OPSWARE INC
Automating IT™

Opsware® SAS 6.0 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

Opsware SAS Version 6.0

Table of contents

Chapter 3: Compliance Dashboard 155

Chapter 4: Reports 175

Chapter 5: Patch Management for Windows 189

Chapter 8: Application Configuration Management 321

Chapter 9: Operating System Provisioning 367

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 Opsware 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: Glossary: Defines terminology and acronyms that are unique to Opsware 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.
`Courier`	Identifies text that is entered or displayed at the command-line prompt, such as Unix commands, Opsware SAS commands, file names, paths, directories, environment variable names, contents of text files that are viewed or edited with a text editor, source code in a programming language, and SQL (database) commands.
Italics	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 iconographic conventions.

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

Guides in the Documentation Set and Associated Users

- The *Opsware® SAS User's Guide: Server Automation* is intended to be read by systems administrators and 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, remediating servers, script execution, configuration tracking, deploying and rolling back code, agent deployment, and using the Opsware Global Shell and opening a Remote Terminal on managed servers. This guide is intended for system administrators who are responsible for all aspects of managing the servers in an operational environment.
- *Opsware® SAS User's Guide: Application Automation* is intended to be read by the system administrator who is responsible for performing the day-to-day functions of managing servers, such as auditing and compliance, software packaging, visual application management, application configuration, and installing software and operating systems on managed servers.
- The *Opsware® SAS Administration Guide* is intended to be read by Opsware administrators who will be responsible for monitoring and diagnosing the health of the Opsware SAS components.
- The *Opsware® SAS Planning and Installation Guide* is intended to be used by advanced system administrators who are responsible for planning all facets of an Opsware SAS installation and for the installation of Opsware SAS in a facility. It documents all the main features of Opsware SAS, scopes out the planning tasks necessary to successfully install Opsware SAS, how to run the Opsware Installer, and how to configure each of the components. It also includes information on system sizing and checklists for installation.
- The *Opsware® SAS Policy Setter's Guide* is intended to be used by system administrators who are responsible for all facets of configuring the Opsware Command Center. It documents how to set up users and groups, how to configure Opsware server management, and how to set up the main Opsware Command Center features, such as patch management, configuration tracking, software repository replicator setup, code deployment, and software provisioning.

Opsware, Inc. Contact Information

The main web site and phone number for Opsware, Inc. are as follows:

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

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

- *https://download.opsware.com/opsw/main.htm*

For troubleshooting information, you can search the Opsware Knowledge Base at:

- *https://download.opsware.com/kb/kbindex.jspa*

The Opsware Customer Support email address and phone number follow:

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

Chapter 1: Visual Application Manager

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Visual Application Manager
- Ways to Use the Visual Application Manager
- Visual Application Manager Display
- Supported Operating Systems
- Launching Visual Application Manager
- Setting Up Application Definitions
- Topology Management
- Analysis and Action
- Error Messages

Overview of Visual Application Manager

The Visual Application Manager is designed to help you understand and manage the operational architecture and behavior of distributed business applications in your IT environment. Business applications are complex collections of services that typically run across many servers and network devices; therefore, it can be difficult to understand or remember what is connected to what, where performance problems originate, how to troubleshoot and resolve problems, and what result would occur if you make a change in your environment.

The Visual Application Manager helps you visualize this type of information in physical and logical drawings. These drawings illustrate how distributed business applications are running across servers and network devices. The drawings include components, connections, and dependencies that show what applications would be affected if you removed a certain server or network device. Since these drawings are created from data collected in real-time, they can help you diagnose and resolve problems. The Visual

Application Manager is integrated with other features in the Opsware SAS Client and Opsware NAS Client to perform change management tasks, such as reconfiguring, patching, auditing, and remediating software and patch policies. This integration means that when you understand your business application and have identified the source of a problem with it, you can take action to remediate the problem.

Visual Application Manager Features

The Visual Application Manager helps you perform the following tasks:

- Discover and draw the components, connections, and dependencies in a Web application.
- Visualize this information in several different layouts, physical and logical, such as in a server view, in a layer 2-based network view, and in an application view.
- Recognize components as known applications and highlight them accordingly.
- Organize recognized components into multi-tier applications to create a logical view that can be analyzed to verify correct operation.
- Manage .vam (Visual Application Manager archive) files, which contain an application definition and topologies.
- Verify the correct operation of Web-based applications.
- Troubleshoot and resolve problems.
- Launch other Opsware SAS Client features, such as Server Explorer, Network Device Explorer, Global Shell, and Remote Terminal, to perform in-depth analysis or to perform actions on the systems under investigation.
- Export drawings to .gif, .jpg, and .svg files for use in other applications, such as Microsoft® Visio or to view scalable drawings in a Web browser.
- Print drawings.

Prerequisites

The Visual Application Manager requires an Opsware Agent that is version 5.1 or higher on managed servers. This enables the Visual Application Manager to scan them.

Ways to Use the Visual Application Manager

The Visual Application Manager collects, models, and displays data about the operational architecture and behavior of applications in your IT environment. The Visual Application Manager also manages this information by using .vam files that contain application definitions and topologies.

The Visual Application Manager uses the information in the Opsware SAS and Opsware 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), and mapping application data to enable you to visualize and analyze your operational environment. Figure 1-1 shows the process you follow to display data about your IT environment.

Figure 1-1: Visual Application Manager Processes

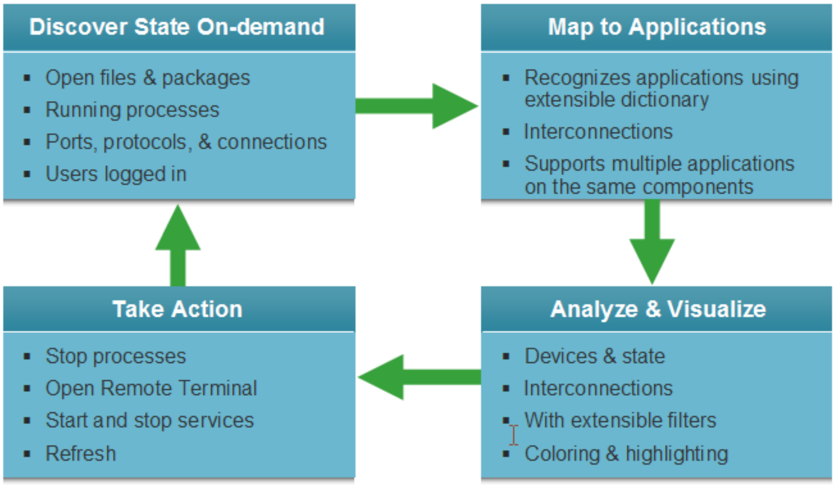

You begin using the Visual Application Manager by selecting servers or server groups in the Opsware SAS Client. The Visual Application Manager scans those servers and draws physical and logical views of their state (as described in the “Discover State On-demand” box in the figure). See “View Tabs” on page 40 for information about how Visual Application Manager draws physical and logical views of your IT environment.

Based on the scan results that are displayed in the views, you can then define applications so that extraneous information is removed and the views show the key information for the application you are viewing (as described in the “Map to Applications” box in the figure). See “Setting Up Application Definitions” on page 50 for information about mapping data to applications.

As you analyze the information about the application you mapped, you can then take action on the servers to troubleshoot and remediate problems (as described in the "Take Action" box in the figure). For example, you can launch other Opsware SAS Client features, such as the Server Explorer, Network Device Explorer, Global Shell, and Remote Terminal. See "Analysis and Action" on page 59 for information about opening the Server Explorer, Network Device Explorer, Global Shell, and Remote Terminal from the Visual Application Manager.

Data Discovery

The Visual Application Manager gathers information about distributed business applications in your environment by scanning managed servers. These business applications consist of complex collections of services that typically run across multiple servers and network devices. In Visual Application Manager, applications and Web-based applications are referred to interchangeably as *applications*.

When you select servers, the Visual Application Manager will scan those servers and then draw physical and logical views. These views do not show servers that are both unscanned and have no IP traffic. Such servers are indicated by an error icon in the Device Tree–where a tooltip for the server (or network device) provides details about the cause of the error.

Data Collection and Display

The Visual Application Manager draws layout views based on data that is collected in real-time snapshots of a scan. Server data is captured directly from servers and then recorded in snapshots. Network device data is scanned and then recorded in snapshots by NAS–where it is retrieved by the Visual Application Manager from the Opsware NAS data model.

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

The Visual Application Manager is a workbench-style interface that displays multiple views of the data in your IT environment. In particular, the Visual Application Manager displays the following information about managed servers and network devices in your environment:

- Processes that are running on managed servers
- TCP and UDP connections between these processes
- Detailed configuration information
- Current runtime state information about servers, connections, and processes
- How servers, interfaces, switches, and switch ports are connected

See “Processes, Process Families, and Extended Process Families” on page 39 and “Symbols Used in Views” on page 45 for an explanation of how the Visual Application Manager interprets this data.

.vam Files

The Visual Application Manager manages data as an application that is represented as a .vam file. A .vam application is analogous to a document managed by Microsoft® Word, which means it can be opened, closed, saved, and edited. A .vam file is a JAR file (which is actually a ZIP file) that can be opened and examined by using WinZip.

Each .vam file contains an application definition (which can be empty) and zero or more topologies. The application definition specifies an application's logical construction in terms of tiers, subtiers and the application components contained in those tiers. Each topology is 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.

A .vam application that does not have any topologies attached is useful for creating an application definition that can be imported (as a template) into other .vam application documents. See “Topology Management” on page 57 for instructions on how to save, open, and edit a .vam file.

Visual Application Manager Display

The Visual Application Manager user interface is a workbench style window that includes Tree tabs, View tabs, a Property Page, a dynamic toolbar, and detailed tooltips for tree and view objects. The trees and the views display physical and logical layouts of the devices in your operational environment. Figure 1-2 shows the types of information that the Visual Application Manager displays.

Figure 1-2: Opsware Visual Application Manager Display

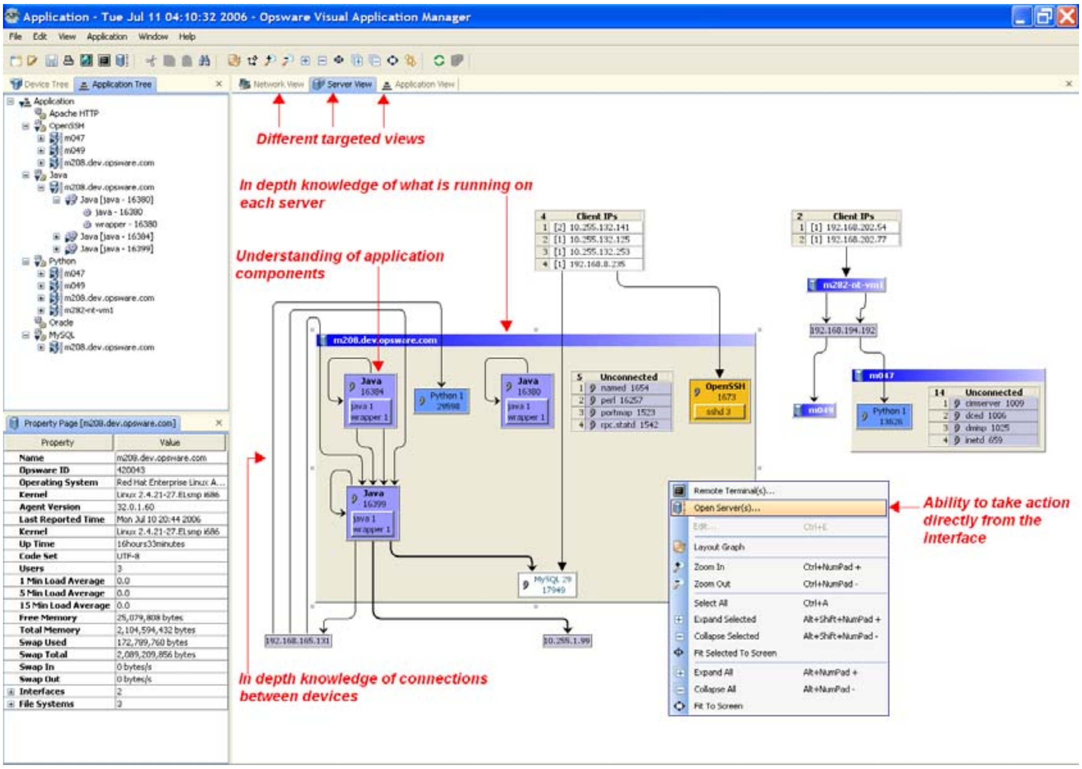

Different targeted views

In depth knowledge of what is running on each server

Understanding of application components

In depth knowledge of connections between devices

Ability to take action directly from the interface

Property Page [m208.dev.opsware.com]

Property	Value
Name	m208.dev.opsware.com
Opsware ID	420043
Operating System	Red Hat Enterprise Linux A...
Kernel	Linux 2.4.21-27.ELsmp i686
Agent Version	32.0.1.60
Last Reported Time	Mon Jul 10 20:44 2006
Kernel	Linux 2.4.21-27.EL smp i686
Up Time	16hours33minutes
Code Set	UTF-8
Users	3
1 Min Load Average	0.0
5 Min Load Average	0.0
15 Min Load Average	0.0
Free Memory	25,079,808 bytes
Total Memory	2,104,594,432 bytes
Swap Used	172,789,760 bytes
Swap Total	2,089,209,856 bytes
Swap In	0 bytes/s
Swap Out	0 bytes/s
⊞ Interfaces	2
⊞ File Systems	3

Supported Operating Systems

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

The Visual Application Manager collects data from managed servers running the following operating systems:

- Windows NT4
- Windows 2000, 2003 Server
- Windows x86_64
- Windows XP x64
- Windows Server 2003 x64
- SLES 9
- Redhat 7.3
- Redhat AS 2.1, 3, 4
- Redhat AS 4 x86_64
- Redhat ES 3, 4
- Redhat WS 3, 4
- Solaris 7, 8, 9, 10
- Solaris 10 x86_64
- AIX 5.1, 5.2, 5.3

Launching Visual Application Manager

You can launch the Visual Application Manager by selecting a server group, by selecting one or more managed servers, or by not selecting any servers.

✓ The Allow Analyze permission is required to use the Visual Application Manager. You also need read access to each managed server you plan to scan. Write access to each managed server is not required to run the Visual Application Manager; however, write access is required to perform any actions on the servers, such as when you use a Remote Terminal. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide*.

To launch the Visual Application Manager, perform the following steps:

1 Launch the Opsware SAS Client.

1 From the Navigation pane, select the Devices tab.

2 Select a server group, or one or more managed servers by performing one of the following steps:

- In the Navigation pane, select Device Groups ➤ Public and then select a public server group. Or, in the Contents pane, select one or more managed servers in that group.
- In the Navigation pane, select All Managed Servers and then select one or more servers in the Contents pane.

3 From the **Actions** menu or context menu, select **Visual Application Manager** to display a Scan Progress message.

After device scanning is completed, the Visual Application Manager window appears containing the Device Tree, Property Page, and Network View. See “Tree Tabs” on page 36, “Property Page” on page 45, and “View Tabs” on page 40.

Scan Time-Out Preference

The Visual Application Manager is optimized to scan a maximum of 50 servers. A number of factors affect the speed in which a scan is completed, including the load on the scanned servers and the load on Opsware. The default scan time-out is set to 200 seconds. You can reset this time-out value to a minimum of 30 seconds or to a maximum of 1,200 seconds.

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

1 From the **Application** menu, select **Set Scan Timeout**.

2 Move the slider to increase or decrease the number of seconds at which you want the scanning process to stop.

3 Click **OK** to save your changes or click **Cancel** to close the window without saving your changes.

Scan time-outs are recorded in .vam files, which means you can have different time-out settings for different applications.

The Toolbar for Visual Application Manager

The Visual Application Manager 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 Opsware Visual Application Manager

TOOLBAR ICON	DESCRIPTION
🗋	Opens a new .vam file (without scanning first).
📂	Opens a previously saved .vam file.
💾	Saves the current application (including topologies) as a .vam file in your local file system or in the Opsware Global File System (OGFS). If the application has not been previously saved, the Save As window displays.
📥	Imports an application definition from a .vam file.
📤	Exports an application definition to a .vam file.
🖨	Prints the selected view, tree, or property page. Displays the Print window where you specify page setup (including printing across multiple pages), a title for the printed view, and so on.
🌐	Opens a Global Shell session.

Table 1-1: Toolbar Icons in Opsware Visual Application Manager (continued)

TOOLBAR ICON	DESCRIPTION
▤	Opens the Terminal Launcher where you select a login ID for a Remote Terminal. When you select one or more managed servers or network devices, a separate dialog is opened for each to allow for different logins.
🖥	Opens the Object Browser Launcher where you select which servers can be viewed in the Server Browser. When you select one or more managed servers, a separate Server Browser is opened for each server.
✂	Removes (cuts) a selected application component or a selected tier in the Application Tree and saves it on the clipboard.
⎘	Copies a selected application component or a selected tier in the Application Tree and saves it on the clipboard.
📋	Pastes a tier or an application component that has been previously cut or copied to the clipboard. See "Pasting an Application Tier" on page 53 and "Pasting an Application Component" on page 56.
🖌	Redraws all components in the selected view. Components that have been manually revised will retain their sizing.
↻	Rotates the selected view, toggling it between a vertical and a horizontal orientation.
🔍+	Zoom in (enlarge the display size of) the selected view.
🔍−	Zoom out (reduce the display size of) of the selected view.

Table 1-1: Toolbar Icons in Opsware Visual Application Manager (continued)

TOOLBAR ICON	DESCRIPTION
⊞	Expands selected tiers in the Application Tree or closed folders in the selected view. Tiers are expanded recursively down to the application component that they contain. Managed servers underneath the application components are not expanded.
⊟	Collapses selected tiers in the Application Tree or closed components in the selected view.
⊡	Resizes the selected components in a currently active view to fit within the screen size.
⊞	Expands all tiers in the Application Tree or expands closed components in the selected view.
⊟	Collapses selected tiers in the Application Tree or opened folders in the selected view.
▭	Resizes all components in the currently active view to fit within the screen size.
⇆	Links trees and views so that elements selected in a view will cause the corresponding element in a tree to also be selected. The Device Tree, Network View, and Server View can be linked together so that selecting an object in one causes it be selected in all three. The Application Tree and Application View can also be linked together so that selecting an object in one causes it to be selected in the other.
🔄	Refreshes the view by collecting and displaying new information gathered.
🖥	Toggles the display of server names on process families in the application view.

Menus and Menu Options

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

View Menu

By default, the **Animate Layout** option is ON (preceded by a check mark). This causes the view 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 view will not be animated (objects are not displayed in motion) each time it is drawn.

Window Menu

The **New View** option enables you to clone a Property Page or any view. See "Property Page" on page 45.

Tree Tabs

The Visual Application Manager displays two tree tabs that show top level information about managed servers, network devices, and applications in your environment. This information is located in the Device Tree tab and in the Application Tree tab.

Device Tree

The Device Tree is a text-based tree view of the 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 View and Server View. The Device Tree contains servers and network devices as its top nodes. Below the servers are process families and extended process families. Below the network devices are VLANs and ports. See Figure 1-4

Figure 1-3: Device Tree

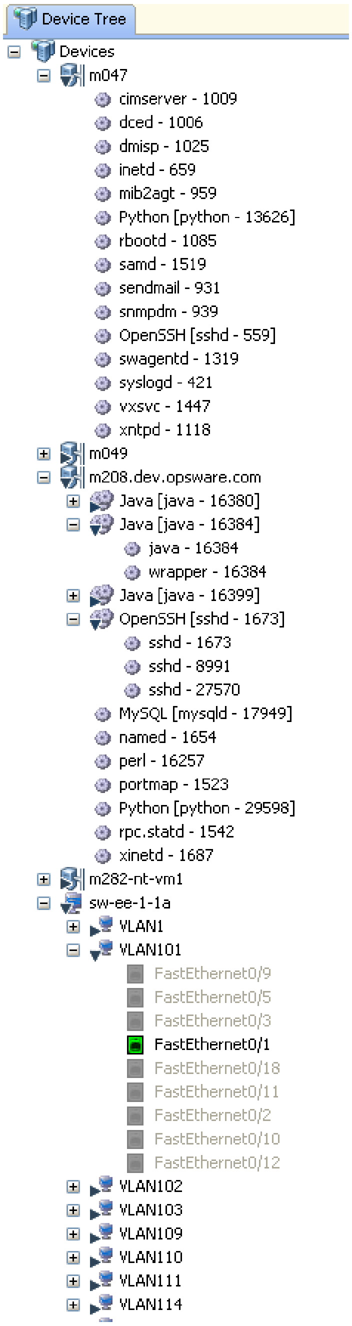

```
Device Tree
Devices
  m047
    cimserver - 1009
    dced - 1006
    dmisp - 1025
    inetd - 659
    mib2agt - 959
    Python [python - 13626]
    rbootd - 1085
    samd - 1519
    sendmail - 931
    snmpdm - 939
    OpenSSH [sshd - 559]
    swagentd - 1319
    syslogd - 421
    vxsvc - 1447
    xntpd - 1118
  m049
  m208.dev.opsware.com
    Java [java - 16380]
    Java [java - 16384]
      java - 16384
      wrapper - 16384
    Java [java - 16399]
    OpenSSH [sshd - 1673]
      sshd - 1673
      sshd - 8991
      sshd - 27570
    MySQL [mysqld - 17949]
    named - 1654
    perl - 16257
    portmap - 1523
    Python [python - 29598]
    rpc.statd - 1542
    xinetd - 1687
  m282-nt-vm1
  sw-ee-1-1a
    VLAN1
    VLAN101
      FastEthernet0/9
      FastEthernet0/5
      FastEthernet0/3
      FastEthernet0/1
      FastEthernet0/18
      FastEthernet0/11
      FastEthernet0/2
      FastEthernet0/10
      FastEthernet0/12
    VLAN102
    VLAN103
    VLAN109
    VLAN110
    VLAN111
    VLAN114
```

When you toggle the Link Views icon , selecting a server or process family in the Device Tree causes the corresponding element in the Network View and Server View to also be selected. Conversely, selecting a list node for a server in the Network View causes all corresponding elements in the Device Tree to also be selected. See Figure 1-4.

Figure 1-4: Device Tree With Link Views

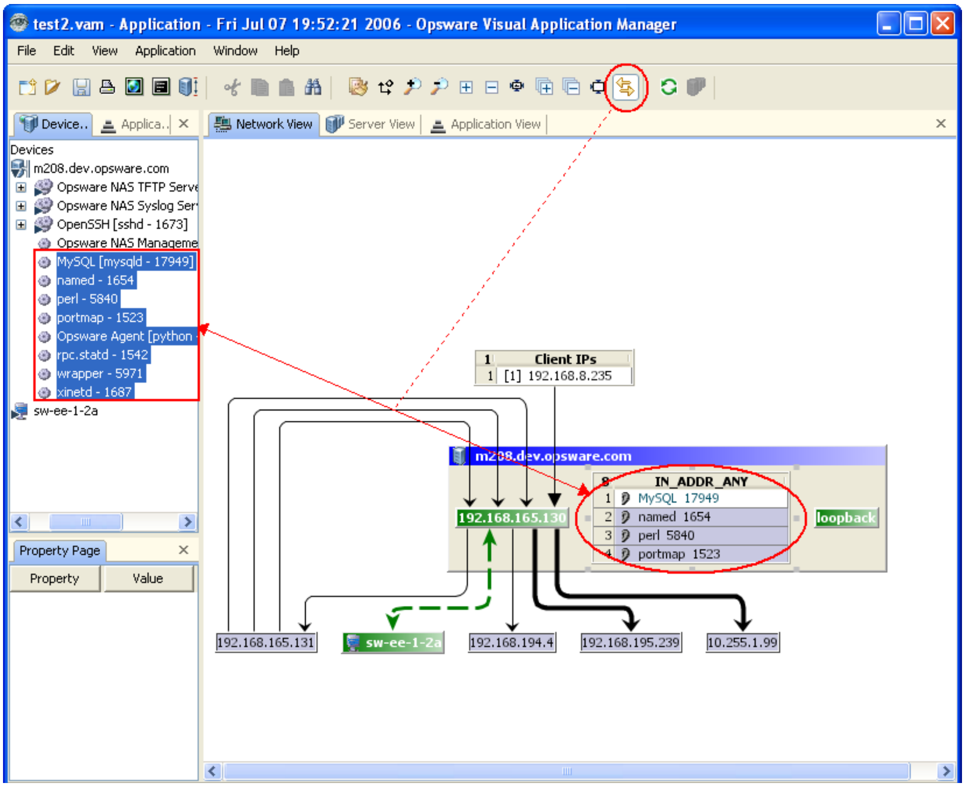

Application Tree

The Application Tree is a text-based tree view of the same information that is displayed in the Application View. The Application Tree contains tiers and subtiers, which in turn contain their associated application components. Below each application component are the process families that match the component and, therefore, appear in the different views. See Figure 1-5.

Application components 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 View–instead, they are highlighted as appropriate in the Network View and Server View.

Figure 1-5: Application Tree

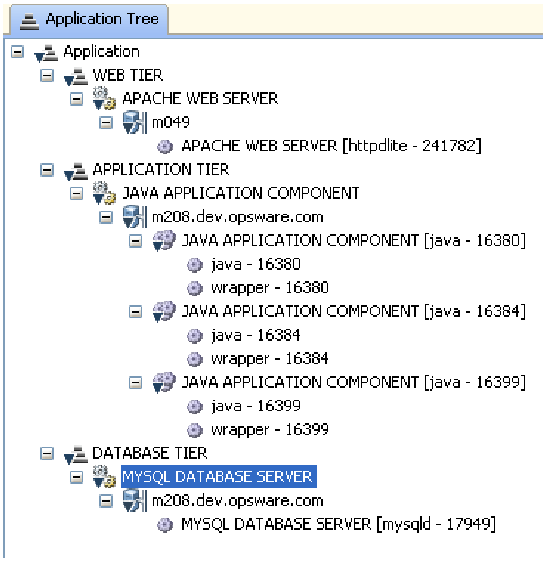

```
Application
  WEB TIER
    APACHE WEB SERVER
      m049
        APACHE WEB SERVER [httpdlite - 241782]
  APPLICATION TIER
    JAVA APPLICATION COMPONENT
      m208.dev.opsware.com
        JAVA APPLICATION COMPONENT [java - 16380]
          java - 16380
          wrapper - 16380
        JAVA APPLICATION COMPONENT [java - 16384]
          java - 16384
          wrapper - 16384
        JAVA APPLICATION COMPONENT [java - 16399]
          java - 16399
          wrapper - 16399
  DATABASE TIER
    MYSQL DATABASE SERVER
      m208.dev.opsware.com
        MYSQL DATABASE SERVER [mysqld - 17949]
```

When you toggle the Link Views icon ⇆, selecting an application tier or application component in the Application Tree causes the corresponding element in the application view to also be selected. See Figure 1-4 for an example of Link Views.

Tree Objects

in the Device Tree and Application Tree, the Visual Application Manager displays objects that represent process families, extended process families, application tiers, and application components.

Processes, Process Families, and Extended Process Families

In the Visual Application Manager, a *process* is a Unix or Windows process that 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 indicated by the ⚙ icon.

An *extended process family* is a set of processes that the Visual Application Manager has heuristically computed to be related, but are not necessarily members of the same process hierarchy. An extended process family is indicated by the icon.

Application Tier

An application tier is 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.

An 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 indicated by the icon.

An application tier object consists of:

- A set of application components
- Optional sub-tiers

Application Component

An application component is an object that represents a a process of running software, such as Apache, Oracle, BEA WebLogic, Microsoft® SQL Server, and so on.

An application component object consists of:

- A signature, which is a set of rules that you provide and that Visual Application Manager uses to identify a process family. These rules use data such as process name, open files, command line, executable path, listen port. See "Adding an Application Component" on page 54.
- Preferences that specify the alias of the application component as displayed in the different views, the background and foreground text color displayed in the different views.

An application component is indicated by the icon.

View Tabs

The Visual Application Manager displays three view tabs that show physical and logical drawings of managed servers, network devices, and connections in your environment. This information is located in the Network View tab, Server View tab, and Application View tab.

You can export a view to a .gif, .jpg, and .svg file for use in other applications where you can annotate the drawing or view the exported file in a Web browser. See “Exporting a View to GIF, JPEG, or SVG” on page 49.

You can also print a view on single and multiple sheets of paper. See “Printing a View” on page 50.

Network View

The Network View is a physical layout that shows how elements of an application interact with the network, including the network interfaces on a server and the layer 2 devices (switches) to which the server is connected. This view shows which process families are connected over which network interfaces in a server, and the switch port, VLANs, and switches to which a server’s network interfaces are connected. A process family that is listening on or connected to more than one network interface appears multiple times in the Network View.

The Network View also shows which external IP addresses (Client IPs) are connected to the application and which external IP addresses (external dependencies) the application connects to and, therefore, depends on.

In the Network View, the Visual Application Manager highlights layer 2 connections that have speed or duplex mismatches. The Network View shows servers that were not scanned but have layer 1 connections. Network devices without connections are not shown. See Figure 1-6.

A network device that only appears in the Device Tree (and not in the Network View) has a warning icon to indicate that it is unknown whether the network device is connected to the servers. See "Error Messages" on page 60.

Figure 1-6: Network View

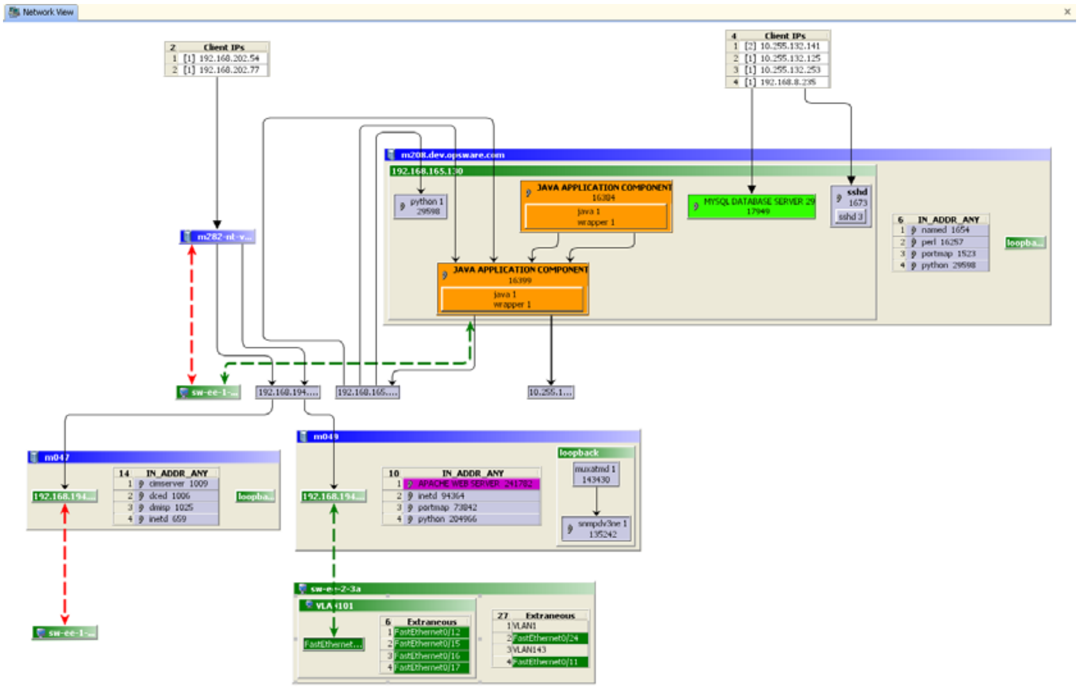

Server View

The Server View is a physical layout that shows how elements of an application map to a set of servers. This view shows which process families are running on which server, and how the processes families are connected to one another. The Server View also shows which external IP addresses (Client IPs) are connected to the application and which external IP addresses (external dependencies) the application connects to and, therefore, depends on. See Figure 1-7.

Figure 1-7: Server View

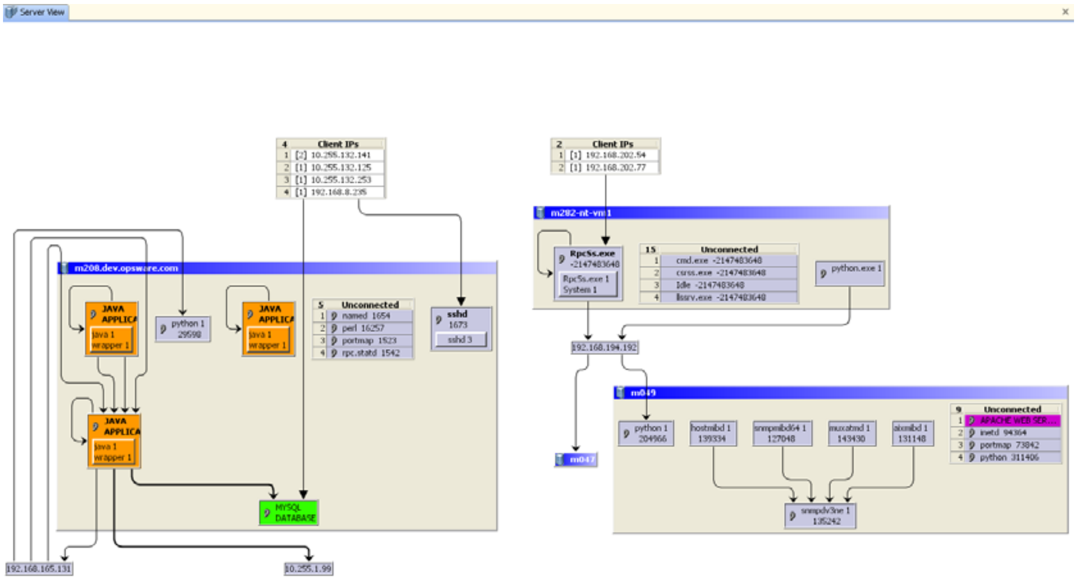

Application View

The Application View is a logical layout that shows the logical structure of an application, including tiers, application components, connections between application components, external clients, and dependencies. You create this view by defining tiers that comprise an application. See "Adding an Application Tier" on page 51.

The Application View also shows which external IP addresses (Client IPs) are connected to the application and which external IP addresses (external dependencies) the application connects to and, therefore, depends on.

Application components can be attached to a tier and used to recognize process families as named elements of an application. 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 View

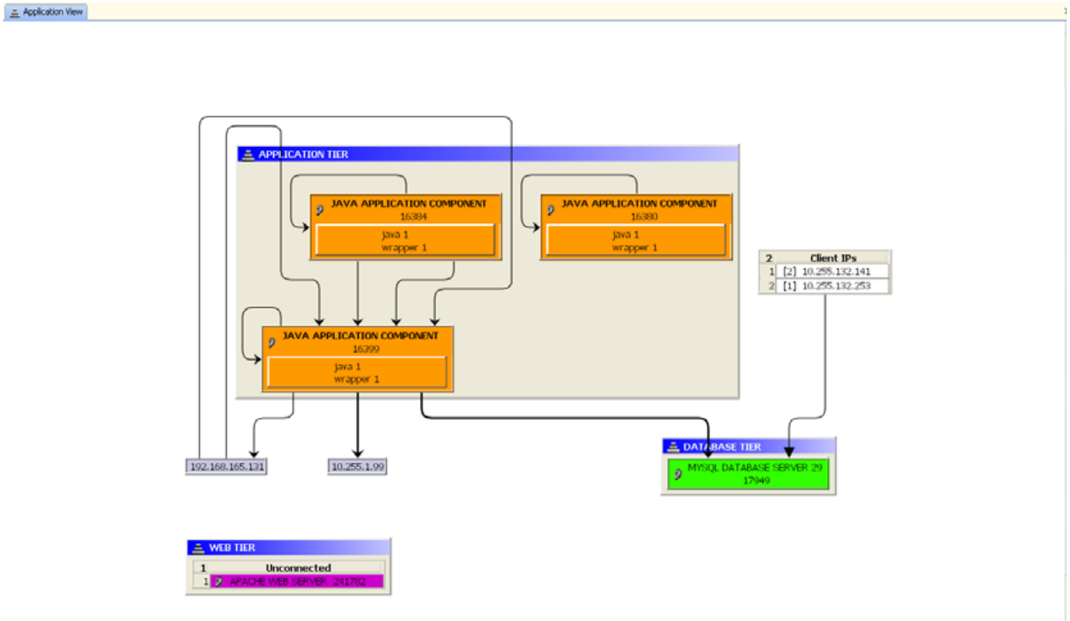

2	Client IPs
1	[2] 10.255.132.141
2	[1] 10.255.132.253

1	Unconnected
1	APACHE WEB SERVER 241782

Symbols Used in Views

The Visual Application Manager uses a variety of symbols in the Network View, Server View, and Application View, such as lines, arrows, diamonds, and so on.

Lines

In the Visual Application Manager views, 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. Processes providing a network service, such as listening for network connections, and processes that have a connection to another process or server are displayed.
- **Layer 2 connection**: A virtual link between servers' network interfaces and switch ports/switches. A layer 2 connection is indicated by a green line in the view.

The thinness or thickness of the line represents the number of connections associated with the link, where a smaller number is indicated by a thinner line and a larger number is indicated by a thicker line.

Arrows

In the Visual Application Manager views, arrows represent the following connections between devices:

- **Solid Arrow**: A solid arrow indicates an inbound or outbound client connection, such as TCP or UDP.
- **Hollow Arrow**: A hollow arrow indicates an inbound or outbound connection.
- **Diamond**: A diamond symbol represents a physical, layer 2 network connection.
- **Color**: You can specify a background or foreground color to identify application components in the views.

Property Page

The Visual Application Manager displays a Property Page for the component selected in the Device Tree, Application Tree, Network View, or Server View. Depending on which object is selected, 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. All MAC addresses seen by each port are also included in the Property Page.

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. See Figure 1-9.

Figure 1-9: Property Page for a Server

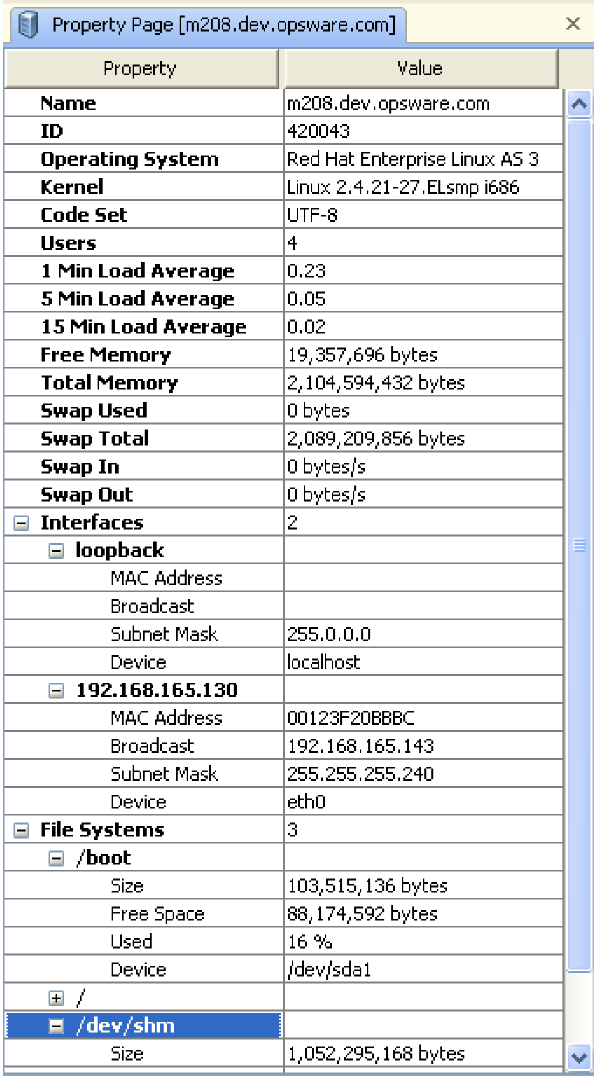

Property Page [m208.dev.opsware.com]

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
Free Space	88,174,592 bytes
Used	16 %
Device	/dev/sda1
⊞ /	
⊟ **/dev/shm**	
Size	1,052,295,168 bytes

Server Property Page

The Property Page for a server displays the following information:

- **Name**: The host name of the server.
- **ID**: The Opsware unique identifier for the server.

- **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 0 (zero) because it is not natively tracked.
- **Memory usage**: The total, free, used memory.
- **Swap usage**: The total, free, 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, bytes free, percent used, and associated device for each file system.

Network Device Property Page

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

- **Speed**: The negotiated speed or configured (if it cannot be collected).
- **Duplex**: The negotiated duplex setting or configured (if it cannot be collected).
- **MAC address**: The Media Access Control ID of the device.

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.
- **Server**: The server on which the process is running.
- **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, 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 .so's on Unix.
- **Open files**: The files the process family currently has open.
- **Packages**: The packages associated with the files the process family has open.
- **Processes**: The number of 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, 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 percentage of CPU consumed by the process since it began.
  - **Environment**: The name and value of each environment variable in the process's 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.
- **Number of subtiers**: The number of subtiers that are currently recognized in the tier.
- **Number of application components**: The number of application components that are currently recognized in the tier.
- **Servers**: The servers that are associated with this tier. Only matching servers are searched for matching application components.

Application Component 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**: The name of the application component as displayed in the different views.

- **Number of process families recognized**: 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.
- **Listen port**: The listen port used to recognize this application component.
- **Connected to port**: The port this server is connected to.
- **Executable path**: The executable path filter used to recognize this application component.
- **Open files**: The open file filter used to recognize this application component.
- **Background color**: The background color displayed in the different views.
- **Foreground color**: The foreground text color displayed in the different views.

Link Property Page

The Property Page for a link displays the following information:

- **Protocol**: 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.

Exporting a View to GIF, JPEG, or SVG

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

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

1. From the **File** menu, select **Export** to display the Export Graph window.
2. Select a file system directory that identifies 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 **Export**.

