5** Select the Content tab.
- 6** To add an application configuration template, click **Add**.
- 7** In the Select Configuration File dialog box, select an Application Configuration template, and then click **OK**.
- 8** If the Application Configuration is run as a script, select the Application Configuration, right-click, and select one of the following menu items: **None** (will not run as script), or **Data-manipulation, Pre-install, Post-install, Post-error**.
- 9** Click **OK** to create the new Application Configuration.

Creating a Configuration Template

An Application Configuration Template is similar to an actual native application configuration file, but one that has had its variable portions marked up with Opware's Configuration Markup Language (CML). (CML is a markup language used for managing configuration files.)

To manage a configuration file with ACM, create an Application Configuration Template. Before a Application Configuration Template can be applied to a server, it needs to be added to an Application Configuration.

An Application Configuration Template can be configured to run as a script, either before all the configurations are made or after. Also, you can set a script to run as a post-error script to rollback all changes if the configuration push fails. See "Setting a Configuration Template to Run as a Script" on page 425 in this chapter for more information.

To create an Application Configuration Template, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2** Select Application Configuration, and then select the Configuration Templates tab.
- 3** From the **Action** menu, select **New**.
- 4** In the Properties tab of the Template Detail dialog box, enter the following information:
 - **Name:** This allows you to enter a name for the Application Configuration or Application Configuration Template. (This is required.)
 - **Description:** This enables you to enter a description.
 - **Version:** This value is set by the person who creates and modifies the Application Configuration/Application Configuration Template. (The version number is not incremented automatically.)
 - **OS:** This allows you to limit the use of the Application Configuration Template to a specific operating system. The Available list of operating systems indicates the operating systems you can associate with the Application Configuration or Application Configuration Template. The Selected list shows the operating systems currently associated with the Application Configuration/Application Configuration Template. Click the arrow to add or remove an operating system to the Application Configuration Template. Once you add an operating system, then only servers using those operating systems can use the Application Configuration Template. If

you do not want this Application Configuration/Application Configuration Template to be associated with an operating system, select the OS Independent option.

- **Customers:** This option allows you to limit the use of the Application Configuration/Application Configuration Template to a specific customer. The Available list of platforms indicates the customers that are currently supported for the Application Configuration or Application Configuration Template. The Selected list shows the customers associated with the Application Configuration/Application Configuration Template. Click the arrow to add or remove customers from the Application Configuration or Application Configuration Template. If you do not want an Application Configuration or Application Configuration Template to be associated with customer, select the OS Independent option.
- **Script Type:** This allows you to set the Application Configuration Template to function as a template, localization file, or script. If the file is a script, you can specify the script language, such as WIndows BAT, JS, VBS, CMD, WSF, and PY; and Unix SH or Other script.
- **Created:** This shows the date that the Application Configuration Template was created.
- **Created By:** This shows the user who created the Application Configuration Template.
- **Last Modified:** This shows the date that the Application Configuration Template was last modified.
- **Modified By:** This shows the user who last modified the Application Configuration Template.
- **Tested:** This option allows you to indicate that the Application Configuration Template has successfully been pushed to a server and that it works.

- 5** Select the Content tab.
- 6** Copy the contents of your CML file here.
- 7** Click **Validate** to validate the CML syntax.
- 8** When you are finished, click **OK**.

Searching for Application Configurations

You can use the SAS Client Search tool to find Application Configurations and Application Configuration Templates in your facility. You can search for Application Configurations by name, by the operating system, and many other criteria.

To search for Application Configurations, perform the following steps:

- 1** From inside the SAS Client, make sure the search pane is activated by selecting View ► Search pane.
- 2** From the top drop down list, select Application Configuration or Application Configuration Templates.
- 3** Click the green arrow button or ENTER to execute the search.
- 4** The results appear in the Contents pane.
- 5** If you want to extend your search criteria, add new criteria in the search parameters section at the top of the Contents pane. You can also save the search by clicking Save, or export the Search results to HTML or CSV.

Viewing Application Configuration Template Sources

In some cases, you will need to examine the contents of your Application Configuration Template and view its CML source, especially if you need to understand which list merging modes have been set in the template before you push the Application Configuration to a server.

For information on Application Configuration sequence merge modes, see “Sequence Merging and Inheritance” on page 412.

To view Application Configuration Template CML source, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2** Select Application Configuration, and then select the Configuration Templates tab.
- 3** To open an Application Configuration Template in the list, double-click it. (Or right-click the template and choose **Open**.)
- 4** Select the Content tab, and you see the CML contents of the Application Configuration Template.

Adding or Removing Configuration Templates

You can add as many Application Configuration Templates to an Application Configuration as you like. If an Application Configuration Template doesn't belong or you no longer need it in an Application Configuration, you can remove it.

To add an Application Configuration Template to an Application Configuration, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2** Select Application Configuration, and then select the Application Configurations tab.
- 3** To open an Application Configuration in the list, double-click it.
- 4** Select the Content tab.
- 5** To add an Application Configuration Template, click **Add**.
- 6** From the Select Configuration dialog box, select the Application Configuration Template, and then click **OK**.

Deleting Application Configurations

If you no longer need an Application Configuration, you can delete it. Once you delete an Application Configuration, you cannot recover it.

To delete an Application Configuration, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2** Select Application Configuration, and then select the Application Configurations tab.
- 3** Select an Application Configuration, right-click, and choose **Delete**. (This will not delete any Application Configuration Templates that belong to the Application Configuration.)
- 4** To delete a Application Configuration Template, select the Configuration Templates tab.
- 5** Select an Application Configuration Template, right-click, and choose **Delete**.

Loading a Template File

If a CML template is already created for use in an Application Configuration, you can upload the template from a local or remote file system.



For configuration files on Windows servers which are encoded in UTF-8, the first three characters of the configuration file might contain a Byte Order Mark (BOM). If you import this file into an Application Configuration Template, the BOM will appear in the template after the file is loaded. If you do not want this BOM to be included in the Application Configuration Template, remove it after you upload the configuration file into the template.

UTF-16 encoding is not supported in the SAS Client.

To load a template file, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2** From the **Action** menu, select **Upload Template**.
- 3** In the Open dialog box, browse to locate the template file (a CML file should have the TPL file extension, but this is not mandatory). If the character encoding of the template file is different than the default encoding of your desktop, select an item from the Encoding drop-down list. (Note that UTF-16 encoding is not supported in the SAS Client.)
- 4** Click **Open**.
- 5** In the Configuration File Upload dialog box, fill out the following information:
 - **Name**: This allows you to enter a name for the Application Configuration or Application Configuration Template. (This is required.)
 - **Description**: This enables you to enter a description.
 - **Version**: This value is set by the person who creates and modifies the Application Configuration/Application Configuration Template. (The version number is not incremented automatically.)
 - **OS**: This allows you to limit the use of the Application Configuration Template to a specific operating system. The Available list of operating systems indicates the operating systems you can associate with the Application Configuration or Application Configuration Template. The Selected list shows the operating systems

currently associated with the Application Configuration/Application Configuration Template. Click the arrow to add or remove an operating system to the Application Configuration Template. Once you add an operating system, then only servers using those operating systems can use the Application Configuration Template. If you do not want this Application Configuration/Application Configuration Template to be associated with an operating system, select the OS Independent option.

- **Customers:** This option allows you to limit the use of the Application Configuration/Application Configuration Template to a specific customer. The Available list of platforms indicates the customers that are currently supported for the Application Configuration or Application Configuration Template. The Selected list shows the customers associated with the Application Configuration/Application Configuration Template. Click the arrow to add or remove customers from the Application Configuration or Application Configuration Template. If you do not want an Application Configuration or Application Configuration Template to be associated with customer, select the OS Independent option.
- **Script Type:** This allows you to set the Application Configuration Template to function as a template, localization file, or script. If the file is a script, you can specify the script language, such as Windows BAT, JS, VBS, CMD, WSF, and PY; and Unix SH or Other script.
- **Created:** This shows the date that the Application Configuration Template was created.
- **Created By:** This shows the user who created the Application Configuration Template.
- **Last Modified:** This shows the date that the Application Configuration Template was last modified.
- **Modified By:** This shows the user who last modified the Application Configuration Template.
- **Tested:** This option allows you to indicate that the Application Configuration Template has successfully been pushed to a server and that it works.

6 Next, select the Content tab.

7 You should see the CML template. Click **Validate** to validate the CML syntax.

8 When you are finished, click **OK**. This will create both the Application Configuration Template and an Application Configuration to house the template.

Setting a Configuration Template to Run as a Script

In addition to using Application Configuration Templates to replace values of actual configuration files, you can also add scripts to an Application Configuration.

For example, you might want to add a post-install script that reboots the server after configuration changes have been made. Or, you might want to use a data-manipulation script to handle certain configuration files which contain unreadable or otherwise unmanageable data before you perform an import, preview, or push the Application Configuration.

If you are configuring an IIS server, you can use a data-manipulation script to read the metabase information into a flat file. When this information gets parsed with the Application Configuration Template, you can run a data-manipulation script to implement the changes in the flat file.

To set an Application Configuration Template as a script, you need to set the Application Configuration Template script type and then specify the type of script execution.

To set a template to run as a script, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2** Select Application Configuration, and then select the Application Configurations tab.
- 3** In the Content pane, double-click the Application Configuration that contains the Application Configuration Template that you want to run as a script.
- 4** In the Configuration Details window, select the Content tab.
- 5** Select the Application Configuration Template in the list, right-click, and choose **Data-manipulation, Pre-install, Post-install, or Post-error** to set the script execution type.



If you would like to change the order in which the Application Configuration Template is run inside the Application Configuration, select the Application Configuration Template, right-click, and select **Move Up** or **Move Down**.

- 6** Select the Application Configuration Template again, right-click, and select **Open Template**.

- 7** In the Template Details window, choose a script type from Type drop-down list. Click **OK**.
- 8** Click **OK** to close the Configuration Details window.



When pushing an application configuration that contains a JScript or VBScript pre- or post-install and post-error scripts, the push succeeds even though the scripts fail. In these cases, the push ignores the scripts errors altogether. The application configuration does not catch the failure of the scripts and allows the push to complete without errors.

If you plan to use these types of scripts, you must make sure that the scripts are free of errors to detect possible failures, and have the script forcibly return a non-zero exit status by invoking `WScript.Quit(<status>)`.

Specifying Template Order

An Application Configuration can contain one or several Application Configuration Templates and scripts. However, you might want to control templates application and script execution order.

For example, you might want to apply changes to certain configuration files before others. Or, you might have a script in the Application Configuration that restarts the server after all the Application Configuration changes have been applied to the application on the server.

To specify template order, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2** Select Application Configuration, and then select the Application Configurations tab.
- 3** To open an Application Configuration in the list, double-click it.
- 4** In the Configuration Detail dialog box, select the Content tab.
- 5** All the Application Configuration Templates and scripts (if there are any) contained within the Application Configuration are displayed. Notice that each Application Configuration Template has a number next to it that indicates the order.
- 6** To reorder the Application Configuration Templates, select one and then click **Move Item Up** or **Move Item Down**.



For better organization, it is useful to position at any pre-install scripts at the top of the list, and position post-install or post-error scripts at the bottom of the list.

- 7 When you are finished, click **OK**.

Editing Default Values for an Application Configuration

Once you have created an Application Configuration, you can edit its default configuration values. An Application Configuration's default values apply to all instances of the application on all attached servers. (An Application Configuration only affects attached servers.)

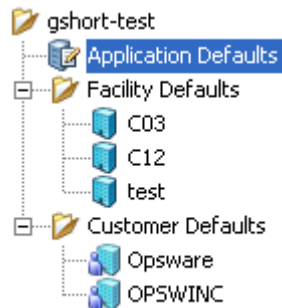
However, you can override the scope of an application configuration's default values by customer or facility. You can also edit specific instances of the application configuration to override the scope of an application configuration's default values. All elements that are required appear in bold font.

To set default values for an application configuration, perform the following steps:

- 1 Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2 Select Application Configuration, and then select the Application Configurations tab.
- 3 In the Content pane, double-click the Application Configuration.
- 4 In the Configuration Details dialog box, select the Default Values tab.
- 5 The left side of the dialog box shows the Application Configuration hierarchy; this allows you to set default values at the application defaults (root) level, the customer level, and the facility level.
- 6 To set default values, select a server in the hierarchy and double-click it. The default values will display.

Figure 8-6 shows an example of the Application Defaults node selected. Any changes to value sets at this level will apply to all facilities and customers – including all applications on all attached servers.

Figure 8-6: Application Configuration Default Values Hierarchy



- 7** Edit the default values for each value set in the Application Configuration Template. The following settings will be displayed:
- **Template:** This enables you to choose the template you want to edit. (Some application configurations can contain multiple Application Configuration Templates.)
 - **Filename:** The name of the configuration file on the managed server that is being managed by the Application Configuration Template. If no name is set, then the file name is inherited from its parent in the inheritance hierarchy. If no file name is set anywhere in the application configuration hierarchy, then the file name listed in the Application Configuration Template is used. This field is set so the Application Configuration knows the configuration file it is supposed to manage. If you have multiple instances of an application on a server, then indicate the full pathname for each configuration file here.
 - **Encoding:** This enables you to choose a character encoding for the source configuration file that the Application Configuration will be managing. The default encoding is the encoding used on the managed server. (Note that UTF-16 encoding is not supported in the SAS Client.)
 - **Preserve Format:** Choose this option if you want to both keep comments and preserve as much of the original ordering and spacing of the actual configuration file on the target server. The Application Configuration feature will attempt to preserve as much of the target source file as possible, but may not be able to preserve all comments and formatting. This options is also required if your

Application Configuration uses the `@!partial-template@` CML tag. For more information on how to use CML, see the CML tutorial located in the *Opsware[®] SAS Policy Setter's Guide*.

- **Preserve Values:** To preserve the values contained in the actual configuration file on the server, choose **Yes** for this option and leave the value blank in all scope levels. With this option selected, the actual file's values will serve as default values for the template, and will be used unless overridden by values at some level of the inheritance hierarchy. By default, this option is turned off.
 - **Show Inherited Values:** This appears only on an Application Configuration instance attached to a server or server group, not at the Application Configuration default values level. Choose this option if you want to show at what values are being inherited from higher levels of inheritance hierarchy. When turned off, you will only see the values set at the current level of the Application Configuration inheritance scope. When turned on, you will see both values set at the current level and those that are inherited. This view is read-only.
 - **Name column:** This is the value set element name from the configuration file. A value set name can consist of a simple type, a list of simple types, or a multidimensional list. Elements that are required appear in bold font. Multidimensional list key names are displayed beneath their parent. Double-click to show or hide multidimensional lists. To add another key name, right-click the parent and select **Add Item**. You can also use the right-click menu to search for other values or keys, copy values, or clear values.
- 8** (Optional) You can copy and paste one value set to another. To do this, select the value set name, right-click, and choose **Copy Values**. Then, paste this value by right-clicking the target value set and choosing **Paste Values**. Copying and pasting will copy the entire value set and will override the old value set.
 - 9** (Optional) You can expand and retract the Application Configuration value set, by right-clicking and choosing **Collapse Subtree**. All name-value hierarchies will be closed. If you would like to find a value set name or value, select the value set, right-click and choose **Find Name** or **Find Value**.
 - 10** When you have finished editing the value sets for the Application Configuration, click **Save Changes**.