Printing a View

You can print a view on single and multiple sheets of paper.

To print a view, 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 view.
3. Click **Print**.

Setting Up 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 data that is relevant to the application of interest. Based on the application tiers and application components you create, Visual Application Manager recognizes and allows you to change the display.

To optimize a meaningful data display, you set up application definitions to:

- Recognize processes as certain application components by giving them meaningful names and appearances (colors).
- Define the logical tiers of an application and display application components according to the tier in which they reside.

See "Evaluation Order" on page 53.

Default Application Definitions

The Visual Application Manager includes a default application that contains no tiers. The default application contains predefined application component signatures that recognize a variety of commonly used application components.

To view the application definition and components, perform the following steps:

1. In the Visual Application Manager window, select the Application Tree to display the default application definition and components.
2. Select a component you want to see (or edit) the information for.
3. Right-click and then select **Edit** to display the Application Component window.

4. View or edit the information and then click **Apply** to save your changes and close the window.

Saving a New Default Application Definition

If you have made changes to the application definition and want to save this as the default, perform the following step:

- From the **File** menu, select **Save As Default**.

Resetting the Default Application Definition

If you have made changes to the application definition and want to restore the previously saved default application, perform the following step:

- From the **File** menu, select **Reset Default Application**.

Application Tiers

Application tiers provide an architectural framework for the application through which application components are organized and displayed. 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 in which process families are assigned to application components as they are recognized. See “Evaluation Order” on page 53.

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

Adding an Application Tier

To add an application tier to 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

6. 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

6. 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**.

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 application components they contain) is significant because it affects the order in which process families are assigned to application components as they are recognized.

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 an application component in the Device Tree, the Paste icon will be disabled.
- Select a tier in the Device Tree and then select the Paste Before action 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, the Paste Before menu action will be disabled.

Application Components

Application components are organized and displayed in application tiers. You can add, edit, delete, cut, copy, and paste application components in the Application Tree. You can paste an application component before or after a selected position in the Application Tree to rearrange the order. The order of application components (and the tiers that contain them) is significant because it affects the order in which process families are assigned to application components as they are recognized by their signature. See “Evaluation Order” on page 53.

Evaluation Order

The order in which application components are recognized is important because a process family is associated with the first application component it matches. Therefore, evaluation order is significant when the recognition criteria for an application component matches the same process family found in multiple application components.

Application components are always evaluated in a predictable, easy-to-understand order. A tier's sub-tiers are evaluated before the tier's application components. The tree of application tiers is searched in this depth-first order by applying application components in the order in which they appear in each tier.

Consider an application definition that has the structure shown in Figure 1-10.

Figure 1-10: Application Component Structure

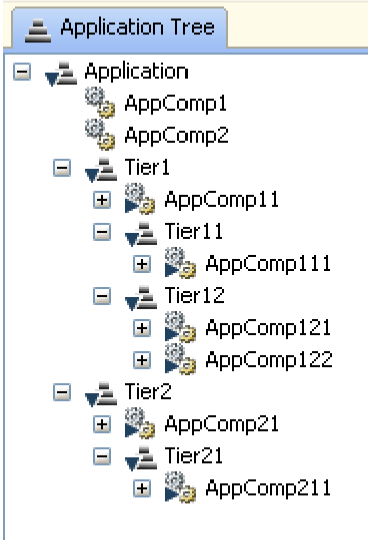

In this application definition example, the application components are evaluated in the following order:

1. Application Component 111

2. Application Component 121

3. Application Component 122

4. Application Component 11

5. Application Component 211

6. Application Component 21

7. Application Component 1

8. Application Component 2

Adding an Application Component

To add an application component to the Application Tree, perform the following steps:

1. Select the Application Tree.

2. From the **Application** menu, select **New Application Component** to display the Application Component window.

3. In the Signature section, enter the following information:
    - **Process Name**: The name of the process family.
    - **Command Line**: The command line that an application component was started with.
    - **Executable Path**: The executable path filter used to recognize 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 .so's on Unix.
- **Connected to Port**: The port the application component is connected to.
- **Listener Port**: The port on which the application component is listening.

4. In the Preferences section, enter the following 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.
5. Click **Apply** to save your changes or click **OK** to save your changes and close the window.

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

Editing an Application Component

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

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

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

Deleting an Application Component

To delete an application component 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 an Application Component

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

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 an Application Component

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

- Select an application component 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 application components. When you select an application component in the Device Tree, the Paste Before menu action will be disabled.
- Select an application component in the Device Tree and then select the Paste Before action from the Edit menu. The application components that you cut or copied to the clipboard will be inserted into the selected application component's parent tier *before* (above) the selected application component.
- Select an application component in the Device Tree and then select the Paste icon 📋. The application components that you cut or copied to the clipboard will be inserted into the selected application component's parent tier *after* (below) the selected application component. 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 application components in the Device Tree, or if you do not select any tiers or applications in the Device Tree, the Paste icon and Past Before menu action will be disabled.

Topology Management

In the Visual Application Manager, a topology is 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. A topology is part of a .vam file. See ".vam Files" on page 29.

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

Opening a .vam File

After you have launched the Visual Application Manager, you can open a previously saved .vam file.

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

1. In the Opsware Visual Application Manager 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 you want to open. In the Application Tiers pane, you can preview the basic tier structure of the application defined.
4. In the Topologies pane, select one or more topologies in the .vam file that you want to open. The default topology is the last one saved.
5. Click **Open**.

Editing a .vam File

After you have launched the Visual Application Manager, you can open a previously saved .vam file and then edit it.

To edit an existing .vam file, perform the following steps:

1. In the Opsware Visual Application Manager 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 or enter the file name.
3. In the left pane, double-click on the .vam file you want to edit. In the Application Tiers pane, you can preview the basic tier structure of the application defined.
4. In the Topologies section, select the topology you want to edit. The default topology is the last one saved.
5. Click **Open**.
6. After you make your changes, select the 💾 toolbar icon or select the **File** menu and then select **Save** to save your changes.

Saving a .vam File

You save a .vam file by using the **Save** and **Save As** options in the **File** menu.

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

1. From the **File** menu, select **Save** or **Save As** to open the Save window.
2. 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 Topologies pane, select the topologies you want to save in the .vam file. Topologies that are not selected will be deleted from the .vam file.

   A topology is identified by its timestamp that was generated when it was recorded. If the last topology 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 Visual Application Manager Archives (*.vam).
6. Click **Save**.

If you exit Visual Application Manager 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.

Analysis and Action

From the Visual Application Manager, you can launch other Opsware SAS Client features, such as Server Explorer, Network Device Explorer, Global Shell, and Remote Terminal to perform in-depth analysis or to perform actions on the devices under investigation. To help with troubleshooting and remediation tasks, you can use these features to stop processes, to start and stop services, and to perform other troubleshooting or automation tasks.

Server Explorer

You can use the Server Explorer to view detailed information about a server.

1. In a view, select one or more servers.
2. Right-click and then select **Open Server** to open the Object Browser Launcher.
3. Select one or more servers and then click **OK** to open a Server Explorer for each server.

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

Global Shell

You can use the Global Shell feature to navigate between servers and connected network devices by tracing their layer 2 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 *Opsware® SAS User's Guide: Server Automation* for information about how to use Global Shell.

Remote Terminal

The Remote Opsware Shell (`rosh`) utility enables you to log in to devices (servers and network devices) and run native commands. You invoke `rosh` from within a Global Shell session. You can run `rosh` and enter native commands interactively, or you can specify the native commands as an option of `rosh`. For example, you can `rosh` on to a switch and run the `show vlan` command to view all VLAN details.

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

- In a view, select one or more servers.
- Right-click and then select **Remote Terminal** to open the Terminal Launcher.
- Select at least one server and login ID and then click **OK** to open a Remote Terminal for each selection.

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

Error Messages

The Visual Application Manager indicates when an error occurred on a managed server by displaying a server error icon ⊗ or a server unreachable error icon ⊗ before the server name in the Device Tree, Network View, and Server View. Move your cursor over the server name to display the detailed error message.

Scan failures and scan time-outs can typically occur when the Opsware 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 Opsware administrator for assistance.

Table 1-2 describes these errors and recommended actions.

Table 1-2: Error Messages in the Visual Application Manager

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 32.
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 Opsware administrator for the required permissions.
Server Capture Failed	The remote capture of data or the transfer of data back to the Opsware core failed.	Review the log file that is in /tmp/.sitemap/<number> for details.
Server ID Invalid	The server's directory was not found in the OGFS, which means that Opsware 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 Opsware Support and provide the log file.
Server Unreachable	The managed server is unreachable by Opsware SAS. This could be caused if the server cannot connect to the Opsware core.	Try again later. If this condition persists, contact your Opsware administrator.
Unknown Scan Error	An unknown error occurred during the scanning process.	Try again later. If this condition persists, contact your Opsware administrator.

Table 1-2: Error Messages in the Visual Application Manager (continued)

ERROR	DESCRIPTION	ACTION
Unsupported Agent for Scan	The Visual Application Manager does not support the Opsware Agent version running on a selected managed server.	Opsware Agent 5.1 or higher is required.
Unsupported OS for Scan	The Visual Application Manager does not support the operating system running on a selected managed server.	See "Supported Operating Systems" on page 31.

Chapter 2: Audit and Remediation

IN THIS CHAPTER

This section discusses the following topics:

Overview of Audit and Remediation

Audit and Remediation allows you 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.

With Audit and Remediation, you can audit a server configuration values based upon a live server (or server snapshot), or based upon your own custom values (or both). 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 create custom audit policies that define company or industry server configuration compliance standards, and which can be used inside of audits or snapshot specifications.

Audit and Remediation Examples

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

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

Capturing a Golden Server Configurations

Sometimes a server becomes configured in such a way that is 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, as a systems administrator you might configure a single server in such a way that represents prefect – a golden server configuration – for a group of Web servers. After you configure this golden server, you would like to duplicate the golden server configuration across a group of servers.

For example, you have a Red Hat Linux server configured with a unique configuration of Apache Web Server, and you wants to duplicate this exact configuration across several other servers. With Audit and Remediation, you could do this by creating an Audit that uses the "golden" server as the source. In the audit, you select those configurations that you would like to be used to audit other servers, such as an application policy and specific application configuration rules. Then you select as the target of the audit those servers to be configured like the golden server. After you runs the audit, you can remediate those target server's whose configurations 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

All IT organizations have security policies they want to enforce, to make sure servers are configured properly and are safe from security attacks. For example, you might want to be sure that a collection of Windows 2003 server have a recent security patch sent out by Microsoft, regardless of the types of applications that are installed on the servers. You can create an audit policy that defines this security configuration, let the systems administrators who directly configure and manage those servers know that this policy exists. The systems administrator would then create an audit and link it to the audit policy you created 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 remediate to have the patch installed. If new patches come out and need to be installed on the target servers, you would update the audit policy with the new patch, and the audit that runs against the target servers would automatically be updated to reflect the new patch definition.

Creating Your Own Ad-Hoc Audit

As a system administrator, your job might be to 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 the run the application, you keep a list that tracks the ideal state of 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 could 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 successful according to your criteria. If something goes wrong and you get paged at four in the morning to fix the problem, you can create an ad hoc audit on the fly to troubleshoot something related to t problem, and 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 desired configuration values for a server, compare expected configurations against live servers, and remediate any differences found by the audit. Through 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 "About Audits" on page 70.

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 which can be saved as an audit policy to help define compliance best practices for use by others to use for running audits. Audit policies can be linked to audits or snapshot specifications, which can maintain the latest changes made to the audit policy it is based upon.

For more information on audit policies, see "Audit Policies" on page 118

Audits and the Compliance Dashboard

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

- A roll up of scheduled audits will appear by default in the Compliance Dashboard in a single column. This status enables you to view at a glance the compliance status of all audits 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 you do not have access to will not be represented in the Compliance Dashboard in the audit roll up.
- Individual audits that have been scheduled can be displayed in 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 on the Compliance Dashboard, see Chapter 3, “Compliance Dashboard”.

Audit Reports

The Opsware SAS Client provides rich reporting capabilities for Audit and Remediation, including such reports as:

- Audit Policy Compliance (All Servers), displaying all servers grouped by their Server Configuration Policy Compliance Level (compliant, non-compliant, scan needed, and so on)
- Audit Policy Compliance by Customer, which displays the Server Configuration Policy Compliance Levels of all servers, grouped by Customer.
- Audit Policy Compliance by Facility, which displays Server Configuration Policy Compliance Level of all servers, grouped by Facility

Reporting for Audit and Remediation also includes several Sarbanes/Oxley (SOX) reports. For more information on how to run and view reports in the SAS Client, see “Reports” on page 175.

Snapshots

snapshots differ from an audit in that the snapshot allows you to take a picture of the current state of configuration of a server. The snapshot is 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 and 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 “About Snapshots” on page 135.

Terms and Concepts

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

- **audit**: A set of rules that express 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 one. 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 one time, 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 rules listed.
  - **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 define 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, which will show how a target server or groups of servers's 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 Audit and Remediation is defined by the audit's or snapshot's rules, which specifies the values expected to be on target servers. An audit result shows differences between what the user has defined as what should be the proper configuration of server objects would be considered out of compliance or non-compliant.
- **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.
- **server object**: An object from a server to which a 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 82.)
- **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 an snapshot specification job is completed, it produces a snapshot.
- **snapshot specification**: And object window that allows a you to define and create a snapshot. In other words, defines 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 be other servers.

About Audits

An audit consists of a collection of rules that enable you to define what should be or what shouldn't be for a server's configuration. And 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 two types of comparisons, based on the source of the audit:

- **Server Comparison**: An audit based upon configuration values from a source server or source snapshot specified at the time 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 upon 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

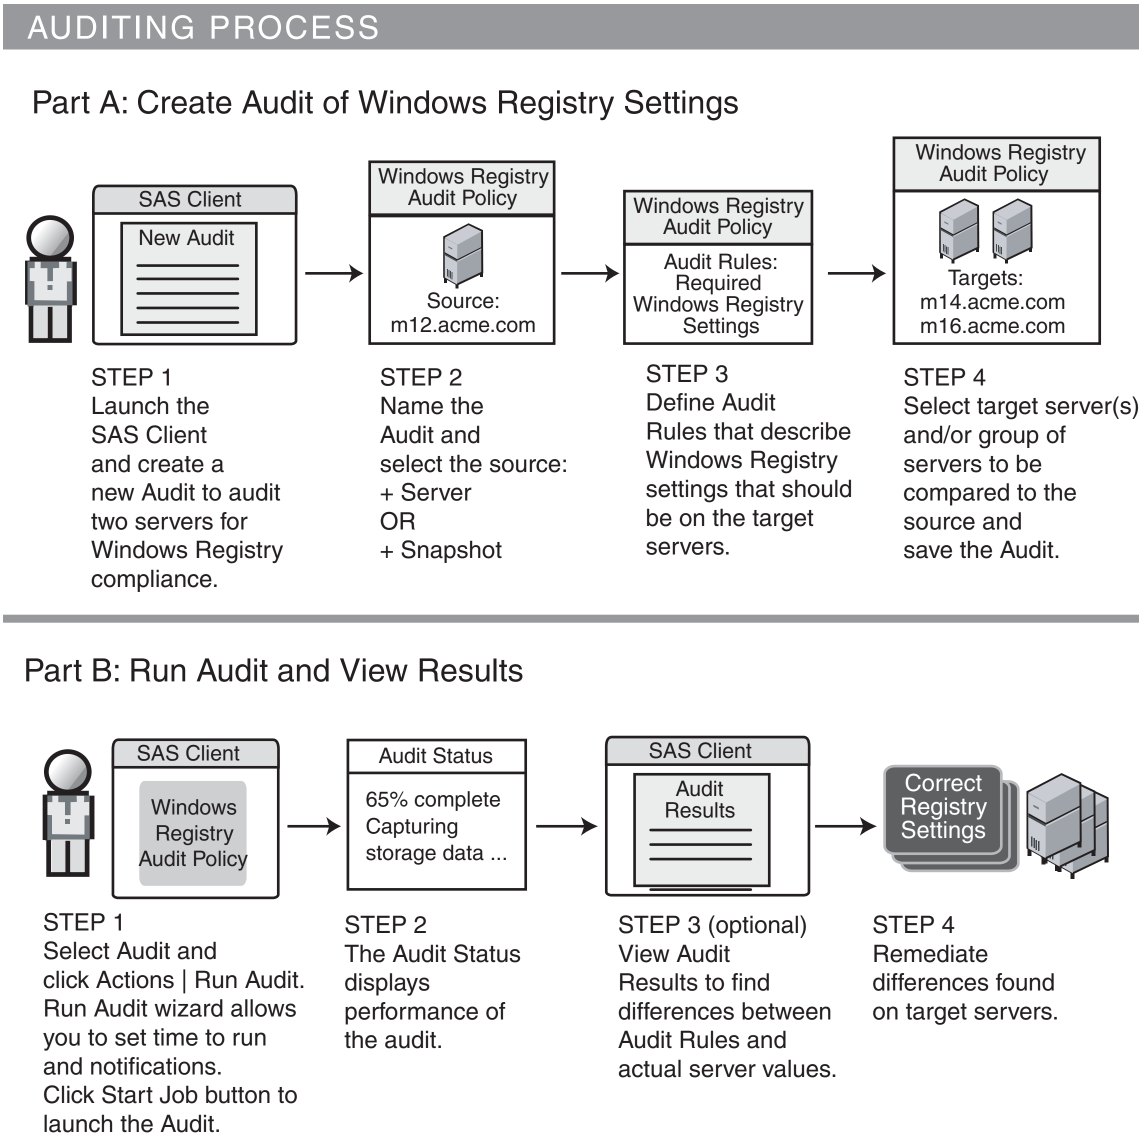

AUDITING PROCESS
Part A: Create Audit of Windows Registry Settings
SAS Client
New Audit
Windows Registry Audit Policy
Source: m12.acme.com
Windows Registry Audit Policy
Audit Rules: Required Windows Registry Settings
Windows Registry Audit Policy
Targets: m14.acme.com m16.acme.com

STEP 1
Launch the SAS Client and create a new Audit to audit two servers for Windows Registry compliance.

STEP 2
Name the Audit and select the source:
+ Server
OR
+ Snapshot

STEP 3
Define Audit Rules that describe Windows Registry settings that should be on the target servers.

STEP 4
Select target server(s) and/or group of servers to be compared to the source and save the Audit.

Part B: Run Audit and View Results
SAS Client
Windows Registry Audit Policy
Audit Status
65% complete Capturing storage data ...
SAS Client
Audit Results
Correct Registry Settings

STEP 1
Select Audit and click Actions | Run Audit. Run Audit wizard allows you to set time to run and notifications. Click Start Job button to launch the Audit.

STEP 2
The Audit Status displays performance of the audit.

STEP 3 (optional)
View Audit Results to find differences between Audit Rules and actual server values.

STEP 4
Remediate differences found on target servers.

Audit Elements

An audit consists of the following elements:

- **Properties**: Name and description of the audit.
- **Source**: The source of an audit can be a server, a snapshot, or a 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: 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 you can define rules for in an audit, see "Server Objects Used in Audits and Snapshots" on page 82.
- **Targets**: The servers to audit – that is, check to see if how the server configuration compares with the configuration defined in the audit. You can choose as many servers and groups of servers to audit or snapshot.
- **Schedule**: You can run the audit on a one time 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 upon success, failure, or simply the completion of the audit Job.

To configure an audit, you select server configuration objects that you want to check the values of, and then apply rules to those objects in order to define their desired configuration state. For example, Figure 2-2 shows an audit that has defined three rules that will determine if any target servers match the rules set for event logging, operating system, and windows services.

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

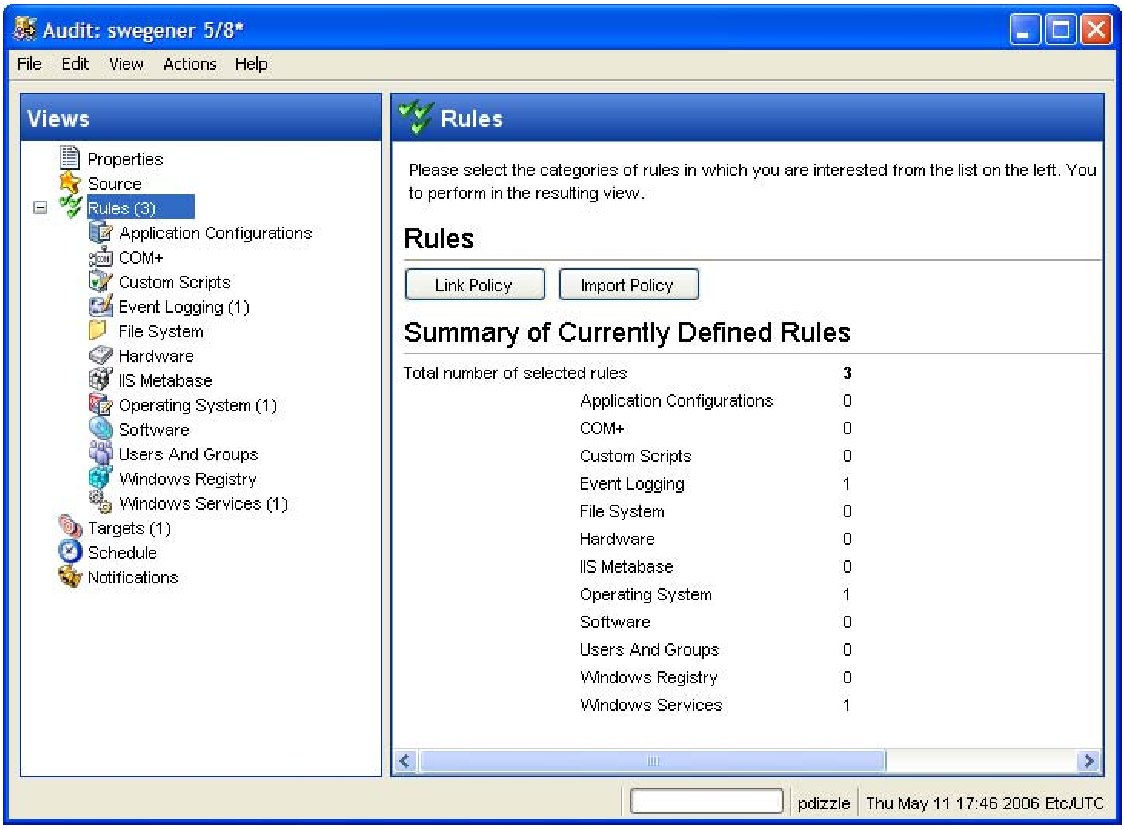

Please select the categories of rules in which you are interested from the list on the left. You to perform in the resulting view.

Rules

Summary of Currently Defined Rules

Total number of selected rules	3
Application Configurations	0
COM+	0
Custom Scripts	0
Event Logging	1
File System	0
Hardware	0
IIS Metabase	0
Operating System	1
Software	0
Users And Groups	0
Windows Registry	0
Windows Services	1

Ways to Create an Audit

You can create an audit from several locations inside the SAS Client, depending upon your purpose: Do you want to audit a specific Server? Do you want to audit a group of servers? Do you want to rerun an audit that has already been run? Do you want to run an audit from a snapshot?

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 to create a new audit
- From a snapshot, to create an audit based upon the snapshot
- From an audit policy, to create an audit based upon 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 wish, 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, 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 the group to which you have access.

To create an audit of 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 to the group of servers hierarchy until you see the group of servers you would like to audit in the Contents pane on the right.
3. Select the group of servers you want to audit 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

If you want to create a new audit and set all your own parameters, 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 Audits, then Windows or Unix.
3. Click once 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 upon the server configuration information 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 once 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 the snapshot you would like 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. This creates a connection between the audit policy and the audit, so if any updates are made to the policy, those changes are automatically reflected in the policy.

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

Viewing Server Audit Usage

Once you have created and run audits, it is useful to understand from a server’s perspective where these audits are being used. 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 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 Contents 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 Contents 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 Contents 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 requires performing the following general steps:

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

To configure an audit, perform the following steps:

1 Create the new audit from anyone of the methods for creating an audit listed at “Ways to Create an Audit” on page 73.

2 In the Audit Window, you can now begin to define the parameters of your audit.

3 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. Since each audit rule is unique and requires its own instructions, for information on how to configure individual audit rules, see "Configuring Audit and Remediation Rules" on page 84.

  If you wish 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 118.

- **Targets**: Choose the Targets of the audit. These are servers, groups of servers, or snapshots 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**.

- **Schedule (Optional)**: Choose when you would like to run the audit, or on a recurring schedule. Choose whether you want to run the to be created 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 one time 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 which day of the week you want the audit to run.

  - **Monthly**: Choose which months you would like to audit run during.

  - **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.

- **Time and Duration:** For each type of schedule, specify the hour, minute, and day of the week, and month for the daily schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose an end date to end the audit schedule, select End and 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 the email addresses of people you would like to receive an email when the audit Job finishes running. You can choose to send the email notification to be sent on both 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.)

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

Audit and Remediation Rules

The purpose of an audits is to enable you to determine how your servers are configured, and if those servers are configured the way you want them to be – that is, configured 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 to 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 82.)

In an audit and audit policy, you can also define what, if any, remediation value you would like the object to have if a server object is found to be different than what you define as 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 you want to check (the specific files or directories you want to check), and the desired state of the object, and a remediation value should the server configuration differ form the audit rule (optional).

An audit rule consists of the following components:

- **Server Object**: 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 thing you can check about it; for example, for the windows services server object, you might be want to know if a specific service exists on a target servers and, and whether or not the service is enabled or disabled.

- **Target Value**: What should be for 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**: The value you want to change for the server object if found to be different from how it is defined in the audit. 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

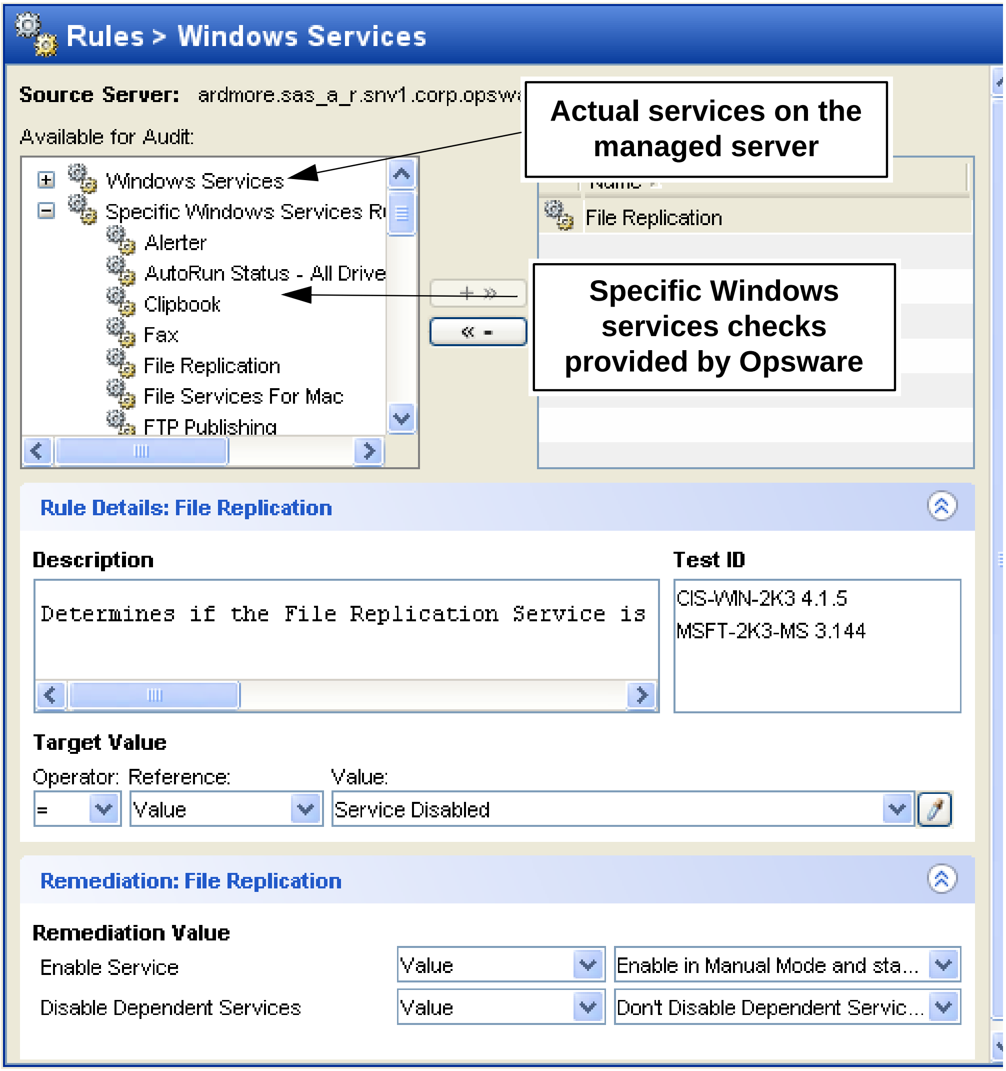

Source Server: ardmore.sas_a_r.snv1.corp.opsw

Rule Details: File Replication

Description

Determines if the File Replication Service is

Test ID

CIS-WIN-2K3 4.1.5
MSFT-2K3-MS 3.144

Target Value

Operator:	Reference:	Value:
=	Value	Service Disabled

Remediation: File Replication

Remediation Value

Enable Service	Value	Enable in Manual Mode and sta...
Disable Dependent Services	Value	Don't Disable Dependent Servic...

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

- **Available for Audit**: Lists all services from the source server available to be added to the audit, plus specific Windows services checks provided by Opsware.
- **Selected for Audit**: The service name File Replication has been chosen.
- **Description**: Describes what is being checked for the rule on the target server. In this case, the audit will check to see if the service is enabled or disabled, as the

description mentions that CIS and Microsoft recommends that this service be disabled.

- **Target Value**: This is the value you want to compare against the target server. For 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 upon 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 was exists on the target server.

- **Remediation Value**: The remediation value determines what action to take if the service on the target server does not match the value you defined in the audit. Remediate values can derive 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 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 about the output against values set in the audit. (Python 1.5.2 only for python scripts.)	Live

Table 2-1: Audit and Remediation Server Objects

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

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

Audit and Remediation does not support device files, and sockets or industry standards.

Configuring Audit and Remediation Rules

Creating an audit (or snapshot specification) requires configuring 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 files on a servers file system, what specifically you want to capture, such as a Windows NT file's Access Level Controls, for an application, which application configuration values you would want to snapshot or audit, plus any remediation values you want to specify if the differences are discovered between the rule and the actual value on the target server.

For example, 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 what the remediation value for the server object should be if the value defined in the audit or snapshot differs from what is found on the server after the audit has run. For example, if a Windows Service is found to be 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 128.

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 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, Audit and Remediation rules for the Software server object evaluate what patches or packages 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

For example, 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 and Event Logging rule that will check the maximum application event log size. The Center for Internet Security (CIS) and Microsoft recommends that this value be set 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 remediate 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

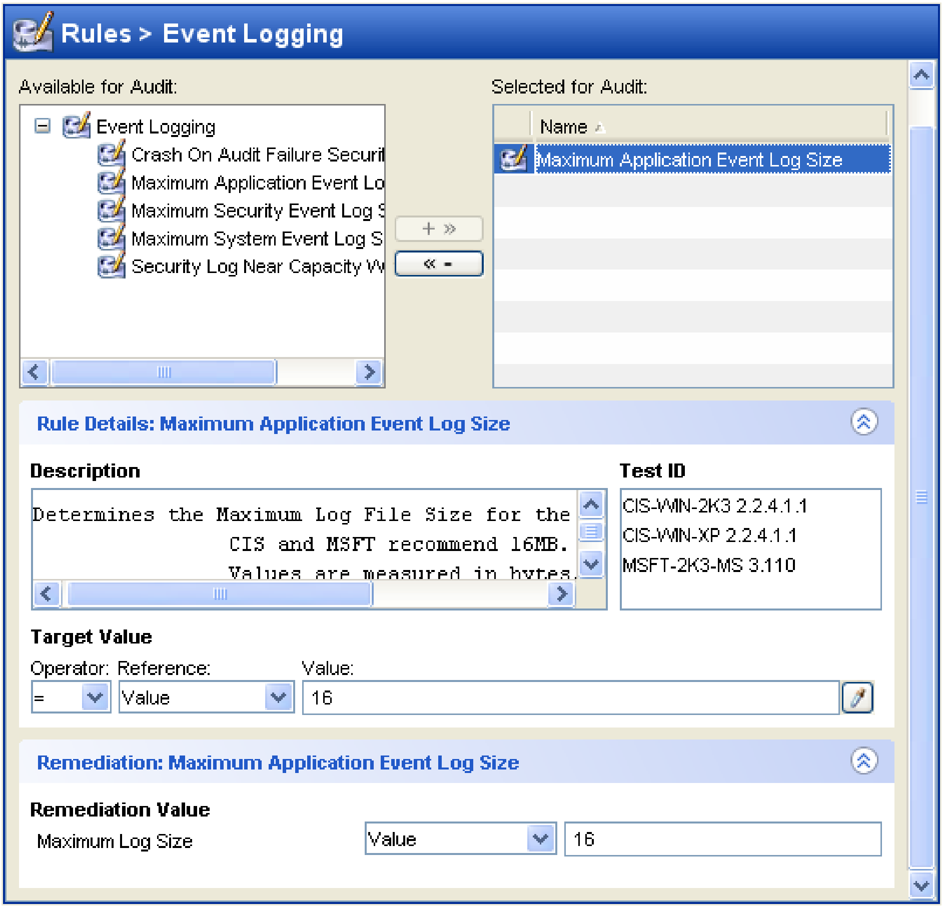

Rules > Event Logging

Available for Audit:

- Event Logging
  - Crash On Audit Failure Securit
  - Maximum Application Event Lo
  - Maximum Security Event Log S
  - Maximum System Event Log S
  - Security Log Near Capacity W

Selected for Audit:

Name
Maximum Application Event Log Size

Rule Details: Maximum Application Event Log Size

Description

```
Determines the Maximum Log File Size for the
         CIS and MSFT recommend 16MB.
         Values are measured in bytes.
```

Test ID

CIS-WIN-2K3 2.2.4.1.1
CIS-WIN-XP 2.2.4.1.1
MSFT-2K3-MS 3.110

Target Value

Operator: = Reference: Value Value: 16

Remediation: Maximum Application Event Log Size

Remediation Value

Maximum Log Size Value 16

In this example, 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 named Available for Audit shows all Event logging objects available from the source server to select to add to the audit, and the top right Selected for audit section shows all Event Logging rules that have been selected for the audit, and from which a rule can be defined.

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. This rule definition includes:
- **Operator**: Uses an operator to set the expression. Operators include equals (=), less than (<), greater than (>), and so on.
- **Reference**: Describes where the source of the value will derive from. You can choose from the following options:
    - **Source**: takes the value of this setting from the source of the audit – 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 the 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.

Once 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 their 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 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 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

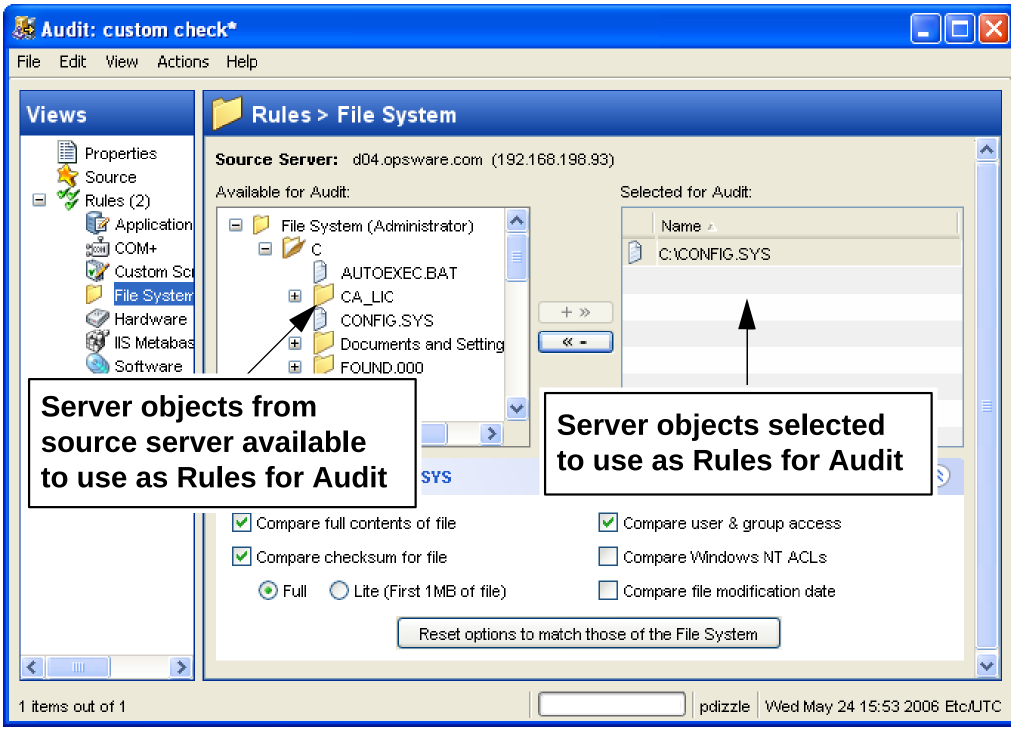

Source Server: d04.opsware.com (192.168.198.93)

Available for Audit:

Selected for Audit:

C:\CONFIG.SYS

Server objects from source server available to use as Rules for Audit

Server objects selected to use as Rules for Audit

☑ Compare full contents of file

☑ Compare checksum for file

● Full ○ Lite (First 1MB of file)

☑ Compare user & group access

☐ Compare Windows NT ACLs

☐ Compare file modification date

1 items out of 1 | pdizzle | Wed May 24 15:53 2006 Etc/UTC

- **Snapshot as Source for an audit or snapshot specification**: Choose a snapshot as the source for your audit or snapshot specification. 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, use a snapshot as the source of a snapshot specification allows you to use the captured server values to take a snapshot of another server. Using a snapshot as the source for a audit or snapshot specification allows you to choose both the results and the rules of the original snapshot specification the snapshot was based upon.

Figure 2-7 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

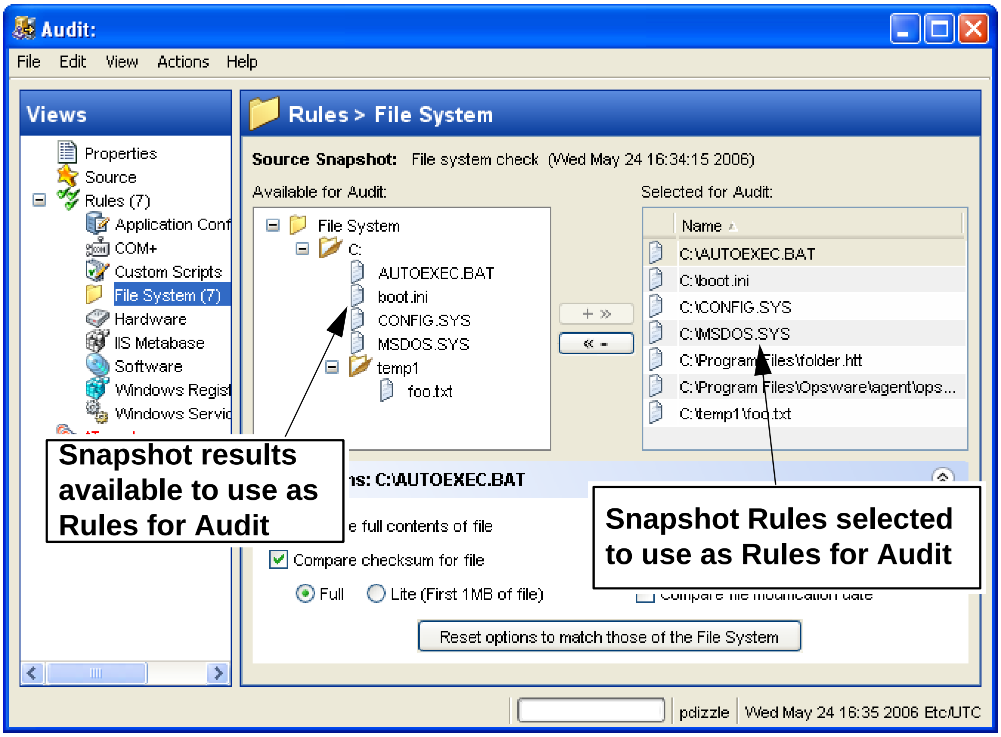

Source Snapshot: File system check (Wed May 24 16:34:15 2006)

Available for Audit:

- File System
  - C:
    - AUTOEXEC.BAT
    - boot.ini
    - CONFIG.SYS
    - MSDOS.SYS
    - temp1
      - foo.txt

Selected for Audit:

- C:\AUTOEXEC.BAT
- C:\boot.ini
- C:\CONFIG.SYS
- C:\MSDOS.SYS
- C:\Program Files\folder.htt
- C:\Program Files\Opsware\agent\ops...
- C:\temp1\foo.txt

ns: C:\AUTOEXEC.BAT

e full contents of file

☑ Compare checksum for file

● Full ○ Lite (First 1MB of file)

□ Compare file modification date

Reset options to match those of the File System

pdizzle | Wed May 24 16:35 2006 Etc/UTC

Snapshot results available to use as Rules for Audit

Snapshot Rules selected to use as Rules for Audit

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

Also important to know is that you cannot save a 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 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

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

Configuring Application Configuration Rules

The application configuration rule inside an audit, snapshot specification, or audit policy allows you configure values for application configuration files on a target server. You can define rules you want to capture and which configuration file values 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, or 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 the View button, you will see the contents of the configuration file from the source of the audit, displaying all key-value pairs you are able to add to the audit rule.

The kind of information you see when you view an application configuration inside the audit windows depends upon the source you choose for 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 a custom one and does not have a name-value pair defined, but the value exists in the source configuration file, you will not see it in the audit or audit policy.