Attaching an Application Configuration to a Server or Group

After you have created an Application Configuration and added all the necessary Application Configuration Templates and scripts and edited its default values, you can add the Application Configuration to a server or public group of servers.


For an Application Configuration to manage an application on a server, it must be added to a server or group of servers. Once you add an Application Configuration to a server or group of servers, the values of the Application Configuration are not applied to the configuration files on the server until you push them to the server. This enables you to add the Application Configuration, edit its values, and then wait until you are ready to apply the changes before pushing them to the server.



You can only add an Application Configuration to a public group of servers.

To attach an Application Configuration to a server or group, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select **Devices**.
- 2** Select a server or group of servers.
- 3** Select a server (or group) from the Content pane, right-click and choose **Configure Applications**.
- 4** You now see the Server Explorer (or Groups Browser), with the Configured Applications folder selected. From the **Action** menu, select **Add Configuration**.
- 5** In the Select Application Configuration dialog box, select an Application Configuration.

Use the search tool  in the upper right corner of the dialog box if the list is large and you want to search by a specific criteria (such as OS, last modified, and so on).

- 6** When you have selected an Application Configuration, click **OK**. The Application Configuration is added to the server or group.

Setting Application Configuration Values on a Server or Group

Once an Application Configuration has been attached to a server or group of servers, you can edit its values. You can also override the default values set at the Application Configuration level. If the server (or group) has multiple instances of an application installed, you can set values for all instances of the application or individual instances.

If you do not edit any values on the Application Configuration at the server or group level, then the values are inherited from the default values set at the Application Configuration level. See “Application Configuration Inheritance” on page 407 in this chapter for more information.

To set Application Configuration values on a server or group, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select a server or group of servers.
- 3** Select a server or group in the Content pane, right-click, and choose **Configure Applications**.
- 4** You now see the Server Explorer (or Groups Browser), with the Configured Applications folder selected. Select the Installed Configurations tab. All Application Configurations that have been attached to the server (or group) will be displayed.



If you do not see an Application Configuration, then none have been attached to the server or group. See “Attaching an Application Configuration to a Server or Group” on page 430 in this chapter for more information.

- 5** From the left side of the Application Configuration hierarchy, select either the top level application folder or an instance of the application, then edit the values of the Application Configuration. Before you start editing values, consider the following about Application Configuration inheritance:
 - If you do not edit any values on the application or application instance level, then all values are inherited from the Application Configuration’s default values. (See “Editing Default Values for an Application Configuration” on page 427 in this chapter for more information.)
 - If you want to see which values are being inherited from a higher level of the Application Configuration hierarchy, select the Show Inherited Values option. Selecting this option will show a read only view of all names and values in the

Application Configuration, and the inherited from column shows where inherited values are derived from.

Once you have selected a level of the Application Configuration to edit, you can now start editing values. Because every configuration file is unique, what you actually see and are able to edit will be different for each Application Configuration.

- 6** Edit the default values for each value set in the Application Configuration Template. The following settings will be displayed:
- **Template:** This enables you to choose the template you want to edit. (Some application configurations can contain multiple Application Configuration Templates.)
 - **Filename:** The name of the configuration file on the managed server that is being managed by the Application Configuration Template. If no name is set, then the file name is inherited from its parent in the inheritance hierarchy. If no file name is set anywhere in the application configuration hierarchy, then the file name listed in the Application Configuration Template is used. This field is set so the Application Configuration knows the configuration file it is supposed to manage. If you have multiple instances of an application on a server, then indicate the full pathname for each configuration file here.
 - **Encoding:** This enables you to choose a character encoding for the source configuration file that the Application Configuration will be managing. The default encoding is the encoding used on the managed server. (Note that UTF-16 encoding is not supported in the SAS Client.)
 - **Preserve Format:** Choose this option if you want to both keep comments and preserve as much of the original ordering and spacing of the actual configuration file on the target server. The Application Configuration feature will attempt to preserve as much of the target source file as possible, but may not be able to preserve all comments and formatting. This options is also required if your Application Configuration uses the `@!partial-template@` CML tag. For more information on how to use CML, see the CML tutorial located in the *Opsware® SAS Policy Setter's Guide*.
 - **Preserve Values:** To preserve the values contained in the actual configuration file on the server, choose **Yes** for this option and leave the value blank in all scope levels. With this option selected, the actual file's values will serve as default values for the template, and will be used unless overridden by values at some level of the inheritance hierarchy. By default, this option is turned off.

- **Show Inherited Values:** This appears only on an Application Configuration instance attached to a server or server group, not at the Application Configuration default values level. Choose this option if you want to show at what values are being inherited from higher levels of inheritance hierarchy. When turned off, you will only see the values set at the current level of the Application Configuration inheritance scope. When turned on, you will see both values set at the current level and those that are inherited. This view is read-only.
 - **Name column:** This is the value set element name from the configuration file. A value set name can consist of a simple type, a list of simple types, or a multidimensional list. Elements that are required appear in bold font. Multidimensional list key names are displayed beneath their parent. Double-click to show or hide multidimensional lists. To add another key name, right-click the parent and select **Add Item**. You can also use the right-click menu to search for other values or keys, copy values, or clear values.
- 7** (Optional) You can copy and paste one value set to another. To do this, select the value set name, right-click, and choose **Copy Values**. Then, paste this value by right-clicking the target value set and choosing **Paste Values**. Copying and pasting will copy the entire value set and will override the old value set.
- 8** (Optional) You can expand and retract the Application Configuration value set, by right-clicking and choosing **Collapse Subtree**. All name-value hierarchies will be closed. If you would like to find a value set name or value, select the value set, right-click and choose **Find Name** or **Find Value**.
- 9** When you have finished editing the Application Configuration values, click **Save Changes**. These changes won't be applied to the configuration files on the server or group until you push the changes. To preview what the changes will look like before you push them, click **Preview**. To push the changes, click **Push**.

Loading Existing Values into a Configuration Template

You might want to import values into the value set editor from a configuration file on a managed server. Selecting the **Import Values** menu item reads the actual existing configuration file on a server, parses the values, and applies them into the instance level value sets for the Application Configuration Template. This shows the values currently in the actual configuration. After you import the values, you can modify some of those values and then push the changes back onto the server.

To load existing values into the value set editor, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select a server or group of servers.
- 3** Select a server or group in the Content pane, right-click, and choose **Configure Applications**.
- 4** You now see the Server Explorer (or Device Groups Browser), with the Installed Configurations tab selected. All Application Configurations that have been attached to the server (or group) will be displayed.
- 5** From the View pane, click the plus (+) symbol to expand Application Configuration folder and select an Application Configuration instance to edit.
- 6** From the Content pane, choose an Application Configuration Template from the Template drop-down list.
- 7** In the File name field, enter the absolute file name of the configuration file that contains the values that you want to import.
- 8** Next, right-click in the Name column and choose **Import Values**. A confirmation message appears, warning you that proceeding with this operation will overwrite any current values. Click **Yes** to proceed.
- 9** All of the values for the Application Configuration Template are replaced with the values from the actual configuration file.
- 10** Click **Save Changes**.

Pushing Changes to a Server or Group

After you have edited Application Configuration values in the Value Set Editor, you must apply them to the application on the server. To do so, you need to perform a push operation. Performing a push operation applies modifications to the actual configuration files on the server (or group).



The way in which sequences (of lists and scalars) are merged when you push depends upon how values have been set in the Application Configuration inheritance hierarchy and what sequence merge modes have been configured in the CML template for the

Application Configuration. For more information about sequence merging, see “Sequence Merging and Inheritance” on page 412.



If your push times out before the push succeeds (default is ten minutes), it could be that the default timeout value set in the Application Configuration is less than the time it takes to push the Application Configuration. See your Opware administrator for help in extending the duration allowed for an Application Configuration push to occur.

To push Application Configuration changes to a server or group, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select a server or group of servers.
- 3** Select a server (or group) from the Content pane, right-click and choose **Configure Applications**.
- 4** You now see the Server Explorer (or Device Groups Browser), with the Configured Applications folder selected. Select the Installed Configurations tab. All Application Configurations that have been attached to the server (or group) will be displayed.
- 5** From the Views pane of the Server Explorer (or Device Groups Browser), select an Application Configuration instance to edit.
- 6** If you wish, make edits to the Application Configuration. (See “Setting Application Configuration Values on a Server or Group” on page 431 in this chapter for more information.)
- 7** To preview the changes and see how they differ from the configuration file on the server, click **Preview**. The Comparison dialog box opens and shows any differences. Click **Close** when you are finished.
- 8** When are ready to apply the changes to the server, click **Push**.

Scheduling an Application Configuration Push

You can schedule an Application Configuration push to run a single time, or on a recurring schedule, such as daily, weekly, or monthly.

To schedule an Application Configuration push, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select a server or group of servers.
- 3** Select a server or group in the Content pane, right-click, and choose **Configure Applications**.
- 4** You now see the Server Explorer (or Device Groups Browser), with the Installed Configurations tab selected. All Application Configurations that have been attached to the server (or group) will be displayed.
- 5** From the View pane, click the plus (+) symbol to expand Application Configuration folder and select an Application Configuration instance.
- 6** Click **Schedule**.
- 7** In the Schedule Job dialog box, set the following parameters:
 - **Schedule:** Choose to Run Once, Daily, Weekly, Monthly, or Custom. By default, the Schedule is set to Weekly.
 - **Crontab String:** (This field appears only if you chose a custom schedule. If you did not choose Custom, then skip to the Start Time field below.) Enter a crontab string for date in this order:
 - Minute (0-59), Hour (0-23)
 - Day of the month (1-31)
 - Month of the year (1-12)
 - Day of the week (0-6 with 0=Sunday)

Any of these fields can contain an asterisk * standing for all possible values. For example, the following crontab string runs the job at midnight every weekday:

```
0 0 * * 1-5
```

The crontab string can also handle serial (1,2,3,4) as well as range (1-5) values. For more information about using crontab strings, consult the crontab man pages on a Unix computer.

- **Start Time:** Select a time for the job to begin running. For one-time jobs, enter the full date and time. For weekly and monthly jobs, enter the time of day. You can enter the values by typing directly into the field using up or down arrows.
 - **Time Zone:** Select a default time zone for the job execution time, or accept the default time zone. The default time shown converts the scheduled time to the time zone set in your user preferences. If you do not set a preferred time zone, the time zone is derived from the Opsware SAS core server (typically UTC).
 - **Day** (Monthly only): Choose the day of the month to run this job.
 - **Days To Run** (Weekly only): Choose the days of the week you want the job to run at the specified time.
 - **Months to Run** (Monthly only): Choose the months during which you want the job to run.
- 8** In the Run Jobs between these Dates section, select a date range during which you would like the job to run.
- **Start:** Choose a start date for the date range.
 - **End:** Choose an end date for the date range.
 - **No End Date:** Choose if you want the job to run indefinitely.
- 9** In the Job Run Notification Email section, enter an email address to receive the results of the job. You can enter multiple email addresses separated by commas or spaces.
- **On Success:** Enter email addresses that will receive notifications of jobs that complete successfully.
 - **On Failure:** Email addresses that will receive notifications of jobs that failed to complete.
- 10** In the Ticket Tracking section, enter a ticket ID from your own job trackins system here.
- 11** When you have finished setting the parameters, click **OK**.

Comparing Two Configuration Templates

To show the difference between two Application Configuration Templates, you can perform a compare operation between them.

To compare two Application Configuration Templates, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 2** Select Application Configuration, and then select the Configuration Templates tab.
- 3** Hold down the CTRL key and select two Application Configuration Templates, right-click, and choose **Compare**.
- 4** The Comparison dialog box displays the difference between the two files. Use the arrows in the upper right of the dialog box to navigate through the two files. To indicate the differences, the Comparison feature uses the following colors:
 - **Green**: This indicates that new information has been added.
 - **Blue**: This indicates that information has been modified.
 - **Red**: This indicates that information has been deleted.
 - **Black**: This indicates no changes.
- 5** When you are finished viewing the differences, click **Close**.

Comparing a Template Against an Actual Configuration File

To show the difference between an Application Configuration Template and the actual file on the server (or group), perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select a server or group of servers.
- 3** In the Content pane, select a server or group, right-click, and choose **Configure Applications**.
- 4** The Installed Configurations tab will be selected. From the Views pane of the Server Explorer (or Device Groups Browser), select an Application Configuration instance.
- 5** If the Application Configuration contains more than one Application Configuration Template, then from the Template drop-down list in the Content pane, choose a Application Configuration Template to compare.

- 6** To preview the differences between the Application Configuration Template and the actual configuration file on the server, click **Preview**. The Comparison dialog box shows the differences between the Application Configuration Template and the actual configuration file. Use the arrow keys in the upper right of the dialog box to navigate through the two files. To illustrate the differences, the Comparison feature uses the following color scheme:
 - **Green**: This indicates that new information has been added.
 - **Blue**: This indicates that information has been modified.
 - **Red**: This indicates that information has been deleted.
 - **Black**: This indicates no changes.
- 7** When you are finished viewing the differences, click **Close** to close the Comparison dialog box.

Scanning Configuration Compliance

After an Application Configuration has been pushed to a server, it is possible that the configuration file on the server becomes changed or altered, either intentionally or by accident. You can scan for configuration compliance on a server to determine if any of the configuration files on the server are out of sync with the values stored in your templates.

If an Application Configuration on a server is out of sync, the server that the Application Configuration is attached to will show the following icon in the server list inside the SAS Client:




For example, open the SAS Client and select the Servers feature icon. A list of all managed servers in your environment is displayed and you can see if any servers show the out of sync icon.

If a server shows this icon, scan for configuration compliance to find out which configuration files on the server are out of sync with the Application Configuration.

For information on how to schedule a configuration compliance scan, see “Scheduling a Configuration Compliance Scan” on page 440.

To scan for configuration compliance, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select a server or group of servers from the Navigation pane, and then select a server that shows the out of sync icon  from the Content pane. From the **Actions** menu, select **Scan** ► **Configuration Compliance**. (You can also multiple-select and scan more than one out of sync server.)
- 3** You will be asked if you are sure you want to scan the Application Configuration on the selected managed server. Click **Yes** to run the scan.
- 4** The Job dialog box appears, showing the details of the scan. Make sure to deselect the Close when finished option at the bottom of the dialog box so the Job dialog box remains open after the scan job has run. Once the job has finished, look in the Completed Status section, and select the Success text. You see a list of servers in the Servers section to the right.
- 5** To view the configuration compliance scan details for a server, in the Servers section, select a server. Below in the Server Detail section, a list of all discrepancies shows which files are out of sync with Application Configuration Templates on the server. To view the Application Configuration, click **Configurations**. The Server Browser appears.
- 6** To troubleshoot the discrepancies, select the out of sync Application Configuration and its templates and click **Preview**. This will show you where the configuration file on the server differs from the values defined in the Application Configuration. Once you have found the discrepancies, you can modify them as needed in the Value Set Editor, and then push the changes to the server. See “Comparing a Template Against an Actual Configuration File” on page 438 in this chapter for more information

Scheduling a Configuration Compliance Scan

You can schedule an configuration compliance scan to run a single time, or on a recurring schedule, such as daily, weekly, or monthly.

To schedule a configuration compliance scan, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select a server or group of servers from the Navigation pane, and then select a server from the Content pane.