Once you view the contents of the source application configuration file, you can select values and define target values – what values the audit should look for on the target server – and also a remediation value 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, which enables you to use the hosts file on that server as the basis of your application configuration audit rule.

You could define an application configuration rule in the audit that expresses the IP address and hostname in the hosts file you want to determine whether or not it exists on the target servers you audit.

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

1. Create the new audit from anyone of the methods for creating an audit listed at "Ways to Create an Audit" on page 73.
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 you want to create a rule from.
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 specificationsection, 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 selected to view a UNIX hosts file, you would see something similar to the following in the Rule Details, File View tab, shown in Windows Services Audit Rule as shown in Figure 2-3:

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

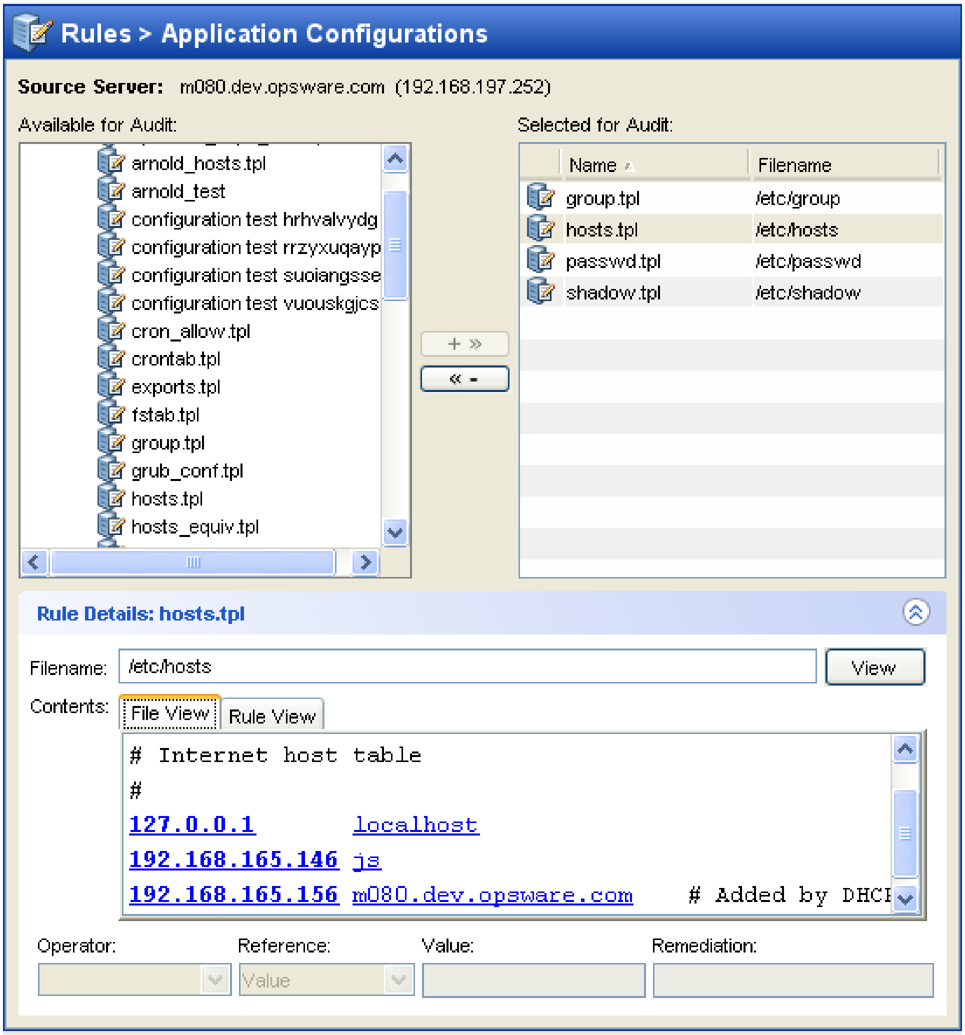

Source Server: m080.dev.opsware.com (192.168.197.252)

Available for Audit:

- arnold_hosts.tpl
- arnold_test
- configuration test hrhvalvydg
- configuration test rrzyxuqayp
- configuration test suoiangsse
- configuration test vuouskgjcs
- cron_allow.tpl
- crontab.tpl
- exports.tpl
- fstab.tpl
- group.tpl
- grub_conf.tpl
- hosts.tpl
- hosts_equiv.tpl

Selected for Audit:

Name	Filename
group.tpl	/etc/group
hosts.tpl	/etc/hosts
passwd.tpl	/etc/passwd
shadow.tpl	/etc/shadow

Rule Details: hosts.tpl

Filename: /etc/hosts

```
# Internet host table
#
127.0.0.1         localhost
192.168.165.146 js
192.168.165.156 m080.dev.opsware.com     # Added by DHCF
```

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. Once you select the IP address, notice that 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 95.

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. That means 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**: = (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**: Value
- **Value**: 127.0.0.1
- **Remediation**: 127.0.0.1

This expresses that you want to look for an IP address with the value of 100.10.10.10 and, if this is not found, then the remediation should be 100.10.10.10 so you can add this to any host files on the target servers that do not contain this IP address.

10 Next, select the hostname in the File View tab area, and in the Rule section, set the following parameters:

- **Operator**: = (equals)
- **Reference**: 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**: host
- **Remediation**: 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, you will see the rule 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 any other rules you wish 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 119.

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. Once 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 as shown in the Rules Details ➤ Contents section, File View tab.

Table 2-2: Application Configuration Audit Rule Color Scheme

TEXT COLOR	DESCRIPTION
Blue underlined	All elements in the source configuration file that can be used in a rule will show with a blue underline.
Dark blue	An element is selected but has no rule has been associated with it.
Light blue	Once you add an element to a rule, it will show with a light blue coloring.
Medium blue	Element is both selected and has a rule associated with it.

Table 2-2: Application Configuration Audit Rule Color Scheme

TEXT COLOR	DESCRIPTION
Green	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. 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 "=". An example of this would be: `127.0.0.1   localhost` 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 like: There is an entry where ip is equal to `127.0.0.1` AND `hostname` is equal to `localhost`.
Bold	Represents a primary key.
Italicized	Custom attribute or Opsware attribute.

Configuring COM+ Rules

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

The SAS Client categorizes the COM+ objects based on an attribute of the object, where the COM object specifies zero or more categories. The Opsware SAS Client displays all COM+ objects in one node in the Rules section of the COM+ object tree in case you do not know the category or the object did not specify one.

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

1. Create the new audit using one of the methods for creating an audit listed at "Ways to Create an Audit" on page 73.
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 you want to create a rule from.
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.
6. To finish configuring the audit, define any other COM+ object or object category rules you wish 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 119.

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**: At audit runtime, the script is executed on the source server, and the return values from the script (exit code or standard output) is compare 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, which 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 that 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, which you can use to force remediation on target servers that you wish to conform to the values returned from the script after it has run. 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 listed at "Ways to Create an Audit" on page 73.
2. Select an audit Source: Server, snapshot, or None.
3. To a script and define the audit rule, you can choose the following to build the custom script:

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

– Select an output, either Exit Code or Standard Output.

– 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 one.

– Reference: Choose where you want the of the script output to derive from.

- Source: This will run the script on the source when you run the audit, get the value the script requests, and then use that value to compare with the value gotten from the script that is run the target server.

  If you choose Source 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 compare 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: Will compare a server attribute found on the source server with the output from the script that is run on the target server.

- Custom Attribute: Will compare a custom attribute found on the target server with the output from the script that is run on the target server. 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.

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. You can also add a remediation script to run if the audit comparison fails from the rule and what is found on the target server. The remediation wont 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 119.

Configuring Event Logging Rules

Event logging rules allows 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 for creating an audit listed at "Ways to Create an Audit" on page 73.

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 Event Logging 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. For each event logging rule, you can define to set slightly different rule parameters:
   - Crash on Audit Failure Security Log Status
     - Target Value: Choose to set as 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 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)
   - 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 119.

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: This setting will perform a Checksum on the contents of the selected file or files in a directory. You can choose to audit the entire contents of the the file, or just the first 1MB of the file.
- User and group access: audits user and group access related to the file and directories.
- Windows NT ACLs (Access Control List): audits 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 112.

There are two categories of file system rules you can define in an audit or snapshot specification, that appear in the Available for Audit section of the Audit window:

- **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 compared these with what is found on the target servers. The purpose for this kind of rule is to determine is the file or directory exists and what its properties are. 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 Opsware SAS Client and 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 for creating an audit listed at "Ways to Create an Audit" on page 73.
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 you select will be used to audit or snapshot on 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 the file system rule options you would like to apply to the selected folder or file. If you would like to reset the original settings of the source file system, click the Reset options to match those of the File System button.

8 For folders, you can also select a File/directory Wildcard option to specify files and directories you want to include or exclude from the audit.

Click the plus button 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 112.

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 119.

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 at “Ways to Create an Audit” on page 73.

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 you select will be used to audit or snapshot on the target server.
6. To finish configuring the audit, 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 119.

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 such IIS Metabase object property information 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 "Ways to Create an Audit" on page 73.
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 on the target server.
6. To finish configuring the audit, 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 119.

Configuring Operating System Rules

Operating system rules allow you to check dozens of different operating system properties and settings, for such things as domain controller settings (for example, LDAP Server Signing Requirements), network configurations (for example, IP Source Routing Protection Level), to such general settings as 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 for creating an audit listed at “Ways to Create an Audit” on page 73.
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 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 you select will be audited on the target servers or snapshot specification.
6. For each operating system rule, you can define to 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

   Specify the value you expect to be on the target server or servers of the audit, or the value you wish to capture on in a snapshot:

   – 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 where you want the of the script output to derive from.
  - 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 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: Will compare a server attribute located on the source server.
  - Custom Attribute: Will compare a custom attribute found on the target server.

Remediation Value

Each remediation value setting will be different depending upon the type of rule.

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 119.

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 listed at "Ways to Create an Audit" on page 73.
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 you select will be used to audit or snapshot on the target server.
6. To finish configuring the audit, 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 119.

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 behavior requiring a user to select CTRL-ALT-DEL to logon is enabled or disabled.
- Display Last Username at Login: Determines if display of the last User to logon is enabled or disabled.
- Message Text for Users Attempting to Log On: Specifies expected text shown to users who attempt to log on to the server.
- Message Title for Users Attempting to Log On: Specifies 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 for creating an audit listed at “Ways to Create an Audit” on page 73.
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 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: This will 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: Will compare a server attribute located on the source server.
   - Custom Attribute: Will compare 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, 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 119.

Configuring Windows Registry Rules

The windows registry rule allows you to select windows registry folders and keys to compare in your audit. The audit compare the selected registry folders and keys and determine if these keys and folders exist on the target servers.

There are two categories of windows registry rules you can define in an audit or snapshot specification, that 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 compared these with what is found on the target servers. The purpose for this kind of rule is to determine is the windows registry key or folder exists and what its properties are. 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. Audit and Remediation 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-#xFFFD], and [#x10000-#x10FFFF]. Invalid control characters cannot be stored by Opsware SAS Client and will be converted to XML entities and will display as &#;. For example, if the data value is 00 00 (in bytes), &#x00; 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 Client and 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 at “Ways to Create an Audit” on page 73.

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 you select will be used to audit or snapshot on the target server.

6. To finish configuring the audit, 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 119.

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 and 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 you can define in an audit or snapshot specification, that appear in the Available for Audit section of the Audit window:

- **Windows Services:** These are comparison-based rules, which enable you to select a service from the source of the audit or snapshot specification and compared these with what is found on the target servers. The purpose for this kind of windows services rule is to determine is the service exists and what its settings are. You cannot set a target or remediation value in this type of rule.

- **Other Windows Services Rules**: These are value-based windows services rules prebuilt into the Opsware SAS Client and 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 at "Ways to Create an Audit" on page 73.

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 you select will be used to audit or snapshot on the target server.

6. To finish configuring the audit, 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 119

File System Inclusion and Exclusion Rules

When configuring a file system rule inside an audit, audit policy, or snapshot specification, you are able to 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

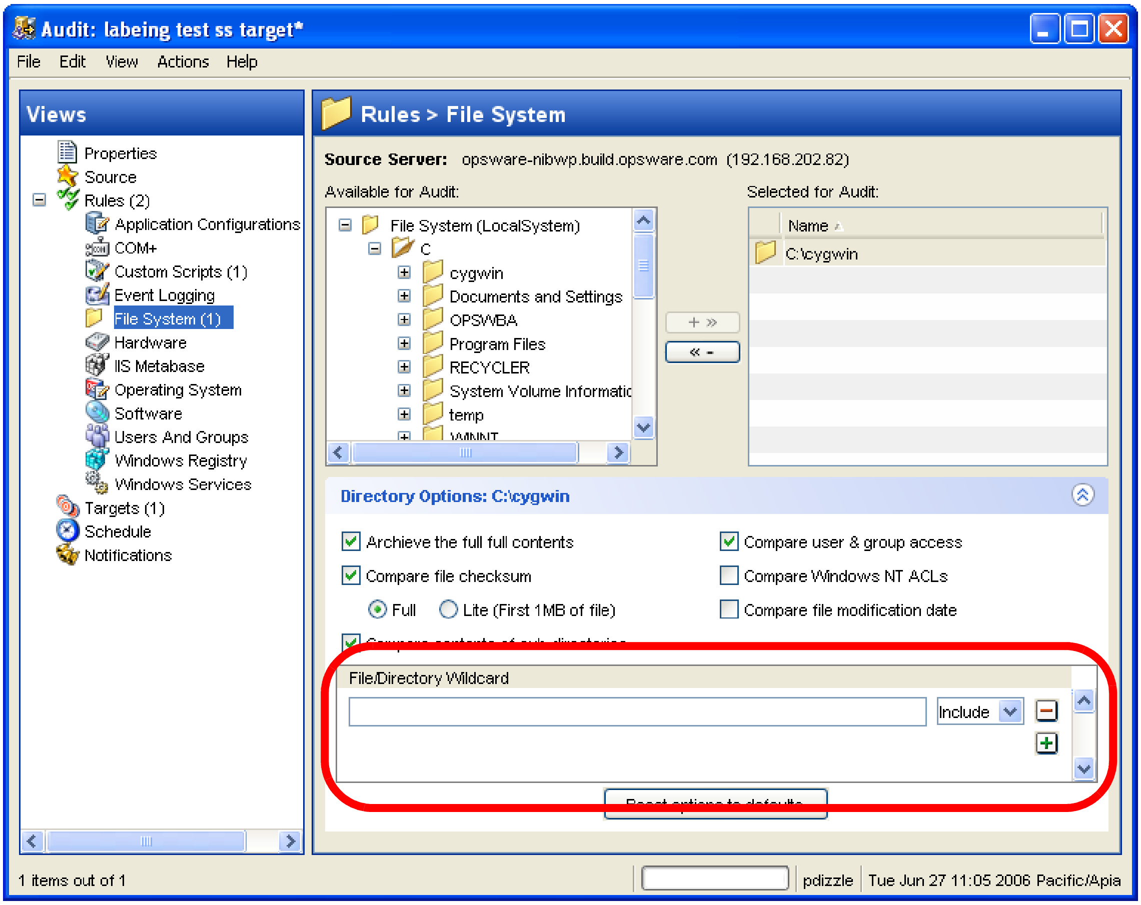

Source Server: opsware-nibwp.build.opsware.com (192.168.202.82)

Directory Options: C:\cygwin

☑ Archieve the full full contents

☑ Compare file checksum

● Full ○ Lite (First 1MB of file)

☑ Compare contents of sub-directories

☑ Compare user & group access

☐ Compare Windows NT ACLs

☐ Compare file modification date

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 102.

Inclusion and Exclusion Rule Types

Audit and Remediation provides the following types of inclusion and exclusion rules when you configure 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, Audit and Remediation will convert it to a “\” (backslash) to be able to use it as a valid path.

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 to 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 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 would 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 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

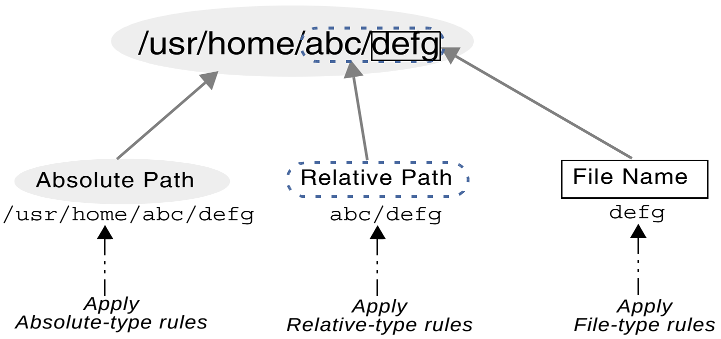

To best explain how these rules are applied, the following examples are provided.

The sample file system structure used in "Example: Including all Files With the .txt Extension in Your Snapshot or Audit", and "Example: Including the last temp.txt file and excluding everything else" is:

```
/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 Files With the .txt Extension in Your 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 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 the File a in Your Snapshot or Audit

If you want to include only the file a 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 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 the last temp.txt file and excluding everything 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 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. How audit rules for a parent and child directory will overlap is best explained by the following examples.

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 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 of 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. In other words, 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. Or, an audit policy can 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.

For information on creating rules for an audit policy, see “Configuring Audit and Remediation Rules” on page 84.

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 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, then Windows or Unix.
3. Click once inside the Content pane and from the **Actions** menu, select **New**.
4. In the Contents 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 upon which you can base the audit policy.
6. From the Contents pane, to select a source for the audit policy, 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 the rules you would like to configure. For more information on how to configure specific rules, see “Configuring Specific Rules” on page 90.

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 his audit policy. If the policy setter makes any changes to the "source" audit or snapshot specification, then changes will be reflected where the policy is linked to.

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 by selecting one from the Library, from the Navigation pane, select Library ➤ Audit and Remediation ➤ Audits, and then double-click the audit you want to open. Or, 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, you will see a message that explains by linking the audit policy into the audit or snapshot specification, you will overwrite any existing rule definitions. Click **Yes** to import the audit policy.

1. 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.

Once 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 by selecting one from the Library, from the Navigation pane, select Library ➤ Audit and Remediation ➤ Audits, and then double-click the audit you want to open. Or, 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 you are importing the audit policy into already has rules defined, you have a choice if you would like to overwrite the existing rules, or merger the audit policy rule 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 To open an existing audit by selecting one from the Library, from the Navigation pane, select Library ➤ Audit and Remediation ➤ Audits, and then double-click the audit you want to open. Or, 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 or servers or snapshot of the audit, evaluating the targets according to the rules defined in the audit. You can run an audit from these 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 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), or you can use the search tool (upper right of the window) 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 tole number of rules defined in the audit (click **View Rule Details** to view the rule defintions), and all targets of the audit (servers and snapshot).
5. Click **Next**.
6. In the Scheduling page, choose if you want the audit to run immediatley, 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 Notifcations page, by default your user will be selected to have a notifcation 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. You can specifiy if you want to the email to be sent upon success of the audit Job job (✓) or failure of the audit job (✗).
10. You can also 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 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. and from the Contents pane, View drop-down list, select Audit and Remediation. You see the Details pane area below the Contents 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 tole number of rules defined in the audit (click **View Rule Details** to view the rule defintions), and all targets of the audit (servers and snapshot).
6. Click **Next**.
7. In the Scheduling page, choose if you want the audit to run immediatley, 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 Notifcations page, by default your user will be selected to have a notifcation 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. You can specifiy if you want to the email to be sent upon success of the audit Job job (✓) or failure of the audit job (✗).
11. You can also 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.
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 that has already been run by starting the audit to run again.

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 tole number of rules defined in the audit (click **View Rule Details** to view the rule defintions), and all targets of the audit (servers and snapshot).
5. Click **Next**.
6. In the Scheduling page, choose if you want the audit to run immediatley, 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 Notifcations page, by default your user will be selected to have a notifcation 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. You can specifiy if you want to the email to be sent upon success of the audit Job job (✓) or failure of the audit job (✗).
10. You can also 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.

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 that are also associated with that audit will be deleted.

You must have permissions to create, view, edit, and delete audit schedules. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration 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. Once 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, and Audit and Remediation, and then select Audits.
3. Select an OS (Windows or UNIX) and then double-click an audit to open it.
4. In the Audit window, in the Views pane, select the schedule object.
5. In the Schedule section, choose whether you want to run the to be created once, daily, weekly, monthly, or on a custom schedule. Parameters include:
   - **None**: No schedule will be set. If you want to run the audit, 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 which day or days of the week you want the audit to run.
   - **Monthly**: Choose which months you would like to audit run during, and which days of the month to run.
   - **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 audit at 1:00 a.m. every weekday:

```
0 1 * * 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.

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 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.

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 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 Contents pane, select Run Audit Task.

4. Double-click the scheduled audit job to open the Audit Window.

5. Select the Schedule object in the Views pane to view the audit schedule.

6. To edit the audit Schedule, modify the following parameters:

   – **Schedule**: Choose whether you want to run the to be created once, daily, weekly, monthly, or on a custom schedule. Parameters include:

   – **None**: No schedule will be set. If you want to run the audit only once, 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 you want the audit to run.

   – **Monthly**: Choose which months you would like to audit run during and the day of the month to run.

   – **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 audit at 1:00 a.m. every weekday:

```
0 1 * * 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 and minute you want the daily schedule to start. Unless you specify an end time, the audit 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 audit schedule to run indefinitely.

7. 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 Contents pane displays all jobs that have been run in on this Opsware core. To display only audit jobs, from the drop-down list at the top of the Contents pane, select Run Audit Task. If you wish to see only those audits that you have scheduled to run, enter your user ID in the User ID field at the top of the Contents pane.
4. To open a completed audit job, double-click it. If you wish to view the audit results, click **View Results**.

Viewing and Remediating Audit Results

An audit defines what server object configurations 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 shows 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 you can perform depends upon what kind of rule was set in the audit: server comparison or value-based.

Viewing Comparison-Based Audit Results: "Copy To" Remediation

Audit results based upon 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 – to remediate, you can copy the rule values of the source objects in the audit to overwrite the values on the target (or add values that exist on the source but do not exist on the target.)

For example, Figure 2-10 shows audit Results for a windows services rule where 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 remediate on the target server.

Figure 2-10: Audit Results For a Comparison-Based Audit Rule

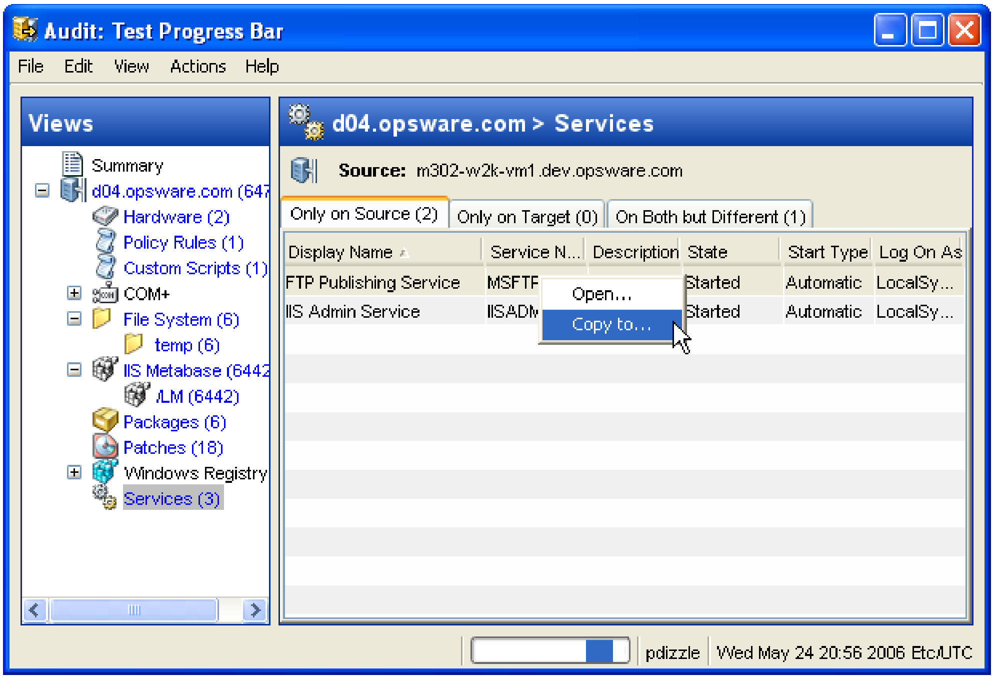

d04.opsware.com > Services

Source: m302-w2k-vm1.dev.opsware.com

Display Name	Service N...	Description	State	Start Type	Log On As
FTP Publishing Service	MSFTP		Started	Automatic	LocalSy...
IIS Admin Service	IISADM		Started	Automatic	LocalSy...

The Audit Results window shows you all objects defined in the audit in the Views pane, and those audit results that failed – differences were 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: audit Results for rules that are value-based present you with the results of the value returned form 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. Remediate will copy the values specified in the audit rule's Remediation values on the server object on the target server.

For example, shows a value-based audit rule, a windows services rule that 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 the rule checked to see if the service was running on the target, and was found to be disabled. You can choose to Remediate the service (click the **Remediate** button), which in this example is will enable the service.

Figure 2-11: Audit Results for a Value-Based Audit rule

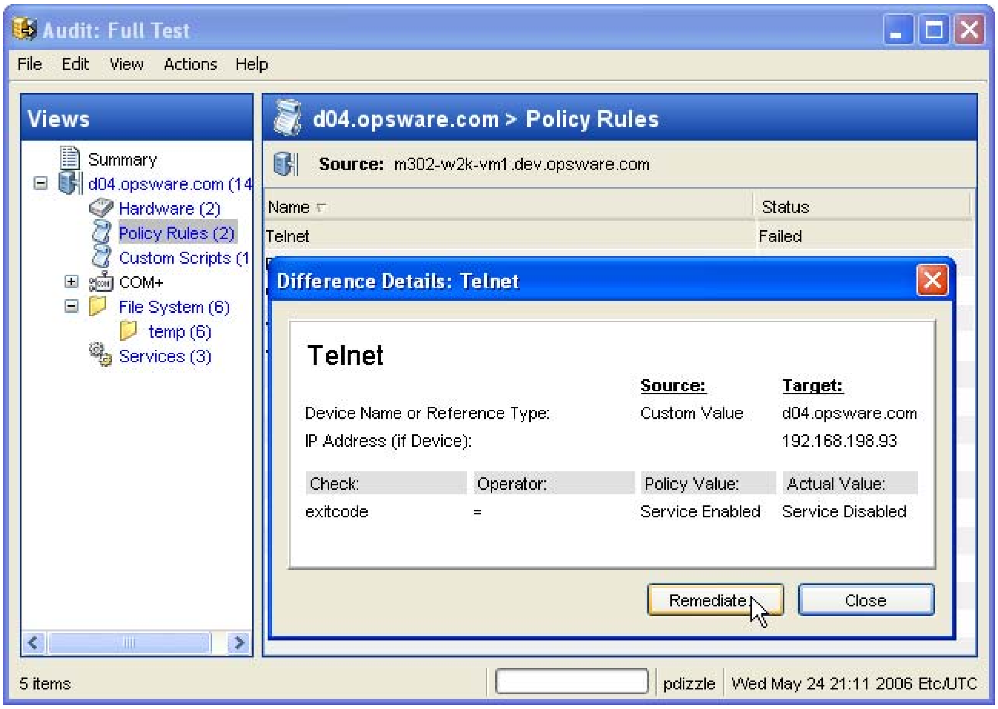

Source: m302-w2k-vm1.dev.opsware.com

Name	Status
Telnet	Failed

Difference Details: Telnet

Telnet

	Source:	Target:
Device Name or Reference Type:	Custom Value	d04.opsware.com
IP Address (if Device):		192.168.198.93

Check:	Operator:	Policy Value:	Actual Value:
exitcode	=	Service Enabled	Service Disabled

pdizzle | Wed May 24 21:11 2006 Etc/UTC

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 to 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 differences of the file’s contents.

With Audit and Remediation, you can view differences of the following server objects in an audit Results window:

- Viewing and Remediating File Audit Results Differences
- Viewing and Remediating Server Object Audit Results Differences

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 which lines in a file 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 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 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, 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 in the file that is in the source or in the file that is in a target. 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 – in other words the audit has "failures" – you can view differences in object properties side by side. Each server object will show different windows, depending upon 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 objects, right-click, and select Open. You see a window that shows the differences between the object as defined the audit and the object on the target server. For example, 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 2-12: Comparison-Based Audit Results Difference: CPU

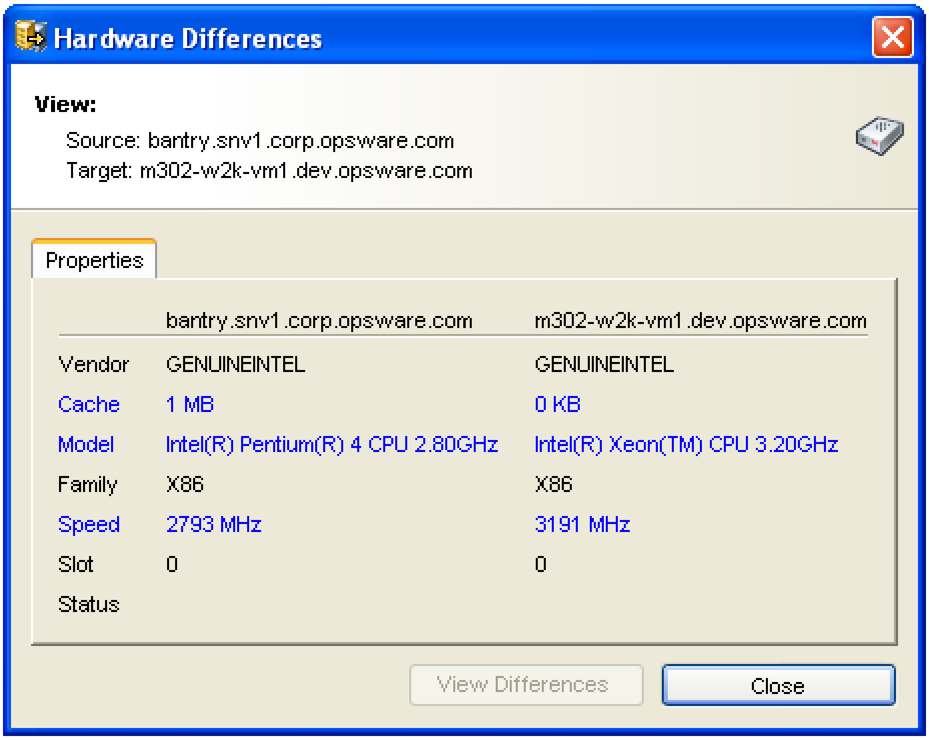

View:

Source: bantry.snv1.corp.opsware.com

Target: m302-w2k-vm1.dev.opsware.com

	bantry.snv1.corp.opsware.com	m302-w2k-vm1.dev.opsware.com
Vendor	GENUINEINTEL	GENUINEINTEL
Cache	1 MB	0 KB
Model	Intel(R) Pentium(R) 4 CPU 2.80GHz	Intel(R) Xeon(TM) CPU 3.20GHz
Family	X86	X86
Speed	2793 MHz	3191 MHz
Slot	0	0
Status		

The window shows the source server information on the left, and the target server information on the right, with all differences between the two CPUs listed in blue font. This type of audit rule cannot be remediated.

7 For a value-based rule, the difference window will be slightly different and will also include a Remediate option, if remediation is possible. For example, shows the differences for a value-based windows registry rule.

Figure 2-13: rule-Based Audit Results Difference: Windows Registry

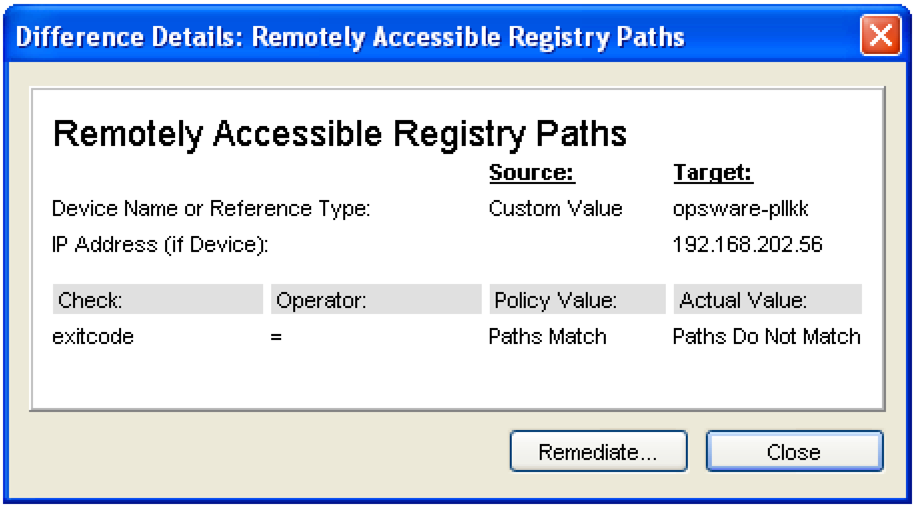

Difference Details: Remotely Accessible Registry Paths

Remotely Accessible Registry Paths

	Source:	Target:
Device Name or Reference Type:	Custom Value	opsware-pllkk
IP Address (if Device):		192.168.202.56

Check:	Operator:	Policy Value:	Actual Value:
exitcode	=	Paths Match	Paths Do Not Match

This difference window displays the audit rule, including the policy value (what should be) and the actual value found on the target server. Since there is a difference between the two, you can click **Remediate** to make the target's actual value match the policy value.

8 To remediate the difference, click **Remediate**.

9 In the Remediate window, enter an optional Job ID, and then click **Remediate**.

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, make sure 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 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. For more information on how to use search, see “SAS Client Search” on page 78.

About Snapshots

A snapshot captures (takes a picture) of how a managed server is configured 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 much in the same 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 rule values 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 to determine how they differ. 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). You can also define rules for server objects in addition to the values defined in the snapshot, to further extend the audit definition.

When a snapshot is used as the source for an audit, all server configuration values captured in the snapshot results are used available to use as rules for the audit. For more information about using a snapshot in an audit, see "Configuring an Audit" on page 77.

Audit Policies and Snapshot Specification

An audit policies collection of rules that define a desired state of a server's configuration. An audit policy can be used inside of an snapshot specification, either through linking or importing. 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 a snapshot specification.

An audit policy can be linked into an snapshot specification so whenever a change is made the audit using the policy will have the latest changes. Or, an audit policy can be imported into an snapshot specification, without keeping the link to the source audit policy. When you import an audit policy into an 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 into a snapshot specification, see "Linking and Importing Audit Policies" on page 119.

Snapshot Specification Elements

An snapshot specification consists of the following elements:

- **Properties**: Name and description of the snapshot specification.
- **Targets**: The servers which 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 wish as targets.
- **Source**: Choosing a server as the source for an snapshot specification allows you to 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 you can define 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: check to see if a server contains a specific Windows Service, and if found, determine if the service is turned off. For a description of server objects you can define rules for in a snapshot specification, see “Configuring Audit and Remediation Rules” on page 84.
- **Schedule**: You can run the snapshot specification as a job on a one time basis, or on a recurring schedule.
- **Notifications**: You can send emails when the snapshot specification job has finished running, and base the notification upon success, failure, or simply the completion of the snapshot specification Job.

When you set up an snapshot specification, you select which objects to check for on the target server. You can also apply rules to these objects that define their desired configuration state, and for some rules, you can define remediation values, in the even that the resulting snapshot is used as the source for an audit.

For example, Figure 2-6 shows snapshot specification that has defined on the target and three rules that will capture configuration information about the target server for event logging, operating system, and windows services.

Figure 2-14: Snapshot Specification Elements

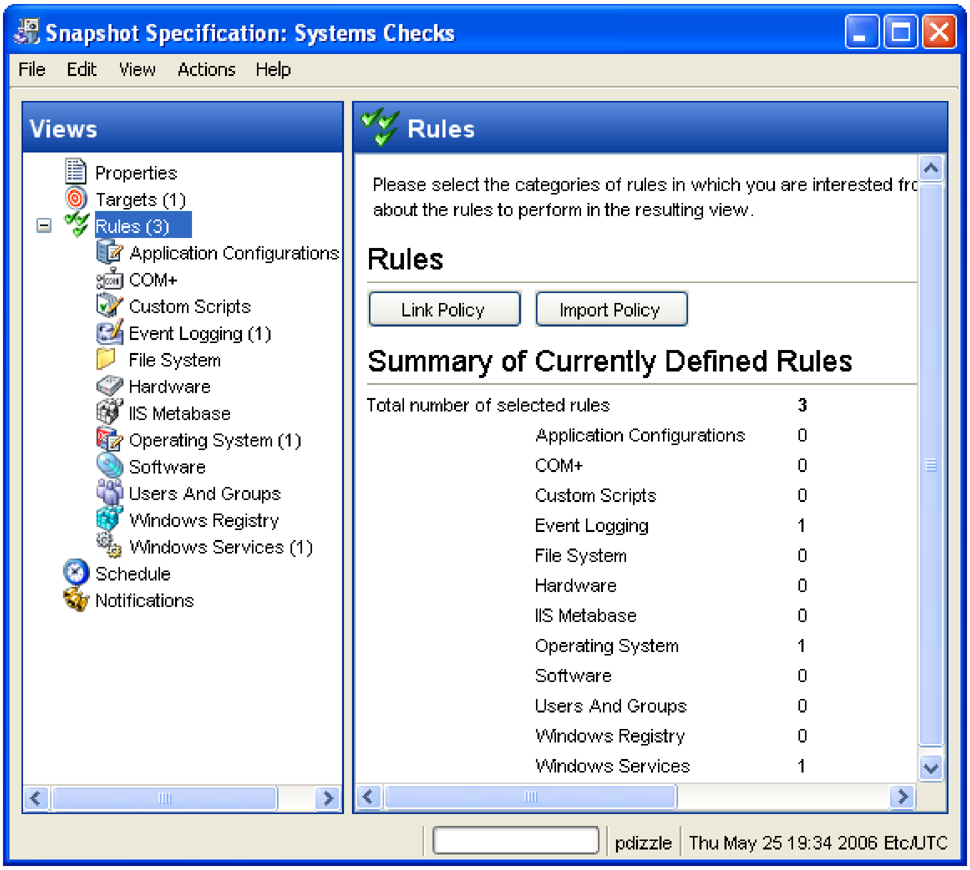

Please select the categories of rules in which you are interested fro about the rules to perform in the resulting view.

Rules

Summary of Currently Defined Rules

Total number of selected rules		3
	Application Configurations	0
	COM+	0
	Custom Scripts	0
	Event Logging	1
	File System	0
	Hardware	0
	IIS Metabase	0
	Operating System	1
	Software	0
	Users And Groups	0
	Windows Registry	0
	Windows Services	1

The Snapshot Process

Taking a snapshot of a server configuration requires two basic steps:

- Creating a snapshot specification, which is a template that defines the configuration parameters you want to capture on a target server.
- Running the snapshot specification job that results in a snapshot.

Figure 2-15 illustrates an example of the snapshot process.

Figure 2-15: Snapshot Process

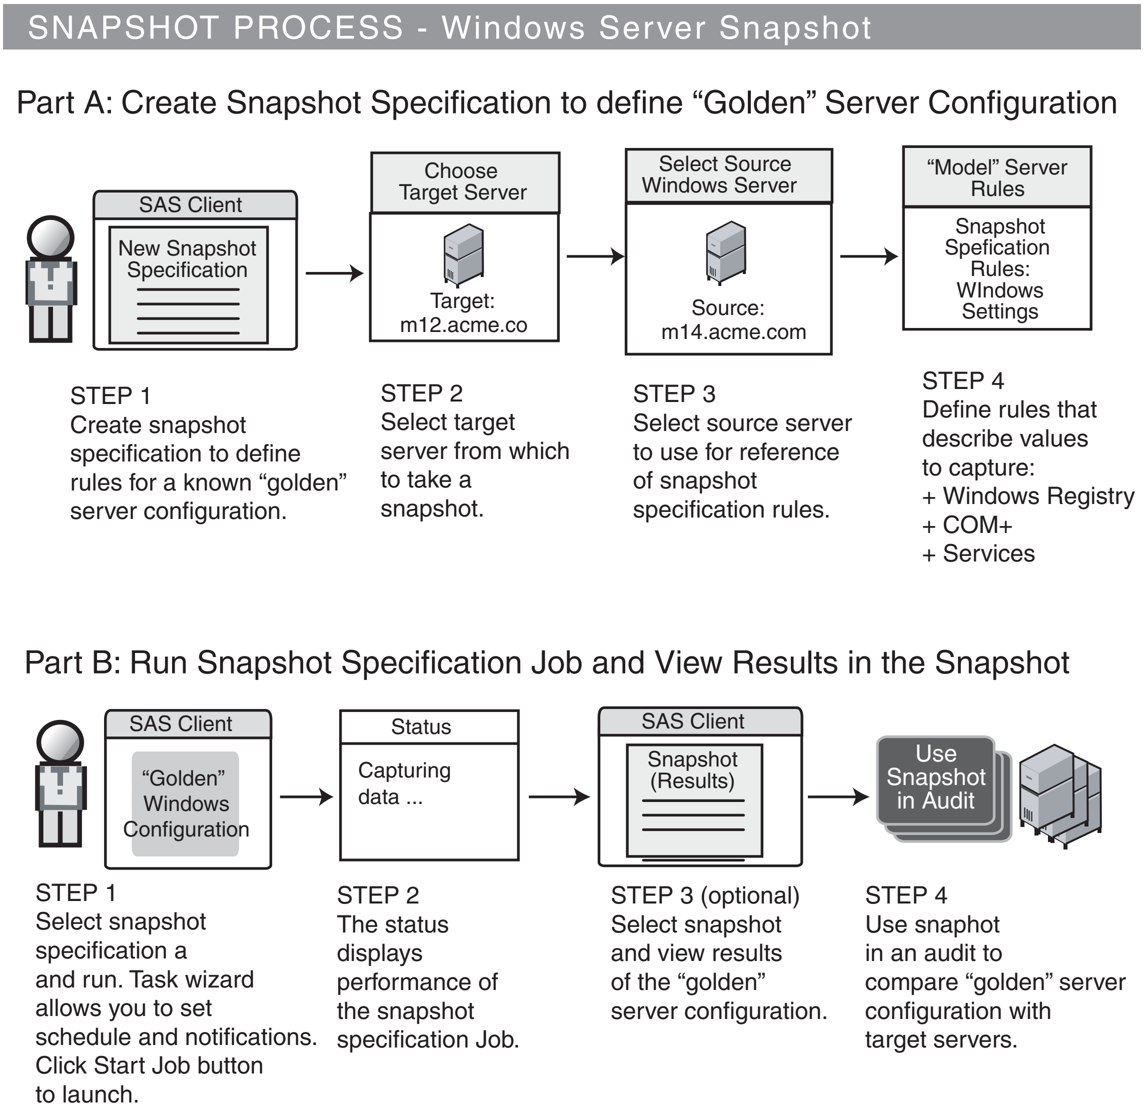

SNAPSHOT PROCESS - Windows Server Snapshot

Part A: Create Snapshot Specification to define “Golden” Server Configuration

SAS Client
New Snapshot Specification
Choose Target Server
Target: m12.acme.co
Select Source Windows Server
Source: m14.acme.com
“Model” Server Rules
Snapshot Spefication Rules: WIndows Settings

STEP 1
Create snapshot specification to define rules for a known “golden” server configuration.

STEP 2
Select target server from which to take a snapshot.

STEP 3
Select source server to use for reference of snapshot specification rules.

STEP 4
Define rules that describe values to capture:
+ Windows Registry
+ COM+
+ Services

Part B: Run Snapshot Specification Job and View Results in the Snapshot

STEP 1
Select snapshot specification a and run. Task wizard allows you to set schedule and notifications. Click Start Job button to launch.

STEP 2
The status displays performance of the snapshot specification Job.

STEP 3 (optional)
Select snapshot and view results of the “golden” server configuration.

STEP 4
Use snaphot in an audit to compare “golden” server configuration with target servers.

Creating a Snapshot Specification

You can create an snapshot specification from two different locations inside the SAS Client, depending upon your purpose: Do you want to snapshot a specific server? Or, do you want to create a snapshot specification from scratch?

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 Configuration 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 difference 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.

3. Click once inside the Content pane and from the **Actions** menu, select **New**.

Configuring a Snapshot Specification

Configuring a snapshot specification requires performing the following tasks:

- Name and describe the snapshot specification
- Choose targets: Choose the servers you want to take a snapshot of. You can choose to snapshot multiple servers or groups of servers.
- Configure the rules: You can define your own custom rules, or you can choose a source server to serve as the basis for the snapshot specification
- Schedule the snapshot specification job: You can schedule the snapshot specification to run once or on a recurring schedule.
- Set up email notifications: You can email selected users if 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 at "Creating a Snapshot Specification" on page 140.
2. In the snapshot specification Window, you can now begin to define the parameters. Enter the following information:
   - **Properties**: Enter a name and description for the snapshot specification.
   - **Source**: By default, the source server for the snapshot specification will be the managed server you chose will be used as the source for the snapshot specification. This means you can 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 Audit and Remediation Rules" on

page 84.

If you wish 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 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 an into a snapshot specification, see "Linking and Importing Audit Policies" on page 119.

- **Targets**: Choose the Targets of the snapshot specification. These are servers groups of servers that you want the configured snapshot specification rules to capture. To add a server or group of servers, click **Add**. To add a snapshot target, in the Snapshot Targets section, click **Add**.
- **Schedule**: Choose when you would like to run the snapshot specification immediately, or on a recurring schedule. Choose whether you want to run the to be created once, daily, weekly, monthly, or on a custom schedule. Parameters include:
- **None**: No schedule will be set. If you want to run the snapshot specification, you have to 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 which day of the week you want the snapshot specification to run.
- **Monthly**: Choose which months you would like to snapshot specification run during.
- **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.
- **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 an end 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 people you would like to receive an email when the snapshot specification Job finishes running. You can choose to send the email notification to be sent 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 1 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 Audit and Remediation Rules" on page 84.

Saving Snapshot As Audit Policy

You can save selection criteria and use it in other snapshot specifications.

To save your 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 that you own (created), that are also associated with that snapshot specification, will be deleted. See "Scheduling Snapshot Jobs" on page 144 in this chapter for more information.

Scheduling Snapshot Jobs

A snapshot specification job schedule specifies when you want Opsware 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 that you own (created), that are also 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 set up a schedule that specifies when you want the snapshot specification to run on a recurring basis. Once 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 snapshot specification, select an OS (Windows or UNIX), and then double-click an snapshot specification to open it.
3. In the snapshot specification window, in the Views pane, select the schedule object.
4. In the Schedule section, choose when you would like to run the snapshot job immediately, or on a recurring schedule. Choose whether you want to run the to be created once, daily, weekly, monthly, or on a custom schedule. Parameters include:
   - **None**: No schedule will be set. If you want to run the snapshot job, you have to select the snapshot specification, right-click, and select **Run Audit**.
   - **Daily**: Choose this option to run the snapshot job on a daily basis.
   - **Weekly**: Choose which day of the week you want the snapshot specification job to run.
   - **Monthly**: Choose which months you would like to snapshot specification job to run during.
   - **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.

5. 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.
6. (Optional) Deselect the End option if you want the snapshot specification job to run indefinitely.
7. 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 Contents pane, select Run Snapshot Task.
4. Double-click the scheduled snapshot specification job to open the snapshot specification Window.
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 when you would like to run the snapshot specification, immediately, or on a recurring schedule. Choose whether you want to run the to be created once, daily, weekly, monthly, or on a custom schedule. Parameters include:
   - **None**: No schedule will be set. If you want to run the snapshot specification, you have to select the snapshot specification, right-click, and select **Run snapshot specification**.
   - **Daily**: Choose this option to run the snapshot job on a daily basis.
   - **Weekly**: Choose which day of the week you want the snapshot job to run.
   - **Monthly**: Choose which months you would like to snapshot specification job run during.

- **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.

- **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 an end date to end the snapshot specification job 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 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 Contents pane displays all jobs that have been run in on this Opsware core. To display only snapshot specification jobs, from the drop-down list at the top of the Contents pane, select Run Snapshot Task. If you wish 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 Contents pane.

4. To open a completed snapshot job, double-click it. If you wish 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 Contents pane displays all jobs that have been run in on this Opsware core. To display only snapshot specification jobs, from the drop-down list at the top of the Contents pane, select Run Snapshot Task. If you wish 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 Contents 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 double-click it to 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 double-click it to 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 Audit and Remediation.
5. In the Contents 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, double-click 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, make sure 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.
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. For more information on how to use search, see “Content Pane” on page 74.

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 148, open a snapshot.

Figure 2-16: Sample Snapshot Browser of a Windows Server

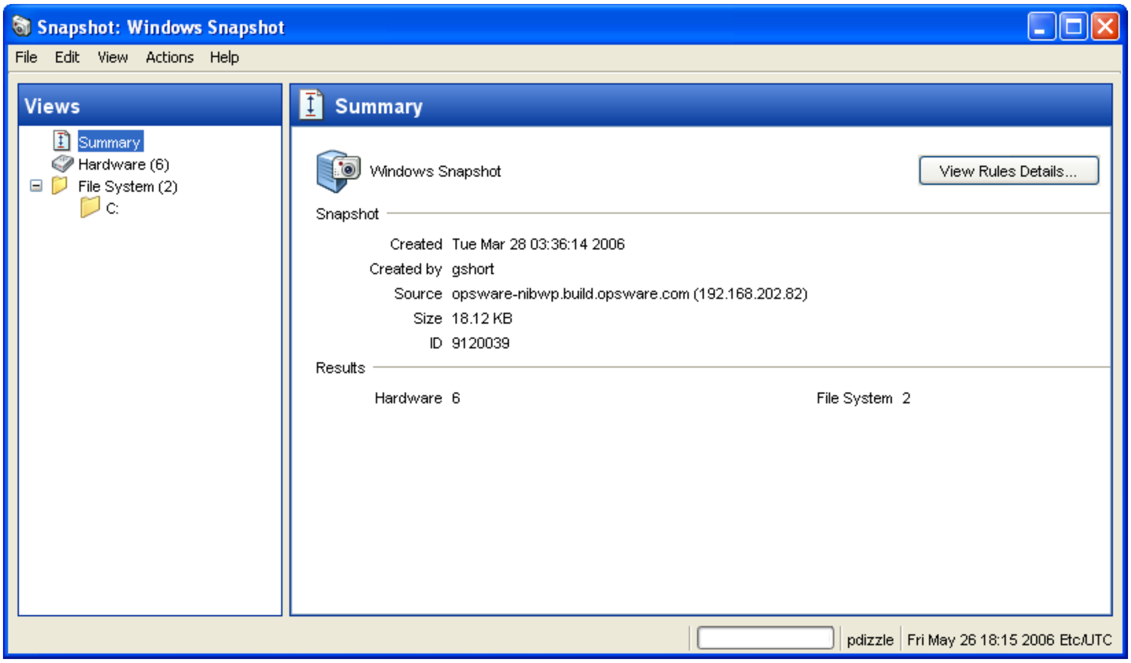

Windows Snapshot

Snapshot

Created Tue Mar 28 03:36:14 2006
Created by gshort
Source opsware-nibwp.build.opsware.com (192.168.202.82)
Size 18.12 KB
ID 9120039

Results

Hardware 6
File System 2

pdizzle | Fri May 26 18:15 2006 Etc/UTC

2. In the snapshot window, you can select:

   – **Summary**: 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 upon.

   – **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**: View information about the installed patches that were recorded in the snapshot, such as the patch type.

- **Installed Packages**: View information about the installed packages that were recorded in the snapshot, such as package type, package version, and release number.
- **File System**: View the directories, file properties and attributes, and contents of files recorded in the snapshot.

If a file in the snapshot exceeds 2MB in file size, Audit and Remediation cannot display the file contents.

- **Windows Services**: View information about the running services recorded in a snapshot, such as the name, description, startup state, startup type, and log on account.
- **Windows Registry**: View 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. Audit and Remediation supports the following windows registry keys: HKEY_CLASSES_ROOT, HKEY_CURRENT_CONFIG, HKEY_LOCAL_MACHINE, and HKEY_USERS.
- **COM+**: View 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.

  Opsware SAS provides warning messages that explain how Windows COM folders were processed. The following scenarios apply:

  When you create a snapshot where you selected 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 where you selected 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 where you selected a Windows COM+ folder that does not contain any objects and a Windows COM+ folder that does contain objects, the Snapshot window displays the folder. Opsware SAS displays a warning that the folder is empty.

– **Metabase**: View information about IIS Metabase objects in the snapshot, such as the ID, name, path, attributes, and data of the object.

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 decommision 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 148, 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 Opsware administrator. See the Opsware® 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 144 in this chapter for more information.

Copying Objects from a Snapshot to a Server

After viewing snapshot contents, you might want to copy certain objects to a target server. Audit and Remediation allows you to copy the following objects to a managed server: directories, files, windows services (state only), and windows registry keys.

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, 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, 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, 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.

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 Configuration 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 148, 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 193 in Chapter 22 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 displays compliance tests for software policies, application configurations, audits, patches, and duplex status – and you can also create your own individual audit tests as well. Each of these compliance tests is based upon an Opsware 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 your IT environment is configured the way you want it to 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 Opsware Command Center. For example, a policy setter can create an 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. In some cases, you can get relevant information 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 make sure they are in compliance with the policy that has been 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 take the necessary action to fix any problems. For example, you might use the Compliance Dashboard to check the status of an individually scheduled audit that makes sure a web application's configuration (such as Apache's http.conf file) is configured according to the standards set by your group. In other words, you want to make sure no one has changed the application's configuration. To verify that no unwanted changes have been made, you could regularly check the Compliance Dashboard for this scheduled audit and see if its compliance status has changed to red (non-compliant), and if it did, 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 and specific problem. For example, if a particular scheduled audit defines security standards for a set of servers – that all Windows 2003 servers contain a specific patch. But recently a security patch has been released and you have to find out if any of the servers contain the patch of not. After the audit has been updated, you can browse to selected Windows 2003 servers in the Compliance Dashboard, rerun the audit, and remediate those server which do not have the 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

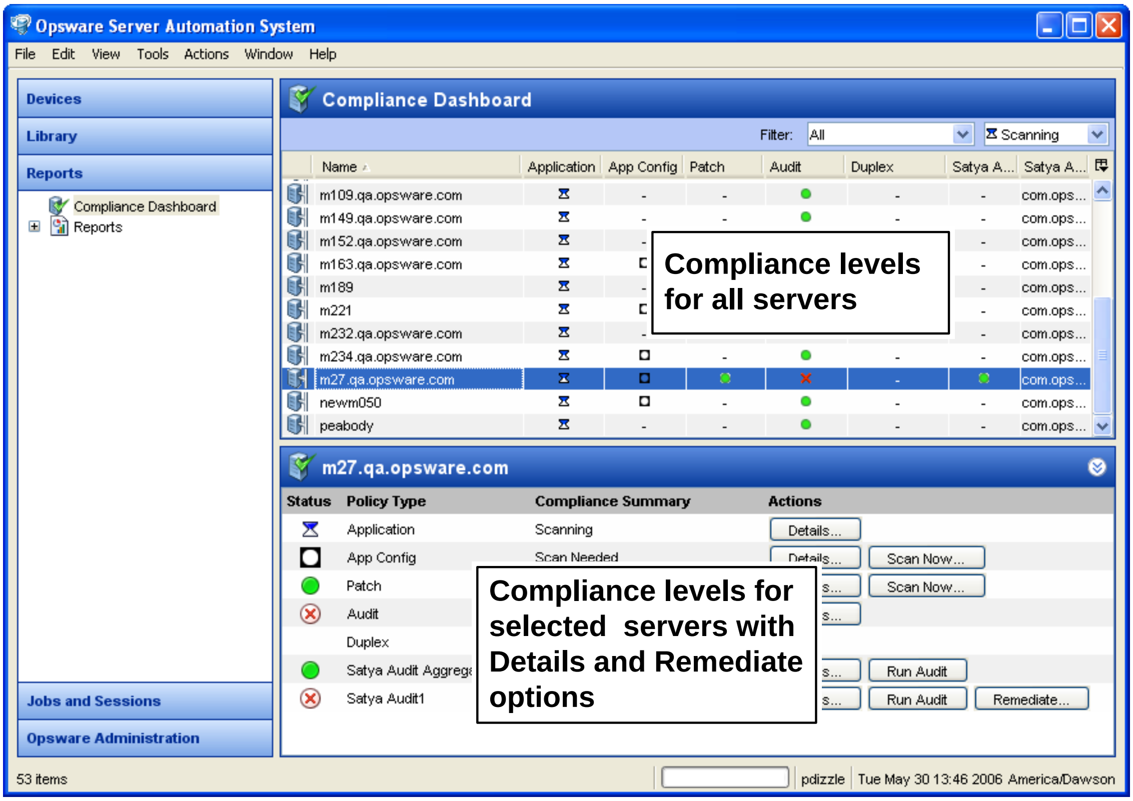

Name	Application	App Config	Patch	Audit	Duplex	Satya A...	Satya A...
m109.qa.opsware.com	⧗	-	-	●	-	-	com.ops...
m149.qa.opsware.com	⧗	-	-	●	-	-	com.ops...
m152.qa.opsware.com	⧗	-				-	com.ops...
m163.qa.opsware.com	⧗	□				-	com.ops...
m189	⧗	-				-	com.ops...
m221	⧗	□				-	com.ops...
m232.qa.opsware.com	⧗	-				-	com.ops...
m234.qa.opsware.com	⧗	□	-	●	-	-	com.ops...
m27.qa.opsware.com	⧗	□	●	✗	-	●	com.ops...
newm050	⧗	□	-	●	-	-	com.ops...
peabody	⧗	-	-	●	-	-	com.ops...

Compliance levels for all servers

Status	Policy Type	Compliance Summary	Actions
⧗	Application	Scanning	Details...
□	App Config	Scan Needed	Details... Scan Now...
●	Patch		s... Scan Now...
⊗	Audit		s...
	Duplex		
●	Satya Audit Aggrega		s... Run Audit
⊗	Satya Audit1		s... Run Audit Remediate...

Compliance levels for selected servers with Details and Remediate options

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 or not the software policy definition matches what is installed on the server. An 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 patched 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 to the roll up of all scheduled audits, 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 Compliance 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
🟢	**Compliant**: Compliance scan ran successfully and the actual server configuration matches the compliance criteria defined in the policy. 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.)
⚠	**Partial**: Compliance scan ran successfully but server configuration did not fully pass the compliance criteria defined in the policy. You will see this status for patch policies if the patch policy has exceptions defined in it.
⊗	**Failed**: Compliance scan ran and the actual server configuration did not match the criteria defined in the policy.
❗	**Scan Failure**: Compliance scan was not able to run.
□	**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
⌛	**Scanning**: Compliance scan currently being run.
–	**Not Applicable**: No policies of this type are attached to the server. Or, in the case of duplex, compliance scan is unable to determined duplex status on the device.

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 make sure 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 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, and helps you answer the question: What is the state of server configuration in my 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 would display as compliant if all configurations defined in the policy matches the same configurations on the server to which the software policy is attached. Compliance statuses include: Compliant, Partial, Noncompliant, Scan Failure, Scan Needed, Scanning, Not Applicable. For more information that explains what these statuses mean, see "Compliance Dashboard Compliance Statuses" on page 159.
- **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 be as simple as, 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 importantly, to be able 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

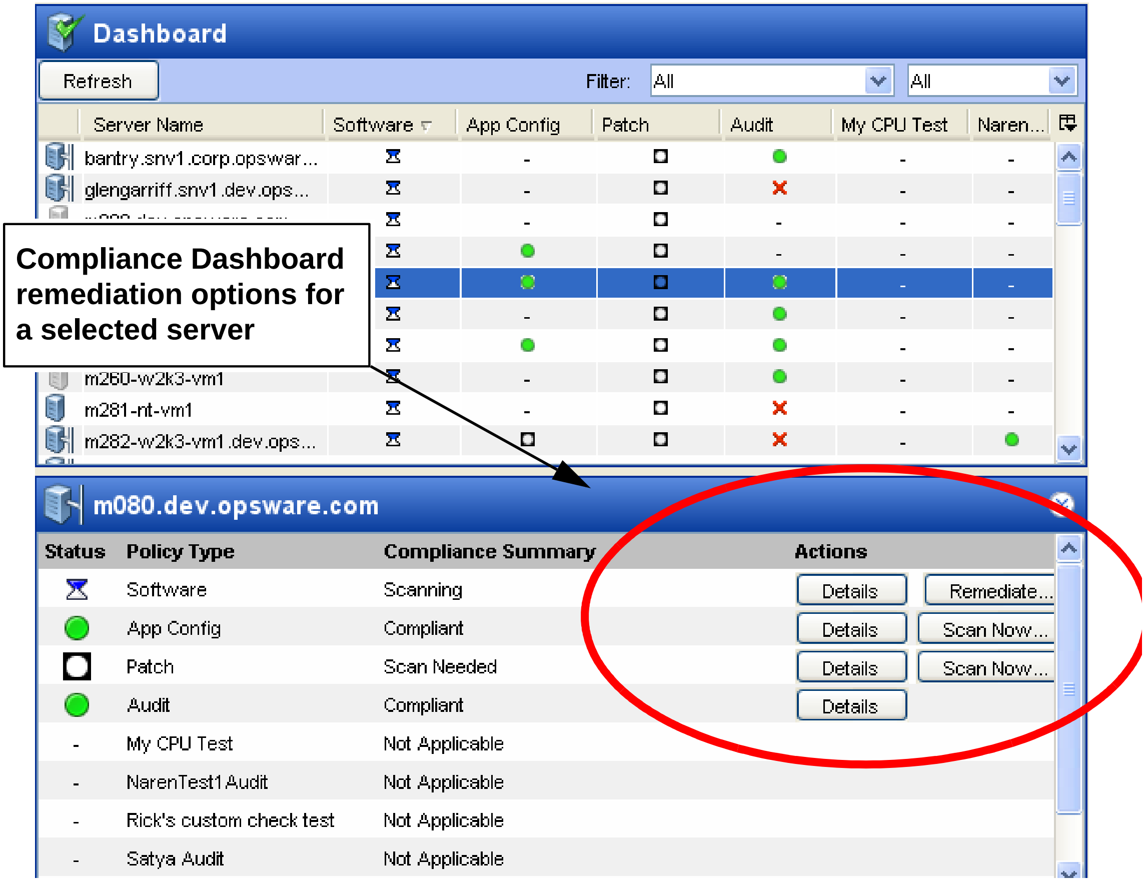

Server Name	Software	App Config	Patch	Audit	My CPU Test	Naren...
bantry.snv1.corp.opswar...	⧗	-	□	●	-	-
glengarriff.snv1.dev.ops...	⧗	-	□	✗	-	-
	⧗	-	□	-	-	-
	⧗	●	□	-	-	-
	⧗	●	□	●	-	-
	⧗	-	□	●	-	-
	⧗	●	□	●	-	-
m260-w2k3-vm1	⧗	-	□	●	-	-
m281-nt-vm1	⧗	-	□	✗	-	-
m282-w2k3-vm1.dev.ops...	⧗	□	□	✗	-	●

Compliance Dashboard remediation options for a selected server

m080.dev.opsware.com

Status	Policy Type	Compliance Summary	Actions
⧗	Software	Scanning	Details Remediate...
●	App Config	Compliant	Details Scan Now...
□	Patch	Scan Needed	Details Scan Now...
●	Audit	Compliant	Details
-	My CPU Test	Not Applicable	
-	NarenTest1 Audit	Not Applicable	
-	Rick's custom check test	Not Applicable	
-	Satya Audit	Not Applicable	

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 display 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 the **Run Audit** button 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 make sure the server conforms to the compliance policy. This might included 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 Chapter 7, “Software Management”.

Understanding Software Compliance Status

An 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 does 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 to with respect to 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 remediate actions:

- **Details**: Opens the Server Explorer showing the software policies view. This shows the software 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 software policy.

- **Remediate**: Opens the Remediate wizard for the server, with available software policies listed. For more information on how to run the Remediate wizard, see “Remediating Software Policies” on page 299.

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 Chapter 8, “Application Configuration Management”.

Understanding Application Configuration Compliance Status

An Application Configuration is either compliant or non-compliant. What this means is, 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 respect to 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 application configuration files definitions 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 Configuration in the Compliance Dashboard, you can perform the following remediate actions:

- **Details**: Opens the Server Explorer for the selected server to Installed Applications view. This view shows you all Application Configurations that have been attached to server. To browser specific instances of the applications, from the View pane in the Server 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 Chapter 5, "Patch Management for Windows" or Chapter 6, "Patch Management for Unix".

Understanding Patch Compliance Status

A patch policy is either compliant or non-compliant. What this means is, if any of the patches installed on the server (excluding patch policy exceptions) do not match the patch policy definitions, then the server is considered out of compliance with its attached patch policies. Specifically, patch policy compliance is defined as:

- **Compliant**: If a server is compliant to with respect to 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 remediate options:

- **Details**: Opens the Server Explorer showing the Patch Policies view. This shows you what patch policies are 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 make sure that it has all the Patches defined in the policy installed.

Audit Compliance

The Compliance Dashboard displays these two types of compliance for audits:

- All roll up of scheduled audits will appear by default in the Compliance Dashboard in a single column named audit. This status enables you to view at a glance the total compliance status of all audits you have scheduled to run on a regular basis 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 that have been scheduled to run on the selected server can be displayed in 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 Chapter 2, "Audit and Remediation" on page 63 of this guide.

Understanding 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 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. Conversely, 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**: All of audit rules for a single audits did not match the actual server configuration values. Failed compliance is represented with 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 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 that will display all Compliance Categories, as shown in Figure 3-3.

Figure 3-3: Compliance Dashboard Category Selector

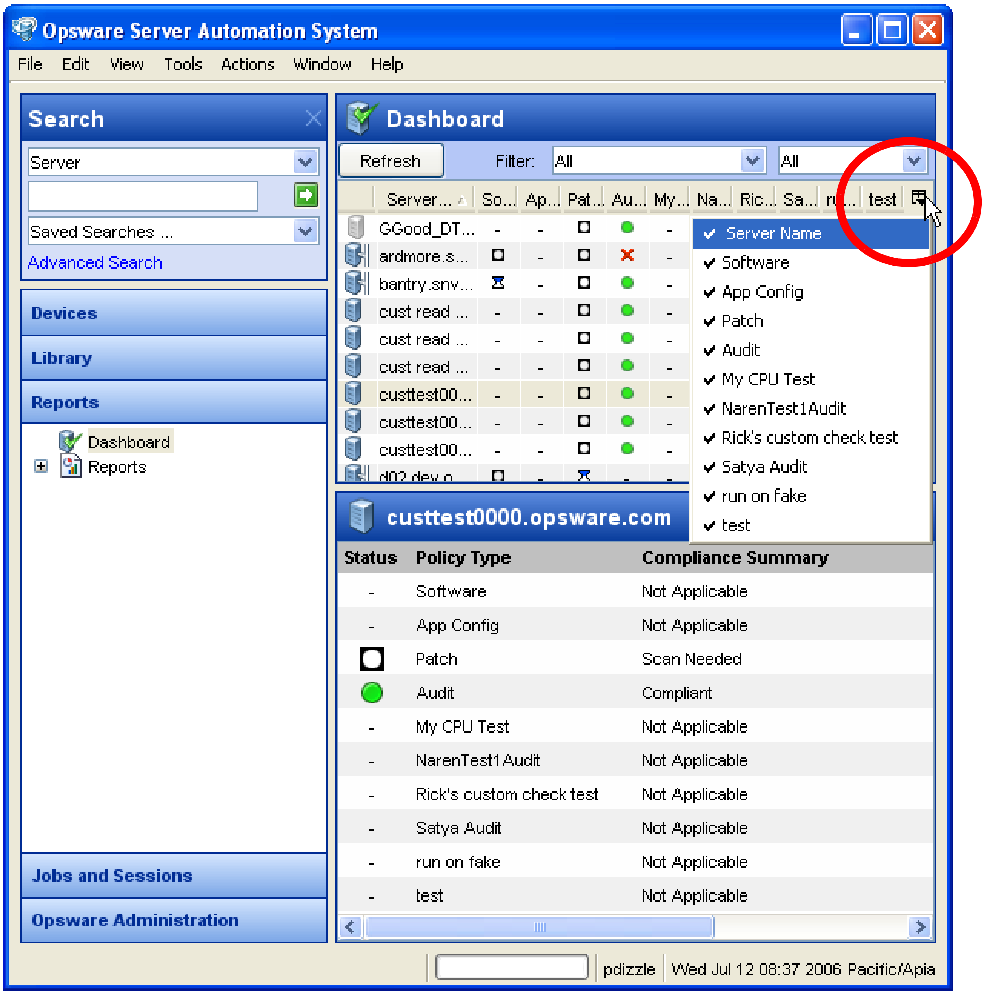

Server...	So...	Ap...	Pat...	Au...	My...	Na...	Ric...	Sa...	ru...	test
GGood_DT...	-	-	□	●	-					
ardmore.s...	□	-	□	✗	-					
bantry.snv...	⧗	-	□	●	-					
cust read ...	-	-	□	●	-					
cust read ...	-	-	□	●	-					
cust read ...	-	-	□	●	-					
custtest00...	-	-	□	●	-					
custtest00...	-	-	□	●	-					
custtest00...	-	-	□	●	-					
d02.dev.o...	□	-	⧗	-	-					

Status	Policy Type	Compliance Summary
-	Software	Not Applicable
-	App Config	Not Applicable
□	Patch	Scan Needed
●	Audit	Compliant
-	My CPU Test	Not Applicable
-	NarenTest1Audit	Not Applicable
-	Rick's custom check test	Not Applicable
-	Satya Audit	Not Applicable
-	run on fake	Not Applicable
-	test	Not Applicable

3. Select an individual audit by name that you want to display in the Compliance Dashboard.

Audit Remediate Options

The types of remediate options you can perform with audits in the Compliance Dashboard depend on if you are working with all scheduled audits or individual audits.

For all scheduled audits (listed as audit in the Compliance Dashboard) run against a selected server, you can perform the following remediate action:

- **Details**: Opens the Server Explorer window showing the Audit & Remediation view, listing 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 displaying 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 Chapter 2, “Viewing and Remediating Audit Results” on page 128 of this guide.

Duplex Compliance

Duplex compliance determines whether or not the duplex of all a server's active interfaces match their corresponding switch ports. For example, if one of a server's network interfaces is set to full duplex and the switch port that interface is connected to is set to half duplex, then the server would be out of compliance.

Duplex compliance is determined by the following criteria:

- **Compliant**: A server will be compliant if the duplex of all of its active network interfaces matches the corresponding switch ports that the interfaces are connected to, the duplex compliance state will be compliant, and 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 that the interface is connected to, 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 of 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 to server where the duplex mismatch occurs so you can troubleshoot problem.

For more information on duplex mismatch, see “Duplex Mismatch” on page 293.

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 to show only the compliance tests you want to see.

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.

You can also choose to show or hide any of the compliance tests, including individually scheduled audits. For example, you might want to only see software compliance test. Or, you might want to see all scheduled audits and specific individually scheduled audits.

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 **Devices ➤ 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 that will 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 an encoding if you want the saved file to use a specific encoding scheme.
4. Click **Export**.

Chapter 4: Reports

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Reports
- Reports Features
- Opsware SAS Client Reports
- User Permissions
- Launching the Reports Feature
- Reports Display
- Running a Report
- Report Results

Overview of 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 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.

Reports Features

The 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.
- Export reports in .html and .xls formats to your local file system for use within your organization.

Opsware 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
Compliance Reports	Software Policy Compliance (All Servers)
	Software Policy Compliance by Customer
	Software Policy Compliance by Facility
	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

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

Table 4-1: SAS Client *Reports (continued)*

REPORT FOLDER	REPORT TITLE
	Users Created in the Last 30 Days User Groups Membership
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 289
- “Audit and Remediation” on page 63
- “Exploring Servers and Jobs in SAS Client” on page 129
- “Patch Management for Windows” on page 189
- “Patch Management for Unix” on page 257
- “NAS Integration” in the *Opsware® SAS User’s Guide: Server Automation*
- “Server Management in Opsware Command Center” in the *Opsware® 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. 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 *Opsware® 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

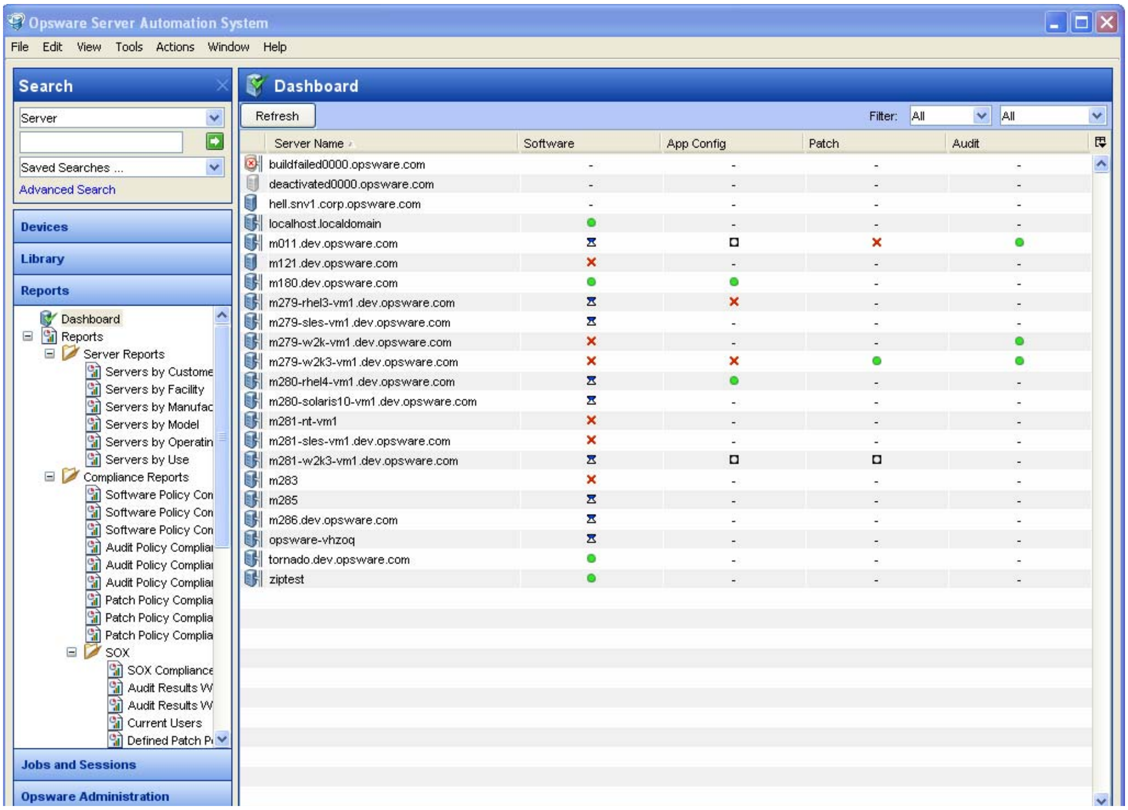

Server Name	Software	App Config	Patch	Audit
buildfailed0000.opsware.com	-	-	-	-
deactivated0000.opsware.com	-	-	-	-
hell.snv1.corp.opsware.com	-	-	-	-
localhost.localdomain	●	-	-	-
m011.dev.opsware.com	⧗	□	✗	●
m121.dev.opsware.com	✗	-	-	-
m180.dev.opsware.com	●	●	-	-
m279-rhel3-vm1.dev.opsware.com	⧗	✗	-	-
m279-sles-vm1.dev.opsware.com	⧗	-	-	-
m279-w2k-vm1.dev.opsware.com	✗	-	-	●
m279-w2k3-vm1.dev.opsware.com	✗	✗	●	●
m280-rhel4-vm1.dev.opsware.com	⧗	●	-	-
m280-solaris10-vm1.dev.opsware.com	⧗	-	-	-
m281-nt-vm1	✗	-	-	-
m281-sles-vm1.dev.opsware.com	✗	-	-	-
m281-w2k3-vm1.dev.opsware.com	⧗	□	□	-
m283	✗	-	-	-
m285	⧗	-	-	-
m286.dev.opsware.com	⧗	-	-	-
opsware-vhzoq	⧗	-	-	-
tornado.dev.opsware.com	●	-	-	-
ziptest	●	-	-	-

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 *Opsware*® *SAS User's Guide: Server Automation*.

Dashboard

The dashboard is the default display when you select Reports from the Navigation pane. See "Compliance Dashboard" on page 155.

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 Opsware 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 reports you create to support the needs of your environment. For information about how to create a custom report, contact Opsware, Inc.