- 3** From the **Actions** menu, select **Schedule Configuration Compliance Scan**.
- 4** In the Schedule Job dialog box, set the following parameters:
 - **Schedule:** Choose to Run Once, Daily, Weekly, Monthly, or Custom. By default, the Schedule is set to Weekly.
 - **Crontab String:** (This field appears only if you chose a custom schedule. If you did not choose Custom, then skip to the Start Time field below.) Enter a crontab string for date in this order:
 - Minute (0-59), Hour (0-23)
 - Day of the month (1-31)
 - Month of the year (1-12)
 - Day of the week (0-6 with 0=Sunday)

Any of these fields can contain an asterisk * standing for all possible values. For example, the following crontab string runs the job at midnight every weekday:

```
0 0 * * 1-5
```

The crontab string can also handle serial (1,2,3,4) as well as range (1-5) values. For more information about using crontab strings, consult the crontab man pages on a Unix computer.

 - **Start Time:** Select a time for the job to begin running. For one-time jobs, enter the full date and time. For weekly and monthly jobs, enter the time of day. You can enter the values by typing directly into the field using up or down arrows.
 - **Time Zone:** Select a default time zone for the job execution time, or accept the default time zone. The default time shown converts the scheduled time to the time zone set in your user preferences. If you do not set a preferred time zone, the time zone is derived from the Opware SAS core server (typically UTC).
 - **Day** (Monthly only): Choose the day of the month to run this job.
 - **Days To Run** (Weekly only): Choose the days of the week you want the job to run at the specified time.
 - **Months to Run** (Monthly only): Choose the months during which you want the job to run.
- 5** In the Run Jobs between these Dates section, select a date range during which you would like the job to run.

- **Start:** Choose a start date for the date range.
 - **End:** Choose an end date for the date range.
 - **No End Date:** Choose if you want the job to run indefinitely.
- 6** In the Job Run Notification Email section, enter an email address to receive the results of the job. You can enter multiple email addresses separated by commas or spaces.
 - **On Success:** Enter email addresses that will receive notifications of jobs that complete successfully.
 - **On Failure:** Email addresses that will receive notifications of jobs that failed to complete.
 - 7** In the Ticket Tracking section, enter a ticket ID from your own job tracking system here.
 - 8** When you have finished setting the parameters, click **OK**.

Rolling Back to a Previous State

Every time you push an Application Configuration to a server, that push is saved in a configuration backup list. At any time, you can revert or rollback to a previous state of an Application Configuration in this list. This enables you to go back to a known configuration state for a specific Application Configuration.

To roll back an Application Configuration to a previous state, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select a server or group of servers.
- 3** Select a server or group of servers from the Content pane, right-click and choose **Configure Applications**.
- 4** You now see the Server Explorer (or Device Groups Browser), with the Installed Configurations tab selected.
- 5** Select the Configuration History tab. A list of all Application Configuration pushes will display. You can sort this list by application name, configuration backup name, date created (when the Application Configuration was pushed), and by user.



If the list is empty, the Application Configuration has never been pushed to the server.

- 6** To roll back to a saved configuration, select a item in the list, and click **Revert**. This restores all configuration files to the state immediately after this backup was made. The original configuration files are also restored and suffixed with “_opsware_backup”.

Chapter 9: Operating System Provisioning

IN THIS CHAPTER

This section discusses the following topics:

- Supported Operating Systems for OS Provisioning
- OS Provisioning
- Hardware Preparation
- New Server Booting
- OS Installation with the SAS Client

Using server build policies, Opware SAS OS Provisioning allows you to provision operating systems onto bare-metal servers quickly and consistently. It also helps ensure that each operating system installed has the right build information.

OS Provisioning supports a large variety of hardware models from different manufacturers out of the box, and it can be configured to support additional hardware models. See the *Opware® SAS Policy Setter's Guide* for more information.



Before you can install operating systems on servers with the OS Provisioning feature, the operating systems must be defined and the OS media must be made available in Opware SAS. Additionally, OS installation profiles can be created in the Opware SAS Web Client and in the Opware Server Automation System Client. Please refer to the *Opware® SAS Policy Setter's Guide*.

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 Provisioning Overview

In Opware SAS, OS Provisioning is installation-based instead of image-based. The OS Provisioning feature uses Red Hat Linux Kickstart (used for Linux and ESX), Sun Solaris JumpStart, and Microsoft Windows unattended installation to install operating systems on servers.

The OS Provisioning feature is fully integrated with Opware SAS; you can install an OS on the following types of servers:

- A bare metal server that does not have an OS installed
- A server that Opware SAS already manages
- A server that is running in the environment but Opware SAS does not manage it

The OS Provisioning feature helps you install operating systems on servers in the following ways:

- Each OS installation profile in the OS Provisioning feature contains all the information necessary to build and maintain a server with that OS.
- When installing an OS on a server, the OS Provisioning feature displays information about server hardware and the operating systems that are compatible with that hardware architecture.



You need a specific set of feature permissions for OS Provisioning. You'll also need permissions to access the servers associated with customers, facilities, or groups of servers. To obtain these permissions, contact your Opware administrator. For more information, see the Permissions Reference appendix in the *Opware® SAS Administration Guide*.

Server Lifecycle for OS Provisioning

Opware SAS enables multiple teams to work together and provision servers. The OS Provisioning feature allows IT teams to separate the tasks of readying servers for provisioning (such as racking servers or connecting them to power and a network) from provisioning the servers with operating systems.

Someone mounts a new server in a rack and connects it to the Opware build network. Then 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 in the Opware SAS Web Client or in the Unprovisioned server list in the SAS Client – and provision it with an OS. In the available state, servers do not have an OS installed and might not have access to disk resources.

During OS provisioning, servers progress through Planned, Unprovisioned, Available, Installing OS, and Managed states. See Table 9-1 for more information on these lifecycle values.

Table 9-1: Opware SAS Lifecycle Values for Servers

OPSWARE LIFECYCLE VALUE	DESCRIPTION
Server Pool Values	
Planned	<p>Indicates that a device record has been created for the server, but an Opware OS Build Agent has not yet been installed. (The OS Build Agent is a small agent that can run in the memory of the bare metal server.)</p> <p>Servers in this stage cannot be provisioned until an OS Build agent is installed.</p> <p>For more information on this unprovisioned lifecycle value, see your Opware Administrator.</p>
Available	<p>Indicates a server on which the OS Build Agent was installed and is running, but that does not have an OS installed.</p> <p>See “Installing OS Build Agents” on page 460 in this chapter for more information.</p>
Installing OS	<p>Indicates that a user is installing an OS on the server. The server stays in the Server Pool list until the installation process finishes successfully, then, the server moves to the Managed Server list.</p>
Build Failed	<p>Indicates a server on which the OS Build Agent was installed and is running, but the installation of an OS failed. The server will remain in the Server Pool list with this status for seven days before Opware SAS deletes the entry.</p> <p>See “Recovering when an OS Build Agent Fails to Install” on page 462 in this chapter for more information.</p>
Managed Server Values	

Table 9-1: Opware SAS Lifecycle Values for Servers (continued)

OPSWARE LIFECYCLE VALUE	DESCRIPTION
Managed	<p>Indicates a server that Opware SAS is managing. Opware SAS performs reachability checks for managed servers.</p> <p>After a server reaches this lifecycle state, the entry for the server moves from the Server Pool list to the Managed Servers list.</p>
Deactivated	<p>Indicates a server that was previously managed by Opware SAS, but is no longer managed by Opware SAS. However, the server's history still exists in Opware SAS. Deactivated servers are not reachable.</p>

OS Provisioning

This section provides information about the OS provisioning process within Opware SAS and contains the following topics:

- Overview of OS Provisioning
- Solaris OS Provisioning
- Linux or VMware ESX OS Provisioning
- Windows OS Provisioning

Overview of OS Provisioning

The process for provisioning new servers on supported operating systems includes the following steps:

- 1** A system administrator unpacks a server, mounts it in a rack, and attaches the server to power and a network that can reach Opware SAS.
- 2** The system administrator prepares the hardware for OS provisioning.
See "Hardware Preparation" on page 453 in this chapter for more information.
- 3** If necessary, the system administrator inserts a bootable floppy or CD provided with Opware SAS. (Using a bootable floppy or CD is not necessary for Intel-based servers that support PXE or Unix servers that support DHCP. This is because these types of servers are capable of booting over a network.)

See “New Server Booting” on page 454 in this chapter for more information.

- 4** The system administrator turns the server on.

For servers capable of booting over the network, powering the server on causes the server to initiate its network boot process. For example, the server sends a boot request to a PXE server.

The Opsware OS Build Manager responds to this network boot request by delivering the Opsware OS Build Agent, a small agent that can run in the memory of the bare metal server. (For servers not capable of booting over the network, the Opsware OS Build Agent is on the bootable floppy or CD.)

The Opsware OS Build Agent constructs an inventory of the server (including server manufacturer, server model, MAC address, available memory, and available storage) and delivers that information to the Opsware OS Build Manager.

- 5** In the SAS Web Client, the system administrator sees this server and its hardware inventory in a list of available servers ready to be provisioned.

See “Verifying Installation of an OS Build Agent” on page 461 in this chapter for more information.

- 6** The system administrator selects the OS or a complete server baseline (which can include a base OS, a set of OS patches, system utilities, and middleware software) to provision.

The system administrator installs the OS or a complete server baseline on the server at that time or schedules the installation for some time in the future.

The OS Provisioning feature installs the selected software onto the server.

- 7** The system administrator uses Opsware SAS to configure networking for the newly provisioned server.

See *Opsware® SAS User's Guide: Server Automation* for more information.

Additionally, the system administrator might choose to configure servers running Red Hat Linux or Sun Solaris operating systems by using the OS Provisioning feature.

See “Reprovisioning a Managed Server” on page 472 in this chapter for more information.

Solaris OS Provisioning

The OS Provisioning feature includes a DHCP-based JumpStart configuration that hides the complexity of JumpStart from the end user. Unlike typical JumpStart systems, the OS Provisioning feature does not require configuration updates to the JumpStart server for each installation that you provision.

Instead, an OS installation profile exists in the OS Provisioning feature for each version of the Solaris OS that you want to install on servers in your environment.

The process for Solaris OS provisioning follows the general OS provisioning process that the OS Provisioning feature established.

See the *Opware® SAS Policy Setter's Guide* for more information on the Solaris build process.

Linux or VMware ESX OS Provisioning

OS Provisioning for Linux or VMware ESX includes a Kickstart or YaST2 system that hides the complexity of Kickstart or YAST2 from the end user.

Mapping a specific installation client to a particular Kickstart or YaST2 configuration is a simple procedure in the OS Provisioning feature. The OS Provisioning feature allows you to easily choose a particular Kickstart or YaST2 configuration through the SAS Web Client at installation time.

The process for Linux or VMware ESX OS provisioning follows the general OS provisioning process that the OS Provisioning feature established.

See the *Opware® SAS Policy Setter's Guide* for more information on the Linux build process.

Windows OS Provisioning

In the OS Provisioning feature, system administrators can perform unattended, scripted installations as well as WinPE-based image installations of Windows NT, Windows 2000, Windows 2003, and Windows XP Professional on bare metal servers.

The installation-based approach allows system administrators to adapt to variations in hardware. The OS Provisioning feature can be set up to install Windows operating systems on known hardware in the managed environment. At build time, the OS Provisioning feature provisions the server with the correct hardware-specific software and drivers based on the hardware signature of the server about to be provisioned.

See the *Opware[®] SAS Policy Setter's Guide* for more information on the Windows build process.

Hardware Preparation

Before you use OS Provisioning to install an OS on a server, the server must meet certain requirements, or the hardware must be prepared in certain ways.

Windows Hardware Preparation Requirements

Before you install Windows on a server, you need to prepare the hardware by performing the following tasks:

- If the hardware has a RAID controller, you may have to extend the Windows OS media distribution based on vendor-specific requirements. The Microsoft Windows OS media might not (depending on the version of Windows) include the necessary drivers for many RAID controllers. Also, certain newer types of SATA controllers may require additional drivers.
- When using the DOS-based PXE or floppy boot images to install the Windows OS, you need to create a FAT16 or FAT32 partition on the primary boot (hard) drive to install the Windows OS on. The boot images contain the necessary functionality to create the required partition.
- If you use the WinPE-based PXE or CD-ROM boot images to install the Windows OS, disk partitioning is performed during the OS installation. You can control the disk partitioning by editing the OS installation profile in Opware. (For more information, see the *Opware[®] SAS Policy Setter's Guide*.)
- If you use the WinPE-based PXE or CD-ROM boot images provided by Opware, and you are using a RAID or SATA controller, you might need to supply a build customization. This will enable you to load the necessary drivers before the OS installation can commence. It is unnecessary to do the same in a DOS-based installation environment, since all disk access is done via the BIOS until the Windows installation is in progress. For more information, see the *Opware[®] SAS Policy Setter's Guide*.

Sun Solaris Hardware Preparation Requirements

To install Solaris on a server, the hardware must meet the following requirements:

- Have a DHCP-capable PROM (older servers can be upgraded to DHCP-capable PROM).
- Be part of the sun4u system architecture (platform group).

You do not need to perform any Opware SAS-specific preparation of the hardware before you install Solaris on a server.

Linux and VMware ESX Hardware Preparation Requirements

Before you install Linux on a server, you need to prepare the hardware by configuring valid, logical drives for RAID.

VMware ESX hardware requirements are the same as Linux.

New Server Booting

This section provides information on booting new servers with Opware SAS and contains the following topics:

- Booting New Servers with Different Operating Systems
- OS Build Agent
- Booting a Windows (DOS), Linux, or VMware ESX Server with PXE
- Booting a Windows Server with PXE Using WinPE
- Booting a Solaris Server Over the Network
- Installing OS Build Agents
- Verifying Installation of an OS Build Agent
- Recovering when an OS Build Agent Fails to Install

Bootting New Servers with Different Operating Systems

On Intel-based servers, you can boot a new server over a network in a hands-off fashion by using PXE. For environments with servers that do not support network boot technology, Opware SAS supports floppy or CD booting.

For Windows and Linux or VMware ESX servers, the Opware Boot Floppy and CD respectively contain a small operating system, network drivers, the software required to mount a network drive, and the Opware OS Build Agent. The Opware Boot Floppy or CD has the software that is otherwise delivered over the network as part of the network boot process.

For Solaris servers, you can provision an OS over the network by using DHCP, but you cannot boot new Solaris servers using a floppy or CD.



To boot servers over the network, the installation client must be able to reach the Opsware DHCP server on the Opsware core network. If the installation client is running on a different network than the Opsware core network, your environment must have a DHCP proxy (IP helper). Alternatively, for Linux and Windows installation clients, you can boot the servers by using an Opsware Boot CD or Floppy instead of booting the servers over the network.

OS Build Agent

The OS Provisioning feature de-couples the task of preparing a server for provisioning from the task of provisioning the server with an OS. This de-coupling of tasks is made possible by the OS Build Agent.

Booting a new server for the first time installs an OS Build Agent on the server; however, the server does not have the target OS installed and might not have access to disk resources. Opsware SAS can still communicate with the server and perform commands on it remotely, because the OS Build Agent is running an OS installed in memory.

The OS Build Agent performs the following functions:

- Registers the server with Opsware SAS when the OS Build Agent starts.
- Listens for command requests from Opsware SAS and performs them.

The OS Build Agent can perform commands even though the target OS is not installed.

Booting a Windows (DOS), Linux, or VMware ESX Server with PXE

The following instructions show you how to boot a Windows (DOS), Linux, or VMware ESX Server with PXE.

For information on how to boot a server with WinPE, see “Booting a Windows Server with PXE Using WinPE” on page 458.

When booting a Linux, VMware ESX, or Windows server by using PXE, the DHCP relay must be running on the router of the build network for PXE to function properly.

Alternatively, if the build script is plugged directly into the boot server providing DHCP service, a DHCP relay is not necessary.

- 1** After you mount the new server in a rack and connect it to the Opsware build network, set up the server to boot by using PXE.

See the hardware vendor's documentation on how to prepare a server to boot by using PXE.

- 2** Power on the server and select the option to boot the server with PXE.
- 3** Choose an Opsware boot image by entering the appropriate text at the boot prompt.

```
windows    - Windows Build Agent (DOS 7.01)
undi       - Windows Build Agent (DOS 7.01 + UNDI)
winpe      - Windows Build Agent (WINPE based)
win-bcom   - Windows Build Agent (DOS 7.01 w/Broadcom driver
v9.07)
linux      - Linux Build Agent (RHEL 3.0-based)
linux4     - Linux Build Agent (RHEL 4.0-based)
solaris    - Solaris x86 Build Agent
localdisk  - Normal boot from localdisk (default after 10
second
```



If you are booting an VMware ESX server, select one of the Linux options. If you are booting an VMware ESX server, select one of the Linux options. If you are provisioning a machine into NT 4, at the boot prompt (DOS menu), enter "windows-old" to continue. This option is hidden on the onscreen menu, but is still accessible by typing it in yourself.

If you are booting a Windows server, the type of hardware being provisioned determines the version of the DOS Windows OS Build Agent. The images for the DOS Windows OS Build Agents vary in terms of the memory management software, disk partitioning capabilities, and network drivers.

If an incompatible boot image is selected for the hardware, an error message might appear at the console during the provisioning process; for example, it might appear when the Windows OS Build Agent is booting and DOS is loading or it might appear later in the process when the Windows Installer is running.

(For more information on WinPE booting, see "Bootting a Windows Server with PXE Using WinPE" on page 458).

See Table 9-2 for the differences between images for the Windows OS Build Agents.

Table 9-2: Differences Between Images for the Windows OS Build Agents

BOOT IMAGE	NETWORK DRIVERS	PREINSTALLATION ENVIRONMENT	DISK PARTITIONING CAPABILITIES
windows	Native DOS	DOS 7.0.1	FAT16 or FAT32
undi	UNDI	DOS 7.0.1	FAT16 or FAT32
winPE	Windows XP/2003/ Vista drivers	WinPE 2.0	NTFS or FAT32
winbcom	Native DOS with Broadcom v9.07	DOS 7.0.1	FAT16 or FAT32

For Windows, if you select a DOS option to boot the server, an additional set of Opware SAS menus appear on the console. This enables you to partition the hardware disk. (If you boot using WinPE, you can specify disk partitioning later, in the OS Installation Profile properties, which will create the partitions when the OS is installed.)

If you do not select an option within 10 seconds, the server defaults to booting from the local disk. To stop the default selection, type something (other than ENTER) at the command line.

- 4** For Windows servers only, select the menu choices to partition the disk based on your specifications.

After the booting process finishes successfully, a message appears on the console that indicates that the server is ready for OS provisioning. An OS Build Agent was installed on the server and the server appears in the Server Pool list in the SAS Web Client.

- 5** (Optional) Record the MAC address and/or the serial number of the server so that you can locate the server in the Server Pool list in the SAS Web Client or in the Unprovisioned Servers list in the SAS Client.

You should verify that the newly racked server shows up in the SAS Web Client and is ready to hand off for OS installation.

See “Verifying Installation of an OS Build Agent” on page 461 in this chapter for more information.

Booting a Windows Server with PXE Using WinPE

Opware OS Provisioning now supports booting a bare metal server with PXE into a WinPE preinstallation environment. You can choose between either a WinPE x86 32 bit environment or a WinPE x64 64 bit environment.

WinPE provides greater flexibility than DOS, because it does not require that you format your hard drive during the boot process. You can define the disk partition configuration later in an OS installation profile, when you create or edit the OS installation profile.

For more information on how to create and edit an OS installation profile, see “Creating an OS Installation Profile” on page 463.

WinPE also allows WIM-based image installation, as an alternative to unattended Windows installations.



When booting a Windows server by using PXE, the DHCP relay must be running on the router of the build network for PXE to function properly. Alternatively, if the build script is plugged directly into the boot server providing DHCP service, a DHCP relay is not necessary.

To boot a bare metal server with PXE into a WinPE preinstallation environment, perform the following steps:

- 1** Mount the new server in a rack and connect it to the Opware build network.
- 2** Set up the server to boot by using PXE.

See the hardware vendor's documentation on how to prepare a server to boot using PXE.

- 3** Power on the server and select the option to boot the server with PXE.

The Opware SAS menu appears and prompts you to select the type of OS Build Agent to install on the server.

- 4** You will see all the Opware boot image options on the screen.

```
Windows - Windows Build Agent (DOS 7.01)
undi    - Windows Build Agent (DOS 7.01 + UNDI)
winpe   - Windows Build Agent (WINPE based)
win-bcom - Windows Build Agent (DOS 7.01 w/Broadcom driver
v9.07)
linux   - Linux Build Agent (RHEL 3.0-based)
```

```
linux4      - Linux Build Agent (RHEL 4.0-based)
solaris     - Solaris x86 Build Agent
localdisk  - Normal boot from localdisk (default after 10
second
```

- 5** At the boot prompt enter: winpe.



If you do not select an option after 10 seconds, the server defaults to booting from the local disk. To prevent the default selection, type in the command line.

- 6** A new menu displays the option to boot a WinPE x86 32 bit environment or a Windows x64 64 bit environment. Make a selection by using the arrow keys to highlight your choice, and then press ENTER.

The server will now be booted with the WinPE preinstallation environment. This may take a few minutes to complete, depending upon the speed of the network and the machine.

Once the booting has finished, a new window will appear indicating that the server has had an Opware Build Agent installed and registered with the Opware core.

- 7** (Optional) Record the MAC address and/or serial number of the server so that you can locate the server in the Server Pool list in the SAS Web Client or in the Unprovisioned Servers list in the SAS Client.
- 8** Verify that the newly racked server shows up in the SAS Web Client and is ready for OS installation. See “Verifying Installation of an OS Build Agent” on page 461 in this chapter for more information.

Booting a Solaris Server Over the Network

When Opware SAS was installed in your facility, the OS Provisioning feature was set up so that the Opware Boot Server listens for broadcast requests from new servers and it responds by using DHCP.

Perform the following steps to boot a Solaris server over the network:

- 1** Mount the new Solaris server in a rack and connect it to the network.

The installation client on this network must be able to reach the Opware DHCP server on the Opware core network. If the installation client is running on a different network than the Opware core network, your environment must have a DHCP proxy (IP helper).

- 2** Enter one of the following commands at the prompt:

```
ok boot net:dhcp - install
```

Or

```
ok boot net:dhcp - install <interface_setting>  
<buildmgr=hostname|IP_address>
```

Where *<interface_setting>* is one of the following options:

```
autoneg, 100fdx, 100hdx, 10fdx, 10hdx
```

You can include an interface setting with the boot command to set the network interface to a specific speed and duplex during OS provisioning. When Opware SAS was installed in the local facility, a default value was provided for this interface setting. Specifying this boot argument allows you to override the default interface setting.

To continue setting the network interface with a specific speed and duplex, you can use a variety of methods, including using a Solaris build customization script or specifying the values in a Solaris Package or RPM in the OS media.

See the *Opware® SAS Policy Setter's Guide* for more information.

Ways that the OS Build Agent Locates the Opware Build Manager

For Solaris OS provisioning, the JumpStart build script runs the OS Build Agent, which contacts the Opware Build Manager (via the Agent Gateway in the core). The Solaris `begin` script attempts to locate the Opware Build Manager in the following ways:

- By using information that the Opware DHCP server provided
- By looking for the host name `buildmgr` in DNS as configured by the DHCP server

You can override the way that the OS Build Agent contacts the Opware Build Manager by specifying a boot argument at the prompt when you boot a new Solaris server:

```
ok boot net:dhcp - install [buildmgr=hostname|IP_address]:port
```

Installing OS Build Agents

You can install an OS Build Agent on a server by booting the server with PXE or an Opware Boot Image (Windows, Linux, or VMware ESX), or by using the network (Solaris). After a successful installation, the server appears in the Server Pool list.

You should verify that the newly racked server shows up in the SAS Web Client and is ready to hand off for OS installation.

The Server Pool list displays the servers that have registered their presence with Opware SAS, but do not have the target OS installed on them. From here, you can install an OS by selecting the server and clicking **Install OS**.

Verifying Installation of an OS Build Agent

Perform the following steps to verify the installation of an OS build agent:

- 1** Log into the SAS Web Client.
- 2** From the Navigation pane, select Servers ► Server Pool. The Server Pool page appears, as Figure 9-1 shows.

Figure 9-1: Server Pool List in the SAS Web Client

The following servers have registered their presence with Opware but do not have a full operating system installed.

All Manufacturers		All Models		Update							
Delete...		Install OS...									1 Total
	Name	MAC Address	Manufacturer	Model	Reported OS	Registered	Lifecycle	Facility	Customer		
<input type="checkbox"/>	m101.tr3.opsware.com	00:11:43:CE:19:4A	DELL COMPUTER CORPORATION	POWEREDGE 750	DOS	05-25-2005	Available	TR3	Not Assigned		

- 3** (Optional) From the drop-down lists, select the manufacturer, model, or facility of the server and click **Update**.
- 4** For Intel x86 and Sun SPARC processor-based servers, locate the MAC address and Host ID of the server that you just booted.

The Lifecycle column indicates the progress or success of the OS Build Agent installation. If the OS Build Agent was successfully installed, the Lifecycle column indicates that the server is available for OS provisioning.

See “Server Lifecycle for OS Provisioning” on page 448 in this chapter for more information.

To obtain information on a server in the SAS Web client, click on the server name. If you are viewing an unprovisioned server in the Unprovisioned Servers list in the SAS Client, double-click on the server to open the Device Explorer. This will display detailed information.

Recovering when an OS Build Agent Fails to Install

When an OS Build Agent fails to install on a server, the server does not appear in the Server Pool list.

You can check the server console for error messages and try to boot the server again with PXE or by using the Opware Boot Floppy or CD.

If all errors were successfully resolved, the initial boot occurs, the OS Build Agent is installed on the server, the server appears in the Server Pool list, and the Lifecycle column indicates that the server is available.

If you are unable to resolve the error condition and install the OS Build Agent on the server so that it appears in the Server Pool list, contact your Opware administrator for troubleshooting assistance.

OS Installation with the SAS Client

This section describes how to install an operating system on an unprovisioned server using the SAS Client. To install an operating system on an unprovisioned server, the operating system must already have its OS media prepared and made available, it must have an OS build customization script created for it, and it must have an OS installation profile created for it. Once you provision a server and install an OS using an OS sequence, it becomes an Opware managed server.



For information on how to set up OS provisioning, see “Operating System Provisioning Setup” on page 131 in the *Opware® SAS Policy Setter's Guide*.

In order to install an OS and provision a server using the SAS Client, you need to create, define, and run an OS sequence. An OS sequence defines what to install on an unprovisioned server, including OS build information from the OS installation profile, selected software and patch policies, and remediation settings. An OS sequence represents a server build policy, and it defines how a server should be provisioned, affecting its software and operating systems. When the OS sequence is defined, it can be used to provision additional servers with the same OS and software.

To install an OS on an unprovisioned server using the SAS Client, you need to perform the following tasks:

- **Creating an OS Installation Profile:** Define the OS, configuration or response file, build customization scripts, OS media, customer association, and packages.
- **Creating an OS Sequence:** Choose the OS installation profile, software policies, patch policies, and remediation policies.
- **Selecting Servers in the Unprovisioned Servers List:** Choose the server or servers which you would like to install the OS and provision.
- **Running An OS Sequence:** Launch the OS sequence to provision the selected unprovisioned server and run the job.

Creating an OS Installation Profile

An OS installation profile defines all necessary parameters of an OS, including the OS type and version, the OS Media Resource Locator (MRL), the configuration or response file, the build customization script, and the packages related to the OS installation.

An OS installation profile is created by a policy setter, and so is beyond the scope of this book. For information on how to create an OS installation profile, see “Operating System Provisioning Setup” on page 131 in the *Opware® SAS Policy Setter's Guide*.

Creating an OS Sequence

An OS sequence defines what to install on an unprovisioned server, including OS build information from the installation profile, selected application and patch policies, and the target servers you want to install the OS on to.



When you create an OS sequence, it will be saved into the Folder list in the Library. You must have permissions to the folder where you want to save the OS sequence. For more information on how folder permissions work, see User and Group Setup in the *Opware® SAS Administration Guide*.

Elements of an OS Sequence

An OS sequence consists of the following components that must be configured before you run the OS sequence:

- **Properties:** Allows you to name the OS sequence and choose a location to save it in a library folder. You must have permissions to write to the folder where you save the OS sequence, otherwise you will be unable to save it in the selected location in the library.

- **Install OS:** Allows you to choose an OS installation profile. If the OS installation profile already has a customer associated with it, you will be unable to select a customer for the OS sequence. If the OS installation profile does not have a customer associated with it, then you can select one here. Once you choose a customer, then all servers on which you install the OS using this OS sequence will be associated with that customer.
- **Attach Software Policies:** Allows you to add a software policy to the OS sequence. When the OS sequence is run, if the remediate is enabled (in Remediate Policies), then all the software in the software policy will be installed on the server during OS installation. If the remediation option is disabled, then none of the software will be installed on the server.

The software policies that you can attach to an OS sequence are restricted by the OS type. In other words, you can only attach software policies when their OS matches the OS installation profile chosen for the OS sequence.

For more information on software policies, see Chapter 7, "Software Management".

- **Attach Patch Policies:** Allows you to select a patch policy to attach to the OS sequence. When run OS sequence is run, if the remediate option is enabled (in Remediate Policies), then all the patches in the patch policy will be installed on the server. If the remediate option is disabled, then none of the patches will be installed on the server.

The patch policies that you can attach to an OS sequence are restricted by the OS type and the patch policy. In other words, you can only attach patch policies when their OS matches the OS installation profile chosen for the OS sequence.

For more information on patches, see Chapter 6, "Patch Management for Unix" or Chapter 5, "Patch Management for Windows".

- **Attach Device Group:** Allows you to select a device group (group of servers) for a the server once the OS sequence has been run. You can select any public static group to attach to the OS sequence. Also, a group of servers can have software and patch policies associated with it. If you enable remediation in the OS sequence (in Remediate Policies), then all software and patches associated with the group of servers will also be installed on the server when you run the OS sequence. If you disable remediation, then none of the software or patches in the policies attached to the group of servers will be installed on the server.

For information on groups of servers, see Server Management in the *Opware® SAS User's Guide: Server Automation*.