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

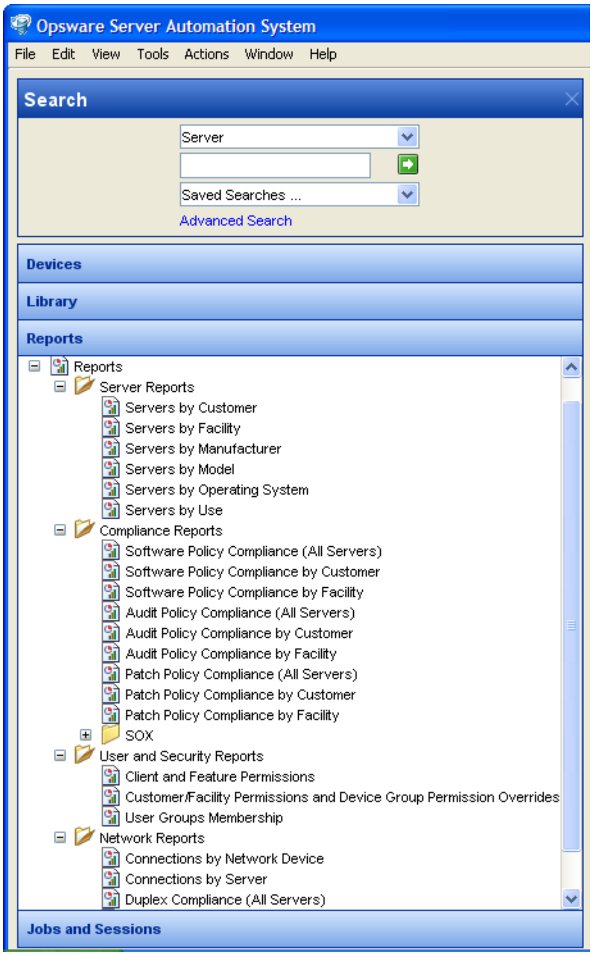

- Reports
  - Server Reports
    - Servers by Customer
    - Servers by Facility
    - Servers by Manufacturer
    - Servers by Model
    - Servers by Operating System
    - Servers by Use
  - Compliance Reports
    - Software Policy Compliance (All Servers)
    - Software Policy Compliance by Customer
    - Software Policy Compliance by Facility
    - 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
    - SOX
  - User and Security Reports
    - Client and Feature Permissions
    - Customer/Facility Permissions and Device Group Permission Overrides
    - User Groups Membership
  - Network Reports
    - Connections by Network Device
    - Connections by Server
    - Duplex Compliance (All Servers)

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 184.

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 a folder within it.
3. Select a report listed in the folder.
4. If there are no report parameters in the Content pane, click **Run**.
5. If there are report parameter 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 184.

Modifying Report Parameters

To modify the default parameters to 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 view 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 View 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

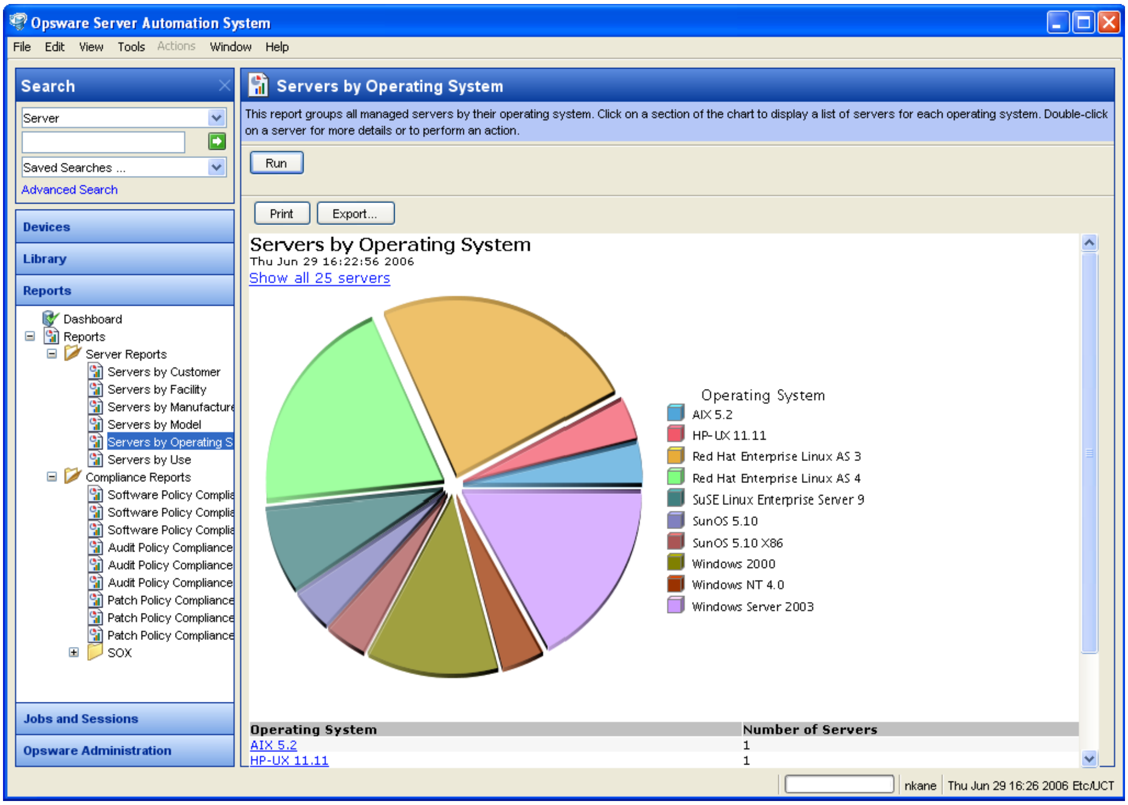

This report groups all managed servers by their operating system. Click on a section of the chart to display a list of servers for each operating system. Double-click on a server for more details or to perform an action.

Servers by Operating System
Thu Jun 29 16:22:56 2006
Show all 25 servers

Operating System	Number of Servers
AIX 5.2	1
HP-UX 11.11	1

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

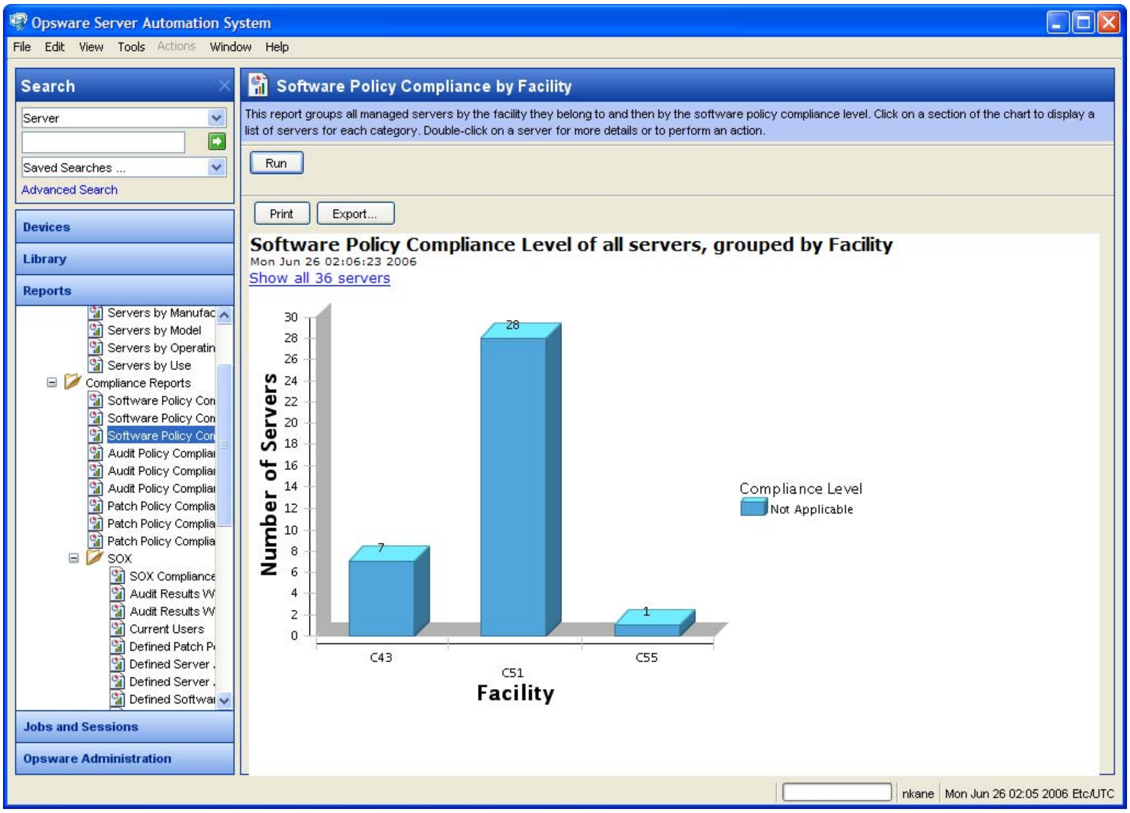

This report groups all managed servers by the facility they belong to and then by the software policy compliance level. Click on a section of the chart to display a list of servers for each category. Double-click on a server for more details or to perform an action.

Software Policy Compliance Level of all servers, grouped by Facility

Mon Jun 26 02:06:23 2006

Show all 36 servers

List View 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

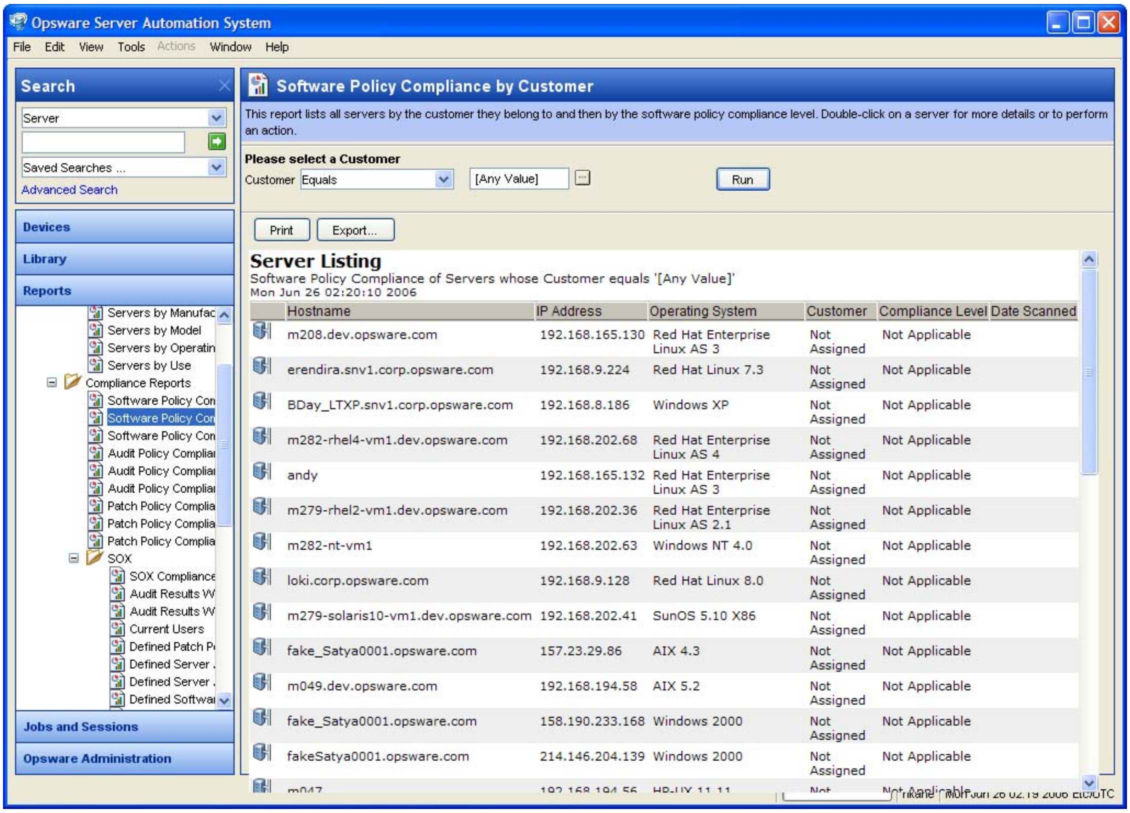

Software Policy Compliance by Customer

This report lists all servers by the customer they belong to and then by the software policy compliance level. Double-click on a server for more details or to perform an action.

Please select a Customer

Customer: Equals [Any Value]

Server Listing

Software Policy Compliance of Servers whose Customer equals '[Any Value]'
Mon Jun 26 02:20:10 2006

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	
erendira.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	Not Applicable	

Exporting a Report

You can export a report for use in other applications in your environment and to attach to email for 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

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 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.

The Opsware 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.

In Patch Management, you can separately schedule when you want patches imported from Microsoft (either automatically or on demand) into Opsware 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 your reboot preferences to adopt, override, postpone, or suppress the vendor's reboot options.

To provide flexibility in how you identify and distribute patches on managed servers or server groups, Patch Management allows you to create patch policies that define groups of patches you need to install. By creating a patch policy and attaching it to a server or server group, you can effectively manage which patches get installed where in your organization. In case you need 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.

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 you need 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 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 install, if a compliance scan has not been run or the installed patch has not been registered, Opsware SAS does not know about it. The preview process provides 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 obsolete other patches or are obsoleted by other patches.

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 install, you can also preview a patch uninstall.

To help you track the patch state of servers or server groups, 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.

Patch Management for Windows Features

The Patch Management for Windows feature is fully integrated with Opsware SAS. This feature leverages from the Opsware SAS server automation features. For example, Opsware SAS maintains a central database (the Model Repository) that includes detailed information about every server under management, the patches and software (applications) installed on managed servers, and the patches and software available for installation. You can use this data to determine the severity of a server’s exposure to a newly discovered vulnerability, and to evaluate the benefits of rolling out a patch versus the costs that might be incurred during downtime and testing activities.

By automating the patching procedure, the Patch Management can reduce the amount of downtime required for patching. Opsware SAS also allows you to separately schedule patch downloads and patch installations so that patching occurs during off-peak hours.

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 identify servers that require patching
- Auditing abilities so that security personnel can track the deployment of important patches

Locales Supported

The Patch Management for Windows feature supports English, Japanese, and Korean locales with the following configuration:

- One locale per mesh (Patches from different locales cannot co-exist in the same mesh.)
- Locale setting of all managed servers is the same as the locale of all patches in a mesh
- No detection that a managed server's locale does not match the current locale in a mesh

You can use the windows-patch-locale script to configure a locale that is not one of the four locales supported by MBSA 1.2.1. See the *Opsware® SAS Administration Guide*.

Opsware SAS Client 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 which columns of patch metadata data to display, depending on what you find useful for any given patch.

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 Opsware SAS Client Library

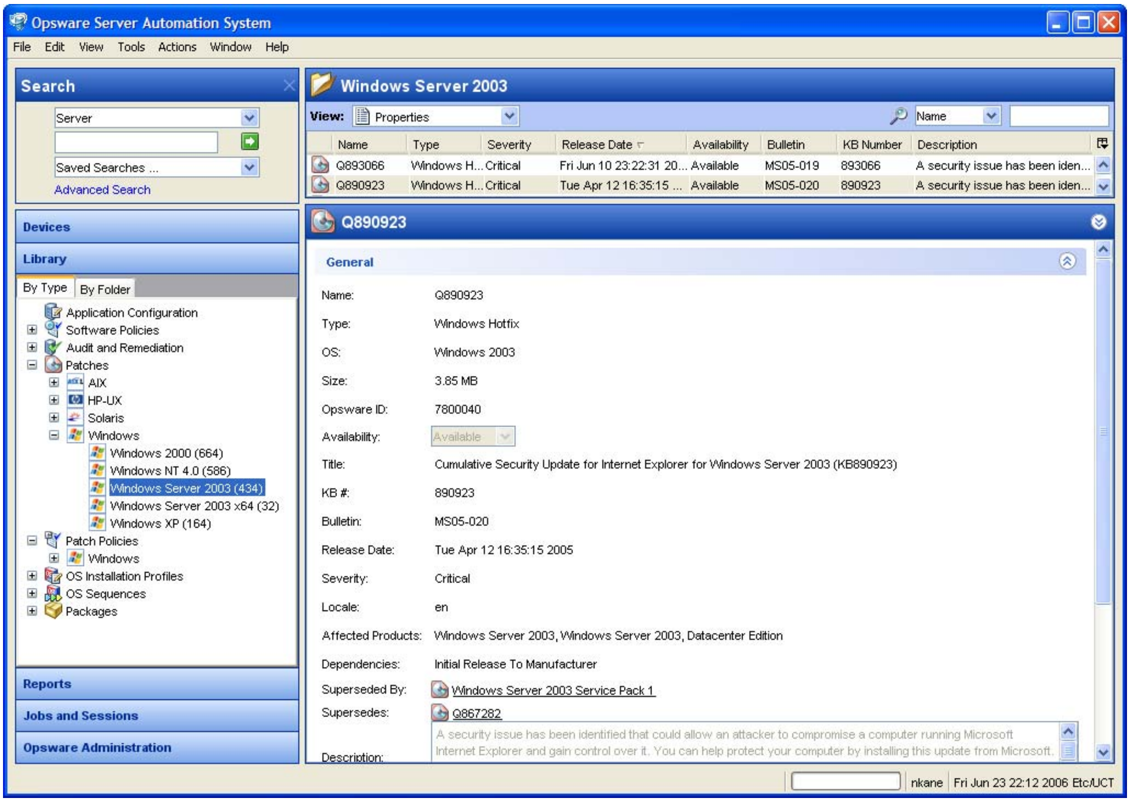

Name	Type	Severity	Release Date	Availability	Bulletin	KB Number	Description
Q893066	Windows H...	Critical	Fri Jun 10 23:22:31 20...	Available	MS05-019	893066	A security issue has been iden...
Q890923	Windows H...	Critical	Tue Apr 12 16:35:15 ...	Available	MS05-020	890923	A security issue has been iden...

Name: Q890923

Type: Windows Hotfix

OS: Windows 2003

Size: 3.85 MB

Opsware ID: 7800040

Availability: Available

Title: Cumulative Security Update for Internet Explorer for Windows Server 2003 (KB890923)

KB #: 890923

Bulletin: MS05-020

Release Date: Tue Apr 12 16:35:15 2005

Severity: Critical

Locale: en

Affected Products: Windows Server 2003, Windows Server 2003, Datacenter Edition

Dependencies: Initial Release To Manufacturer

Superseded By: Windows Server 2003 Service Pack 1

Supersedes: Q867282

Description: A security issue has been identified that could allow an attacker to compromise a computer running Microsoft Internet Explorer and gain control over it. You can help protect your computer by installing this update from Microsoft.

nkane Fri Jun 23 22:12 2006 Etc/UCT

Patch Management for Windows Prerequisites

You must have Internet Explorer 6.0 or later installed on a managed server to support Patch Management. This version of Internet Explorer supports the Microsoft XML parser and related DLL files that are required for its native Microsoft Baseline Security Analyzer (MBSA) tool (mbsacli.exe). Opsware SAS uses version 1.2.1 and 2.0 of the MBSA tool for patch management. Vendor-recommended patches that are installed during the patch remediate process are based on MBSA 2.0.

You must also have the following versions of Opsware Agent to support certain functionality in Patch Management:

Table 5-2: Opsware Agent Requirements

PATCH FUNCTIONALITY	OPSWARE AGENT VERSION
Install Patch task window	Opsware Agent 4.5 or later
Uninstall Patch task window	Opsware Agent 4.5 or later
Remediate task window	Opsware Agent 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 which patches 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. See the *Opsware® SAS Administration Guide* for information about automatically importing the Microsoft patch database.

Once every twenty-four hours, the Opsware Agent on a Windows server compares the server's current state against the Microsoft patch database (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 the data is stored 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 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, instructions on when to reboot, 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 5-1: 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. Opsware SAS uses these tools behind the scenes, which 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. Detects which application files 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: Microsoft framework for patch update.

See “Importing Windows Utility Files” on page 236.

Windows Hotfixes

After a Microsoft Windows hotfix is imported into Opsware 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 Opsware 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 normally require that the server is rebooted, the reboot can sometimes be postponed until all hotfixes have been installed. The user performing the install must first run Qchain.exe before doing the reboot, in order to be sure 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, they will require a reboot and 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 could be safely postponed until the end of the chained installation.

Opsware SAS will *not* reorder the install 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 install operation.

Opsware SAS runs Qchain.exe on the managed server after the install of each Windows hotfix and before any associated reboot to guard 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 any information about your operational environment that is available in Opsware SAS 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: 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 from operating system vendors and determines which vendor patches 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 are able to assess the necessity of applying patches issued by vendors. A policy setter is 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, 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 the 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 server groups. 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 or create new policies, mark them 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

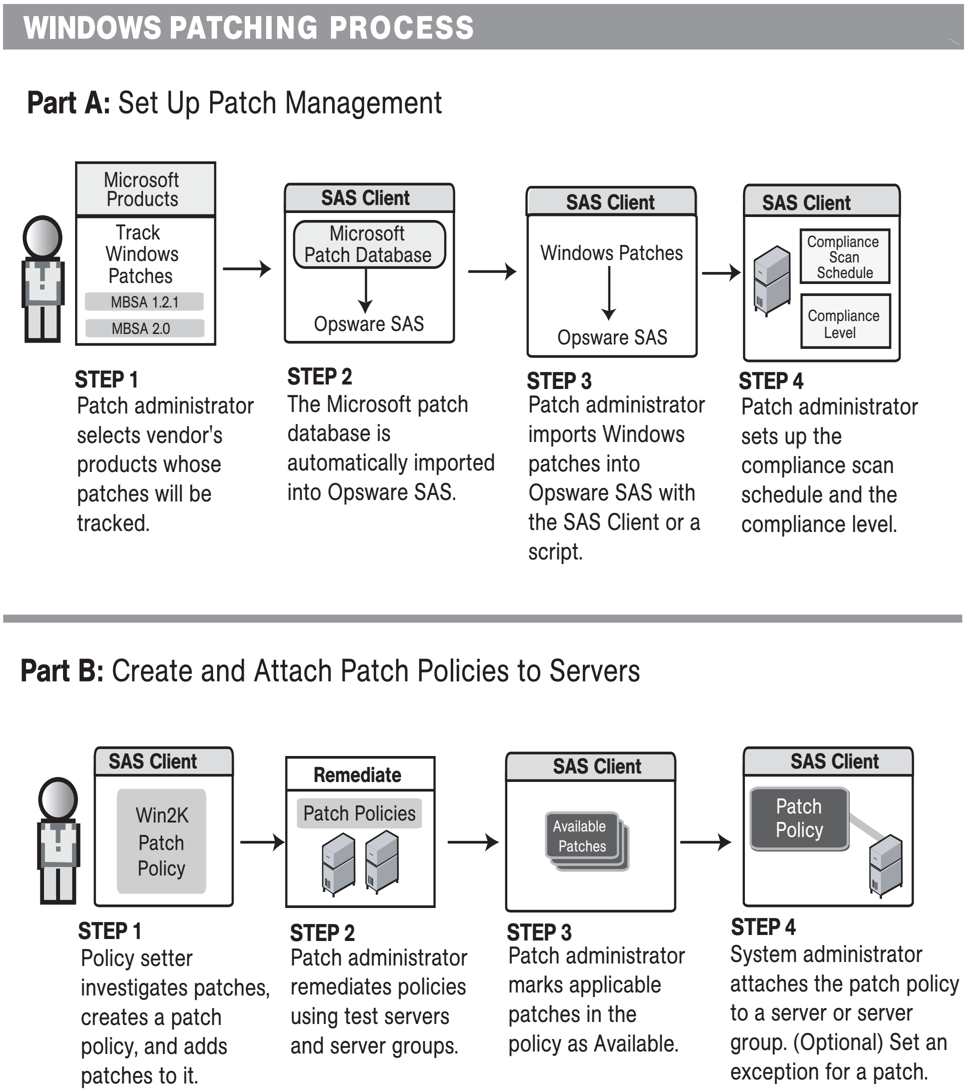

STEP 1
Patch administrator selects vendor's products whose patches will be tracked.

STEP 2
The Microsoft patch database is automatically imported into Opsware SAS.

STEP 3
Patch administrator imports Windows patches into Opsware SAS with the SAS Client or a script.

STEP 4
Patch administrator sets up the compliance scan schedule and the compliance level.

STEP 1
Policy setter investigates patches, creates a patch policy, and adds patches to it.

STEP 2
Patch administrator remediates policies using test servers and server groups.

STEP 3
Patch administrator marks applicable patches in the policy as Available.

STEP 4
System administrator attaches the patch policy to a server or server group. (Optional) Set an exception for a patch.

Figure 5-3: Windows Patching Process: Part C and Part D

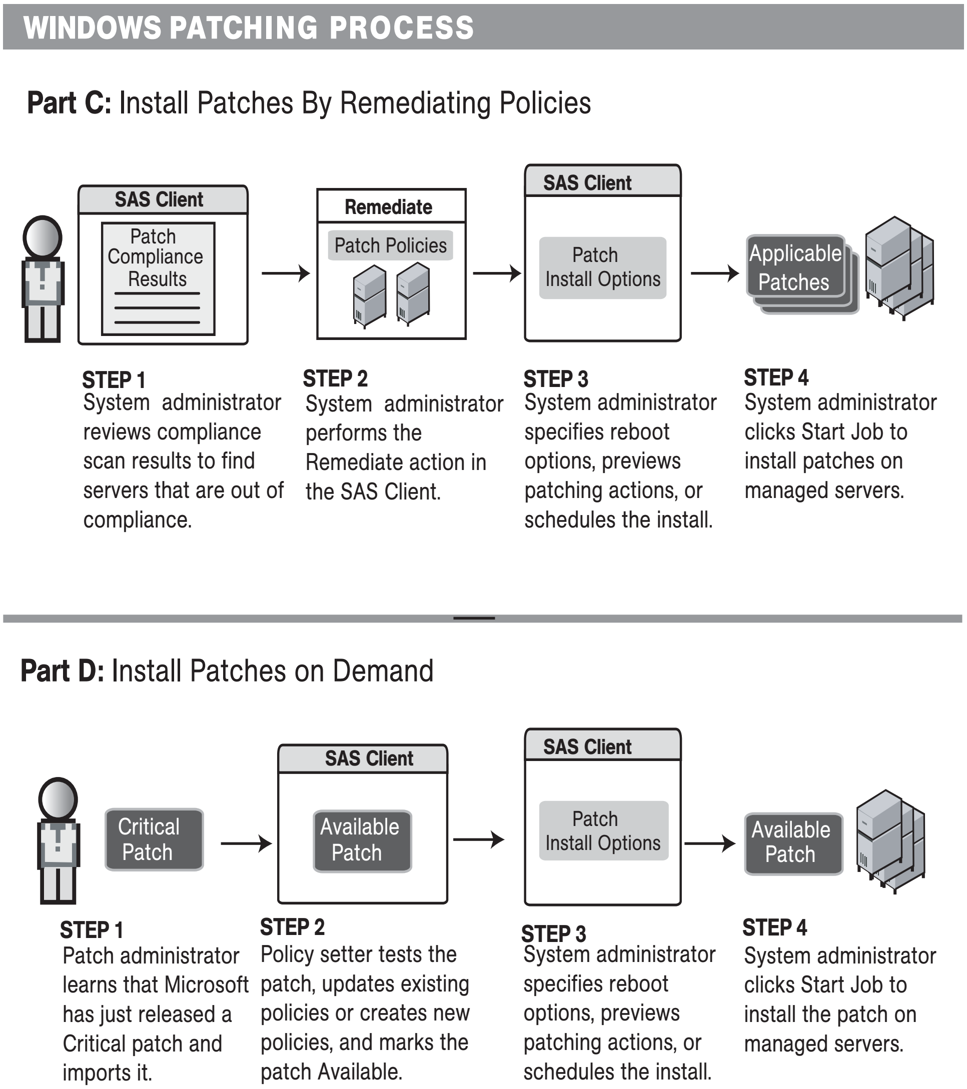

WINDOWS PATCHING PROCESS

Part C: Install Patches By Remediating Policies

STEP 1 System administrator reviews compliance scan results to find servers that are out of compliance.

STEP 2 System administrator performs the Remediate action in the SAS Client.

STEP 3 System administrator specifies reboot options, previews patching actions, or schedules the install.

STEP 4 System administrator clicks Start Job to install patches on managed servers.

Part D: Install Patches on Demand

STEP 1 Patch administrator learns that Microsoft has just released a Critical patch and imports it.

STEP 2 Policy setter tests the patch, updates existing policies or creates new policies, and marks the patch Available.

STEP 3 System administrator specifies reboot options, previews patching actions, or schedules the install.

STEP 4 System administrator clicks Start Job to install the patch on managed servers.

Patch Properties

A patch is a piece of object code (binaries) that is inserted into (patched into) an executable program to temporarily fix a known defect. Patch Management displays detailed information (properties) about a patch.

Figure 5-4: Windows Patch Properties

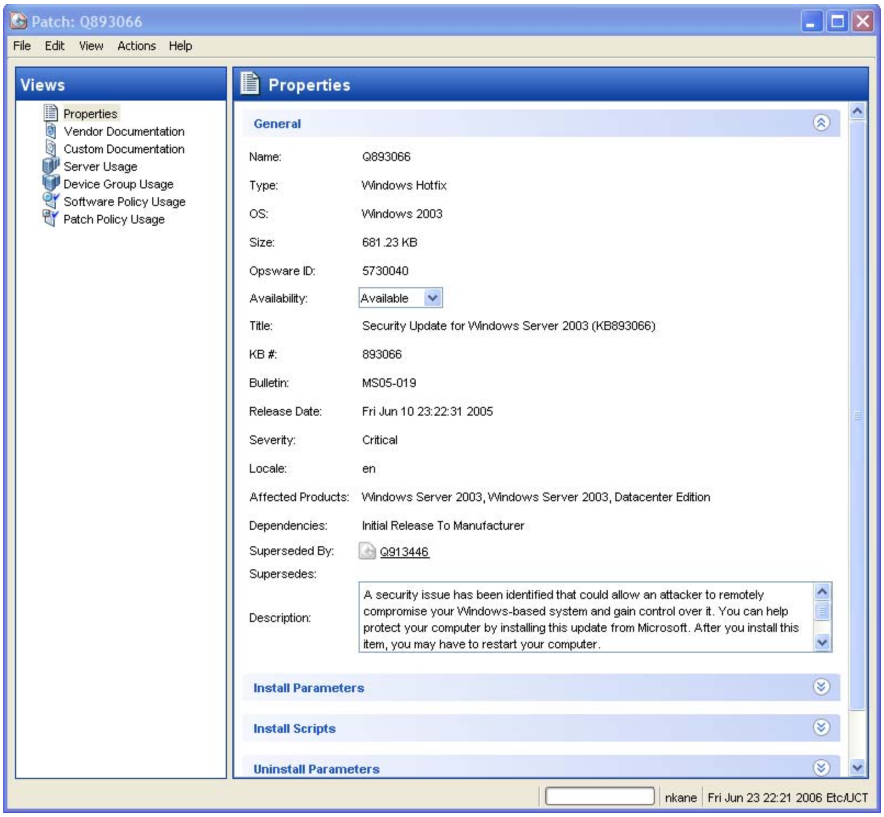

Name: Q893066
Type: Windows Hotfix
OS: Windows 2003
Size: 681.23 KB
Opsware ID: 5730040
Availability: Available
Title: Security Update for Windows Server 2003 (KB893066)
KB #: 893066
Bulletin: MS05-019
Release Date: Fri Jun 10 23:22:31 2005
Severity: Critical
Locale: en
Affected Products: Windows Server 2003, Windows Server 2003, Datacenter Edition
Dependencies: Initial Release To Manufacturer
Superseded By: Q913446
Supersedes:
Description: A security issue has been identified that could allow an attacker to remotely compromise your Windows-based system and gain control over it. You can help protect your computer by installing this update from Microsoft. After you install this item, you may have to restart your computer.

Patch properties includes 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).
- **Opsware ID**: The Opsware SAS unique ID for the patch.

- **Availability**: The status of a patch within Opsware 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 Opsware SAS.
  - **Limited**: The patch has been imported into Opsware SAS but cannot be installed. This is the default patch availability.
  - **Available**: The patch has been imported into Opsware 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): One of following Microsoft severity ratings for this patch:
  - **Critical**: A patch whose exploitation could allow the propagation of an internet worm, without user action.
  - **Important**: A patch whose exploitation could result in a compromise of the confidentiality, integrity, or availability of user data, or of the integrity or availability of processing resources.
  - **Moderate**: Exploitability is mitigated to a significant degree by certain factors, such as default configuration, auditing, or difficulty of exploitation.
  - **Low**: A patch whose exploitation is extremely difficult, or whose impact is minimal.
- **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 obsoleted (replaced) by. This relationship does not apply to MBSA 1.2.1 patches.
- **Supersedes** (Optional): A list of patches that this patch obsoletes (replaces). 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 obsolete (supersede) or are obsoleted (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 obsoleted 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—where 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 Opsware 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 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 Opsware 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 View 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**.

Finding Vendor-Recommended Patches

To find out which patches 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 Opsware 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 in the drop-down list to display these types of patches for the selected server.

Finding Servers That Have a 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 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 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 ➤ 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 Opsware 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. See "Automatically Importing Windows Patches" in the *Opsware® SAS Administration Guide* for information about the script.

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**.

Exporting a 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 Patch Information

You can export the following information about patches that are installed on a server, patches that are recommended by the vendor, and patches with model information on the selected server (such as patch policies or patch policy exceptions) *and* are also recommended by the vendor into a comma-separated value (CSV) file:

- Server Name
- OS
- Service Pack (This is the service pack level of the server being reported, such as Service Pack 0, Service Pack 1, and so on.)
- KB#
- Bulletin (This is 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
- Time Queried (This is the last software registration by the Agent.)
- Time Installed (This is the time the patch was installed.)
- Type (This is the patch type.)
- Compliance Level (This is an integer that represents the compliance level.)
- Compliance (This is a text description that displays when you place your cursor over the Compliance column in the Patch Preview pane.)

- Exception Type
- Exception Reason

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 **Actions** menu, select **Export Patch Info to CSV**.
4. In the Export to CSV window, navigate to a folder and enter the file name.
5. 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.
6. 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

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 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.

1 From the Navigation pane, select Library ➤ By Type ➤ Patches.

1 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.

2 From the Content pane, select a patch.

3 From the **Actions** menu, select **Delete Patch**.

4 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 which Windows patches 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 leverage from the patch policies and exceptions and you perform ad hoc patch installs, you need to remediate to get the applicable patches installed on servers.

Patch Policy

A patch policy is a group of patches that you want to install on Opsware SAS managed servers. All patches in a patch policy must apply to the same Windows operating system.

A patch policy provides broad flexibility in how you distribute 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) to install 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 server groups. 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 task 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–to explain how patch policies compare with corresponding patch policy exceptions.

Patch Management supports the following types of patch policies:

- **User-defined patch policy**: This allows an Opsware SAS user to specify which patches 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) to apply only those patches that their environment needs.

- **Vendor-recommended patch policy**: Membership of patches is defined by what MBSA recommends 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 it is associated with. See the *Opsware® 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 server groups.
- More than one patch policy can be attached to a server or public server group.
- A patch policy can be created, edited, and deleted by users who have permissions. 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 in 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 server group) by deselecting or adding individual patches for 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—to 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 detailed 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 server 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 server group), you control which patches should be installed or not installed on a server. A precedence hierarchy in Patch Management delineates how a policy or an exception, whether the policy or exception is attached at the server or server group level, is applied to a patch installation.

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 server group.
- Patch policy exceptions that are attached to a public server group take precedence over patch policies that are attached to a public server group.
- If a server is in multiple public server groups, a Never Installed patch policy exception type always take precedence over an Always Installed patch policy exception type for the same patch.

Remediate Process

To ensure patch compliance, Patch Management identifies vulnerable managed servers and simultaneously deploys patches to many servers when a remediate process is performed. The remediate process applies an entire patch policy and even multiple policies to an operation that examines the patch policy and the managed servers that it is attached to. (A policy must be attached to a server or a server group before you can remediate the policy with that server or server group.)

The remediate process requires that the selected managed server is running Opsware Agent 5.5 and a Windows 2000 Service Pack 3 (or higher) operating system or a Windows 2003 operating system. You cannot use the remediate 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 Opsware Agent 5.5. Use the Install Patch task window to install patches on servers that are running these operating systems or Opsware Agents 4.5 or earlier.

As a best practice, each time a policy setter reviews the latest Microsoft patch releases and subsequently updates a patch policy (by adding new patches to a policy), a remediate should be performed. In these situations, a remediate process provides demand forecasting information that allows you to determine how patch policy changes would impact servers that this policy is attached to.

If the remediate 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 remediate 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 remediate process will uninstall it.

After Opsware SAS determines what packages need to be installed to complete the remediate operation, remediate uses a set of standard system utilities to complete the operation. See “Supporting Technologies for Patch Management” on page 197.

To help you optimally manage the conditions under which patch policies are remediated, Patch Management allows you to specify remediate options and pre and post remediate scripts, and set up ticket IDs and email notifications to alert you about the status of the remediate process. The Remediate task window guides you through setting up these conditions.

Figure 5-5: Remediate Task Window

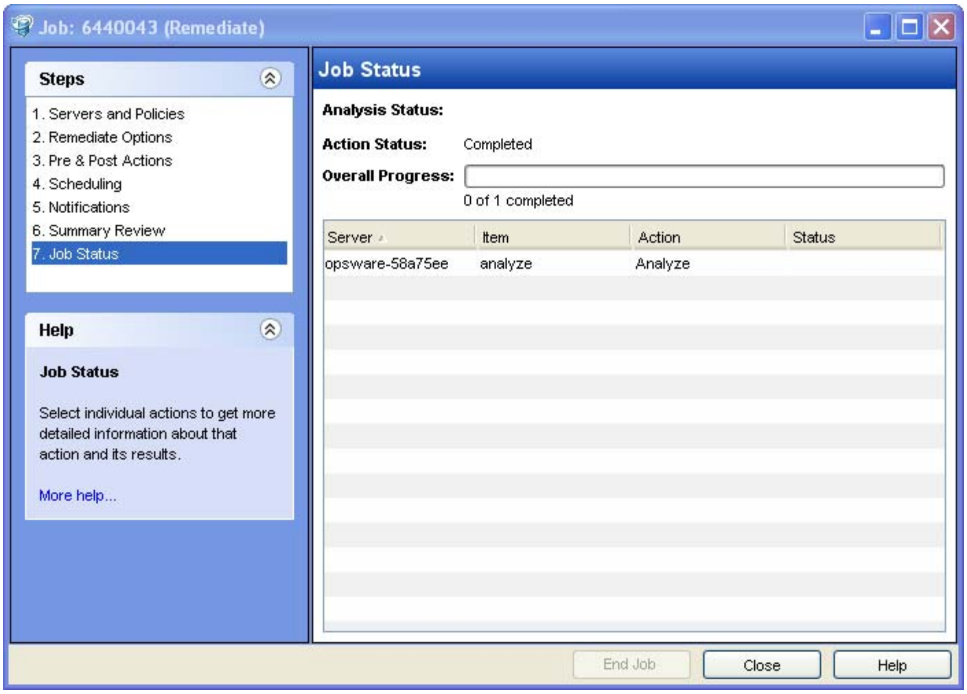

Job: 6440043 (Remediate)

Steps

1. Servers and Policies
2. Remediate Options
3. Pre & Post Actions
4. Scheduling
5. Notifications
6. Summary Review
7. Job Status

Help

Job Status

Select individual actions to get more detailed information about that action and its results.

More help...

Job Status

Analysis Status:

Action Status: Completed

Overall Progress:

0 of 1 completed

Server	Item	Action	Status
opsware-58a75ee	analyze	Analyze	

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 remediate operation for a server group, patches will only be remediated if any server in the server 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 Servers.

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 task window appears: Servers and Server Groups.
For instructions on each step, see the following sections:

- Setting Remediate Options
- Setting Reboot Options for Remediate
- Specifying Pre and Post Install Scripts for Remediate
- Scheduling a Patch Installation for Remediate
- Setting Up Email Notifications for Remediate
- Previewing a Remediate

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 remediate 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 Remediate task 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 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 these options, perform the following steps:

1. From the Remediate task 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 install to be scheduled separately.
3. Select the Error Options check box if you want the remediate 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 task window.

Setting Reboot Options for Remediate

To minimize the downtime that server reboots can cause, you can control when servers reboot during a patch installation.

You 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 task window

When you are selecting reboot options in the Remediate task window, Opsware, Inc. recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the task window. 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 in the task window. 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 Opsware control).

The following options in the Remediate task window determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Remediate task 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 task 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 task window.

Specifying Pre and Post Install Scripts for Remediate

For each patch remediate, you can specify a command or script to run before remediate or after remediate. 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 remediate process:

- **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 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-Instal**l: 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 task 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 Command Center. 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 task window.

Scheduling a Patch Installation for Remediate

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 Remediate task 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 install 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 install immediately.

   - **Run Task At**: This enables you to specify a date and time that you want the download or install performed.

3. Click **Next** to go to the next step or click **Cancel** to close the Remediate task window.

Setting Up Email Notifications for Remediate

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 task 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 task window.

If you previously selected Staged in the Remediate Options step, the Notifications pane displays notification options for both the download and install phases.

Previewing a Remediate

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 what patches 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, server groups, and patches that are listed in the Summary Step window will be submitted to remediate 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 enables you to see what patches go on what servers (regardless of any patch policy and server group membership changes that may have occurred). If you Preview a remediate that is scheduled for a future time, this same list of servers, server 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 task 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 server groups, you must click Preview again for results that include your changes.

The remediate preview does not report on the behavior of the server as though the patches have been applied.

To preview a policy remediate, perform the following steps:

1. From the Remediate task 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. 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 task 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 install, 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 install.
   - **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 task window. You can stop a job only if it is scheduled.

Verifying Patch Policy Compliance

A patch policy identifies patches that should be installed on a managed server. A patch policy exception identifies a patch that should or should not be installed.

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 and then select a server.
3. The status of patches will display. The Compliance column shows whether a patch is Compliant, Partial, or Non-Compliant.

   - **Non-Compliant** The patch is installed on the server, but is not in the policy, or the patch is not installed on the server but is in the policy.
   - **Partial**: The policy and exception do not agree, and the exception does not have data in the Reason field.
   - **Compliant**: This status indicates one of the following conditions:
     - A patch is installed on the server and is in a policy, or a patch is not installed on the server and is not in a policy.
     - A patch is installed on the server and there are additional patches with the same QNumber in a patch policy or exception. In this case, all patches with the same QNumber are considered installed when Patch Management calculates patch compliance.
     - A patch is not installed on the server and is in a patch policy or has an always install exception, and is not recommended by the vendor. In this case, the patch is considered as if it has a never install exception because it is not recommended by the vendor.

   If the icon is Compliant, no further action is required. If the icon is Partial or Non-Compliant, perform the following steps.

   - **No Indicator:** The patch is installed on the server and it is not in a policy and there is no exception.

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 server groups. You must first detach the policy from the servers or server groups 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.
1. Select a specific Windows operating system.
2. From the Content pane of the main window, select a policy.
3. 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.
1. Select a specific Windows operating system and view the list of Windows patches.
2. From the Content pane, select the patch.
3. From the View drop-down list, select Patch Policies.
4. From the Show drop-down list, select Policies without Patch Added.
5. Select a policy. From the **Actions** menu, select **Add to Patch Policy**.
6. 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.
1. Select a specific Windows operating system and view the list of Windows patches.
2. 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 server group). You must perform this action before you remediate a policy with a server (or server group).

To attach the policy, perform the following steps:

1. From the Navigation pane, select Library ➤ By Type ➤ Patch Policies.
1. Select a specific Windows operating system and view the list of Windows patch policies.
2. From the Content pane, select a patch policy.
3. From the View drop-down list, select Servers (or Server Groups).
4. From the Show drop-down list, select Servers with Policy Not Attached (or Server Groups with Policy Not Attached).
5. From the Preview pane, select one or more servers.
6. From the **Actions** menu, select **Attach Server**.
7. 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.
1. Select a specific Windows operating system and view the list of Windows patch policies.
2. From the Content pane, select a patch policy.
3. From the View drop-down list, select Servers (or Server Groups).
4. From the Show drop-down list, select Servers with Policy Attached (or Server Groups with Policy Attached).
5. From the Preview pane, select one or more servers.
6. From the **Actions** menu, select **Detach Server**.
7. Click **Detach**.

Setting a Patch Policy Exception

A patch policy exception indicates whether the patch is installed during the Remediate action. (The Install Patch and Uninstall Patch actions ignore patch policy exceptions.) A patch policy exception overrides the policy. You specify an exception for a particular patch and server (or server group), 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 patch.
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:

- **Always Install**: The patch should be installed on the server even if the patch is not in the policy.

- **Never Install**: The patch should not be installed on the server, even if the patch is 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 display of patches with exceptions.

9. Click **OK**.

Finding an Existing Patch Policy Exception

You can search for managed servers 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 server group/patch association.

   A never install exception on a patch/server association.

   A never install exception inherited to a server from a server group/patch association.

Copying a Patch Policy Exception

To copy an exception between servers or server groups, perform the following steps:

1. From the Navigation pane, select Library ➤ By Type ➤ Patches.

2. Expand the Patches and select a specific Window 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 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 server 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.

1 Expand the Patches and select a specific Window operating system.

2 From the Content pane, select a patch.

3 From the View drop-down list, select Servers.

4 From the Show drop-down list, select Servers with Exception.

5 From the Preview pane, select a server.

6 From the **Actions** menu, select **Remove Exception.**

Patch Compliance

Patch Management performs conformance tests (compliance checks) against managed servers and public server 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, based on compliance levels and compliance rules. 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 tell you which servers are in compliance (have all required patches installed) and which servers are out of complies (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 as Scan Needed in the GUI. In this case, instead of waiting for the scan schedule to iterate, you can request the scan to run 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 which compliance icons 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 that 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. Double-click on a policy to perform an action.

Patch Policy Compliance (All Servers)

This report groups all managed servers by their patch policy compliance level to show compliant and non-compliant servers. Click on a section of the chart to display a list of servers for a certain compliance level. Double-click on a server for more details or to perform an action.

Patch Policy Compliance by Customer

This report lists all servers by the customer they belong to and then by the patch policy compliance level. Double-click on a server for more details or to perform an action.

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. Click on a section of the chart to display a list of servers for each category. Double-click on a server for more details or to perform an action.

Servers in Compliance With Their Patch Policies

This report lists all managed servers that are in compliance with all of their attached patch policies. Double-click on a server for more details or to perform an action.

Servers Not in Compliance With Their Patch Policies

This report lists all managed servers that are not in compliance with their attached patch policies. Double-click on a server for more details or to perform an action.

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. Double-click on a server for more details or to perform an action.

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. Double click on a server for more details or to perform an action.

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.
- 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. See “Automatically Importing Windows Patches” in the *Opsware® 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 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. See "Automatically Importing Windows Patches" in the *Opsware® SAS Administration Guide* for information about the script.

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 this 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 `mssecure.cab` (MBSA 1.2) or `wsusscan.cab` (MBSA 2.0) file and click **Import**. These files 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 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. See "Automatically Importing Windows Patches" in the *Opsware® SAS Administration Guide* for information about the script.

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. Click **Edit**.

4. In the Edit Patch Properties window, use the include and exclude arrows to select the products whose patches you want to track (from either the Windows MBSA 2.0 tab or the Windows MBSA 1.2.1 tab) and then click **Select**.

   If you select the MBSA 1.2.1 tab and this is a fresh install, the list of Products in Patch Database is empty. Click **Edit** to select the products you want to track patches for.

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 Rules.
3. Select one of the following compliance levels: Policy and Exception, Policy Only, or Customized.

Importing Windows Utility Files

You can import the following Windows Utility files:

- mbsacli20.exe
- mbsacli.exe
- parsembsacli20.exe
- qchain.exe
- WindowsUpdateAgent20-x86.exe
- WindowsUpdateAgent20-x64.exe
- wusscan.dll

To import a Windows Utility file, 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**.