- **Remediate Policies:** Allows you to choose to enable or disable remediation when the server is provisioned with the OS sequence. Default is enabled. If you disable remediation, then when you run the OS sequence, the OS will be installed but no policies in the OS sequence will be remediated – this means that no software or patches in any of the policies attached to the OS sequence will be installed when the sequence is run.

If you enable remediation, then all software and patches in all policies attached to the server will be installed when the OS sequence is run. This is also true for any policies attached to the group of servers selected for the OS sequence. You can also set reboot and pre and post installation script options.

Creating an OS Sequence

To create an OS Sequence, perform the following steps:

- 1** In the SAS Client, from the Navigation pane, select Library and then select OS Sequences.
- 2** Choose an OS folder.
- 3** From the **Actions** menu, select **Create New**.
- 4** In the Views pane of the OS Sequence window, select Properties and enter a name for the OS sequence.
- 5** Click **Change** in the Content pane to choose a location in the folder library to save the OS sequence. You must have permissions to write to the folder where you save the OS sequence.
- 6** Next, from the Views pane, click **Select Install** to choose an OS installation definition.
- 7** If the OS installation profile does not have a customer associated with it, then select a customer from the Assign Customer drop-down list. If the OS installation profile already has a customer associated with it, you will be unable to select a customer for the OS sequence. All servers provisioned with this OS installation profile will be associated with the specified customer (if a customer has been assigned).
- 8** From the Views pane, select Attach Software Policy.

- 9** At the bottom of the Content pane, click **Add** and select a software policy to add to the OS sequence.
- 10** From the Views pane, select Attach Patch Policies.
- 11** At the bottom of the Content pane, click **Add** and select a patch policy to add to the OS sequence.
- 12** From the Views pane, select Attach Device Group.
- 13** At the bottom of the Content pane, click **Add**. Select a device group to place the server into, after the OS sequence has been run. You can only select a public static group for this option.
- 14** From the Views pane, select Remediate Policies.
- 15** In the Content pane, choose to enable or disable remediation when the server is provisioned with the OS sequence. If you select Disable Remediation, then when you run the OS sequence, the OS will be installed but no policies in the OS sequence will be remediated – this means that no software in any of the policies attached to the OS sequence will be installed when the sequence is run.
- 16** If you select Enable Remediation, then you will need to configure the Rebooting and Scripts parameters. For the rebooting options, you can select one of the following:
 - **Reboot servers as dictated by properties on each installed item:** Selecting this option will allow any reboot settings to run that might be set in any software or patch policies attached to the OS sequence.
 - **Hold all server reboots until after all items are installed:** This option will override any pre-install reboot options that might be set in any software or patch policies attached to the OS sequence. If any post-install reboots have been set, then they will execute after the OS has been installed.
 - **Suppress all server reboots:** This option will override reboot options set in any software or patch policies attached to the OS sequence.
- 17** Next, in the Scripts section, select either a Pre-Install/Post-Install Script. These tabs allow you to set a pre- or post-install script to be executed before the OS sequence has been run and after the OS has been installed. Click **Enable Script** to enable a the script parameters.
- 18** From the Select drop-down list, select either Saved Script or Ad Hoc Script. Each script type has its own settings:

Saved Script

- **Command:** Add any commands or arguments to be executed here.
- **Script Timeout:** Enter a numerical value for the number of minutes to pass until the script will timeout.
- **User:** Enter a user name and password, or choose to run the script as Local System. (If using a Unix OS, choose root as the user.)
- **Error:** Select if you want the OS sequence job to stop if the script returns an error.

Ad Hoc Script


- **Type:** Choose UNIX shell for Unix systems, or for Windows, select BAT or VBSCRIPT.
- **Script:** Enter the text of the script. An Ad-Hoc script runs only for this operation and is not saved in Opsware SAS. In the Script box, enter the contents of the script.
- **Command:** If the script requires command-line flags, enter the flags here.
- **Script Timeout:** Enter a numerical value for the number of minutes to pass until the script will timeout.
- **User:** Enter a user name and password, or choose to run the script as Local System account. (If using a Unix OS, choose root as the user.)
- **Error:** Select if you want the OS sequence job to stop if the script returns an error.

- 19** When you have finished making your selections, from the **File** menu, select **Save** to save the OS Sequence.

Selecting Servers in the Unprovisioned Servers List

To provision a server and install an OS, select an unprovisioned server from the Unprovisioned Servers list in the SAS Client. Servers in the Unprovisioned Servers list have registered their presence, but do not have an OS installed. From this location, you can install an OS by selecting an unprovisioned server. See Figure 9-2.

Figure 9-2: Unprovisioned Servers in the SAS Client




Select an unprovisioned server in the list and the Content pane will display detailed information about the unprovisioned server that was gathered by the OS Build Agent after a network boot.

The View drop-down list enables you to view the server in the following ways:

- **Summary:** Provides information about the host name set by booting the server the first time over the network or by using an Opware Boot Floppy or CD. It also displays the OS of the OS Build Agent (Windows, Red Hat Linux, or Solaris), processor type, manufacturer and model of the server, and Opware registration information.

- **Properties:** Displays placeholders for various management and reported information which will be filled in later once the server is provisioned.
- **Hardware:** Displays details about the hardware on the server, such as a processor type, physical and virtual memory, storage and network interfaces.
- **Custom Attributes:** Allows you to read and manage custom attributes.
- **History:** Indicates the first event associated with the server.


You can also search for an unprovisioned server using the search tool  in the upper right corner of the Content pane. You can choose a filter, then enter text to search for the server.



You also have the option of running an OS sequence from the Library and then selecting a server or servers as you configure the Run OS Sequence window.



Some servers in the Unprovisioned Servers list will be in a server lifecycle state called Planned, which means the server has been partially prepared for OS provisioning. (It has a device record created for it, but no Opsware mini agent installed yet.) Servers in this state will not be able to have an OS sequence run on them.

To display the server lifecycle stage value in the Unprovisioned Servers list, in the upper right corner of the Content pane, select the column selector  and from the list select Lifecycle. For more information, see your Opsware administrator.

Before Running an OS Sequence – Firewall Considerations

The following operating systems come with default firewall settings that must be modified during the OS installation process in order to allow the Opsware agent to be properly installed and configured on the target server.

- VMware ESX Server 3.0
- Windows 2003 x64 and Windows 2003 R2
- Windows XP SP2

OS Provisioning will make minor modifications to the firewall configurations on the managed server such that communication between the Opsware core and the Opsware agent will succeed.

VMware ESX 3.0 Firewall Settings

VMware ESX 3.0 ships by default with an iptables firewall that will block communication between the core and the mini-agent or agent. In order for communications to and from the Opsware core to succeed, rules will be added to the VMware ESX firewall by the build scripts and the Opsware agent.

Windows 2003 SP1 and Windows XP SP2 Firewall Settings

For Windows 2003 SP1 and Windows XP SP2, in order for OS provisioning and ongoing management to succeed, Opsware must ensure that the Windows firewall settings are configured to bypass the default "Security Out Of the Box" experience and allow communication over the Opsware ports. Thus the Opsware OS provisioning process will update the WindowsFirewall settings in the unattend.txt answer file as necessary for provisioning and management to work.

OS provisioning will look for the following Windows Firewall configurations in unattend.txt:

- There is no WindowsFirewall configuration in unattend.txt.
- There is a WindowsFirewall configuration, but it does not allow the ports needed by Opsware.
- There is a WindowsFirewall configuration that does allow the ports needed by Opsware (no changes will be made).

In any of the cases, after running an OS sequence and installing the OS (and agent), any predefined firewall settings will remain in tact, with the exception that the Opsware agent will be installed and all of its necessary ports will have been opened.

Running An OS Sequence

To install an OS on an unprovisioned server, select a server from the Unprovisioned Servers list and run an OS sequence, or start an OS sequence and choose a target server in the Run OS Sequence window.




After you run an OS sequence job, if the OS Sequence does not have remediation enabled, the newly provisioned servers will not immediately perform a full software registration. Full software registration will occur after a small variable delay usually less



than one hour. Thus when provisioning without remediation, the server's installed software packages and patches might not be listed immediately after the OS Sequence job completes. If this occurs, check again after one hour.

To run an OS Sequence and install an OS on an unprovisioned server, perform the following steps:

- 1** Choose a way to install an OS on an unprovisioned server:
 - From the Navigation pane, select Devices ► Unprovisioned Servers. Select a server and from the **Actions** menu, select **Run OS Sequence**.
 - Or
 - From the Navigation pane, select Library ► OS Sequences. Select the OS of the OS sequence, then select the OS sequence that you want to run and from the **Actions** menu, select **Run OS Sequence**.



If the **Run OS Sequence** menu item is grayed out, one or more of the unprovisioned servers is in a server lifecycle stage of Pre-Undefined. Servers in this stage cannot be provisioned. You can display the server lifecycle stage value in the Unprovisioned Servers list. In the upper right corner of the Content pane, select the column selector . From the list, select Lifecycle. For more information, see your Opsware administrator.

- 2** In the Select OS Sequence pane, click **Add** to add an OS sequence or click **Next** if OS Sequence is already listed.
- 3** In the Run OS Sequence window, step one requires that you add an unprovisioned server or servers to provision. To add a server, click **Add**.
- 4** Click **Next**, and in the Scheduling pane choose if you want to run the OS sequence, immediately, or at a later date and time.
- 5** Click **Next** and in the Notifications pane, select an email notifier. Click **Add Notifier** and enter an email address.
- 6** You can specify if you want the email to be sent upon success of the OS sequence job () or failure of the OS sequence job (.

- 7** The ticket ID field is only used when Opware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 8** Click **Next**, and review the OS sequence information before you run the job.
- 9** Click **Start Job** to run the OS sequence. When the OS sequence job begins to run, click on the Job in the Job Status window or click **Close** to exit the Job Status window. You can also check the status of the Job by clicking on Job Logs under Jobs and Sessions in Navigation Pane.
- 10** When the OS sequence job has completed successfully, you can check the Devices
 - ▶ All Managed Servers list to see the newly provisioned server.



If you scheduled the OS sequence job to run at a later date and would like to cancel it, from the Navigation pane, select Jobs and Sessions ▶ Recurring Schedules. Then, select the job, right-click and select **Stop**.

Reprovisioning a Managed Server

You have the option of reprovisioning a managed server, but keep in mind that reprovisioning a server completely removes all data on the server.

While all data will be lost when you reprovision a server, you have the option of preserving the network configuration of the server. Also, some attributes will be saved when you reprovision the server, which are defined in the build script for each OS. For more information on OS provisioning build scripts, see OS Provisioning Setup in the *Opware® SAS Administration Guide*.



You can only reprovision a server that runs the Solaris or Linux operating system (but not Solaris x86).

To reprovision a managed server, perform the following steps:

- 1** From the Navigation pane, select Devices ▶ All Managed Servers.
- 2** Select a managed server to reprovision and from the **Actions** menu, select **Run OS Sequence**.
- 3** You will be shown a warning message that you are about to reprovision a managed server. By doing so, you will lose all data on the server. Click **Yes** to proceed.

- 4** In the Run OS Sequence window, please select the appropriate option before you begin the reprovisioning:
 - Yes, I understand the OS installation process will erase all data on the selected servers. (Mandatory. You must select this option in order to proceed.)
 - Please preserve the network configuration for the selected servers. (Optional)
- 5** Click **Next**. In the Run OS Sequence window, select an unprovisioned server or servers to provision. To add a server, click **Add**.
- 6** Click **Next**. In the Select OS Sequence pane, click **Add** to add an OS sequence.
- 7** Click **Next**, and in the Scheduling pane, choose if you want to run the OS sequence, immediately, or at a later date and time.
- 8** Click **Next** and in the Notifications pane, select an email notifier. Click **Add Notifier** and enter an email address.
- 9** (Optional) Specify if you want the email to be sent upon the success of the OS sequence job or failure of the OS sequence job.
- 10** You can also specify a Ticket Tracking ID in the Ticket ID field.
- 11** Click **Next**, and review the OS sequence information before you run the job.
- 12** Click **Start Job** to run the OS sequence. When the OS sequence has run, click **View Results** to view the results of the OS sequence job.
- 13** When the OS sequence job has been run, you can check the Devices ► All Managed Servers list to see the newly reprovisioned server.

Appendix A: VAM Platform Support

IN THIS APPENDIX

This section discusses the following topics:

- Supported Platforms in VAM

This appendix provides information about the operating system platforms and architecture that VAM supports scanning and displaying application (process families), server, and device information.



This list of operating system support for VAM is a subset of the supported platforms for the Opware Agent, since in order for VAM to be able to fully scan a server it must be under Opware management with an Opware agent. For more information on supported platforms for the Opware Agent, see the chapter on server asset tracking in the *Opware[®] SAS User's Guide: Server Automation*

Supported Platforms in VAM

For non-Linux and non-VMware operating systems, VAM supports each operating systems kernel out of the box, and assumes no customizations have been made. For a list of non-Linux and non-VMware operating systems and kernels listed supported by VAM, see Table A-1.

For Linux and VMware ESX 3 operating systems, there are certain out of the box kernel versions that VAM supports. For information on Linux and VMware operating systems and kernels supported by VAM, see Table A-2.

Table A-1: VAM Supported Operating Systems - Non-Linux/VMware

VAM SUPPORTED OPERATING SYSTEMS	OS VERSIONS	ARCHITECTURE
AIX		
	AIX 4.3 AIX 5.1 AIX 5.2 AIX 5.3	POWER
HP-UX		
	HP-UX 10.20 HP-UX 11.00 HP-UX 11.11	PA-RISC
	HP-UX 11i v2	PA-RISC and Itanium
Sun Solaris		
	Sun Solaris 6 Sun Solaris 7 Sun Solaris 8 Sun Solaris 9	Sun SPARC
	Solaris 10, Update 2 and Update 3	Sun SPARC, 32 bit x86, 64 bit x86 and Niagara
Fujitsu Solaris		

Table A-1: VAM Supported Operating Systems - Non-Linux/VMware (continued)

VAM SUPPORTED OPERATING SYSTEMS	OS VERSIONS	ARCHITECTURE
	Fujitsu Solaris 8 Fujitsu Solaris 9 Fujitsu Solaris 10	Fujitsu SPARC
Windows		
	Windows NT 4.0 Windows 2000 Server Family Windows Server 2003	32 bit x86
	Windows Server 2003 x64	64 bit x86 (not Itanium)
	Windows XP Professional	32 bit x86
	Windows XP Professional x64	64 bit x86 (not Itanium)

Table A-2: VAM Supported Operating Systems - Linux and VMware

VAM SUPPORTED OPERATING SYSTEMS	VERSIONS	KERNEL	ARCHITECTURE
Red Hat Linux			
	Red Hat Linux 7.2	2.4.7-10	32 bit x86
	Red Hat Linux 7.3	2.4.18-3	32 bit x86
	Red Hat Linux 8.0	2.4.18-14 2.4.18-17.8.0	32 bit x86
	Red Hat Enterprise Linux 2.1 AS	2.4.9	32 bit x86
	Red Hat Enterprise Linux 3 AS Red Hat Enterprise Linux 3 ES Red Hat Enterprise Linux 3 WS	2.4.21-x.EL	32 bit x86 and Itanium
	Red Hat Enterprise Linux 4 AS Red Hat Enterprise Linux 4 ES Red Hat Enterprise Linux 4 WS	2.6.9-x.EL	32 bit x86 64 bit x64 (VAM does not support Itanium for these Red Hat OS and kernel versions)
SUSE Linux			
	SUSE Linux Standard Server 8	2.4.18	32 bit x86
	SUSE Linux Enterprise Server 9	2.4.21, 2.6.5	32 bit x86 64 bit x64
	SUSE Linux Enterprise Server 10	2.6.16.13, 2.6.16.21	32 bit x86 64 bit x64
VMware			
	VMware ESX Server 3	2.4.21- 37.0.2.ELvnmix	32 bit x86 64 bit x64

Appendix B: Glossary

IN THIS APPENDIX

This appendix describes terminology and acronyms used in Opsware SAS.

ACM See Application Configuration Management.

Ad-Hoc scripts A script that is created (or uploaded) and then immediately executed by a user. The script is intended for one-time use and is not stored in Opsware SAS.

administrator See Opsware administrator.

Agent See Opsware Agent.

Agent Installer An application that installs the Opsware Agent on a server.

Agent Uninstaller An application that uninstalls the Opsware Agent on a server.

application A set of tiers and sub-tiers, such as a Web tier running Apache on Linux, an application tier running WebLogic on Windows, and a database tier running Oracle on Solaris.

application component An object that represents a process of running software, such as Apache, Oracle, BEA WebLogic, Microsoft® SQL Server, and so on.

application configuration Contains application configuration templates associated with an application.

Application Configuration Management (ACM) An Opsware feature that enables you to manage and modify configuration files for applications on managed servers.

application definition Specifies an application's logical construction in terms of tiers, subtiers and the application components contained in those tiers. Allows you to transform a data display that contains extraneous and hard-to-understand information into a focused and easy-to-understand view of the data that is relevant to the application of interest.

application tier A set of process families that you create to organize the Application View so that you can see a compelling diagram of multiple servers, the connections among them, clients connecting to them, and dependencies to which they connect.

attribute A single property of a configuration item, the value of which describes the behavior of the configuration item. Configuration items can have one or more attributes. For example, some of the attributes for the configuration item SAS Server includes hostname, life cycle, agent status, and management IP.

audit A set of rules that express the desired state of a managed server's configuration objects – for example, a server's file system directory structure or files, a server's Windows Registry, application configuration, and so on.

audit job The process that performs an audit.

audit policy A collection of rules that define a desired state of configuration for a server.

audit result The results of running an audit, which will show how a target server or groups of servers's configuration object values match or mismatch the values.

Automated Configuration Tracking An Opsware feature that allows users 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.

available patch A patch that the patch administrator has tested and marked as available. Only patches that have been marked as available can be installed by anyone other than a patch administrator. (The patch administrator can install an unavailable patch in order to test it.)

available server A reserve of new, unconfigured Opsware-enabled servers ready for quick deployment. The provisioned server can be moved into the Live environment to replace existing servers, add capacity, or support new applications. While optional, this provides faster recovery options in cases of hardware failure.

backup A feature in Automated Configuration Tracking that performs a backup of a file or database when it detects a change to a tracked configuration file or database. This action is performed only if the backup action is selected in the configuration tracking policy for the file or database.

backup (CDR) The process of saving the entire contents of the current Live directory for a specific service to the Backup directory. Code Deployment & Rollback (CDR) saves the backup copy to the local disk for the host on which the Backup operation was run. Only one backup copy is maintained at any time for a service.

backup event An event that causes configuration files or configuration databases to be backed up. Types of backup events include manual, full, and triggered.

Boot Server A part of the OS Provisioning feature that supports network booting of Sun and x86 systems with inetboot and PXE, respectively. The processes used to provide this

support include the Internet Software Consortium DHCP server, and Sun Solaris TFTP and NFS.

Build Manager A part of the OS Provisioning feature that facilitates communication between the OS Boot Agent and the Command Engine for OS provisioning.

CDR See Code Deployment & Rollback (CDR).

change log An audit trail of changes made to a node (read-only). The log identifies who has recently modified the node and who has added or removed software packages, operating systems, and servers. It also tracks who has created or removed subordinate nodes.

Code Deployment & Rollback (CDR) An Opsware feature used to push updated code and content to staging host servers.

Code Deployment Role A specific role that authorizes access to capabilities and functions with the Opsware Code Deployment & Rollback feature.

combined event history log A record of events performed on servers and network devices in your environment. These events are recorded in detail as actions performed on a certain date, by a certain user, on a certain server, or on a certain network device.

Command Engine The Opsware SAS component that enables distributed programs to run across many servers. The Command Engine handles the entry of scripts into the Opsware Model Repository (the script storage location in Opsware SAS) and the versioning of stored scripts. Command Engine scripts are written in Python and run on the Command Engine server.

Communication Test A feature that helps in identifying managed servers with unreachable Opsware Agents. A Communication Test lists all servers with unreachable agents, returns specific errors associated with each unreachable agent, and provides troubleshooting information to resolve the error. To determine if an Opsware Agent is reachable, the Communication Test runs the Command Engine to Agent Communication, Crypto Match, Agent to Command Engine Communication, Agent to Data Access Engine, Agent to Software Repository Communication, and Machine ID mismatch.

compliance The degree to which a server object conforms to a test.

configuration item (CI) An object which can be viewed and managed with Opsware OMDB. For example SAS Server, NAS Device, SAS Patch Policy.

configuration template A set of values that represent the configuration file of an application.

configuration tracking policy The configuration tracking policy defines the set of files or configuration databases to be monitored, and the actions to be taken when change is detected to a tracked file.

configuration tracking reconcile The process by which new configuration tracking policies or changes to existing configuration tracking policies are deployed on servers.

core See Opsware core.

custom attributes Attributes such as miscellaneous parameters and named data values that users can set for servers. This is used when performing a variety of Opsware functions, including network and server configuration, notifications, and CRON script configurations.

custom extension Developed and deployed by Opsware Inc. support representatives, a custom extension is a Command Engine script that extends Opsware SAS functionality. A custom extension is often tailored to meet the specific needs of a customer.

customer An account within Opsware SAS that has access to designated resources, such as servers and software.

cutover A feature in CDR, that causes the Update directory and current Live directory to be identical. The different files from the Update directory are automatically synchronized with the current Live directory.

Data Access Engine The XML-RPC interface to the Model Repository that simplifies interaction with various clients such as the SAS Web Client, system data collection, and monitoring agents on servers.

data center Legacy term. See facility.

data link connection A layer 2 connection that includes switches that are directly connected to a managed server and switches that are indirectly connected through other switches.

DCML Exchange Utility (DET) A command-line tool (cbt) for importing and exporting Opsware SAS content. The primary function of this tool is to inject a newly-installed Opsware core with content from an existing core.

deactivated server A server removed from Opsware management even though its history still exists.

deployment Within CDR, automatically pushes code and content from a staging server to a live network server.

deprecated A possible state of a package or patch in Opsware SAS. A deprecated package or patch can no longer be installed on a managed server, but might still be installed on a server before the patch or package was deprecated.

device A managed server or network device.

device group In SAS, a device group can contain managed servers and network devices, or only managed servers. In NAS, a device group contains only network devices. A device group helps you categorize your devices (servers and network devices) in ways that make sense for your organization.

Distributed Scripts An Opsware feature that allows you to manage scripts in your managed environment.

Dormant Opsware Agent An Opsware Agent that runs in the dormant mode after its installation when the Opsware SAS core is not available on the network. The dormant agent periodically attempts to contact the core and when the core is available, it performs the initialization tasks to complete its installation.

duplex mismatch A configuration mismatch between the speed and duplex of a managed server and a connected network device.

dynamic group A server group that contains servers added to or removed from the group based on a set of user-defined rules.

email notification list An the Automated Configuration Tracking feature that sends an email to those listed in an email notification list whenever a change to a tracked file or configuration database is detected.

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

facility The collection of servers that a single Opsware core manages. A facility can be all or part of a data center, server room, or computer lab.

folder A hierarchal file-system type organization structure that allows you to organize and manage your software resources and allows you to define security boundaries to control access to its content across user groups.

full backup During a full backup, all tracked configuration files that were selected to be backed up are backed up (and not just the files that have changed). Full backup is performed if you select backup as the action for a tracked configuration file.

gateway See Opware Gateway.

Global Shell A Unix shell that accesses the Opware Global File System (OGFS) and managed servers in a secure environment.

group See device group.

Health Check Monitor (HCM) The command-line scripts that check the status of Opware SAS core components.

IDK Intelligent Software Module (ISM) Development Kit. The command-line tools (primarily ismtool) that build and upload ISMs into Opware SAS.

Import Media tool A utility script included with Opware SAS that is used to import OS media from the Media Server to Opware SAS.

inclusions/exclusions criteria Specifies how to include and exclude directories and files during the snapshot or audit process.

incremental backup During an incremental backup, only targets that have changed since the last backup (and that have been selected to be backed up) are backed up. Incremental backup is performed if you select backup as the action for a tracked configuration file.

initialization Legacy term. See OS Provisioning.

IP Range Groups A designated set of servers assigned to a customer account, grouped by either a physical or a logical list.

IP Ranges A designated grouping of servers.

ISM Intelligent Software Module. A set of file and directories that include application bits, installation scripts, and control scripts for applications that are to be installed on managed servers.

ISM control A script within an ISM package that can be run on a managed server.

job Any major process run by the SAS Web Client or the SAS Client such as a communication test or a software installation.

layer 2 connection See data link connection.

Library In the SAS Client, the Library provides a way to display features such as application configurations, software policies, patches, patch policies, packages, OS sequences, and OS profiles managed by Opware.

Live directory In CDR, the directory that stores the actual code and content required to run a live site.

MAC See Media Access Control Address (MAC).

Machine ID (MID) A unique identifier that Opware SAS uses to identify the server. Opware SAS assigns a unique number to the server when it first registers and stores the Machine ID and uses it to identify each server.

managed server A server that has an Opware Agent installed on it and is under the control of a particular Opware core.

management IP The IP address that Opware SAS uses to communicate with the Opware Agent on the server.

manifest Within CDR, a list of files that preview the results of an update to be performed. Each entry in the list specifies the file size, the last-modified date, the time stamp, and the full directory path to the listed file.

Media Access Control Address (MAC) The network interface card's unique hardware number. The MAC is used as the server's physical address on the network.

Media Resource Locator (MRL) A network path in URL format that is registered with Opware SAS. The path defines the installation media for an OS.

Media Server Contains the vendor-supplied OS media used during OS provisioning over the network. The OS media on the Media Server is accessed over the network by using NFS for Linux and Solaris OS provisioning, and by using SMB for Windows OS provisioning.

MID See Machine ID.

Model Repository The Opware database that stores information about managed server configurations within Opware SAS. It contains all information necessary to build, operate, and maintain an Opware-managed site. It includes data on all managed servers, the hardware associated with those servers, their memory, CPUs, storage capacity, configuration, IP addresses, DNS configuration, and so on.

Model Repository Multimaster Component The application that propagates and synchronizes changes from each Model Repository database to all other Model Repository databases.

Modeling and Change Simulation Engine Opware SAS enables users to first model and simulate proposed operational changes to their environment before propagating these changes to production servers and applications. Using the information contained in the Software and Environment Trees, the Modeling and Change Simulation Engine maintains a model of the various hardware and software configurations and other

customer characteristics associated with each of the production servers under Opware SAS control.

MRL See Media Resource Locator (MRL).

multimaster core An Opware core that belongs to a multimaster mesh.

multimaster infrastructure component See Model Repository Multimaster Component.

multimaster mesh A set of two or more Opware cores that are linked by synchronizing the data in the Model Repositories at each of the cores. The Model Repositories in each of the cores are continually updated so that they are exact duplicates of each other. All the Opware cores in a multimaster mesh can be managed through a single SAS Web Client.

My Jobs A page in the SAS Web Client that displays a list of jobs from the Model Repository such as software installation or server provisioning.

My Scripts Private scripts that can only be executed by the user who created the script. My Scripts are created for personal use.

name-value pairs Legacy term. See custom attributes.

NAS See Network Automation System (NAS).

NAS Integration An Opware feature that enables you to closely examine detailed information about managed servers and the network devices connected to them so that you can determine how they are related and then, subsequently, coordinate and implement those changes.

OCLI See Opware command Line Interface (OCLI).

OGFS See Opware Global File System.

OMDB See Opware Operational Management Database.

Opware administrator The person responsible for overall administration, policy, and practices for accessing Opware SAS. The administrator can add users, define access to specific Opware SAS features, allow users to view site information, and deploy new code and content to their site.

Opware Agent Intelligent software on Opware-managed servers that is used to make changes to the servers. Depending on the request, the agent might use Global Opware services. Some functions supported include software installation and removal, software and hardware configuration, server status reporting, and auditing.

Opware API An API that accesses the functionality and data model of Opware SAS. The Opware API supports Web Services, Java RMI, and other types of clients. The Opware API enables you to extend, integrate, and customize Opware SAS.

Opware Automation features Opware SAS is made up of a set of Opware Automation features. Opware Automation features are the components that automate particular IT processes. The Opware Automation features include the following functions: Configuration Tracking, Code Deployment and Rollback, Script Execution, and Data Center Intelligence Reporting.

Opware Automation Platform The Opware API and runtime environment that facilitate the integration and extension of Opware SAS. The Opware APIs expose core services such as audit compliance, Windows patch management, and OS provisioning. The runtime environment (hub) executes scripts that can access the Opware Global File System (OGFS).

Opware Discovery and Agent Deployment A feature that helps deploy Opware Agents to a large number of servers through the SAS Client.

Opware Operational Management Database (OMDB) A configuration management database designed to create and maintain a record of the infrastructure data such as applications, servers, networks, and storage in your IT environment from various sources such as Opware SAS, Opware NAS, and other third party systems.

OMDB Dashboard The OMDB Dashboard displays either a default or user-configured collection of OMDB reports in an attempt to provide useful summary information on one window.

Opware SAS Web Client A web-based user interface for managing the Opware environment.

Opware Command Line Interface (OCLI) OCLI 1.0 has two commands, upload and download, which import and export files between the local file system and the Software Repository. OCLI 2.0 is intended for use only by Opware Inc. support representatives.

Opware core The server side of Opware SAS server-agent architecture. A core consists of the Opware components (such as the Model Repository, the Software Repository, the Data Access Engine, and the Command Engine) for a particular installation.

Opware Gateway Provides connectivity with an Opware core either directly or through a network of gateways. Opware Agents communicate with the core through Opware Gateways. All traffic between the servers in the Satellite and the core that manages them is routed through Opware Gateways.

Opware Global File System (OGFS) The Opware Global File System is a single, unified file system view of all file systems for all managed servers in Opware SAS.

Opsware installation Either a standalone core, multimaster core, or Opsware Satellite.

Opsware model space The Opsware Global File System (OGFS) structure that is derived from the Model Repository.

Opsware Network Automation System (NAS) The network device management application that tracks, validates, and automates network infrastructure changes. It enables compliance and enforces operational best practices across globally distributed, multi-vendor networks.

Opsware Process Automation System (PAS) The management application that automates incident resolution, maintenance tasks, and process integration in the data center. Opsware PAS is a run book automation platform.

Opsware Satellite Installed in a remote facility, an Opsware satellite provides a network connection and bandwidth management for a core that manages remote servers. A Satellite must be linked to at least one core, which may be either standalone or multimaster.