4. In the Import Patch Utility window, specify where you want the file to be imported to in your local file system, the file name, and the file type, such as .exe or .dll. The Importing Utility Update window displays during the process.

Exporting Windows Utility Files

You can export the following Windows Utility files:

- mbsacli20.exe
- mbsacli.exe
- parsembsacli20.exe
- qchain.exe
- WindowsUpdateAgent20-x86.exe
- WindowsUpdateAgent20-x64.exe
- wusscan.dll

To export a Windows Utility file, 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**.
4. In the Export Patch Utility window, specify where you want the file to be imported to in your local file system, the file name, and the file type.

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**.

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 the default command-line arguments that you want to perform the installation.

To help you optimally manage the conditions under which Windows patches are installed, 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 task window guides you through setting up these conditions.

Figure 5-6: Install Patch Task Window

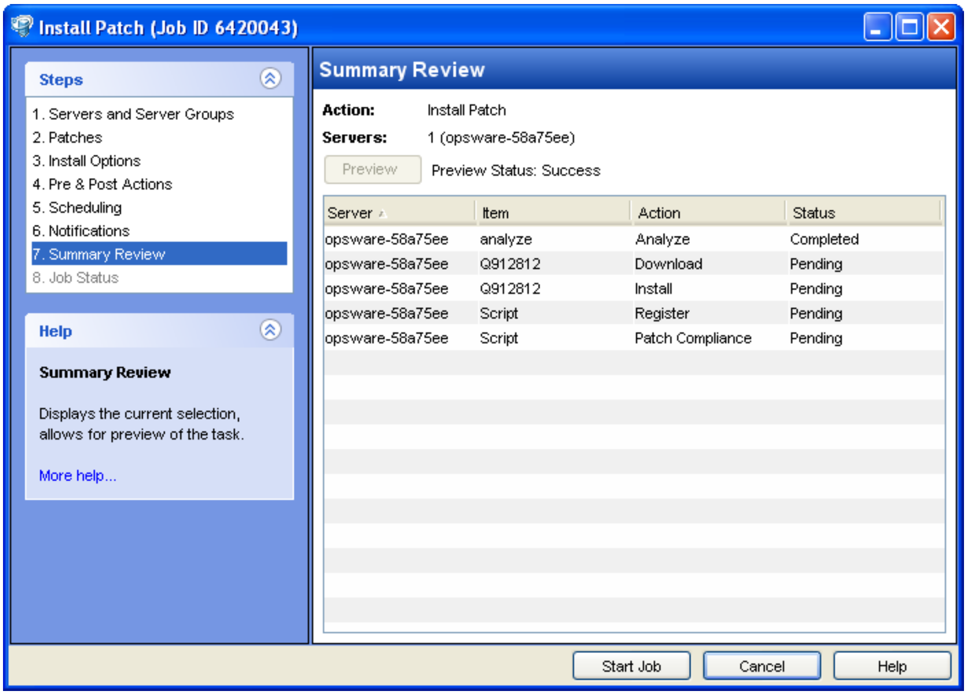

Summary Review

Displays the current selection, allows for preview of the task.

More help...

Summary Review

Action: Install Patch

Servers: 1 (opsware-58a75ee)

Preview Status: Success

Server	Item	Action	Status
opsware-58a75ee	analyze	Analyze	Completed
opsware-58a75ee	Q912812	Download	Pending
opsware-58a75ee	Q912812	Install	Pending
opsware-58a75ee	Script	Register	Pending
opsware-58a75ee	Script	Patch Compliance	Pending

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 in by Opsware SAS. See “Setting Install Options” on page 242 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, to prevent the Patch Management feature from passing in the -z or

-q or -z -q flag. By default, Opsware 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 Opsware SAS uses.

Table 5-2: Default Installation Flags

WINDOWS PATCH TYPE	FLAGS
Windows Hotfix	`-q -z`
Windows Security Rollup Package (treated identically to a Hotfix by the Patch Management feature)	`-q -z`
Windows OS Service Pack	`-u -n -o -q -z`

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 install or uninstall process if you do not click OK in the confirmation dialog. For Hotfixes, the Agent will time out if 5 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 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 Window 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 task appears: Servers and Server Groups.

For instructions on each step, see the following sections:

– Setting Install Options

– Setting Reboot Options for a Patch Install

– Specifying Pre and Post Install Scripts for a Patch Install

– Scheduling a Patch Installation

– Setting Up Email Notifications for a Patch Install

– Previewing a Patch Installation

– Viewing Job Progress of a Patch Install

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 task 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 216 for another method of installing a patch.

Setting 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 task 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 install 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, Opsware 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 task window.

Setting Reboot Options for a Patch Install

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 task window, Opsware, Inc. recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the task window. 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 in the task window. 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 Install Patch task 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 task 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 task window.

Specifying Pre and Post Install Scripts for a Patch Install

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 task 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 Command Center. 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 task window.

Scheduling a 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 task 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 install or download performed.

3. Click **Next** to go to the next step or click **Cancel** to close the Install Patch task 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 Patch Install

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 task 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 task window.

If you previously selected Staged in the Install Options step, the Notifications pane displays notification options for both the download and install phases.

Previewing a 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 (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 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. For example, if a managed server is running Windows 2000 Service Pack 3 (or higher) or Windows 2003, and an Opsware 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 task 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 task 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 task window without launching the install.

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 Install

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 task window, click **Next** to advance to the Job Progress step. This will start the install 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**: Opsware SAS examines the patches needed for the install, 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.
   - **Pre/Post Install/Download Script**: If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstall.
   - **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 *Opsware® 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 task 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 Opsware 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 task window guides you through setting up these conditions.

Figure 5-7: Uninstall Patch Task Window

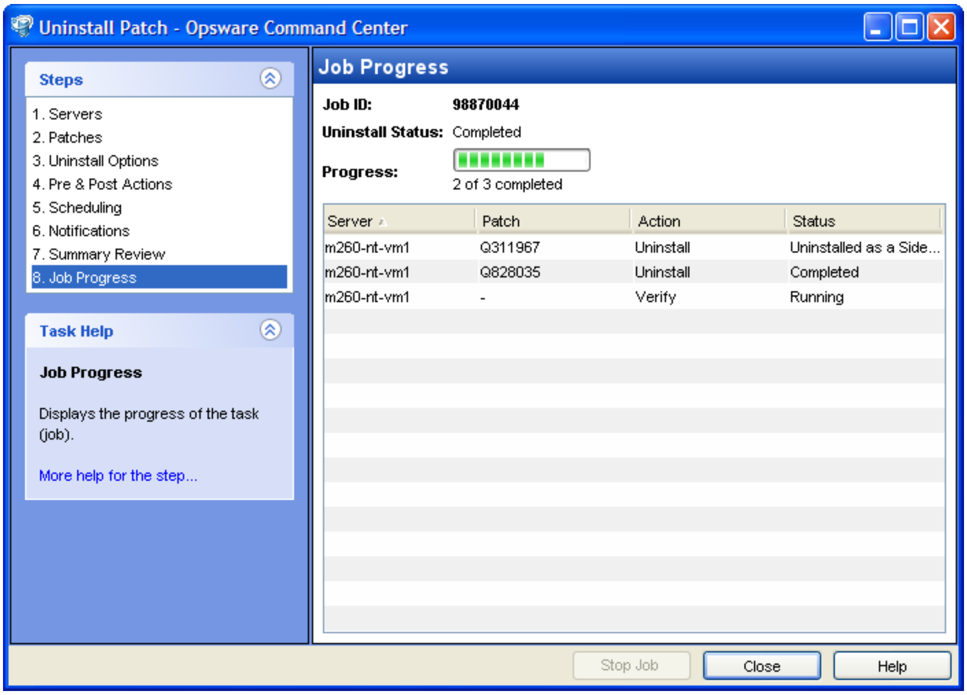

Job Progress

Displays the progress of the task (job).

Job Progress

Job ID: 98870044

Uninstall Status: Completed

Progress: 2 of 3 completed

Server	Patch	Action	Status
m260-nt-vm1	Q311967	Uninstall	Uninstalled as a Side...
m260-nt-vm1	Q828035	Uninstall	Completed
m260-nt-vm1	-	Verify	Running

Uninstallation Flags

You can specify uninstallation flags that are applied whenever a Windows patch is uninstalled. However, Opsware 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 in by Opsware 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, Opsware 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 5-3: Default Uninstallation Flags

WINDOWS PATCH TYPES	FLAGS
Windows Hotfix	`-q -z`
Security Rollup Package	`-q -z`
Windows OS Service Pack	Not uninstallable

Uninstalling a 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 Window 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 task appears: Servers.
For instructions on each step, see the following sections:

- Setting Reboot Options for a Patch Uninstall
- Specifying Pre and Post Install Scripts for a Patch Uninstall
- Scheduling a Patch Uninstallation
- Setting Up Email Notifications for a Patch Uninstall
- Viewing Job Progress of a Patch Uninstall

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 uninstall 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 task 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 uninstall.

To set these options, perform the following steps:

1 From the Uninstall Patch task 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, Opsware 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 task window.

Setting Reboot Options for a Patch Uninstall

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 task window, Opsware, Inc. recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the task window. 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 in the task window. 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 task 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 task 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 task window.

Specifying Pre and Post Install Scripts for a Patch Uninstall

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 uninstall:

- **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 task 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 Opsware Command Center. 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 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 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 uninstall, perform the following steps:

1. From the Uninstall Patch task 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 uninstall in the Summary Review step.
   - **Run Task At**: This enables you to specify a later date and time that you want the uninstall performed.
3. Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch task window.

Setting Up Email Notifications for a Patch Uninstall

You can set up email notifications to alert users when the patch uninstall operation completes successfully or with errors.

To set up email notifications, perform the following steps:

1. From the Uninstall Patch task 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 task window.

Previewing a Patch Uninstallation

The uninstallation preview process provides an up-to-date report about the patch state of servers. The uninstall 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 uninstall have that patch installed (based on the MBSA).

The uninstall preview process does not report or simulate the behavior of a system with patches removed from the server.

To preview a patch uninstall, perform the following steps:

1. From the Uninstall Patch task 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 task window without launching the uninstall.

   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 Uninstall

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 task 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 uninstall, 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 uninstall.
   - **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 task 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
- Policy Management
- Patch Administration for Unix
- Patch Installation
- Patch Uninstallation

Overview of Patch Management for Unix

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 later 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.

After the patch is integrated into your environment, you can make it part of your standard builds with Opsware templates.

Patch Management for Unix 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 identify servers that require patching
- Auditing abilities so that security personnel can track the deployment of important patches

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 you first provision a server with Opsware SAS, the same data is immediately recorded.

When a new patch is issued, you can use the Opsware Command Center to immediately identify which servers require patching. Opsware SAS provides a Software Repository where you upload patches and other software. Users access this software from the Opsware Command Center to install patches on the appropriate servers.

After a server is brought under management, you should install all patches 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, and instructions on when to reboot and how to handle error codes from the pre-install and post-install scripts.

Opsware SAS Client 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, depending on what you find useful for any given patch.

Figure 6-1: Unix Patches in the Opsware SAS Client Library

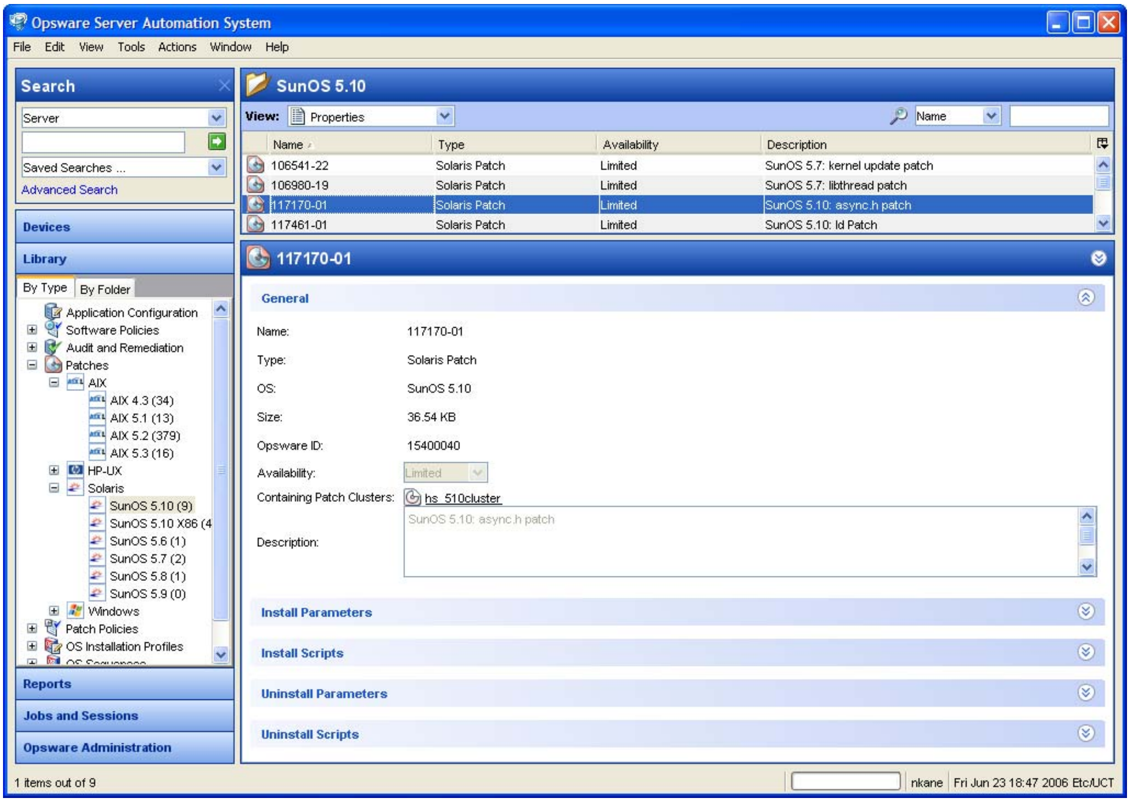

Name	Type	Availability	Description
106541-22	Solaris Patch	Limited	SunOS 5.7: kernel update patch
106980-19	Solaris Patch	Limited	SunOS 5.7: libthread patch
117170-01	Solaris Patch	Limited	SunOS 5.10: async.h patch
117461-01	Solaris Patch	Limited	SunOS 5.10: ld Patch

Name: 117170-01
Type: Solaris Patch
OS: SunOS 5.10
Size: 36.54 KB
Opsware ID: 15400040
Availability: Limited
Containing Patch Clusters: hs_510cluster
Description: SunOS 5.10: async.h patch

Searching for Patches and Policies

In the SAS Client, you can search for any information about your operational environment that is available in Opsware 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 *Opsware® SAS User's Guide: Server Automation*.

Patch Management Roles for Unix

Opsware SAS provides support for rigorous change management by assigning the functions of patch management to two different types of administrators:

- The patch administrator (often referred to as the security administrator), who has the authority to upload and test, and edit patch options
- The system administrator, who applies the patches (that have been approved for use) uniformly and automatically according to the options that the patch administrator specifies

✓ Only the patch administrator should have the Patches permission, which gives access to advanced features. To obtain these permissions, contact your Opsware administrator. See the Permissions Reference appendix in the *Opsware® 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 Opsware SAS, patch administrators are granted specific permissions that allow them to upload patches into Opsware SAS, 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 available in the Opsware Command Center, 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 though 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

Table 6-1: Supported Unix Versions and Patch Types

UNIX VERSIONS	PATCH TYPES
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
Solaris 10	Solaris Patch Solaris Patch Cluster

Underlying Technologies for Patch Management on Unix

Behind the scenes, the Patch Management feature uses utilities and technologies that are specific for a particular operating system. Although the utilities vary, Opsware SAS enables you to perform patch management through a single interface, without having to worry about invoking a number of different patching utilities.

Opsware SAS models the way it treats patches on the way the underlying utility treats a patch. For example, if the Solaris patchadd utility is not able to install one patch contained in a patch cluster, the Solaris utility continues to install the remaining patches in the patch cluster. Opsware 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	HU-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		
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, Opsware 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, Opsware SAS still associates the earlier version of the fileset with the APAR.

When uploading an LPP, Opsware 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 be able 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 Opsware 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 *Opsware® 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, Opsware 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. You can also uninstall the patch cluster if you installed it with the Patch Management feature.

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 Opsware 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 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.

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

A patch is a piece of object code (binaries) that is inserted into (patched into) an executable program to temporarily fix a known defect. Patch Management displays detailed information (properties) about a patch.

Figure 6-2: Unix Patch Properties

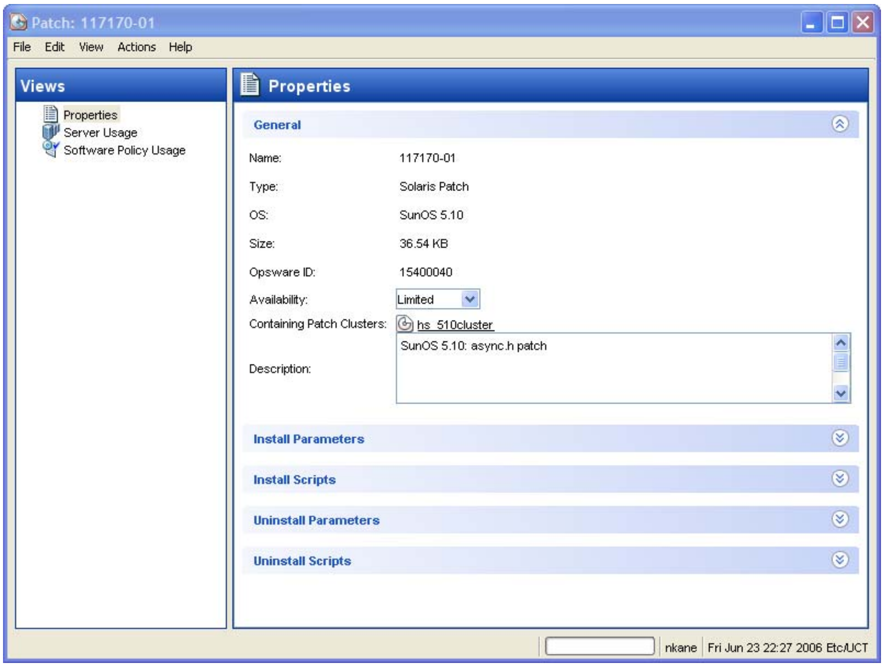

Name: 117170-01

Type: Solaris Patch

OS: SunOS 5.10

Size: 36.54 KB

Opsware ID: 15400040

Availability: Limited

Containing Patch Clusters: hs_510cluster

Description: SunOS 5.10: async.h patch

nkane | Fri Jun 23 22:27 2006 Etc/UCT

Patch properties includes the following information:

- **Name**: The Unix name for the patch.

- **Type**: The type of Unix patch. Table 6-1identifies 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.
- **Opsware ID**: The Opsware SAS unique ID for the patch.
- **Availability**: The status of a patch within Opsware SAS, which can be one of the following:
  - **Limited**: The patch has been imported into Opsware SAS but cannot be installed. This is the default patch availability.
  - **Available**: The patch has been imported into Opsware 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 Opsware 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 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 Opsware 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 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 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 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 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.

1 Expand Patches and select a specific Unix operating system. The Content pane will display all patches associated with that operating system.

2 From the Content pane, select a patch.

3 From the **Actions** menu, select **Delete Patch**.

4 In the Delete Patches windows, click **Delete**.

Policy Management

In Patch Management, 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.

You can choose to have Unix patching in your server environment comply to the model that these software policies define. If you choose to leverage from the software policies and you perform ad hoc patch installs, you need to remediate to get the applicable patches installed on servers.

See “Software Management” on page 289 for more information about creating and remediating software policies to install Unix patches.

Patch Compliance Reports

To help troubleshoot problems that involve patch compliance, you can run and examine several patch compliance reports. By using the Reports feature in the SAS Client you can produce the following patch compliance reports that 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. Click on a section of the chart to display a list of servers for a certain compliance level. Double-click on a server for more details or to perform an action.

Patch Policy Compliance by Customer

This report lists all servers by the customer they belong to and then by the patch policy compliance level. Double-click on a server for more details or to perform an action.

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. Click on a section of the chart to display a list of servers for each category. Double-click on a server for more details or to perform an action.

See the *Opsware® 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 *Opsware® 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 Opsware 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 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*.

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 that the Opsware Agent runs on the managed server to install the patch. You can override the default command-line arguments that you want to perform the installation.

To help you optimally manage the conditions under which Unix patches are installed, 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 task window guides you through setting up these conditions.

Figure 6-3: Install Patch Task Window

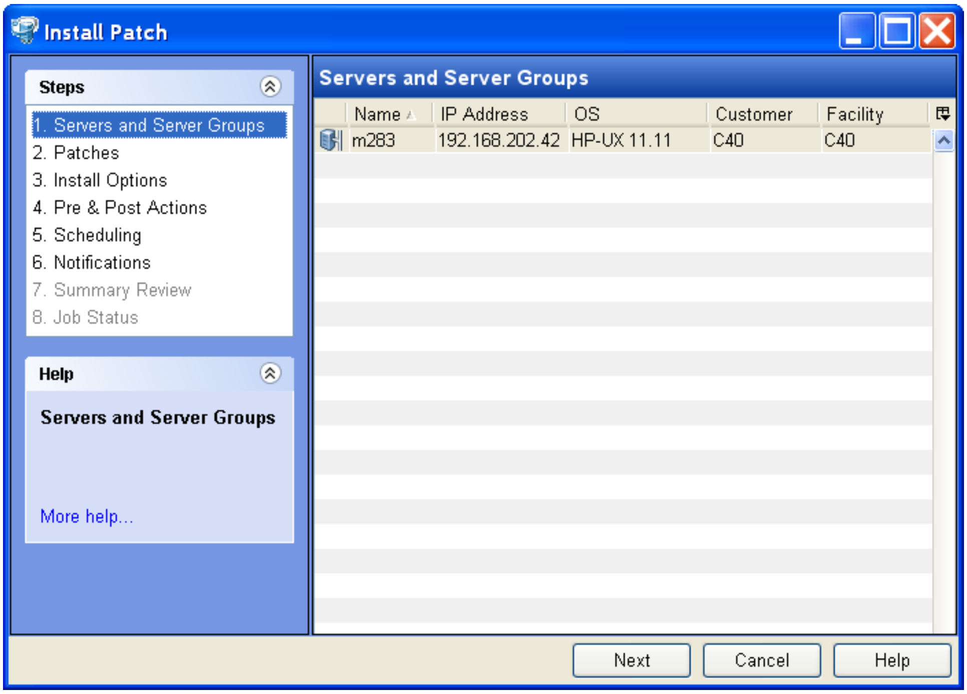

Name	IP Address	OS	Customer	Facility
m283	192.168.202.42	HP-UX 11.11	C40	C40

Installation Flags

You can specify installation flags that are applied whenever a Unix 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 in by Opsware SAS. See “Setting Install Options” on page 276 for information about how to specify commands and flags.

The following table lists the default installation flags that Opsware 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, however, the Patch Management feature does not automatically filter out servers that do not have the application installed that the patch is intended for. Although the Patch Management feature does not prevent you from doing so, you must not attempt to apply application patches to servers that do not have the necessary applications installed.

If an application patch is intended for an application that is running on more than one version of the same operating system, you cannot apply the patch to all of the servers at the same time. An application patch is associated with only one operating system version. You must first select the patch for one operating system, select the servers where the application is installed, and apply the patch. You must repeat this process for each version of the operating system where the application is installed.

Similarly, when uninstalling application patches that are installed on multiple versions of the same operating system, you cannot uninstall all of the patches at the same time. You must repeat the uninstallation process for each version of the operating system where the patch is installed.

Installing a 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.

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 task appears: Servers and Server Groups.
   For instructions on each step, see the following sections:

   - Setting Install Options
   - Setting Reboot Options for a Patch Install
   - Specifying Pre and Post Install Scripts for a Patch Install
   - Scheduling a Patch Installation
   - Setting Up Email Notifications for a Patch Install
   - Previewing a Patch Installation
   - Viewing Job Progress of a Patch Install

   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 task 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 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 task 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 install 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 task window.

Setting Reboot Options for a Patch Install

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 task window, Opsware, Inc. recommends that you use the Unix reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the task 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 task 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 Install Patch task 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 task 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 task window.

Specifying Pre and Post Install Scripts for a Patch Install

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-Instal**l: 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 task 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 Command Center. 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 task window.

Scheduling a 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 task 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 install or download performed.

3. Click **Next** to go to the next step or click **Cancel** to close the Install Patch task 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 Patch Install

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 task 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 task window.

If you previously selected Staged in the Install Options step, the Notifications pane displays notification options for both the download and install phases.

Previewing a 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 task 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 task window without launching the install.

   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 Install

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 task window, click **Next** to advance to the Job Progress step. This will start the install 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**: Opsware SAS examines the patches needed for the install, 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.
   - **Pre/Post Install/Download Script**: If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstall.
   - **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 *Opsware® 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 task 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 Opsware 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 task window guides you through setting up these conditions.

Figure 6-4: Uninstall Patch Task Window

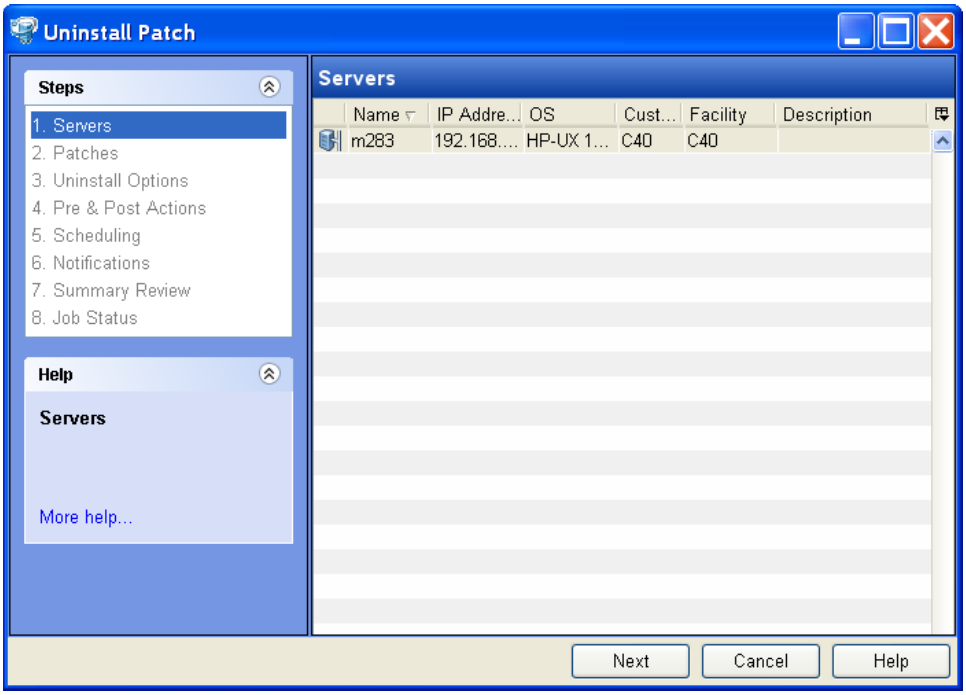

Uninstall Patch

Steps

1. Servers
2. Patches
3. Uninstall Options
4. Pre & Post Actions
5. Scheduling
6. Notifications
7. Summary Review
8. Job Status

Help

Servers

More help...

Servers

Name	IP Addre...	OS	Cust...	Facility	Description
m283	192.168....	HP-UX 1...	C40	C40	

Uninstallation Flags

You can specify uninstallation flags that are applied whenever a Unix patch is uninstalled. However, Opsware 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 in by Opsware SAS.

The following table lists the default uninstallation flags that Opsware 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 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 task appears: Servers.
   For instructions on each step, see the following sections:

   – Setting Reboot Options for a Patch Uninstall

   – Specifying Pre and Post Install Scripts for a Patch Uninstall

   – Scheduling a Patch Uninstallation

   – Setting Up Email Notifications for a Patch Uninstall

   – Viewing Job Progress of a Patch Uninstall

   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 uninstall 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 task 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 uninstall.

To set these options, perform the following steps:

1. From the Uninstall Patch task 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 task window.

Setting Reboot Options for a Patch Uninstall

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 task window, Opsware, Inc. recommends that you use the Unix reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the task 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 task 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 task 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 task 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 task window.

Specifying Pre and Post Install Scripts for a Patch Uninstall

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 uninstall:

- **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 task 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 Opsware Command Center. 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 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 uninstall, perform the following steps:

1. From the Uninstall Patch task 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 uninstall in the Summary Review step.
   - **Run Task At**: This enables you to specify a later date and time that you want the uninstall performed.
3. Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch task window.

Setting Up Email Notifications for a Patch Uninstall

You can set up email notifications to alert users when the patch uninstall operation completes successfully or with errors.

To set up email notifications, perform the following steps:

1. From the Uninstall Patch task 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 task window.

Previewing a Patch Uninstallation

The uninstallation preview process provides an up-to-date report about the patch state of servers. The uninstall 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 uninstall have that patch installed.

The uninstall preview process does not report or simulate the behavior of a system with patches removed from the server.

To preview a patch uninstall, perform the following steps:

1. From the Uninstall Patch task 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 task window without launching the uninstall.

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 Uninstall

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 task 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 uninstall, 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 uninstall.
   - **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 task window.

Chapter 7: Software Management

IN THIS CHAPTER

This section contains the following topics:

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 server, 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 297 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.

The Software Management feature also 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 317 in this chapter for more information.

The Reporting feature in Opsware SAS allows you 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 318 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 *Opsware® SAS Policy Setter's Guide* for information about uploading packages, and creating and managing software policies.

Software Installation Process

The software installation process 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 and remediating non-compliant servers, and generating software compliance reports across servers.

Figure 7-1: Software Deployment Process

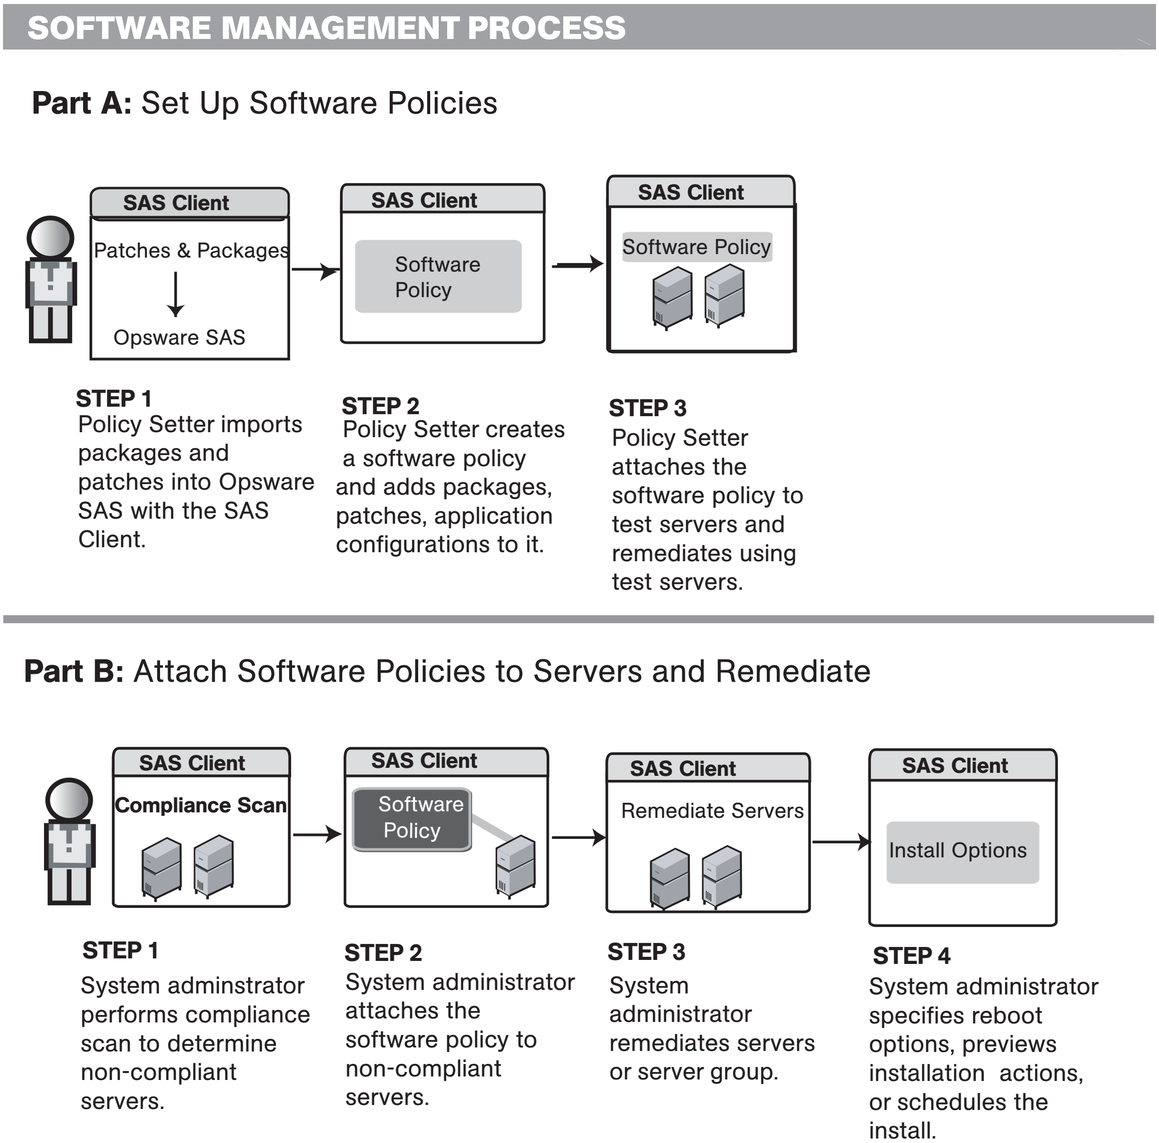

SOFTWARE MANAGEMENT PROCESS

Part A: Set Up Software Policies

STEP 1
Policy Setter imports packages and patches into Opsware SAS with the SAS Client.

STEP 2
Policy Setter creates a software policy and adds packages, patches, application configurations to it.

STEP 3
Policy Setter attaches the software policy to test servers and remediates using test servers.

Part B: Attach Software Policies to Servers and Remediate

STEP 1
System adminstrator performs compliance scan to determine non-compliant servers.

STEP 2
System administrator attaches the software policy to non-compliant servers.

STEP 3
System administrator remediates servers or server group.

STEP 4
System administrator specifies reboot options, previews installation actions, or schedules the install.

Ways to Install Software in Opsware SAS

Opsware 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 292 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 189 in this chapter for more information.
- Use Visual Packager to prepare installable software packages. See the *Opsware® 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 321 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 299 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 Opsware administrator. See the *Opsware® 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

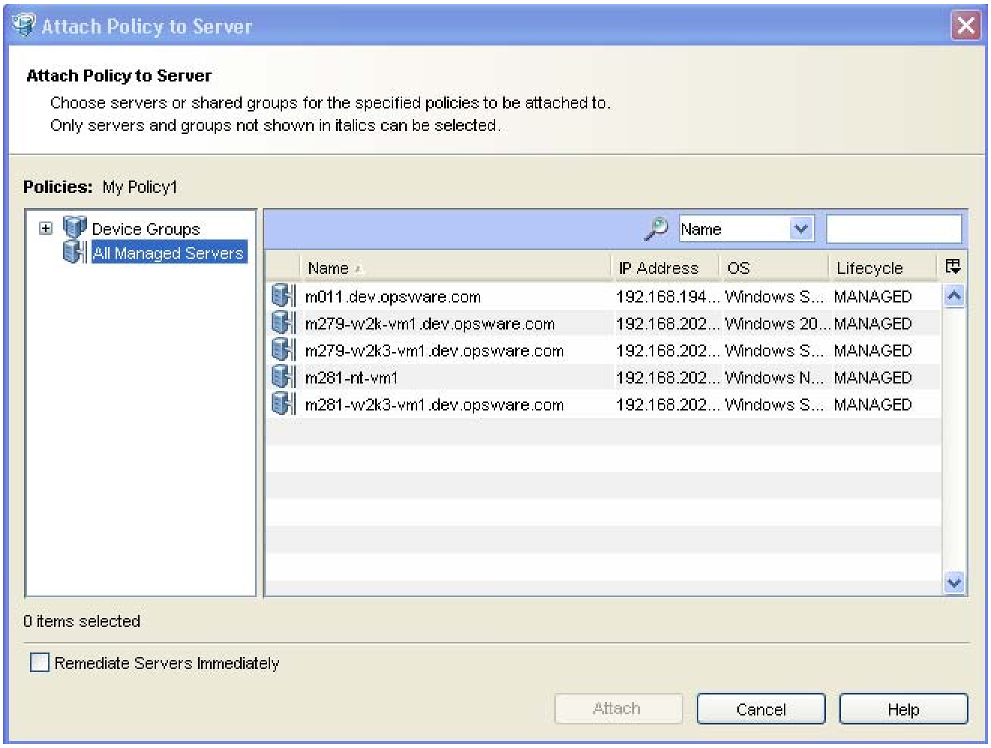

Attach Policy to Server

Attach Policy to Server

Choose servers or shared groups for the specified policies to be attached to.
Only servers and groups not shown in italics can be selected.

Policies: My Policy1

Name	IP Address	OS	Lifecycle
m011.dev.opsware.com	192.168.194...	Windows S...	MANAGED
m279-w2k-vm1.dev.opsware.com	192.168.202...	Windows 20...	MANAGED
m279-w2k3-vm1.dev.opsware.com	192.168.202...	Windows S...	MANAGED
m281-nt-vm1	192.168.202...	Windows N...	MANAGED
m281-w2k3-vm1.dev.opsware.com	192.168.202...	Windows S...	MANAGED

0 items selected

□ Remediate Servers Immediately

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 299 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 299 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** ➤ **Attach Software Policy**. The Attach Server to Policy window appears as shown in Figure 7-3.

Figure 7-3: The Attach Server to Policy Window in the SAS Client

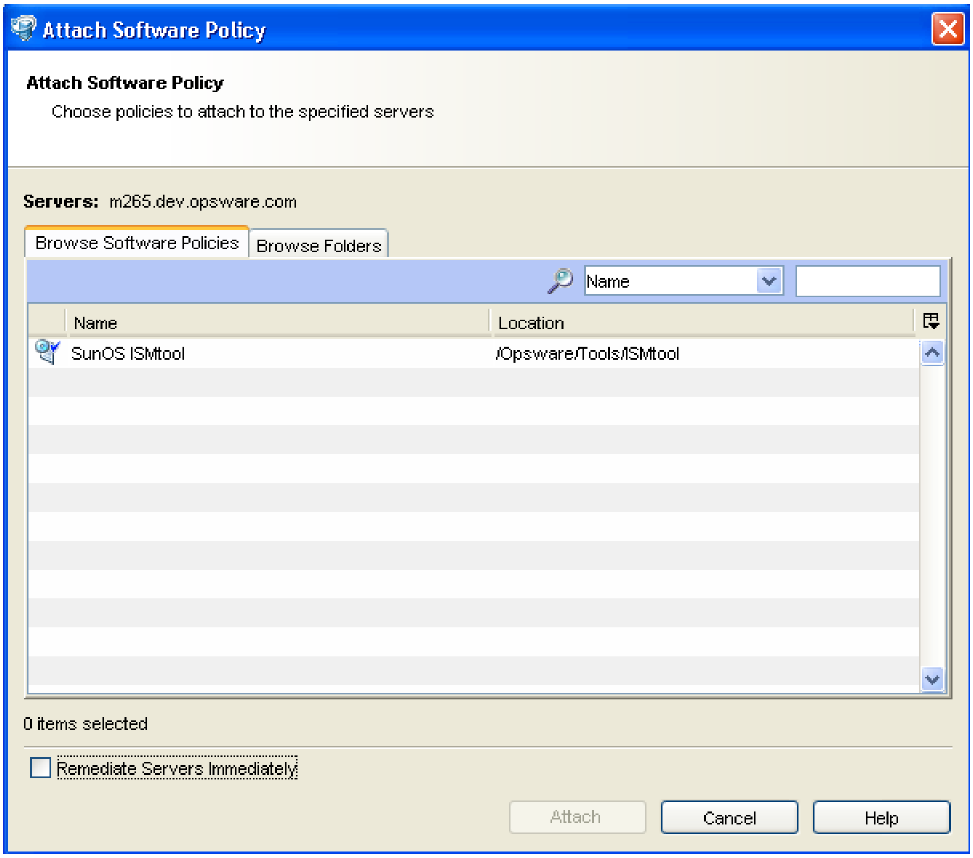

Attach Software Policy

Attach Software Policy

Choose policies to attach to the specified servers

Servers: m265.dev.opsware.com

Name	Location
SunOS ISMtool	/Opsware/Tools/ISMtool

0 items selected

☐ Remediate Servers Immediately

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 299 in this chapter for more information.

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 299 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 Opsware administrator. See the *Opsware® 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 7-4.

Figure 7-4: The Detach Software Policy Window in the SAS Client

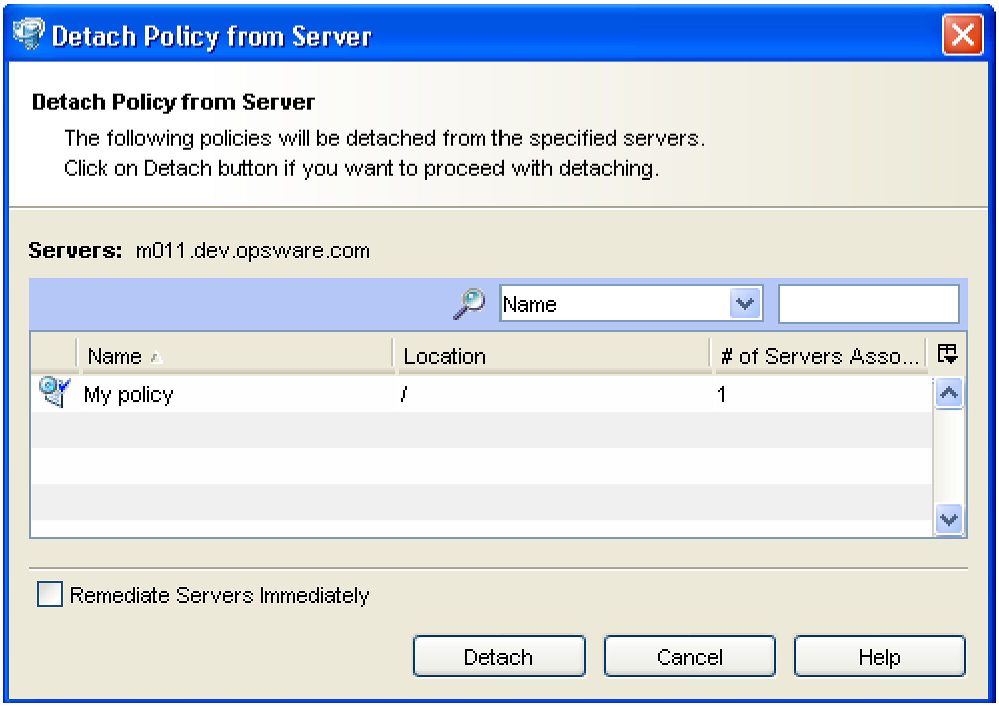

Detach Policy from Server

The following policies will be detached from the specified servers.
Click on Detach button if you want to proceed with detaching.

Servers: m011.dev.opsware.com

Name	Location	# of Servers Asso...
My policy	/	1

□ Remediate Servers Immediately

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 299 in this chapter for more information.

Overview of Software Policies Remediation

The remediate 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 remediate process uninstalls the software in a software policy.

The remediate process allows you to specify remediation options and pre and post installation scripts required for the remediate process, schedule the download and the installation phase of the remediation process, set up email notifications to alert you about the status of the remediate process, and associate a Ticket ID with each remediate process.

You must have a set of permissions to remediate polices. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide* for more information.

The Remediate window allows to you remediate 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 allows to you remediate software policies and consists of the following steps:

Figure 7-5: The Remediate Window in the SAS Client

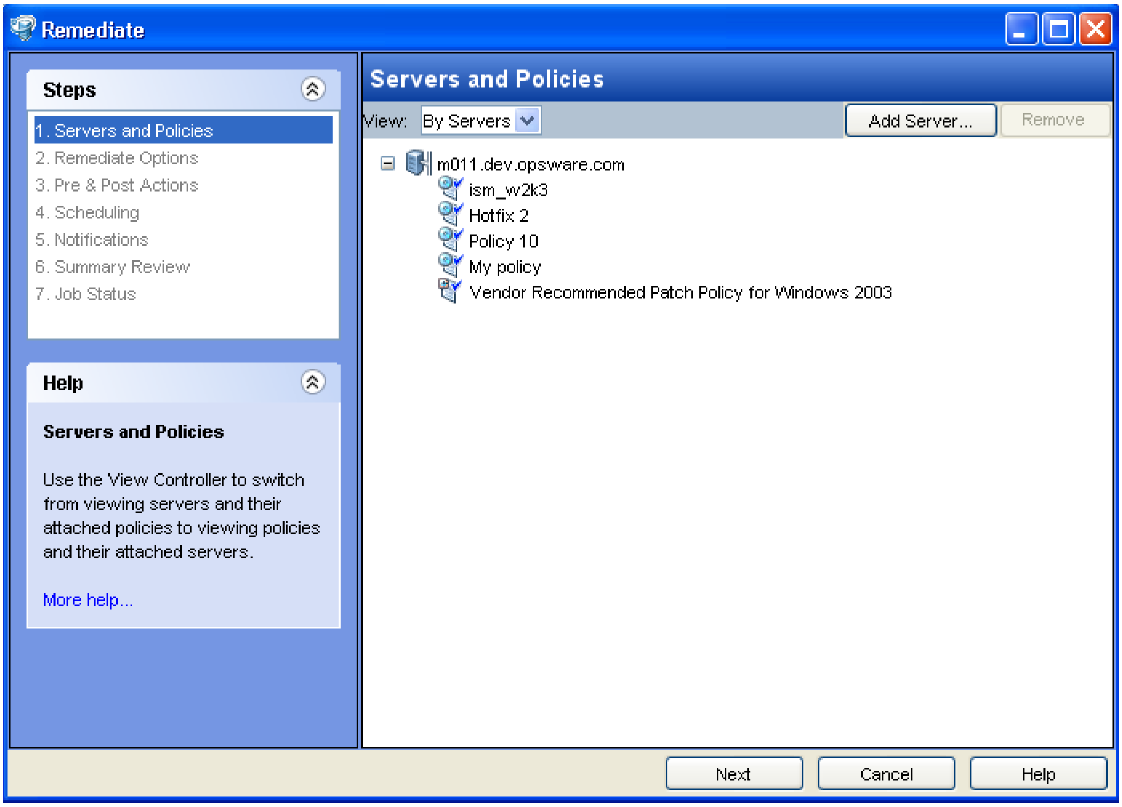

m011.dev.opsware.com
ism_w2k3
Hotfix 2
Policy 10
My policy
Vendor Recommended Patch Policy for Windows 2003

Servers and Policies

Use the View Controller to switch from viewing servers and their attached policies to viewing policies and their attached servers.

- 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 polices 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 298.
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 **Add Server** to add servers to the list or select a server and click **Remove** to remove servers from the list. You can also select software policies attached to a server and click **Remove** to remove the software policies attached to a server.
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 to 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 remediate 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 Opsware 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 Opsware 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.

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 Opsware SAS after you upload the script to Opsware SAS. An Ad-Hoc script is intended only for one operation and is not stored in Opsware SAS.

6 If you selected Saved Scrip 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 install ion and download stage to be run immediately or at a specified date and time.

Perform the following steps to schedule the remediate 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 remediate process. You can associate a Ticket ID with the remediate policy 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 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 remediate process and have the option preview the remediate 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 remediate process, and the scripts that will be executed during the remediate process.

Perform the following steps to preview the remediate process:

1. From the Remediate window, click **Next** to advance to the Summary Review step.
2. Verify the summary information displayed for the remediate 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 I 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 jogs in the SAS Client. See the*Opsware® SAS User's Guide: Server Automation* for information about job logs.

Overview of Software Policy Template

Opsware SAS allows you to install software by using a software policy template. A software policy template can only contain other software policies. A software policy template is not persistently associated with a server or group of servers. When you install a software policy template to a server or group of servers, the software policies specified in the software policy template are installed. If you update a software policy template, servers that already had the software policy template applied are not automatically modified to match the updated software policy template. You must install the software policy template again to reflect the changes made to the software policy template on the server.

A software policy template has the following features:

- A software policy template is not associated with a server or group of servers.
- A software policy template contains other software policies.
- A software policy template is associated with an operating system family.
- Software policy templates are located in folders.
- Custom attributes can be set on a software policy template.

Installing software on a server by using a software policy template consists of the following steps:

- Creating a software policy template

  See the *Opsware® SAS Policy Setter's Guide* for information about creating a software policy template.

- Adding software policies to a software policy template

  See the *Opsware® SAS Policy Setter's Guide* for information about adding software policies.

- Installing the software policy template See "Installing Software Using a Software Policy Template" on page 306 in this chapter for more information.

Installing Software Using a Software Policy Template

You must have a set of permissions to install to software policy template. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide* for more information.

The Install Software Policy Templates window allows to you install a software policy template on a server.

Ways to Open the Install Software Policy 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 Software Policy Template**. The Install Software Policy 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 policy template.
3. From the **Actions** menu select **Install Software Policy Template**. The Install software policy Templates window appears.

The Install Software Policy Templates allows you to install the software on a server and consists of the following steps:

Figure 7-6: The Install Software Policy Templates Window in the SAS Client