Opsware Server Automation System (SAS) The server management application that preserves the knowledge of system administrators, network engineers, and database administrators in a centralized knowledgebase. It automates previously manual tasks associated with the deployment, support, and the growth of a data center infrastructure.

OS Build Agent A part of the OS Provisioning feature that is responsible for registering bare metal servers in Opsware SAS and guiding the installation process.

OS media Installation software for an OS from the software vendor that is distributed on a CD-ROM, or DVD, or can be obtained by downloading the software from the vendor's FTP site.

OS Provisioning The process of installing a basic set of software components, including an operating system and an Opsware Agent. After provisioning is complete, the server is ready to be managed by Opsware SAS.

Package Repository Legacy term. See Software Repository.

package A collection of executables, configuration, or script files that are associated with an Opsware-installable application or program. In Opsware SAS, a package contains software package files registered in the Software Repository and software for operating systems, applications (for example, BEA WebLogic, IBM WebSphere), databases, customer code, and software configuration information.

packaging server A managed server that has the IDK installed on it. Visual Packager requires a packaging server for each type of operating system for the packages you plan to create.

PAS See Opsware Process Automation System (PAS).

patch A piece of object code (binaries) that is inserted into (patched into) an executable program to temporarily fix a known defect.

patch management administrator Administrator responsible for testing patches and defining patch options, such as installation and uninstallation scripts. A patch cannot be installed by other personnel until the patch administrator has marked the patch available through the SAS Web Client.

Patch Management An Opsware feature that allows you to upload, test, and deploy patches in a safer and uniform way.

permission A setting within a User Group that allows or disallows access to Opsware SAS features and resources. A resource is usually a set of managed servers. The set of managed servers corresponds to a facility, customer, or server group. A feature is an action, such as provisioning an OS.

physical connection The connection inferred from data link connections and represents direct connections (cables) between server and switches.

platform The name and version of an operating system.

pluggable checks Audit rules developed as scripts in a command-line environment and then imported into Opsware SAS.

post-install script A shell script invoked on a managed server immediately after a software package is installed on a managed server.

post-uninstall script A shell script invoked on a managed server immediately after a software package is removed from the managed server.

pre-install script A shell script invoked on a managed server immediately before a software package is installed on a managed server.

pre-uninstall script A shell script invoked on a managed server immediately before a software package is removed from the managed server.

preview remediate Before Opsware SAS installs software on a server, it performs a preview remediate, and determines what will happen when the actual remediate is performed (for example, what packages will be installed or removed, what server reboots are required, and so forth).

primary IP A locally-configured IP address of the management interface.

private group A type of server group that can be edited, or deleted by the Opsware user who created the server group.

privileges See Permissions and User Group.

public group A type of server group that can be created, edited, or deleted by any Opsware user who has Manage Public Server Groups permissions.

Pytwist A set of Python libraries that provide access to the Opsware API from managed servers and custom extensions.

realm In Opsware SAS, a realm is a routable IP address space, which is serviced by one or more Opsware Gateways. The managed servers that connect to an Opsware core via a Gateway are identified as being in that Gateway's realm.

remediate The process of updating the actual software configuration of a server based on the specified configuration stored in the Model Repository.

remediate output After a remediate operation completes, Opsware SAS displays the remediate output for each server that was remediated. The remediate output aggregates output from the various installation, uninstallation, or post-installation scripts, messages from Opsware SAS, and messages from the system utilities that remediate uses to perform the installation and uninstallation of packages, operating systems, and patches.

reference server A managed server that is compliant (performs as expected) and is also referred to as a known working server or a baseline server.

remote terminal A terminal window for a Unix server or an RDP client window for a Windows server.

Reports An Opsware feature that enables you to create reports that provide comprehensive, real-time information about managed servers, network devices, software, patches, customers, facilities, operating systems, compliance policies, and users and security in your environment.

restore A function of the Automated Configuration feature that allows the user to return the configuration file or database to a previous state, when the backup action for a tracked file or database is selected.

restore Within CDR, the process of restoring the previous Live directory from the Backup directory to the Live directory.

restore queue The queue in which the configuration files are placed before they are restored to a server.

rollback Within CDR, returns a site to the state prior to the last cutover. During rollback, restores the set of modified and deleted files to the Live directory.

rosh The remote Opsware shell is a command that makes a client connection enabling you to remotely run programs on managed servers.

SAS See Opsware Server Automation System (SAS).

Satellite See Opsware Satellite

Script Execution See Distributed Scripts.

selection criteria Rules that instruct Opsware SAS to collect information about server objects and how to collect that information. It optionally can also collect data for file comparison and inclusions/exclusions. Selection criteria is required for the snapshot and audit processes.

sequence The process within CDR that simplifies deployment by grouping a collection of service operations and synchronizations that can be performed as a single task.

Sequence Editor In CDR, a predefined User Group used to create, modify, or delete a sequence definition.

Sequence Performer (Production) In CDR, a predefined User Group used to directly perform or request performance of a sequence action on production hosts.

Sequence Performer (Staging) In CDR, a predefined User Group used to directly perform or request performance of a sequence action on staging hosts.

Sequence Requester (Production) In CDR, a predefined User Group used to request performance of a sequence action on production hosts.

Sequence Requester (Staging) In CDR, a predefined User Group used to request performance of a sequence action on staging hosts.

servers Any specific hardware. Specific nodes are attached to servers that determine the specific software, configuration, and other server attributes.

server assimilation Opsware SAS assimilates servers that are already functioning in the operational environment, which allows users to deploy and manage new applications installed on those servers. Assimilating servers installs Opsware Agents on the servers and registers them with the Model Repository.

server baselines The process of defining and provisioning servers with standard configurations. Opsware templates can be used to automate the building of complete server baselines.

Server Explorer A feature of the SAS Client that allows you to browse and manage servers and server groups in your facility.

Server ID The primary key in the Opsware Model Repository that represents a given server. The Server ID is used internally in Opsware SAS.

server lifecycle The various server states assigned to a server by Opware SAS. Server states include Unprovisioned, Available, Installing OS, and Managed.

server management The process by which users can manage and track servers in an Opware-managed environment. Opware SAS forces changes to the operating environment by first changing the centralized configuration information in the Model Repository and then changing the actual configuration of physical servers.

Server Pool Servers that have registered their presence with Opware SAS, but do not have a full operating system installed.

server provisioning The process of installing a basic set of software components including the operating system, an Opware Agent, and other system utilities and debugging tools to manage the server. Configuration is defined in the Model Repository.

Server Search A feature that allows you to search for servers based on a variety of criteria, including OS version, installed package, customer, and installed patch.

Server Status A feature that defines server availability. The three major status conditions are USE, STAGE, and STATE.

server-based configuration tracking policy A configuration tracking policy that is defined for a particular server or group of servers, rather than for a particular software node (application).

service A host application (for example, BEA WebLogic, Allaire ColdFusion, Microsoft IIS, Apache Web Server, or iPlanet Application Server).

Service Editor In CDR, a predefined User Group used to define and modify or delete service definitions.

Service Performer (Production) In CDR, a predefined User Group used to directly perform or request performance of service operations on production hosts (servers).

Service Performer (Staging) In CDR, a predefined User Group used to directly perform or request performance of service operations on staging hosts.

Service Requester (Production) In CDR, a predefined User Group used to request performance of service operations on production hosts.

Service Requester (Staging) In CDR, a predefined User Group used to request performance of service operations on staging hosts.

service-instance Multiple independent instances of a service running on a host (for example, BEA WebLogic, which can run single or multiple instances).

Service Levels User-defined categories that are used to group servers in an arbitrary way. For example, a user can group servers by functionality, tier, application, or ontology.

Shared Scripts Public scripts that every Opsware SAS user can access.

signature A set of rules that you provide and that VAM uses to identify a process family.

Site Backup directory In CDR, the directory that stores a complete backup of the Live directory when the user issues a Backup service operation.

Site Previous directory In CDR, the directory that stores the files that have changed between the current Live directory and its previous state prior to the last cutover. It holds all the changes necessary to revert the Live directory back to the state that it was in before the last cutover.

snapshot A record of how an Opsware managed server is configured at a particular point in time. Snapshots allow administrators to audit the configuration of servers and deploy files and software to correct discrepancies. A snapshot can be based on specified server objects. Audit and Remediation records one snapshot per server.

snapshot job The process that created a snapshot of a server or server group.

snapshot specification A definition of a target and selection criteria that will be examined during the snapshot process to capture and record information about a managed server.

software compliance Software compliance indicates whether or not the all software policies attached to the selected server are compliant with the actual server configuration.

software policy A software policy contains packages, patches, application configurations, and other software policies and allows you to install software and configure applications on managed servers simultaneously.

Software Repository The central repository for all software managed by Opsware SAS. It contains software packages for operating systems, applications, databases, customer code, and software configuration information.

Software Repository Cache An Opsware Satellite component that contains local copies of files. The Software Repository Cache stores files from the Software Repository of an Opsware core or from another Software Repository Cache, and supplies the cached files to Opsware Agents on managed servers.

Software Repository Replicator A component providing backup functionality for Software Repositories running in a multimaster mesh.

source In the snapshot process, this is the managed server that information is recorded about. In the audit process, this is an existing snapshot or server you are comparing selection criteria *from*.

standalone core An Opware core that manages servers in a single facility. Unlike a multimaster core, a standalone core does not communicate with other cores.

static group A server group in which the servers are added to and removed from the group manually.

synchronization The process within CDR to move modified files from a directory on a source host to a directory on a destination host.

Synchronization Editor In CDR, a predefined User Group used to create, modify, or delete a synchronization definition.

Synchronization Performer In CDR, a predefined User Group used to directly perform or request performance of a synchronization action.

Synchronization Requester In (CDR), a predefined User Group used to request performance of a synchronization action.

target In the snapshot process, this is the managed server or server group you are recording information about. In the audit process, this is an existing snapshot, server, or server group you are comparing selection criteria *to*.

topology A snapshot that represents the state of a set of network devices and managed servers, the process families running on those servers, the connections among those process families, and any external clients and dependencies.

tunnel A TCP connection between two Gateways that carries multiplexed TCP or UDP connections.

Update directory The directory that CDR writes to when synchronizing modified files in source and destination hosts. After synchronization, the Update directory is different from the current Live directories. After cutover, the Update directory and current Live directory are identical.

user An individual with access to the Opware SAS. An Opware user belongs to one or more User Groups, which control the access of its members.

User Group Represents a role played an organization's Opware users. The permissions specified for a user group determine what the group's members can do with Opware SAS.

Value Set Editor Enables you to change the values in a configuration file by editing that file's value set. Each entry configuration file is represented inside the value set editor as a value set, (a key name and a value).

Visual Application Manager An Opware feature that is designed to help you understand and manage the operational architecture and behavior of distributed business applications in your IT environment.

Web Services API A SOAP-based interface that facilitates the integration of operations and business support systems with Opware SAS. The Opware Web Services API 2.2 is intended for Opware SAS 5. A different Web Services API, introduced in Opware SAS 6.1, accesses the Opware API.

Web Services Data Access Engine A web services interface to the Model Repository that provides increased performance to other Opware SAS components.

Index

A

- ad hoc audit, running170
- adding
 - application configuration, server or group to ..430
 - configuration template421
 - rule exceptions to an audit162
- Ad-Hoc Scripts
 - definition of479
- AIX
 - APARs
 - about339, 340
 - uploading339, 340
 - LPPs, about339
 - APARs. See AIX APARs.
 - application configuration
 - Compliance Dashboard215
 - compliance remediation options216
- Application Configuration Management
 - adding, configuration template421
 - adding, server or group to430
 - application configuration, definition of479
 - application configuration, overview403
 - applying and inheriting values407
 - CML407
 - comparing, configuration templates438
 - configuration out of synch with server439
 - configuration template, overview403
 - configuring audit rules132
 - creating
 - application configuration417
 - configuration template419
 - default values, inheritance408
 - definition of479
 - deleting, configuration template422
 - editing, default values427
 - instance values, inheritance409
 - loading
 - template423
 - values433
 - main components402
 - preserve values406
 - process400
 - pushing changes, servers or groups to434
 - roll back442
 - scheduling
 - push436
 - scan for compliance439
 - sequence merging
 - append mode413
 - overview412
 - prepend mode415
 - primary key option414
 - replace mode413
 - setting values, on431
 - setting, templates to run as scripts425
 - show inherited values406
 - specifying, template order426
 - value set editor
 - definition of495
 - overview403
- Application Map (VAM)54
- attaching
 - software policy to server368
- audit
 - audit job, definition of480
 - audit results
 - viewing and remediating179
 - auditing process111
 - audits and the Compliance Dashboard107
 - Compliance Dashboard217
 - compliance remediation options220
 - configuring, overview118
 - elements of112
 - performing
 - audit results, from171
 - reports107
 - re-running from audit results171
 - results, comparison based remediation176
 - results, value based remediation177
 - running from the library169
 - running on a server170
 - saving as audit policy168
 - schedule configuration compliance scan440
 - scheduling172

scheduling recurring	173	selection criteria	
searching for	182	definition of	491
selection criteria		inclusions/exclusions	155
inclusions/exclusions	155	inclusions/exclusions, definition of	484
snapshot used in	183	snapshot	
sources, server or snapshot	128	definition of	493
viewing completed audit job	175	job, definition of	493
ways to create	113	template, definition of	493
creating from a server	114	source, definition of	494
from a group of servers	114	target, definition of	494
from a snapshot	115	terms and concepts	108
from an audit policy	115	viewing	
from the library	115	and remediating audit results	179
Audit and Remediation		ways to create an audit	113
audit policies	165	audit policy	
audit process overview	111	creating	166
audit results	175	linking and importing	166
capturing golden server configuration	105	overview	165
creating an audit policy	166	saving	168
creating your own ad hoc audit	106	audit template	
deleting		definition of	480
snapshot	491	available	
snapshot specification	191	patches, definition of	480
enforcing security policies	105	servers, definition of	480
examples (use cases)	104	B	
exceptions	161	backups	
adding to an audit	162	backup event, definition of	480
editing	164	full, definition of	483
rules that cannot have exceptions	162	incremental backup, definition of	484
linking and importing audit policies	166	booting	
overview	104	Solaris servers, over network	459
performing audit		Build Manager	
audit results, from	171	definition of	481
rules	120	OS Build Agents, locating	460
configuring	124	C	
configuring, application configuration	131	change logs, definition of	481
configuring, COM+	138	code deployment	
configuring, custom script	139	definition of	480, 481
configuring, even logging	141	deployment, definition of	482
configuring, file system	143	live directory, definition of	484
configuring, hardware	144	restore, definition of	490
configuring, IIS Metabase	145	sequences	
configuring, operating system	146	definition of	491
configuring, software	148	sequence editor, definition of	491
configuring, TON	152	sequence performer, definition of	491
configuring, users and groups	148	sequence requester, definition of	491
configuring, Windows Registry	150	site backup directory, definition of	493
configuring, Windows services	151		
example of configuring	125		
server objects	122		
scheduling audits	172		

- site previous directory, definition of 493
 - user roles, definition of 481
 - COM+ object
 - configuring Audit and Remediation rule 138
 - Command Engine, definition of 481
 - command line interface, definition of 487
 - Communication Test
 - definition of 481
 - comparing scan results (VAM)
 - about 88
 - how to 93
 - heuristics used 100
 - interesting object attribute differences 92
 - object attribute differences 91
 - object existence comparison 90
 - source and target scans 89
 - comparing, configuration templates 438
 - compliance
 - performing software compliance scan 397
 - software 396
 - viewing in VAM for servers 66
 - Compliance Dashboard
 - application configuration 215
 - audit 217
 - compliance statuses 209
 - duplex 220
 - filtering and sorting information 221
 - general categories 208
 - overview 205
 - patch 216
 - refreshing 210
 - remediation overview 212
 - software compliance 214
 - terms and concepts 210
 - usage, proactive or reactive 206
 - viewing 207
 - compliance status for Compliance Dashboard 209
 - component signature (VAM)
 - about 81
 - creating 84
 - cutting and copying 85
 - defined 49
 - deleting 85
 - editing 85
 - evaluation order 82
 - pasting 86
 - properties of 74
 - components
 - Application Configuration Management, of 402
 - Configuration Markup Language 407
 - configuration template
 - adding 421
 - adding, scripts 425
 - comparing 438
 - creating 419
 - definition of 481
 - deleting 422
 - loading
 - CML file 423
 - values into 433
 - overview 403
 - specifying, order 426
 - configuration tracking
 - backups, definition of 480
 - definition of 480
 - email notification lists
 - definition of 483
 - policies
 - definition of 482
 - reconcile, definition of 482
 - restore queue, definition of 490
 - restore, definition of 490
 - server-based policies, definition of 492
 - configuring
 - snapshot specification 188
 - conventions used in the guide 25
 - copying
 - objects to a server from a snapshot 201
 - creating
 - application configuration 417
 - application definition in VAM, overview 77
 - application tier (VAM) 79
 - audit policy 166
 - component signature (VAM) 84
 - configuration template 419
 - OS sequence 463
 - snapshot specification 187
 - snapshot specification from library 187
 - custom attributes
 - definition of 482
 - custom extensions
 - definition of 482
 - custom script
 - configuring Audit and Remediation rule 139
 - customers
 - definition of 482
 - cutting over
 - definition of 482
- D**
- Data Access Engine, definition of 482

data center. See facility.

deactivating
 servers, definition of482

deleting
 configuration template422
 snapshot200, 491
 snapshot job schedule195
 snapshot specification191

depots
 patch management338

deprecated packages, definition of483

detaching
 software policy to server382

device tree (VAM)50

DHCP
 Linux servers, requirements for using455, 458
 servers, booting with450
 Solaris servers, usage of453, 454

directories
 live, definition of484
 site backup directory, definition of493
 Update directory, definition of494

duplex
 Compliance Dashboard220
 compliance remediation options221

E

editing
 audit rule exceptions164
 audit schedule174
 default values, application configuration for427
 snapshot job schedule194

email
 notification lists for configuration tracking, definition
 of483

e-mail notifications276, 277, 304, 313, 354, 362

encoding
 choose for application configuration source406
 uploading template file423

Environment Tree
 definition of483

error messages (VAM)98

evaluation order, VAM component signatures82

event logging
 configuring Audit and Remediation rule141

Exceptions, Audit and Remediation
 about161
 adding to an audit162
 considerations162
 rules that cannot have exceptions162

F

facilities, definition of483

File System
 configuration Audit and Remediation rule143

filtering and sorting Compliance Dashboard
 information221

font295

G

Global Shell
 definition of484

groups, definition of484

H

hardware preparation, overview453

hardware,configuring Audit & Remediation rule ..144

hotfix chaining248

HP-UX
 depots
 patch management338

I

icons, toolbar icons (VAM)41

IIS Metabase
 configuring Audit and Remediation rule145

import media tool, definition of484

importing audit policy167

Install Patch Wizard349, 356

Install Patch wizard297, 348

install scripts,specifying274, 360

installation
 flags, overview297, 308, 348

installing
 Install Patch Wizard349, 356

Opware Agents
 definition of491

OS Build Agents
 verification461
 software using software policy368, 382

Intelligent Software Module (ISM)
 definition of484

IP range groups, definition of484

IP ranges
 definition of484

ISM Control
 running392

ISM control, definition of484

J

Japanese293
 jobs
 definition of484
 snapshot, definition of493

K

Korean293

L

layer 2 connections displayed in VAM 60
 life cycle
 definition of492
 link selection (VAM) 44
 linking an audit policy167
 loading
 template423
 values, configuration template into 433
 local attachment, definition of485
 locales293

M

Machine ID (MID), definition of485
 management IP
 definition of485
 manifests, definition of485
 mbsacl.exe244, 248
 Media Access Control Address, definition of485
 media resource locators (MRLs), definition of ...485
 media server, definition of485
 Microsoft Baseline Security Analyzer244
 Microsoft patch management prerequisites244
 Model Repository 246, 267, 268
 Model Repository Multimaster Component, definition
 of485
 Model Repository, definition of485
 modeling and change simulation engine
 definition of485
 msixexec.exe248
 multimaster
 mesh, definition of486
 multimaster core, definition of486
 My Jobs
 definition of486

N

network

Solaris servers, booting over459
 Network Map (VAM) 57

O

opening a .vam file 87
 operating systems
 configuring Audit and Remediation rule 146
 patch management, supported for336
 provisioning445
 Opware administrator, definition of486
 Opware Agent246
 definition of486
 dormant, definition of483
 registration332
 Opware Agent Installer
 definition of479
 Opware Agent Uninstaller, definition of479
 Opware Discovery and Agent Deployment
 definition of487
 Opware Gateway, definition of487
 Opware Global File System
 definition of487
 Opware guides
 contents of 23
 conventions used 25
 documentation set 26
 icons in guide, explained 25
 Opware model space, definition of488
 Opware SAS
 core, definition of487
 documentation set 26
 features, definition of487
 installation, definition of488
 related documentation 26
 Opware Satellite, definition of488
 OS Build Agents
 Build Manager, locating460
 definition of488
 failure to install, recovering from462
 overview455
 verifying installation461
 OS media, definition of488
 OS provisioning
 definition of488
 hardware preparation453
 import media tool, definition of484
 Linux servers452
 managed server values449
 media resource locators (MRLs), definition of .485
 media server, definition of485

OS Build Agents		patch policy	266
definition of	488	patch policy exception	268
overview	455	patch reboot options	273, 301, 351
using	455	patch uninstallation, previewing	314, 362
SAS Client		patch uninstallation, scheduling	313, 361
creating an OS sequence	463	patches	
overview	462	installation flags	297, 348
reprovisioning a managed server	472	types supported	336
running an OS sequence	470	uninstallation flags	308
select unprovisioned servers	468	performing	
Server Pool values	449	audit	
Solaris servers	452	audit results, from	171
Windows servers	452	policy setter	250
OS sequence		populate-opsware-update-library	260
attach device group	464	post-install script, definition of	489
attach patch policy	464	post-uninstall script, definition of	489
attach software policy	464	pre-install script, definition of	489
creating	463	pre-uninstall script, definition of	489
set remediate policy	465	preview remediate	
		definition of	489
P		primary IP, definition of	489
package types		printing a VAM map	64
AIX APAR	339, 340	private group, definition of	489
HP-UX depots	338, 340	process families (VAM)	
LPP	339	defined	50
RPM	336	properties of	72
Windows Hotfix	247, 298	property page (VAM)	
packages		application tier	73
definition of	488	component signature	74
deprecated, definition of	483	links	69
packaging server		network device	71
definition of	488	network interface	72
patch		port group	71
Compliance Dashboard	216	process family	72
compliance remediate options	217	server	65
patch compliance	267, 268, 270	server compliance	66
patch compliance scan	245	virtual server	70
patch installation, previewing	305, 355	virtual switch	71
patch installation, scheduling	275, 304, 354	public group, definition of	490
patch management		push	
definition of	489	application configuration	
Microsoft patch releases	245	onto a server or group	434
operating systems, supported	336	scheduling	436
patch administrators			
definition of	489	Q	
patch information from Agent	332	qchain.exe	248
patch testing, support for	332	QNumber	248, 277, 279
roles	335		
supported Unix versions	336		
uploading automatically	260		

R

realm, definition of	490
reference server, definition of	490
refreshing	
Compliance Dashboard	210
remediate	
application configuration compliance	216
audit compliance	220
definition of	490
duplex compliance	221
output	
definition of	490
overview	372
patch compliance	217
preview	
definition of	489
software policy	374
remediating	
audit results	178
remote Opsware shell	
definition of	490
remote terminal	
definition of	490
reports	
software policy	397
reprovisioning a managed server, SAS Client	472
rolling back	
definition of	490
RPM	
patching	336
rules	
Audit and Remediation	120
running	
ad hoc audit	170
audit from server	170
audit from the library	169
ISM Controls	392
OS sequence, SAS Client	470
snapshot specification	191

S

SAS Client	
OS installation with	462
SAS Web Client	
My Jobs	
definition of	486
patch administration in	346
saving	
a VAM map to an image file	63
snapshot specification as policy	190

scan	
configuration compliance	439
patch compliance	285
software compliance	397
scan timeout (VAM)	75
scheduling	
audit	172
audit, recurring	173
snapshot job	193
scripts	
Distributed Scripts	
definition of	483
My Scripts, definition of	486
post-install script, definition of	489
post-uninstall script, definition of	489
pre-install script, definition of	489
pre-uninstall script, definition of	489
setting, templates to run as script	425
Shared Scripts, definition of	493
search	
audit	182
searching	
server search, definition of	492
sequence merging, application configuration	
append mode	413
prepend mode	415
primary key option	414
replace mode	413
sequences	
editor for CDR, definition of	491
server	
attaching software policy	368
detaching software policy	382
server baselines, definition of	491
Server Explorer	
definition of	491
server groups	
adding, application configuration	430
definition of	491
dynamic, definition of	483
private, definition of	489
public, definition of	490
pushing changes, application configuration of	434
setting values on	431
static, definition of	494
Server ID, definition of	491
server management	
definition of	492
groups, definition of	484
IP range groups, definition of	484
IP ranges, definition of	484

Server Map (VAM)	56	difference between snapshot specification ...	183
server objects		editing job schedule	194
Audit and Remediation	122	job, definition of	493
Server Pool		locating	196
definition of	492	locating in SAS Client	195
servers		process	186
adding, application configuration	430	saving as audit policy	168
booting		scheduling	193
over network	459	template, definition of	493
configuration policies for, definition of	492	used in an audit	183
deactivating, definition of	482	used with audit policies	183
definition of	491	viewing contents of	197
life cycle		viewing job schedule	195
definition of	492	snapshot specification	187
Machine ID, definition of	485	and audit policies	183
managed, definition of	485	configuring	188
management IP, definition of	485	configuring rules for	190
media access control address, definition of ..	485	creating from library	187
primary IP, definition of	489	creating from server	187
provisioning, definition of	492	deleting	191
pushing changes, application configuration of	434	elements of	184
reference server, definition of	490	relationship to snapshots	183
reprovisioning, SAS Client	472	running	191
Server ID, definition of	491	selection criteria	
Service Levels, definition of	493	inclusions/exclusions	155
setting values on	431	software	
Solaris servers, booting	459	configuring Audit and Remediation rule	148
source, definition of	494	software compliance	
status, definition of	492	Compliance Dashboard	214
Service Levels, definition of	493	compliance remediate options	214
services		software installation	
definition of	492	attaching software policy	368
editor for CDR, definition of	492	detaching software policy	382
performer (production) for CDR, definition of	492	installing using software policy	368, 382
performer (staging) for CDR, definition of ..	492	installing, software policy template	385
requester (production) for CDR, definition of	492	overview	365
requester (staging) for CDR, definition of ..	492	process	366
service-instance for CDR, definition of	492	remediate	
setting		overview	372
application configuration values	431	remediate software policy	374
configuration templates to run as scripts ..	425	software policy template overview	383
Shared Scripts		ways to install	368
definition of	493	software management	
showing		installation, overview	365
individual audits in Compliance Dashboard	219	software policy	
snapshot		attaching to server	368
copying objects to server from	200	compliance overview	396
definition of	493	detaching to server	382
deleting	200, 491	installing software	368, 382
template	191	installing, software policy template	385
deleting job schedule	195	performing software compliance scan	397

remediate374
 remediate overview372
 reports397
 running ISM controls392
 software policy template overview383
 software policy template
 installing385
 overview383
 Software Repository246
 Software Repository Cache, definition of493
 Software Repository Replicator, definition of493
 Software Repository, definition of493
 Solaris
 booting servers over network459
 OS provisioning452
 specifying
 application configuration template order426
 application configuration to run as script425
 standalone core, definition of494
 static group, definition of494
 Summary Review277, 306, 314, 355, 362, 363
 synchronizations
 definition of494
 editor for CDR, definition of494
 performer for CDR, definition of494
 requester for CDR, definition of494

T

target, definition of494
 templates
 local attachment, definition of485
 TON rules
 configuring Audit and Remediation rule152
 troubleshooting
 OS Build Agents
 installation failure462
 verifying installation461
 tunnel, definition of494

U

Uninstall Patch wizard308, 357
 uninstallation
 flags, overview297, 308, 348
 unzip.exe248
 user roles, definition of494
 users and groups, configuring Audit and Remediation
 rules148
 users, definition of494

V

value set editor
 definition of495
 value set editor, application configuration
 editing default values427
 overview403
 preserve values406
 show inherited values406
 verifying
 installation of OS Build Agents461
 viewing
 audit results179
 audit server usage115
 completed audit job175
 Compliance Dashboard207
 snapshot contents197
 snapshot job schedule195
 Virtualization Director
 scan time out preference75
 virtualization settings74
 Virtualization Map (VAM)58
 Visual Application Manager (VAM)
 .vam files
 managing87
 opening87
 overview32
 saving87
 saving as template88
 accessing servers
 Device Explorer76
 remote terminal76
 comparing scan results88
 comparison types89
 how to93
 interesting object attribute differences92
 object attribute difference91
 object existence comparison90
 source and target89
 data collection and display46
 error messages98
 filtering data94
 criteria to use95
 relationships between results96
 types of filters94
 using regular expressions96
 icons in toolbar41
 link selection tool44
 menus44
 new in this release30
 options74
 overview29

prerequisites to run	32	configuring Audit and Remediation rule	150
property pages	64	Windows servers	
application tier	73	OS provisioning	452
component signatures	74	Windows services	
connection links	69	configuring Audit and Remediation rule	151
network devices	71	Windows Update Agent	248
network interfaces	72	wizards	
port groups	71	Install Patch	349, 356
process families	72		
server	65		
server compliance	66		
virtual servers	70		
virtual switches	71		
supported operating systems	32		
symbols used in maps	62		
user interface explained	39		
VAM application	47		
adding a tier	79		
application tier	49		
application tree	47		
component signature	49		
component signatures evaluation order	82		
component signatures explained	81		
copying and cutting a tier	80		
creating definition	77		
deleting a tier	80		
device tree	50		
editing a tier	80		
pasting a tier	81		
process families	50		
templates	78		
tiers	79		
VAM maps	53		
application map	54		
network map	60		
printing	64		
saving as an image file	63		
server map	56		
virtualization map	58		
virtualization settings	74		
Visual Packager			
definition of	488		

W

Web services APIs, definition of	495
Web Services Data Access Engine, definition of	495
Windows Hotfix	
installation flags	298
uploading	298
Windows Registry	