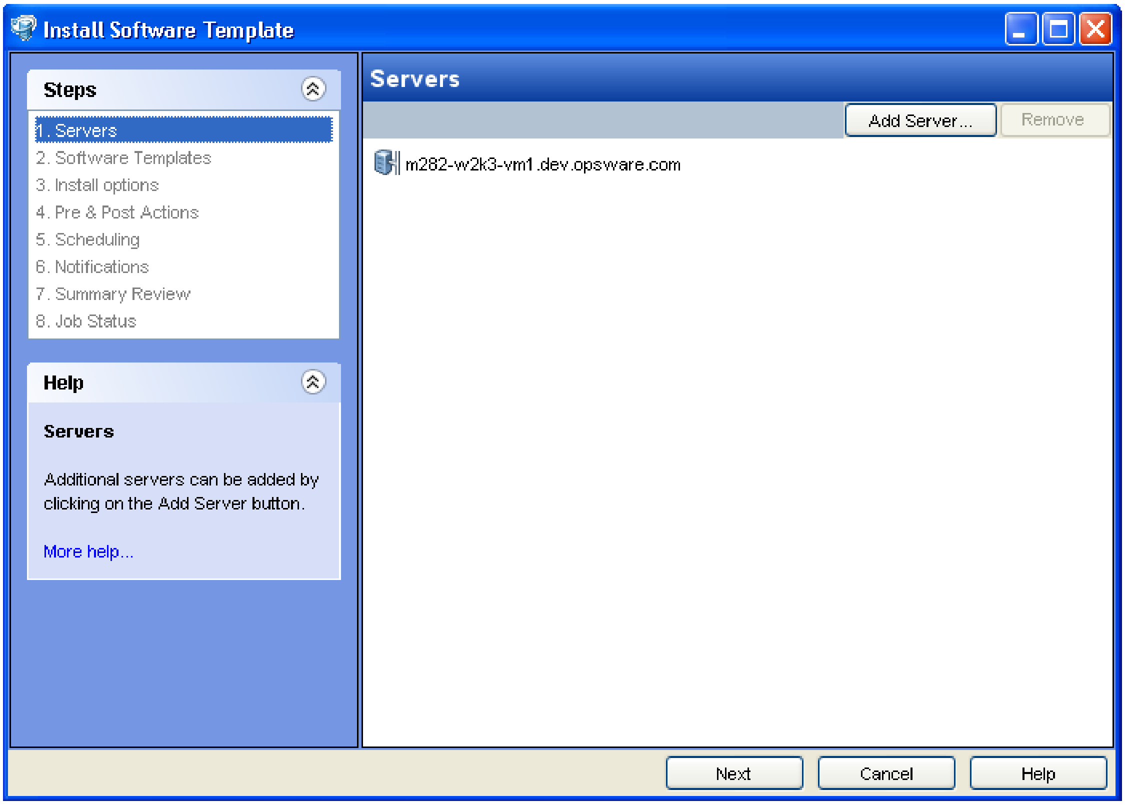

Additional servers can be added by clicking on the Add Server button.

More help...

m282-w2k3-vm1.dev.opsware.com

- Selecting Servers
- Selecting Software Policy 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 Policy Templates window from the server list, the first step in the window is Selecting Servers. If you access the Software Policy Templates window from the software policies list, the first step in the window is Selecting Software Policies Templates.

Selecting Servers

In this step, you can specify the servers for installing the software policy template.

Perform the following steps to select servers:

1. In the Install Software Policy Templates window, select the Servers step.
2. (Optional) Click **Add Server** to add additional servers to the list or click **Remove** to remove servers from the list.
3. Select the servers.
4. Click **Next** to proceed to the Software Policy Templates step.

Selecting Software Policy Templates

In this step you can specify the software policy templates to install on servers.

Perform the following steps to select software policy templates:

1. From the Software Policy Templates window, click **Next** to advance to the Software Policy Template step.
2. In the Software Policy Template window, click **Add Template**. The Attach Software Policy window appears.
3. Select the software policy templates to be installed on the servers.
4. (Optional) Click **Remove** to remove any software policy 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 policy template.

Perform the following steps to set the installation options:

1. From the Software Policy 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.

In this step, 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 Opsware 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 Opsware 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 Policy 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 Opsware SAS after you upload the script to Opsware SAS. An Ad-Hoc script is intended only for one operation and is not stored in Opsware SAS.
6. If you selected Saved Scrip 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 then 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 Policy 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 associate a Ticket ID with the software installation job.

Perform the following steps to set up email notifications:

1. From the Software Policy 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 Policy 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 policy 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 to be performed for the job to be completed.

Perform the following steps to view the job status:

1. From the Software Policy 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 Policy 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 Policy 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 *Opsware® 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 292 in this chapter for more information.

You must have a set of permissions to run an ISM Control. To obtain these permissions, contact your Opsware administrator. See the *Opsware® 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 allows you to run an ISM Control on a server and consists of the following steps:

Figure 7-7: The Run ISM Control Window in the SAS Client

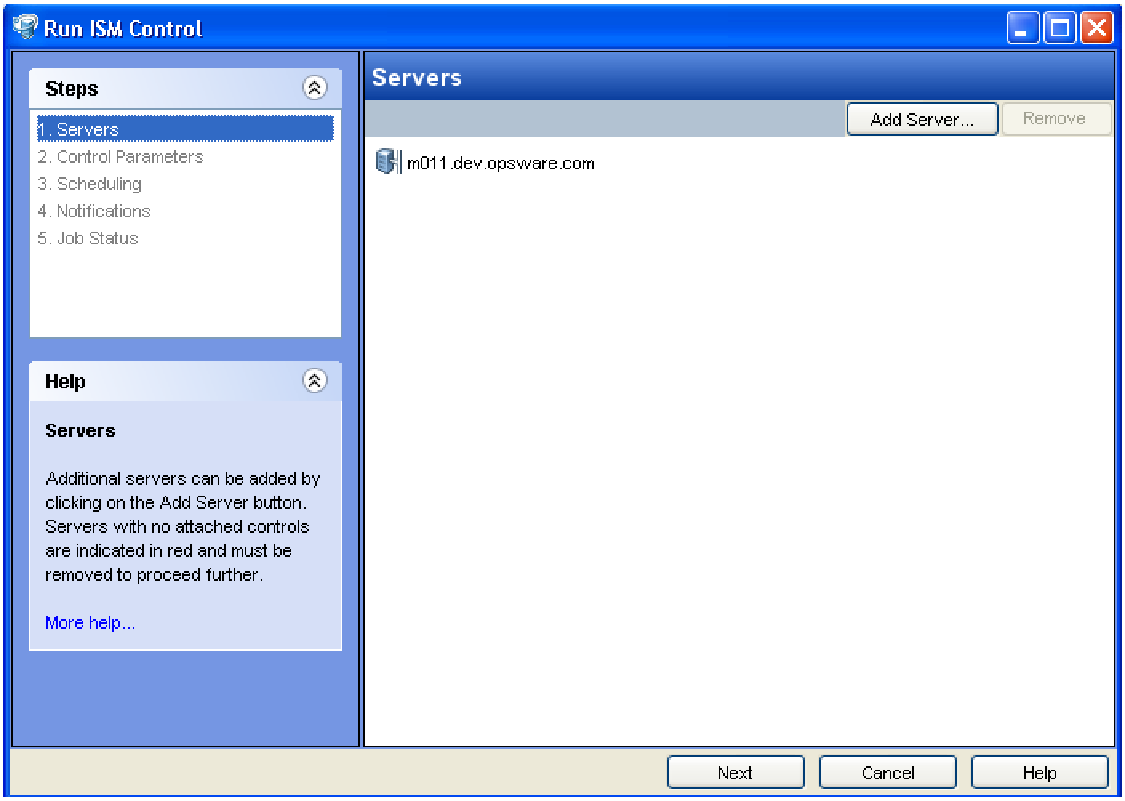

Additional servers can be added by clicking on the Add Server button. Servers with no attached controls are indicated in red and must be removed to proceed further.

m011.dev.opsware.com

- 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 **Add Server** to add additional servers to the list or click **Remove** 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 only 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 to be performed 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 *Opsware® SAS User’s Guide: Server Automation* for information about job logs.

Software Policy Compliance

Software compliance indicates whether or not the all 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 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 one of the software policy attached to the server is not compliant, the server is said to be non-compliant. You can also view the software policy which is not compliant and then 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**: Software compliance information is currently being calculated and is represented by the icon ⧗.
- **Scan Needed**: Software compliance information needs to be calculated or the compliance information might not be accurate 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 155 in Chapter 3 for information about how to perform a software compliance scan.

To perform a compliance scan from the server list See "Performing a Software Compliance Scan" on page 318.

Performing a Software Compliance Scan

You must have a set of permissions to perform a software compliance scan. To obtain these permissions, contact your Opsware administrator. See the *Opsware® 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 299 in this chapter for more information.

Software Policy Reports

The Reporting feature in Opsware 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, export the reports to HTML and XLS, and perform actions on the results.

The Opsware SAS allows you to run the following software policy reports:

- **Software Policy Compliance**: Groups all managed servers by their software policy compliance level to show compliant and non-compliant servers.
- **Software Policy Compliance By Customer**: Lists all servers by the customer they are associated with and then by the software policy compliance level.
- **Software Policy Compliance By Facility**: Displays a chart of all servers by the facility they are associated with and then by the software policy compliance level.
- **Defined Software Policies**: Lists all the software policies by name and their location in the folder hierarchy.
- **Servers With Attached Software Policies**: Lists all servers that have one or more software policies attached.
- **Servers In Compliance With Their Software Policies**: Lists all servers that are in compliance with all of their attached software policies.
- **Servers Not In Compliance with their Software Policies**: Lists all servers that are not in compliance with all of their attached software policies.
- **Servers Without Attached Software Policies**: Lists all servers that have no software policies attached.

See “Reports” on page 175 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)

Opsware 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 Opsware-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 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

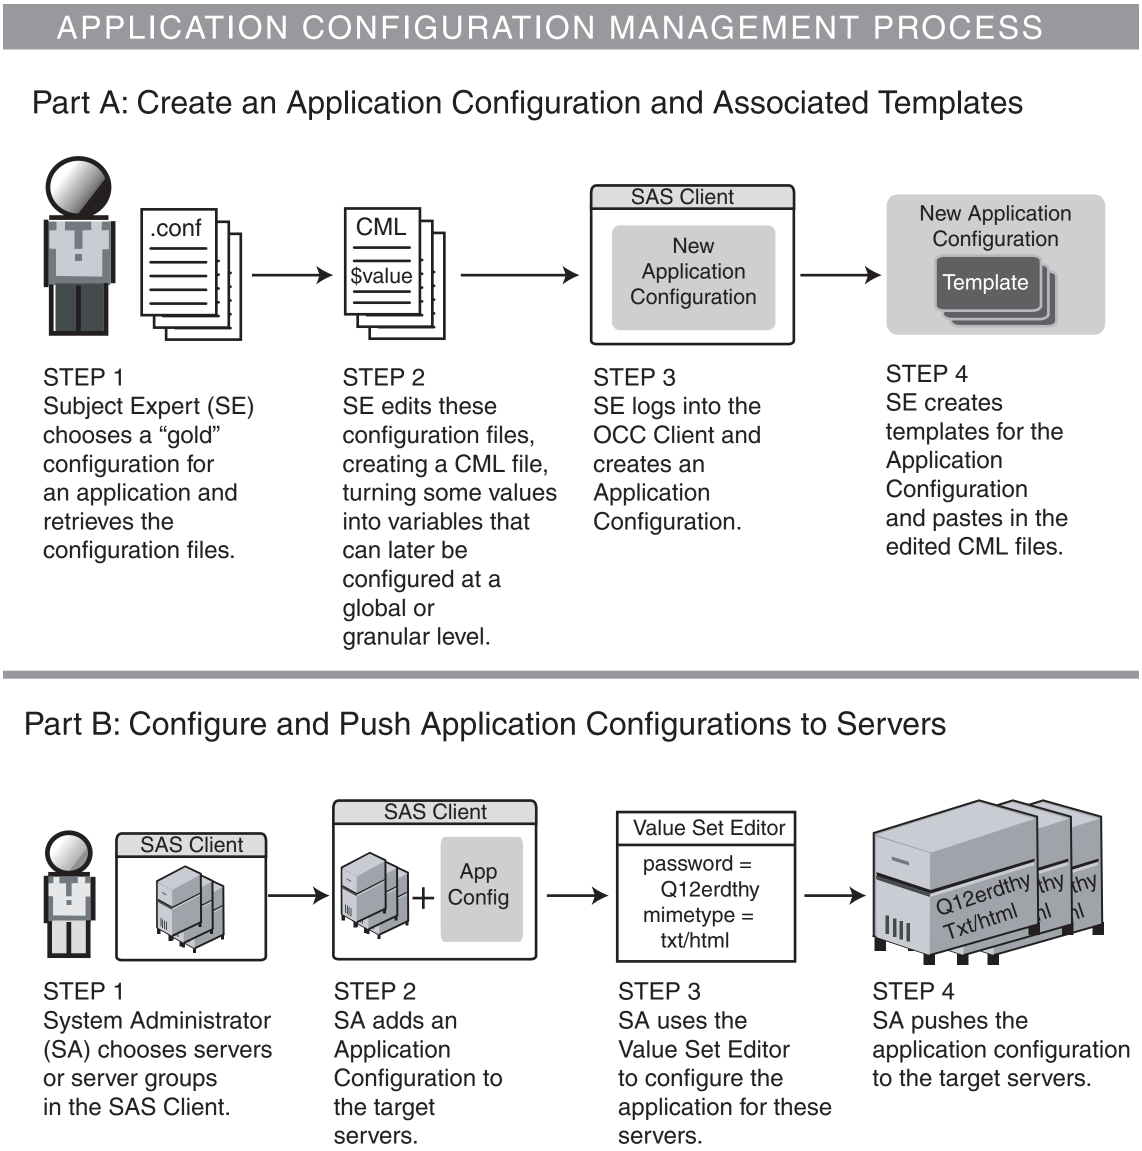

Part A: Create an Application Configuration and Associated Templates

STEP 1
Subject Expert (SE) chooses a "gold" configuration for an application and retrieves the configuration files.

STEP 2
SE edits these configuration files, creating a CML file, turning some values into variables that can later be configured at a global or granular level.

STEP 3
SE logs into the OCC Client and creates an Application Configuration.

STEP 4
SE creates templates for the Application Configuration and pastes in the edited CML files.

Part B: Configure and Push Application Configurations to Servers

STEP 1
System Administrator (SA) chooses servers or server groups in the SAS Client.

STEP 2
SA adds an Application Configuration to the target servers.

STEP 3
SA uses the Value Set Editor to configure the application for these servers.

STEP 4
SA pushes the application configuration to the target 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

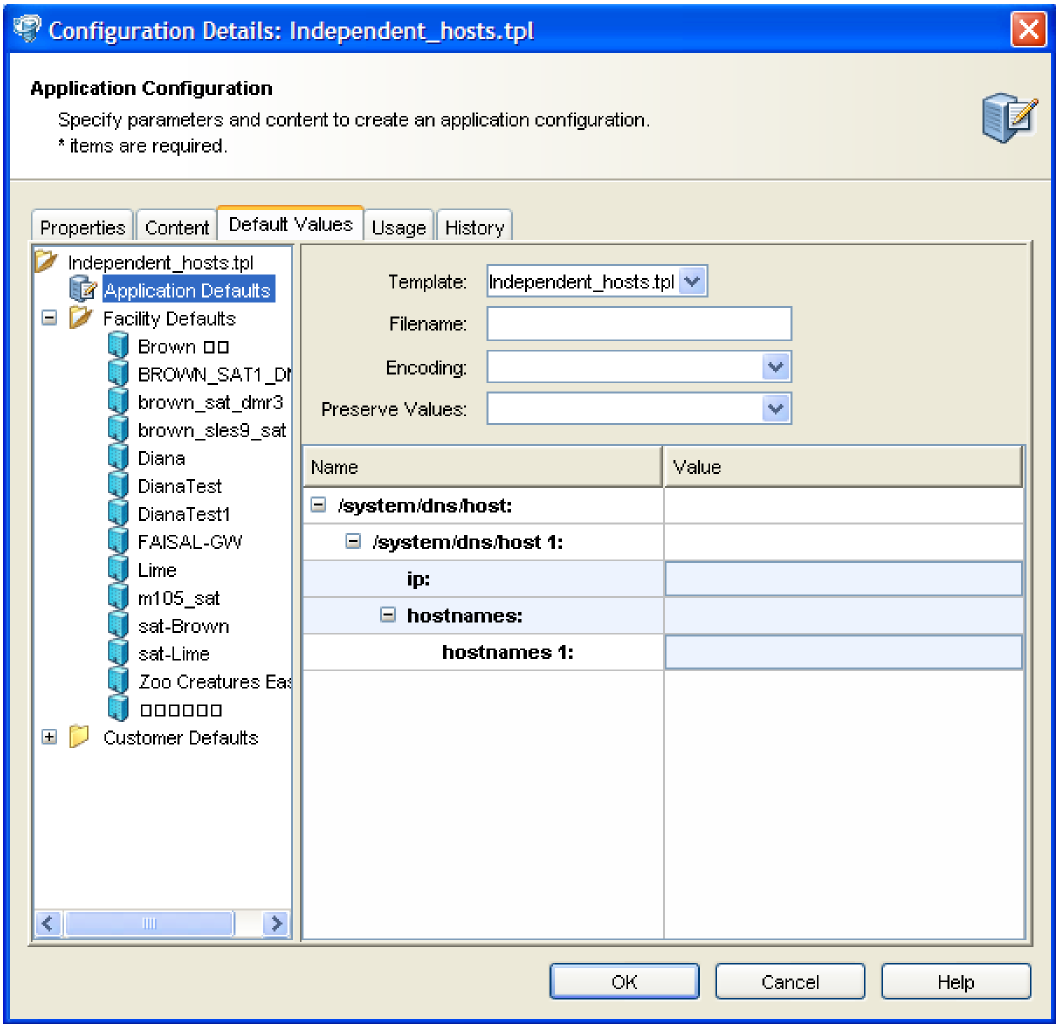

Configuration Details: Independent_hosts.tpl

Application Configuration

Specify parameters and content to create an application configuration.
* items are required.

Template: Independent_hosts.tpl
Filename:
Encoding:
Preserve Values:

Name	Value
/system/dns/host:	
/system/dns/host 1:	
ip:	
hostnames:	
hostnames 1:	

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

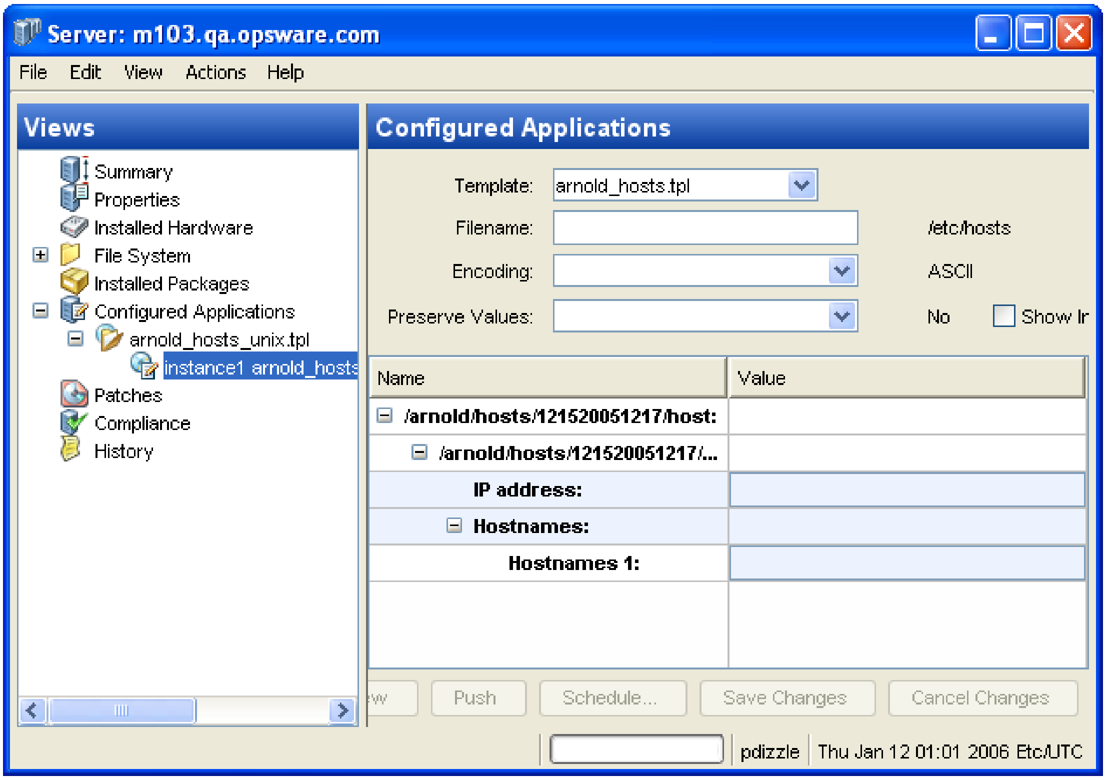

Template: arnold_hosts.tpl

Filename: /etc/hosts

Encoding: ASCII

Preserve Values: No

Name	Value
/arnold/hosts/121520051217/host:	
/arnold/hosts/121520051217/...	
IP address:	
Hostnames:	
Hostnames 1:	

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.
- **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

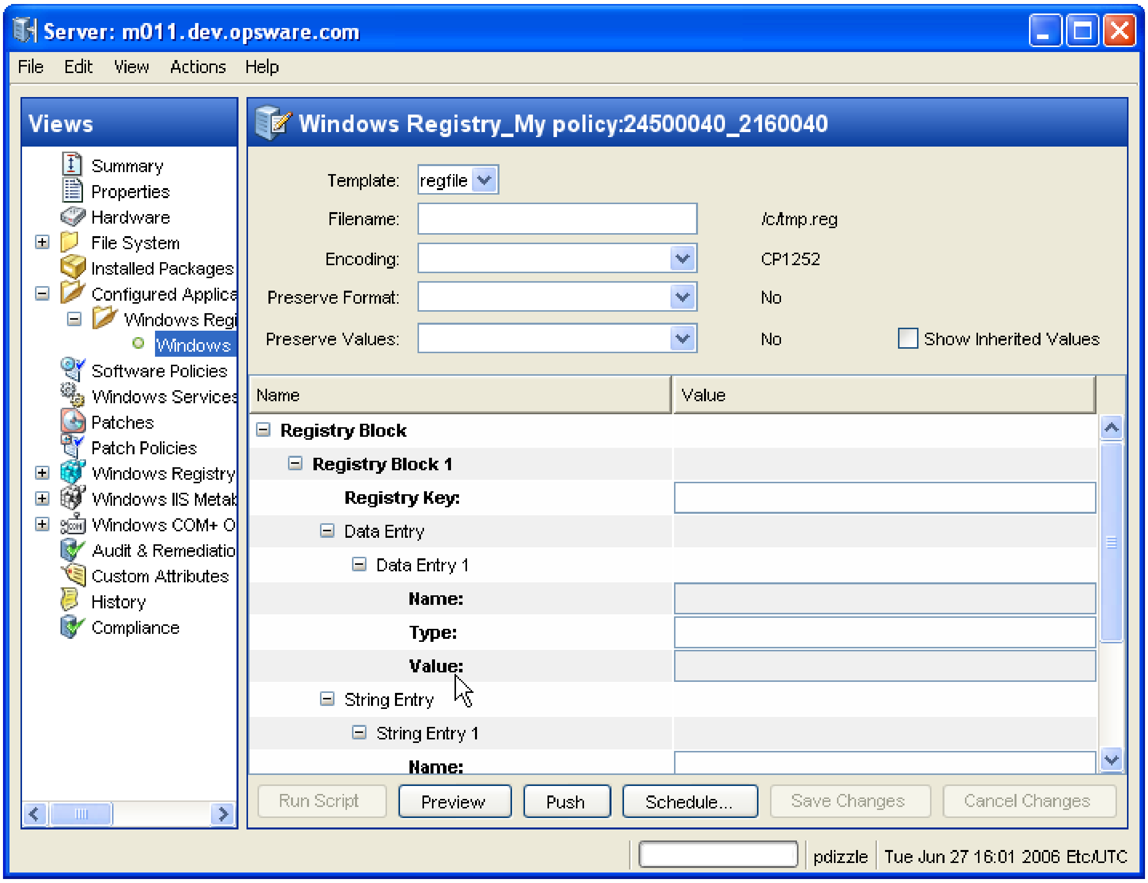

Server: m011.dev.opsware.com

Windows Registry_My policy:24500040_2160040

Template:	regfile	
Filename:		/c/tmp.reg
Encoding:		CP1252
Preserve Format:		No
Preserve Values:		No

□ Show Inherited Values

Name	Value
⊟ Registry Block	
⊟ Registry Block 1	
Registry Key:	
⊟ Data Entry	
⊟ Data Entry 1	
Name:	
Type:	
Value:	
⊟ String Entry	
⊟ String Entry 1	
Name:	

pdizzle | Tue Jun 27 16:01 2006 Etc/UTC

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

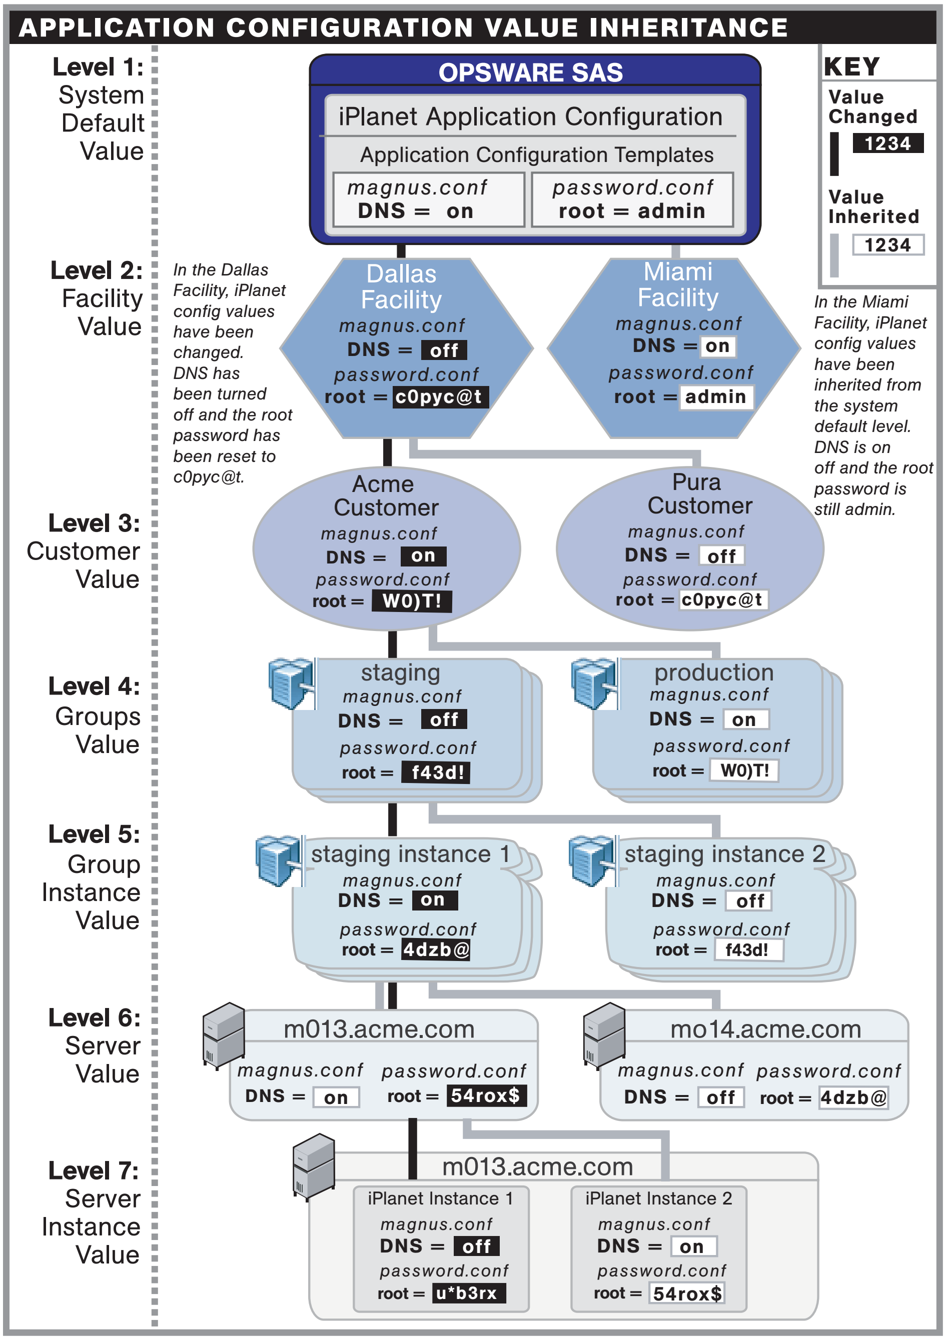
APPLICATION CONFIGURATION VALUE INHERITANCE
KEY
Value Changed
1234
Value Inherited
1234
Level 1: System Default Value
OPSWARE SAS
iPlanet Application Configuration
Application Configuration Templates
magnus.conf DNS = on
password.conf root = admin
Level 2: Facility Value
In the Dallas Facility, iPlanet config values have been changed. DNS has been turned off and the root password has been reset to c0pyc@t.
Dallas Facility
magnus.conf DNS = off
password.conf root = c0pyc@t
Miami Facility
magnus.conf DNS = on
password.conf root = admin
In the Miami Facility, iPlanet config values have been inherited from the system default level. DNS is on off and the root password is still admin.
Level 3: Customer Value
Acme Customer
magnus.conf DNS = on
password.conf root = W0)T!
Pura Customer
magnus.conf DNS = off
password.conf root = c0pyc@t
Level 4: Groups Value
staging
magnus.conf DNS = off
password.conf root = f43d!
production
magnus.conf DNS = on
password.conf root = W0)T!
Level 5: Group Instance Value
staging instance 1
magnus.conf DNS = on
password.conf root = 4dzb@
staging instance 2
magnus.conf DNS = off
password.conf root = f43d!
Level 6: Server Value
m013.acme.com
magnus.conf DNS = on
password.conf root = 54rox$
mo14.acme.com
magnus.conf DNS = off
password.conf root = 4dzb@
Level 7: Server Instance Value
m013.acme.com
iPlanet Instance 1
magnus.conf DNS = off
password.conf root = u*b3rx
iPlanet Instance 2
magnus.conf DNS = on
password.conf root = 54rox$

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 preprended (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 343.

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 361.

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
- Auditing an Application Configuration
- Scheduling an Application Configuration Audit
- 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 Opsware'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 347 in this chapter for more information.

To create a 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 334.

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.

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.
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

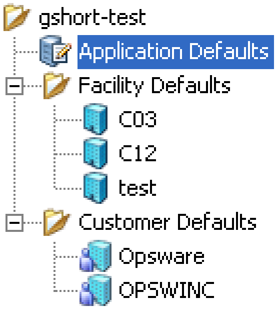

7 Edit the default values for each value set in the Application Configuration Template. The following settings will be displayed:

- **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.
- **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.
- **Value column**: This allows you to 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 (if a parent or ancestor has settings configured). To use an Opsware or custom attribute for the value, click the browse (...) button to access the Set Value dialog box.

8. To edit or change a value, either type a string value directly into the field, or click once in the value field, then click the browse (...) button to access the Set Value window. Choose one of the following options:

   - **No Value**: Choose this to set no value to the value set key.
   - **Block Inheritance**: Choose this option if you do not want to inherit any values from values set at higher levels of the Application Configuration inheritance hierarchy. The effect this has when you push the template to the server depends on if the value is a scalar or a list value:
     - **Scalars**: If the value is a scalar, this key's value will be removed from the configuration file when pushed to the server. Thus, this option is a means of removing a scalar value from a configuration file.
     - **Lists**: If the value is list, then any values from higher levels of the inheritance hierarchy will be blocked, but the current level and any lower levels of the scope will be pushed to the configuration file on the server.
     - **Note**: To block a namespace sequence from inheriting from other scopes, you should add a new namespace sequence that has the a single scalar value or the only entry in a sequence set to <Block Inheritance> with all other fields empty.

- **Any Value**: Enter a value here.
- **Opsware Attribute**: From the drop-down list, choose an Opsware attribute to use, such as customer name, customer ID, chassis ID, device ID, and so on.
- **Custom Attribute**: Enter your own custom attribute here.

9 (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.

10 (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**.

11 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 329 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 352 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 349 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.

- **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.
- **Value column**: This allows you to 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 (if a parent or ancestor has settings configured). To use an Opsware or custom attribute for the value, click the browse (...) button to access the Set Value dialog box.

7. To edit or change a value, either type a string value directly into the field, or click once in the value field, then click the browse (...) button to access the Set Value window. Choose one of the following options:

   - **No Value**: Choose this to set no value to the value set key.
   - **Block Inheritance**: Choose this option if you do not want to inherit any values from values set at higher levels of the Application Configuration inheritance hierarchy. The effect this has when you push the template to the server depends on if the value is a scalar or a list value:

  - **Scalars**: If the value is a scalar, this key's value will be removed from the configuration file when pushed to the server. Thus, this option is a means of removing a scalar value from a configuration file.
  - **Lists**: If the value is list, then any values from higher levels of the inheritance hierarchy will be blocked, but the current level and any lower levels of the scope will be pushed to the configuration file on the server.
  - **Note**: To block a namespace sequence from inheriting from other scopes, you should add a new namespace sequence that has the a single scalar value or the only entry in a sequence set to <Block Inheritance> with all other fields empty.
- **Any Value**: Enter a value here.
- **Opsware Attribute**: From the drop-down list, choose an Opsware attribute to use, such as customer name, customer ID, chassis ID, device ID, and so on.
- **Custom Attribute**: Enter your own custom attribute here.

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 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 334.

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 that the time it takes to push the Application Configuration. See your Opsware 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 353 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 successully.
   - **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.

Auditing an Application Configuration

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 audit the Application Configurations 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. If you scan the list of servers, you can see if any servers show the out of sync icon.

If a server shows this icon, run an Application Configuration audit to find out which configuration files on the server are out of sync with the Application Configuration.

To run an Application Configuration audit, 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 **Audit Application Configurations**. (You can also multiple-select and audit more than one out of sync server.)

3 You will be asked if you are sure you want to audit the Application Configuration on the selected managed server. Click **Yes** to run the audit.

4 The Job dialog box appears, showing the details of the audit. 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 audit 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 audit 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 361 in this chapter for more information

Scheduling an Application Configuration Audit

You can schedule an Application Configuration audit to run a single time, or on a recurring schedule, such as daily, weekly, or monthly.

To schedule an Application Configuration audit, perform the following steps:

1 Launch the SAS Client. From the Navigation pane, select Devices.

1 Select a server or group of servers from the Navigation pane, and then select a server from the Content pane.

2 From the **Actions** menu, select **Schedule Application Configuration Audit**.

3 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.

4 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.

5 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 successully.

- **On Failure**: Email addresses that will receive notifications of jobs that failed to complete.

6. In the Ticket Tracking section, enter a ticket ID from your own job trackins system here.

7. 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 rollback 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 rollback 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

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 Opsware SAS. Additionally, OS installation profiles can be created in the Opsware Command Center and in the Opsware Server Automation System Client. Please refer to the *Opsware® SAS Policy Setter's Guide*.

Supported Operating Systems for OS Provisioning

The OS Provisioning feature supports installation of the following versions of Red Hat Linux, Sun Solaris, and Microsoft Windows operating systems:

- Red Hat Linux 6.2
- Red Hat Linux 7.1
- Red Hat Linux 7.2
- Red Hat Linux 7.3
- Red Hat Linux 8.0
- Red Hat Linux Enterprise Linux 2.1 AS

- Red Hat Linux Enterprise Linux 2.1 ES
- Red Hat Linux Enterprise Linux 2.1 WS
- Red Hat Linux Enterprise Linux 3 AS
- Red Hat Linux Enterprise Linux 3 WS
- Red Hat Linux Enterprise Linux 3 ES
- Red Hat Linux Enterprise Linux 4 AS
- Red Hat Linux Enterprise Linux 4 WS
- Red Hat Linux Enterprise Linux 4 ES
- Sun Solaris 6
- Sun Solaris 7
- Sun Solaris 8
- Sun Solaris 9
- Sun Solaris 10 (See the Note following this list.)
- Fujitsu Solaris 8
- Fujitsu Solaris 9
- Fujitsu Solaris 10
- SUSE Linux Standard Server 8
- SUSE Linux Enterprise Server 8
- SUSE Linux Enterprise Server 9
- Windows NT 4.0
- Windows 2000
- Windows 2003
- Windows 2003 x64
- Windows XP Professional

For Solaris on SPARC, the OS Provisioning feature only supports hardware that is supported by Solaris 10 Update 1, regardless of the version (7, 8, 9, or 10) of Solaris being provisioned.

The OS Provisioning feature works with floppy, CD, or network booting.

The OS Provisioning feature does not provision HP-UX or AIX operating systems. However, you can integrate Opsware SAS with Network Installation Management (NIM) to provision AIX and Ignite-UX to provision HP-UX. See the *Opsware® SAS Administration Guide* for more information about how to integrate Opsware SAS with HP-UX and AIX OS provisioning systems.

The OS Provisioning feature supports a large variety of hardware models from different manufacturers out of the box. You can also configure the OS Provisioning feature to support additional hardware models. See the *Opsware® SAS Policy Setter's Guide* for more information about how to extend the OS Provisioning feature to support new hardware.

OS Provisioning

This section provides information on OS Provisioning within Opsware SAS and contains the following topics:

- Overview of OS Provisioning
- Server Life Cycle for OS Provisioning

Overview of OS Provisioning

In Opsware SAS, OS Provisioning is installation-based instead of image-based. The OS Provisioning feature uses Red Hat Linux Kickstart, Sun Solaris JumpStart, and Microsoft Windows unattended installation to install operating systems on servers.

The OS Provisioning feature is fully integrated with Opsware 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 Opsware SAS already manages

- A server that is running in the environment but Opsware SAS does not manage it

The OS Provisioning feature facilitates installing 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 which operating systems 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 Opsware administrator. For more information, see the Permissions Reference appendix in the *Opsware® SAS Administration Guide*.

Server Life Cycle for OS Provisioning

Opsware SAS is designed to enable multiple teams to work together to provision servers. The OS Provisioning feature allows IT teams to separate the tasks of readying servers for provisioning (such as racking servers, connecting them to power and a network) from provisioning the servers with operating systems.

Someone mounts a new server in a rack and connects it to the Opsware build network. Then they boot the server for the first time by using an Opsware Boot Floppy or CD or by using the network. At a later time, a different system administrator can select the available server from the Server Pool list (in the Opsware Command Center or 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.

See Table 9-1 for the life cycle values for servers. During OS provisioning, servers progress through the following Opsware life cycle state changes:

Unprovisioned ➤ Available ➤ Installing OS ➤ Managed

Table 9-1: Opsware SAS *Life Cycle Values for Servers*

OPSWARE LIFE CYCLE VALUE	DESCRIPTION
Server Pool Values	
Available	Indicates a server on which the OS Build Agent was installed and is running, but an OS has not been installed on the server. The OS Build Agent is a small agent that can run in the memory of the bare metal server. See “Installation of OS Build Agents” on page 381 in this chapter for more information.
Installing OS	Indicates that a user is installing an OS on the server. The server stays in the Server Pool list until the installation process finishes successfully; then, the server moves to the Manage Server list.
Build Failed	Indicates a server on which the OS Build Agent was installed and is running, but the installation of an OS failed. The server will remain in the Server Pool list with this status for 7 days before Opsware SAS deletes the entry. See “Recovering when an OS Build Agent Fails to Install” on page 382 in this chapter for more information.
Managed Server Values	
Managed	Indicates a server that Opsware SAS is managing. Opsware SAS performs reachability checks for managed servers. After a server reaches this life cycle state, the entry for the server moves from the Server Pool list to the Manage Servers list.
Deactivated	Indicates a server was previously managed by Opsware SAS but is no longer managed by Opsware SAS. However, the server’s history still exists in Opsware SAS. Deactivated servers are not reachable.

OS Provisioning

This section provides information about the OS provisioning process within Opsware SAS and contains the following topics:

- Overview of OS Provisioning
- Solaris OS Provisioning
- Linux OS Provisioning
- Windows OS Provisioning

Overview of OS Provisioning

The process for provisioning new servers of all supported operating systems includes the following steps in the OS Provisioning feature:

1. A system administrator unpacks a server, mounts it in a rack, and attaches the server to power and a network that can reach Opsware SAS.

2. The system administrator prepares the hardware for OS provisioning.

   See "Hardware Preparation" on page 375 in this chapter for more information.

3. If necessary, the system administrator inserts a bootable floppy or CD provided with Opsware SAS. Using a bootable floppy or CD is not necessary for Intel-based servers that support PXE or Unix servers that support DHCP because these types of servers are capable of booting over a network.

   See "New Server Booting" on page 376 in this chapter for more information.

4. The system administrator turns the server on.

   For servers capable of booting over the network, powering the server on causes the server to initiate its network boot process. For example, the server sends a boot request to a PXE server.

   The Opsware OS Build Manager responds to this network boot request by delivering the Opsware OS Build Agent, a small agent that can run in the memory of the bare metal server. (For servers not capable of booting over the network, the Opsware OS Build Agent is on the bootable floppy or CD.)

   The Opsware OS Build Agent constructs an inventory of the server (including server manufacturer, server model, MAC address, available memory, and available storage) and delivers that information to the Opsware OS Build Manager.

5 In the Opsware Command Center, the system administrator sees this server and its hardware inventory in a list of available servers ready to be provisioned.

See “Verifying Installation of an OS Build Agent” on page 381 in this chapter for more information.

6 The system administrator selects the OS or a complete server baseline (which can include a base OS, a set of OS patches, system utilities, and middleware software) to provision.

The system administrator selects to install the OS or complete server baseline on the server at that time or schedule the installation for some time in the future. .

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 reprovision servers running Red Hat Linux or Sun Solaris operating systems by using the OS Provisioning feature.

See “Reprovisioning a Managed Server” on page 391 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 *Opsware® SAS Policy Setter’s Guide* for more information about the detailed steps that occur during the Solaris build process.

Linux OS Provisioning

The OS Provisioning feature 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 Opsware Command Center at installation time.

The process for Linux OS provisioning follows the general OS provisioning process that the OS Provisioning feature established.

See the *Opsware® SAS Policy Setter's Guide* for more information about the detailed steps that occur during the Linux build process.

Windows OS Provisioning

In the OS Provisioning feature, system administrators can perform unattended, scripted installations of Windows NT, Windows 2000, and Windows 2003 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 the known hardware makes and models 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.

The process for Windows OS provisioning follows the general OS provisioning process that the OS Provisioning feature established.

See the *Opsware® SAS Policy Setter's Guide* for more information about the steps that occur during the Windows build process.

Hardware Preparation

Before you use the OS Provisioning feature to install an OS on a server, the server must meet certain requirements, or the hardware must be prepared in certain ways, as Table 9-2 shows.

Table 9-2: Required Hardware Preparation for Servers Managed by Opsware SAS

OPERATING SYSTEM	HARDWARE REQUIREMENTS
Microsoft Windows	Before you install Windows on a server, you need to prepare the hardware by performing the following tasks: • If the hardware has a SCSI RAID controller, you must extend the Windows OS media distribution based on vendor specific requirements. The Windows OS media from Microsoft Corporation does not include the necessary drivers for these SCSI-RAID controllers. • Depending on the version of Windows, create a FAT16 partition or FAT32 on which to install the Windows OS You can create this required partition when using the Windows Boot images or PXE to boot a server the first time. The boot image contains the functionality to create the required partition. See “Booting a Windows or Linux Server with PXE” on page 377 in this chapter for more information. See “Booting a Windows or Linux Server” on page 379 in this chapter for more information.
Sun Solaris	To install Solaris on a server, the hardware must meet the following requirements: • Have a DHCP-capable PROM (older servers can be upgraded to DHCP-capable PROM) • Be part of the sun4u system architecture (platform group) You do not need to perform any Opsware SAS-specific preparation of the hardware before you install Solaris on a server.
Linux	Before you install Linux on a server, you need to prepare the hardware by configuring valid, logical drives for RAID.

New Server Booting

This section provides information on booting new servers with Opsware SAS and contains the following topics:

- Booting New Servers with Different Operating Systems
- OS Build Agent
- Booting a Windows or Linux Server with PXE
- Booting a Windows or Linux Server
- Booting a Solaris Server Over the Network
- Installation of OS Build Agents
- Verifying Installation of an OS Build Agent
- Recovering when an OS Build Agent Fails to Install

Booting 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, Opsware SAS supports floppy- or CD-based booting.

For Windows and Linux servers, the Opsware Boot Floppy and CD respectively contains a small operating system, network drivers, software required to mount a network drive, and the Opsware OS Build Agent. The Opsware Boot Floppy or CD has the software that is otherwise delivered over the network as part of the network boot process.

For Solaris servers, you can provision an OS over the network by using DHCP.

To boot servers over the network, the installation client must be able to reach the Opsware DHCP server on the Opsware core network. If the installation client is running on a different network than the Opsware core network, your environment must have a DHCP proxy (IP helper). Alternatively, for Linux and Windows installation clients, you can boot the servers by using an Opsware Boot CD or Floppy instead of booting the servers over the network.

OS Build Agent

The OS Provisioning feature de-couples the task of readying a server for provisioning from provisioning the server with an OS. This de-coupling of tasks is possible because of the OS Build Agent.

Booting a new server for the first time installs an OS Build Agent on the server; however, the server does not have the target OS installed and might not have access to disk resources. 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 or Linux Server with PXE

Perform the following steps to boot a Windows or Linux Server with PXE:

1. After you mount the new server in a rack and connect it to the Opsware build network, set up the server to boot by using PXE.

   See the hardware vendor’s documentation on how to prepare a server to boot by using PXE.

2. Power on the server and select the option to boot the server with PXE.

   The Opsware SAS menu appears and prompts you to select the type of Opsware Build Agent to install on the server.

   ```
   windows   - Windows Build Agent (DOS 6.22)
   undi      - Windows Build Agent (DOS 6.22 + UNDI)
   win98     - Windows Build Agent (DOS 7.01)
   undi98    - Windows Build Agent (DOS 7.01 + UNDI))
   linux     - Linux Build Agent
   localdisk - Normal boot from localdisk (default after 10 sec)
   ```

   Which version of the Windows Build Agent you should select depends on the type of x86 hardware being provisioned. The images for the Windows Build Agents vary in terms of the memory management software, disk partitioning capabilities, and network drivers – DOS or universal network device interface (UNDI) – that they contain.

For example, if you are provisioning a server that has more than 2GB of RAM, you should select the `Win98` or `Undi98` Boot Image. If an incompatible Boot Image is selected for the hardware, an error message appears at the console. The error message can appear at any point during the provisioning process; for example, it might appear when the Windows Build Agent is booting and DOS is loading or it might appear later in the process when the Windows Installer is loading. See Table 9-3 for the differences between images for the Windows Build Agents.

Table 9-3: Differences Between Images for the Windows Build Agents

BOOT IMAGE	NETWORK DRIVERS	MEMORY MANAGEMENT SOFTWARE	DISK PARTITIONING CAPABILITIES
`windows`	DOS	DOS 6.22	FAT16
`undi`	UNDI	DOS 6.22	FAT16
`win98`	DOS	Windows 98	FAT32
`undi98`	UNDI	Windows 98	FAT32

If you do not select an option after 10 seconds, the server defaults to booting from local disk.

If you select Windows as the option for booting the server, an additional set of Opsware SAS menus appear on the console so that you can partition the hardware disk.

3. For Windows servers only, select the menu choices to partition the disk based on your specifications.

   After the booting process finishes successfully, a message appears on the console that indicates that the server is ready for OS provisioning. An OS Build Agent was installed on the server and the server appears in the Server Pool list in the Opsware Command Center.

4. (Optional) Record the MAC address of the server so that you can locate the server in the Server Pool list in the Opsware Command Center.

You should verify that the newly racked server shows up in the Opsware Command Center and is ready to hand off for OS installation.

See "Verifying Installation of an OS Build Agent" on page 381 in this chapter for more information.

When booting a Linux or Windows server by using PXE, the DHCP relay must be running on the router of the build network for PXE to function properly.

Booting a Windows or Linux Server

You can boot different types of x86 hardware by using an Opsware Boot Floppy (Windows, Windows 98) or by using an Opsware Boot CD (Red Hat Linux or SUSE Linux) because a Boot Floppy or CD can contain multiple NIC drivers.

When you boot a Windows server with a boot floppy, select the Windows or Windows 98 boot floppy based on the server's memory and disk partitioning requirements.

See Table 9-3 for the differences between the Windows and Windows 98 OS build images.

Perform the following steps to boot a Windows or Linux server:

1. After you mount the new server in a rack and connect it to the Opsware build network, insert the Windows Boot Floppy or Linux Boot CD (depending on which OS you want to install on the server).

2. Power on the server. A hardware-vendor specific message appears on the console.

   If you selected Windows as the option for booting the server, Opsware menus appear on the console so that you can partition the hardware disk.

3. For Windows servers only, select the menu choices to partition the disk based on your specifications.

   After the booting process finishes successfully, a message appears on the console that indicates that the server is ready for OS provisioning. An OS Build Agent was installed on the server and the server appears in the Server Pool list in the Opsware Command Center.

4. (Optional) Record the MAC address of the server so that you can locate the server in the Server Pool list in the Opsware Command Center.

You should verify that the newly racked server shows up in the Opsware Command Center and is ready to hand off for OS installation.

See "Verifying Installation of an OS Build Agent" on page 381 in this chapter for more information.

Booting a Solaris Server Over the Network

When Opsware SAS was installed in your facility, the OS Provisioning feature was set up so that the Opsware Boot Server listens for broadcast requests from new servers and it responds by using DHCP.

Perform the following steps to boot a Solaris server over the network:

1. Mount the new Solaris server in a rack and connect it to the network.

   The installation client on this network must be able to reach the Opsware DHCP server on the Opsware core network. If the installation client is running on a different network than the Opsware core network, your environment must have a DHCP proxy (`IP helper`).

2. Enter one of the following commands at the prompt:

   ```
   ok boot net:dhcp - install
   ```

   Or

   ```
   ok boot net:dhcp - install <interface_setting>
   <buildmgr=hostname|IP_address>
   ```

   Where `<interface_`*`setting`*`>` is one of the following options:

   `autoneg, 100fdx, 100hdx, 10fdx, 10hdx`

   You can include an interface setting with the boot command to set the network interface to a specific speed and duplex during OS provisioning. When Opsware 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 *Opsware® SAS Policy Setter's Guide* for more information.

Ways that the OS Build Agent Locates the Opsware Build Manager

For Solaris OS provisioning, the JumpStart build script runs the OS Build Agent, which contacts the Opsware Build Manager (via the Agent Gateway in the core). The Solaris `begin` script attempts to locate the Opsware Build Manager in the following ways:

- By using information that the Opsware DHCP server provided
- By looking for the host name `buildmgr` in DNS as configured by the DHCP server

You can override the way that the OS Build Agent contacts the Opsware Build Manager by specifying a boot argument at the prompt when you boot a new Solaris server:

```
ok boot net:dhcp - install [buildmgr=hostname|IP_address]:port
```

Installation of OS Build Agents

After you install an OS Build Agent on a server by booting the server with PXE or an Opsware Boot Image (Windows and Linux) or by using the network (Solaris), the server appears in the Server Pool list.

You should verify that the newly racked server shows up in the Opsware Command Center and is ready to hand off for OS installation.

The Server Pool list displays the servers that have registered their presence with Opsware 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 Opsware Command Center.

2. From the navigation panel, click Servers ➤ Server Pool. The Server Pool page appears, as Figure 9-1 shows.

Figure 9-1: Server Pool List in the Opsware Command Center

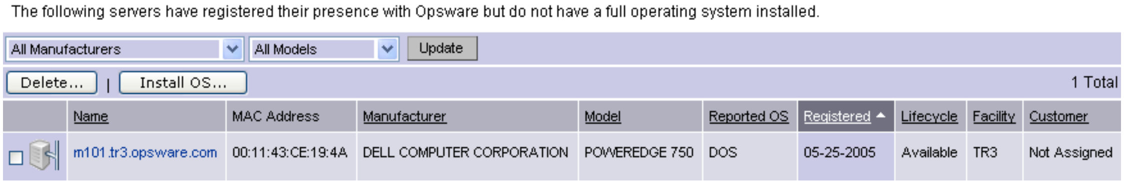

The following servers have registered their presence with Opsware but do not have a full operating system installed.

	Name	MAC Address	Manufacturer	Model	Reported OS	Registered ▲	Lifecycle	Facility	Customer
□	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 that you want to verify and click **Update**.

4. For Intel x86 processor-based servers, locate the MAC address of the server that you just booted.

   Or

   For Sun SPARC processor servers, locate the chassis ID of the server that you just booted.

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

The Life cycle column indicates the progress or success of the OS Build Agent installation. If the OS Build Agent was successfully installed, the Life cycle column indicates that the server is available for OS provisioning.

See "Server Life Cycle for OS Provisioning" on page 370 in this chapter for more information.

Recovering when an OS Build Agent Fails to Install

When an OS Build Agent fails to install on a server, the server does not appear in the Server Pool list.

You can check the server console for error messages and try to boot the server again with PXE or by using the Opsware Boot Floppy or CD.

If all errors were successfully resolved, the initial boot occurs, the OS Build Agent is installed on the server, the server appears in the Server Pool list, and the Life cycle column indicates that the server is available.

If you are unable to resolve the error condition and install the OS Build Agent on the server so that it appears in the Server Pool list, contact your Opsware administrator for troubleshooting assistance.

OS Installation with the SAS Client

This section describes how to install as operating system on an unprovisioned server using the SAS Client. In order to install an operating system on an unprovisioned server, the operating system type you want to install must already have its OS media prepared and made available, have an OS build customization script created for it, and as well must have an OS installation profile created for it. Once you provision a server and install an OS, it becomes an Opsware managed server.

For information on how to set up OS provisioning, see "OS Provisioning Setup" on page 213 in the *Opsware® 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 installation profile, selected software and patch policies, and remediation settings. An OS sequence represents a server build policy that defines how a server should be provisioned, encompassing all types of software, including the operating system. Once you have defined an OS sequence, others can run the OS sequence 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, build customization scripts, OS media, customer association, and packages.

- Creating an OS Sequence

  Choose the OS installation profile, software policies, patch policies, and remediate polices.

- 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 OS type and version, OS media resource locator, OS build customization script, and 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 “OS Provisioning Setup” on page 213 in the *Opsware® 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 in the installation profile, selected application and patch policies, and the target servers you want to install the OS on to. Figure 9-2 shows the OS sequence window.

Figure 9-2: OS Sequence Window

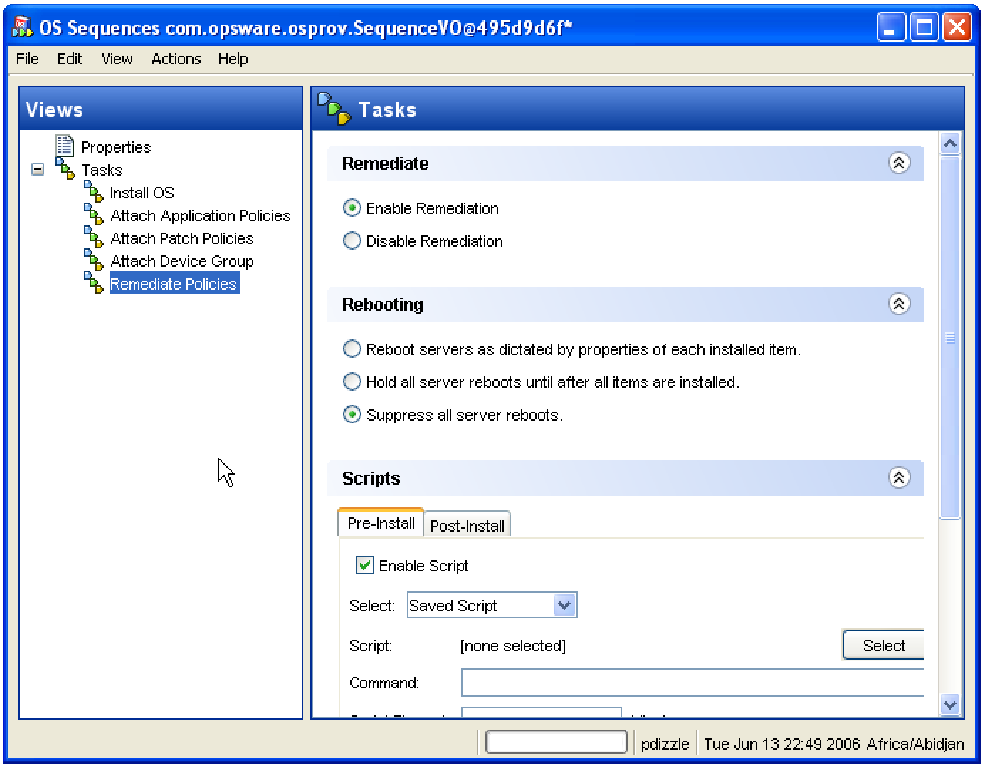

Remediate

● Enable Remediation

○ Disable Remediation

Rebooting

○ Reboot servers as dictated by properties of each installed item.

○ Hold all server reboots until after all items are installed.

● Suppress all server reboots.

Scripts

☑ Enable Script

Select: Saved Script

Script: [none selected]

Command:

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 *Opsware® SAS Administration Guide*.

An OS sequence consists of the following components that must be configured before you run the OS sequence:

- **Properties**: Allows you to name for 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, other wise you will not be able to save it in the selected location in the library.
- **Install**: Allows you to choose an OS installation definition. If the OS installation profile already has a customer associated with it, you will not be able 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 to which you install the OS using this OS sequence will be associated with that customer.
- **Attach Software Policy:** 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 remediate option is disabled, then none of the software will be installed on the server.

  The choice of software policies you can attach to an OS sequence is restricted by the OS type the software policy. In other words, you can only attach software policies whose 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 select a patch policy to attach to the OS sequence. When you 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 when you run the OS sequence. If the remediate option is disabled, then none of the patches will be installed on the server.

  The choice of patch policies you can attach to an OS sequence is restricted by the OS type the patch policy. In other words, you can only attach patch policies whose 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) where you want to places 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 associate with it. If you enable remediate in the OS sequence (in Remediate Policies), then all software and patches associated with the to the group of servers will also be installed on the server when you run the OS sequence. If you disable remediate, 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 *Opsware® SAS User's Guide: Server Automation*.

- **Remediate Polices**: Allows you to choose to enable or disable remediate when the server is provisioned with the OS sequence. Default is enabled. If you disable remediate, then when you run the OS sequence, the OS will be installed but no policies in the OS sequence will be remediated – this mean 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 remediate, then all software and patches in all policies attached to the server in the OS sequence, and any policies attached to the group of servers selected for the OS sequence, will be installed on the server when the OS sequence is run. You can also set reboot and pre and post installation script options.

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 for the OS type you want to create for the OS sequence.
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 of the OS sequence.
5. Click **Change** in the Contents 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 here from the Assign Customer drop-down list. If the OS installation profile already has a customer associated with it, you will not be able 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 Contents 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 Contents 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 Contents pane, click **Add** to add a group of servers where you would like to have the server placed once the OS sequence has been run. You can select any public static group to attach to the OS sequence.

14. From the Views pane, select Remediate Polices.

15. In the Contents pane, choose if you want to enable or disable remediate 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 mean 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 and 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**: Click **Select** to choose a saved script.
- Command: Add any commands or arguments to be executed here.
- Script Timeout: Enter a numerical value for number of minutes to pass until the script will timeout.
- User: Choose if you want the script to be run as either Local System, or enter your own username and password.
- 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 number of minutes to pass until the script will timeout.
- User: Choose if you want the script to be run as either Local System or enter your own username and password.
- 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 Uprovisioned Servers list in the SAS Client. Servers in the Uprovisioned Servers list have registered their presence with Opsware SAS but do not have an OS installed. From this location, you can install an OS by selecting an unprovisioned server. See Figure 9-3.

Figure 9-3: Unprovisioned Servers in the SAS Client

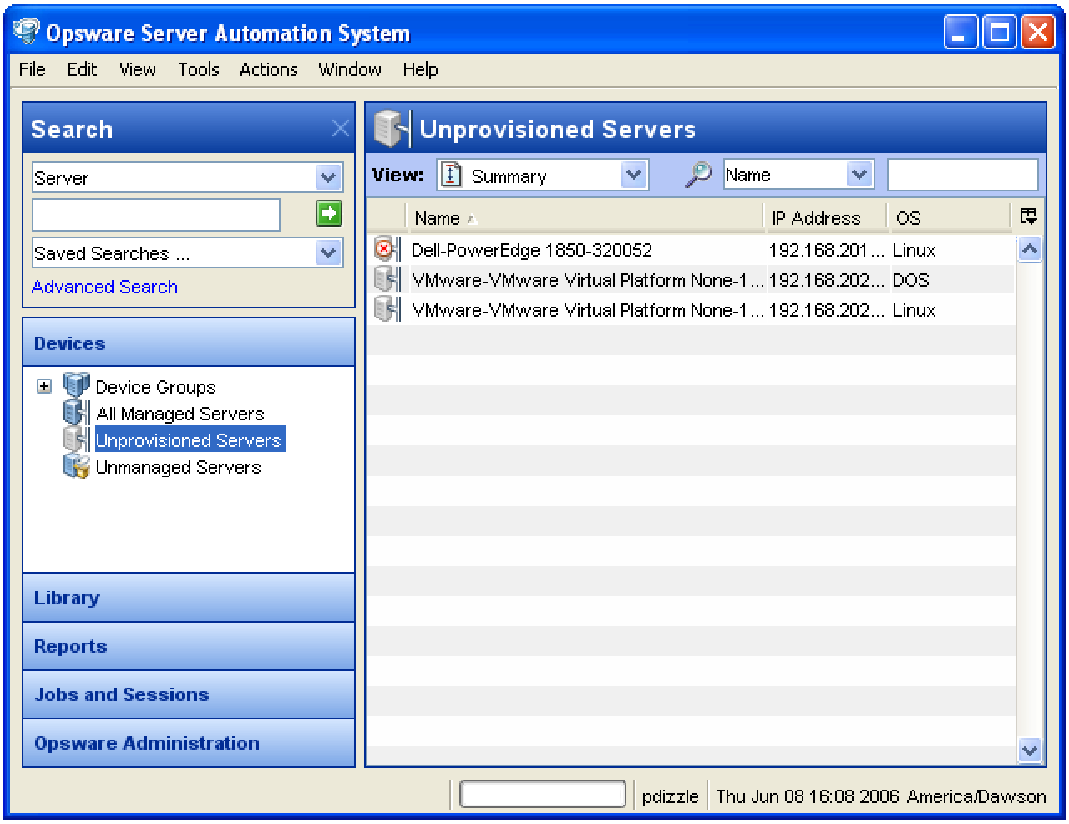

Name	IP Address	OS
Dell-PowerEdge 1850-320052	192.168.201...	Linux
VMware-VMware Virtual Platform None-1...	192.168.202...	DOS
VMware-VMware Virtual Platform None-1...	192.168.202...	Linux

pdizzle | Thu Jun 08 16:08 2006 America/Dawson

Select an unprovisioned server in the list and the Content pane below will display detailed information about the unprovisioned server that was gathered by the Opsware Build Agent after a network boot. Use the View drop-down list to select a view of the server:

- **Summary**: Provides information about the host name set by booting the server the first time over the network or by using an Opsware Boot Floppy or CD, the OS of the OS Build Agent (Windows, Red Hat Linux, or Solaris), processor type, manufacturer and model of the server, Opsware registration information.

- **Properties**: Placeholders for various management and reported information which will be filled in later once the server is provisioned.
- **Hardware**: Details about the hardware on the server, such a processor type, physical and virtual memory, storage and network interfaces.
- **Custom Attributes**: Here you can 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 Contents 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.

Running An OS Sequence

To install an OS on an unprovisioned server, select a server from the Unprovisioned Server list and run an OS sequence, or you can simply start an OS sequence and choose a target server in the Run OS Sequence window.

To run an OS Sequence and install an OS on an unprovisioned server, perform the following steps:

1. There are two ways to install an OS on an unprovisioned server by running an OS sequence:

   – From the Navigation panel, select Devices ➤ Unprovisioned Servers. Select a server and from the **Actions** menu select **Run OS Sequence**.

   OR

   – From the Navigation panel, select Library ➤ OS Sequences. Select the OS of the OS sequence, then select the OS sequence you want to run and from the **Actions** menu, select **Run OS Sequence**.

2. In the Run OS Sequence window, step one requires that you select an unprovisioned server or servers to provision. To add a sever, click the **Add** button and add a server.

3. Click **Next**, and in the Select OS Sequence pane click the **Add** button to add an OS sequence.

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 Opsware 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 has run, click the **View Results** button to view the results of the OS sequence job.

10. When the OS sequence job has been run, 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 reprovision a server, you have the option of preserving the network configuration of the server. Also, some attributes will be saves 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 *Opsware® SAS Administration Guide*.

You can only reprovision a server that runs the Solaris or Linux operating system.

To reprovision a managed server, perform the following steps:

1. From the Navigation panel, 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 and warning message that explains you are about to reprovision a managed server, and 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 sever, click the **Add** button and add a server.

6 Click **Next**, and in the Select OS Sequence pane click the **Add** button 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 You can specific if you want to the email to be sent upon success of the OS sequence job (check mark) or failure of the OS sequence job (X).

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 the **View Results** button 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: 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 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.

audit A set of rules that express 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 as defined in the audit.

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) Process of saving the entire contents of the current Live directory for a specific service to the Backup directory. Code Deployment & Rollback (CDR) saves the backup copy to the local disk for the host on which the Backup operation was run. Only one backup copy is maintained at any time for a service.

backup event An event that causes configuration files or configuration databases to be backed up. Types of backup events include manual, full, and triggered.

blocked attachment An attachment that is not installed when that template is applied. The attachment also does not appear in child templates or folders.

Boot Server A part of the OS Provisioning 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). Tracks changes made to a node. Identifies who has recently modified the node to add or remove software packages, add or remove operating systems, add or move servers, and create or remove subordinate nodes.

Code Deployment & Rollback (CDR) An Opsware 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.

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. The Communication Test runs various tests like Command Engine to Agent Communication, Crypto Match, Agent to Command Engine Communication, Agent to Data Access Engine, Agent to Software Repository Communication, and Machine ID mismatch to determine if an Opsware Agent is reachable.

compliance The degree to which a server object conforms to a test.

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 Process by which new configuration tracking polices or changes to existing configuration tracking polices 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 in the Opsware Command Center. Used when performing a variety of Opsware functions, including network and server configuration, notifications, and CRON script configurations.

custom extension Custom Command Engine scripts that extend Opsware SAS functionality to customers to cover their specific needs.

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. Performed automatically by determining the differences between the Update directory and current the Live directory. The files that are different are synchronized from the Update directory to the current Live directory.

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

data center Legacy term. *See* facility.

deactivated server Server removed from Opsware management even though its history still exists.

deployment Within CDR, automatically pushes code and content from a staging server to a live network server.

deprecated A possible state of a package or patch in 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 Legacy term. *See* server.

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 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.

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 In the Automated Configuration Tracking feature, an email can be sent to the email addresses in the email notification list whenever a change to a tracked file or configuration database is detected.

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

facility 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.

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 Opsware Gateway.

Global Shell A terminal window for the Opsware Global File System (OGFS) in your Opsware SAS.

group See server group.

IDK Intelligent Software Module (ISM) Development Kit. The tools from Opsware Inc. used to build and upload ISMs.

Import Media tool A utility script included with Opsware SAS that is used to import OS media from the Media Server to Opsware 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.

inherited attachment An attachment that is inherited from an ancestor folder or a template.

initialization Legacy term. *See* OS Provisioning.

IP Range Groups A designated set of servers assigned to a customer account, grouped by either a physical or a logical list.

IP Ranges A designated grouping of servers.

ISM Intelligent Software Module. A set of file and directories that include application bits, installation scripts, and control scripts. When an ISM is uploaded into an Opsware core, a node is created for the application and installable packages are attached to the node.

ISM control A script within an ISM package that can be run on a managed server.

job Any major process run by the Opsware Command Center or the Opsware Command Center Client such as Communication Test or Install Software.

Live directory In CDR, the directory that stores the actual code and content required to run a live site.

local attachment An attachment that is attached directly to a folder or a template.

MAC *See* Media Access Control Address (MAC).

Machine ID (MID) A unique identifier that Opsware SAS uses to identify the server. Opsware 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 Opsware Agent installed on it and is under the control of a particular Opsware core.

management IP The IP address that Opsware SAS uses to communicate with the Opsware Agent on the server.

manifest Within CDR, a list of files that indicate the results or preview of an update to be performed. Each entry in the list specifies the file size, last-modified date and timestamp, and the full directory path to the listed file.

Media Access Control Address (MAC) The network interface card's unique hardware number. The MAC is used as the server's physical address on the network.

Media Resource Locator (MRL) A network path in URL format that is registered with Opsware 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 Opsware database that stores information about managed server configurations within Opsware SAS. It contains all information necessary to build, operate, and maintain an Opsware-managed site, including a list of all servers under management, the hardware associated with these servers, including memory, CPUs, storage capacity, and the configuration of these servers, including IP addresses, DNS configuration, and so on.

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

MRL *See* Media Resource Locator (MRL).

multimaster core An Opsware core that belongs to a multimaster mesh.

multimaster infrastructure component See Model Repository Multimaster Component.

multimaster mesh A set of two or more Opsware cores that are linked by synchronizing the data in the Model Repositories at each of the cores. The Model Repositories in each of the cores are continually updated so that they are exact duplicates of each other. All the Opsware cores in a multimaster mesh can be managed through a single Opsware Command Center.

My Jobs A page in the Opsware Command Center that displays a list of jobs from the Model Repository such as software installation or server provisioning.

My Scripts Private scripts that can only be executed by the user who created the script. My Scripts are created for personal use.

name-value pairs Legacy term. See custom attributes.

node A hierarchical set of categories or types that classify hardware, software, configuration, or other components of a site's infrastructure. Simplifies server management (for example, servers within Opsware SAS) and the software applications and configurations associated with those servers.

OCLI *See* Opsware command Line Interface (OCLI).

OGFS See Opsware Global File System.

Opsware administrator Responsible for overall administration, policy, and practices for individuals accessing Opsware SAS. Can add users and define access to specific Opsware SAS features that allow users to view site information and deploy new code and content to their site.

Opsware Agent Intelligent software on Opsware-managed servers that is used to make changes to the servers. Depending on the request, the agent might use Global Opsware services. Some functions supported include software installation and removal, software and hardware configuration, server status reporting, and auditing.

Opsware Automation features Opsware SAS is made up of a set of Opsware Automation features. Opsware Automation features are the components that automate particular IT processes. The Opsware Automation features include the following functions: Software Provisioning, Patch Automation, Configuration Tracking, Code Deployment and Rollback, Script Execution, and Data Center Intelligence Reporting.

Opsware Discovery and Agent Deployment A feature that helps deploy Opsware Agents to a large number of servers through the Opsware Command Center Client.

Opsware Command Center A web-based user interface for managing the Opsware environment.

Opsware Command Line Interface (OCLI) An alternative interface to the Opsware Command Center. The OCLI allows you to perform some actions not possible though the browser-based interface of the Opsware Command Center, such as uploading multiple packages, patches, AIX filesets, and so forth, in a batch operation.

Opsware core The server side of Opsware SAS server-agent architecture. A core consists of the Opsware components (such as the Model Repository, the Software Repository, the Data Access Engine, and the Command Engine) for a particular installation.

Opsware Gateway Provides connectivity with an Opsware core either directly or through a network of gateways. All traffic between the servers in the Satellite and the core that manages them is routed through Opsware Gateways.

Opsware Global File System (OGFS) The Opsware Global File System is a single, unified file system view of all file systems for all managed servers in Opsware SAS.

Opsware installation Either a standalone core, multimaster core, or Opsware Satellite.

Opsware model space The Opsware Global File System (OGFS) file system structure that is derived from the Model Repository.

Opsware Satellite Installed in a remote facility, an Opsware satellite provides 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 SAS The server management application to preserve the knowledge of system administrators, network engineers, and database administrators in a centralized knowledgebase. Automates previously manual tasks associated with the deployment, support, and growth of a data center infrastructure.

OS Build Agent A part of the OS Provisioning 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 Process of installing a basic set of software components, including an operating system and an Opsware Agent to add a server into the Opsware managed environment. After provisioning is complete, the server is ready to be managed by 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. Contains 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.

patch management administrator Administrator responsible for testing patches and defining patch options, such as installation and uninstallation scripts. A patch cannot be installed by other personnel until the patch administrator has marked the patch available through the Opsware Command Center.

Patch Management An Opsware 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 or software nodes. The set of managed servers corresponds to a facility, customer, or server group.

platform The name and version of an operating system.

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.

realm One or more Opsware Gateways that service the managed servers contained within an Opsware realm. In Opsware SAS, a realm is a routable IP address space, which is serviced by one or more gateways. All 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.

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 Queue in which 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.

Satellite See Opsware Satellite

Script Execution See Distributed Scripts.

selection criteria Rules that instruct Opsware SAS what server objects you want to collect information about, how to collect the server objects, and (optionally) file comparison and inclusions/exclusions criteria. Selection criteria is required for the snapshot and audit processes.

sequence Process within CDR that simplifies deployment by grouping a collection of service operations and synchronizations that can be performed as a single task.

Sequence Editor In CDR, a predefined User Group to create, modify, or delete a sequence definition.

Sequence Performer (Production) In CDR, a predefined User Group to directly perform or request performance of a sequence action on production hosts.

Sequence Performer (Staging) In CDR, a predefined User Group to directly perform or request performance of a sequence action on staging hosts.

Sequence Requester (Production) In CDR, a predefined User Group to request performance of a sequence action on production hosts.

Sequence Requester (Staging) In CDR, a predefined User Group 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 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 Opsware Command Center Client that allows you to browse and manage servers and server groups in your facility.

server group A feature used to organize servers into groups in order to perform the same action on all of the servers. Server groups can be comprised of individual servers as well as other server groups.

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 Opsware SAS. Server states include Unprovisioned, Available, Installing OS, and Managed.

server management Process by which users can manage and track servers in an Opsware-managed environment. Opsware 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 Opsware SAS, but do not have a full operating system installed.

server provisioning The process of installing a basic set of software components that include the operating system, an Opsware Agent, and other system utilities and debugging tools to manage the server. Configuration is defined in the Model Repository.

Server 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 to define and modify or delete service definitions.

Service Performer (Production) In CDR, a predefined User Group to directly perform or request performance of service operations on production hosts (servers).

Service Performer (Staging) In CDR, a predefined User Group to directly perform or request performance of service operations on staging hosts.

Service Requester (Production) In CDR, a predefined User Group to request performance of service operations on production hosts.

Service Requester (Staging) In CDR, a predefined User Group 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 ontogeny.

Shared Scripts Public scripts that every Opsware SAS user can access.

Site Backup directory In CDR, the directory that stores a complete backup of the Live directory when the user issues a Backup service operation.

Site Previous directory In CDR, the directory that stores the files that have changed between the current Live directory and its previous state prior to the last cutover. It holds all the changes necessary to revert the Live directory back to the state that it was in before the last cutover.

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. Server Compliance 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 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 Opsware 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 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 to create, modify, or delete a synchronization definition.

Synchronization Performer In CDR, a predefined User Group to directly perform or request performance of a synchronization action.

Synchronization Requester In (CDR), a predefined User Group 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*.

template Used to install a set of (usually related) applications through a single invocation of a wizard.

template inheritance Process by which templates and folders inherit all attachments of the folder they reside in. Inheritance is propagated from parent (folder) to child (template or folder) and to all children of children.

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 Opsware SAS. An Opsware user belongs to one or more User Groups, which control the access of its members.

User Group Represents a role played an organization's Opsware users. The permissions specified for a user group determine what the group's members can do with Opsware 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).

Web Service API A web services interface that facilitates the integration of operations and business support systems with Opsware SAS. The Opsware Web Services APIs allow other IT systems, such as customers' existing monitoring, trouble ticketing, billing, and virtualization technology, to exchange information with Opsware SAS.

Web Services Data Access Engine A web services interface to the Model Repository that provides increased performance to other Opsware SAS components.

Index