HP Server Automation

for the HP-UX, Solaris, Red Hat Enterprise Linux,
VMware, and Windows operating systems

Software Version: 7.50

User's Guide: Server Automation

Document Release Date: September 2008

Software Release Date: September 2008

Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

For information about third party license agreements, see the Third Party and Open Source Notices document in the product installation media directory.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notices

Trademark Notices

Microsoft®, Windows®, Windows Vista®, and Windows® XP are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

`http://h20230.www2.hp.com/selfsolve/manuals`

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

`http://h20229.www2.hp.com/passport-registration.html`

Or click the New users - please register link on the HP Passport login page.

Support

Visit the HP Software Support Online web site at:

`www.hp.com/go/hpsoftwaresupport`

This web site provides contact information and details about the products, services, and support that HP Software offers.

For downloads, see:

`https://h10078.www1.hp.com/cda/hpdc/display/main/index.jsp?zn=bto&cp=54_4012_100__`

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services

- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

`http://h20229.www2.hp.com/passport-registration.html`

To find more information about access levels, go to:

`http://h20230.www2.hp.com/new_access_levels.jsp`

Table of Contents

Chapter 3: Getting Started with SA Client 69

Chapter 4: Agent Management 105

Chapter 6: Virtualization Director 209

Chapter 7: Server Tracking in the SAS Web Client 231

Chapter 8: Server Management in SAS Web Client 275

Chapter 9: NA Integration 353

Chapter 11: Code Deployment and Rollback 401

Chapter 12: Configuration Tracking 429

Appendix B: Agent Utilities 487

Appendix C: Global Shell Utilities Syntax 519

Appendix D: OGFS Directories 531

Preface

Welcome to HP Server Automation (SA) – an enterprise-class software solution that enables customers to get all the benefits of the SA data center automation platform and support services. SA 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 SA, 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: Introduction to HP Server Automation: Provides a high-level overview of HP Server Automation, including the system features, Web Service APIs, and multimaster.

Chapter 2: Getting Started with the SAS Web Client: Includes information about supported operating systems and browsers, navigation of the user interface, and an explanation of each of the features found on the SAS Web Client home page.

Chapter 3: Getting Started with SA Client: Includes information about how to get started using the SA Client, the user interface, the client installation and launch, and SA Client main features.

Chapter 4: Server Agent Management: Includes information about Server Agents on managed servers, the Discovery and Agent Deployment feature, and Server Agent reachability communication test.

Chapter 5: Exploring Servers and Jobs in SA Client: Provides information browsing servers, server groups, and jobs, also includes information about copying files on a managed server's file system.

Chapter 6: Virtualization Director: Provides information on managing your virtual servers with the SA Client, such as viewing virtual servers in your managed environment, provisioning VMware ESX 3, creating and managing Solaris local zones, and searcing for , grouping, and performing operations on virtual servers.

Chapter 7: **Server Tracking in the SAS Web Client**: Provides information about server asset tracking, server lists, server search, server histories and reports.

Chapter 8: **Server Management in SAS Web Client**: Provides information about all aspects of server management including server groups, server life cycle, server locking, and service levels.

Chapter 9: **Integration with NA**: Provides information about how you can monitor managed servers and network devices that are connected to them by using HP Server Automation (SA), Network Automation (NA), Global Shell, and Server Automation Visualizer (SAV).

Chapter 10: **Global Shell**: Provides information about the Global Shell, the Global File System, setting user and group permissions with the Global Shell, accessing the Global Shell and remote terminals on servers, and explains Global Shell commands.

Chapter 12: **Code Deployment and Rollback**: Provides information about uploading code and content to staging, and performing services, synchronizations, and sequences to deploy code and content to managed servers.

Chapter 13: **Configuration Tracking**: Provides information about configuration tracking policies, the supported types of configuration files and databases, how changes are detected, and the tracking policies for a specific server. It also discusses reconciling customized tracking policies, performing manual backups, viewing backup history, restoring backed up files, and enabling and disabling configuration tracking.

Appendix A: **Communication Test Troubleshooting**: Provides troubleshooting information to diagnose Server Agent unreachability problems.

Appendix B: **Server Agent CLI Utilities**: Provides information about installing Server Agents using the Command Line Interface (OCLI) and the Server Agent Upgrade Tool.

Appendix C: **Global Shell Utilities**: Describes the syntax and usage rules for the `aaa`, `rosh`, and `swenc` commands.

Appendix D: **OGFS Directories**: Provides an overview of the directories under `/opsw`, which can be accessed from within a Global Shell session.

Appendix E: **Custom Extensions:** Provides information on running custom extensions using the SAS Web Client.

Appendix F: **Glossary**: Defines terminology and acronyms that are unique to HP Server Automation.

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, HP Server Automation 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 icons.

ICON	DESCRIPTION
📄	This icon represents a note. It identifies especially important concepts that warrant added emphasis.

ICON	DESCRIPTION
☑	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 *User's Guide: Server Automation* is intended for system administrators responsible for all aspects of managing servers in an operational environment. It describes how to use SA, introducing the system and the user interface. It provides information about managing servers, remediating servers, script execution, configuration tracking, deploying and rolling back code, and agent deployment. It also explains how to use the Global Shell and open a Remote Terminal on managed servers.
- The *SA User's Guide: Application Automation* is intended for system administrators responsible for performing the day-to-day functions of managing servers. It reviews auditing and compliance, software packaging, visual application management, application configuration, and software and operating system installation on managed servers.
- The *SA Administration Guide* is intended for administrators responsible for monitoring and diagnosing the health of the SA core components. It also documents how to set up SA user groups and permissions.
- The *SA Planning and Installation Guide* is intended for advanced system administrators responsible for planning all facets of an SA installation. It documents all the main features of SA, scopes out the planning tasks necessary to successfully install SA, explains how to run the BSA Installer, and details how to configure each of the

components. It also includes information on system sizing and checklists for installation.

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

Chapter 1: Introduction to HP Server Automation

IN THIS CHAPTER

This section discusses the following topics:

- Overview of HP Server Automation (SA)
- Types of SA Users
- SA Features
- HP Server Automation Terms and Concepts

Overview of HP Server Automation (SA)

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

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

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

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

The following figure provides an overview of how SA automates server and application operations across all major platforms and a wide range of applications. Each feature that is shown in the diagram is discussed in the following sections.

Figure 1-1: Overview of SA Features

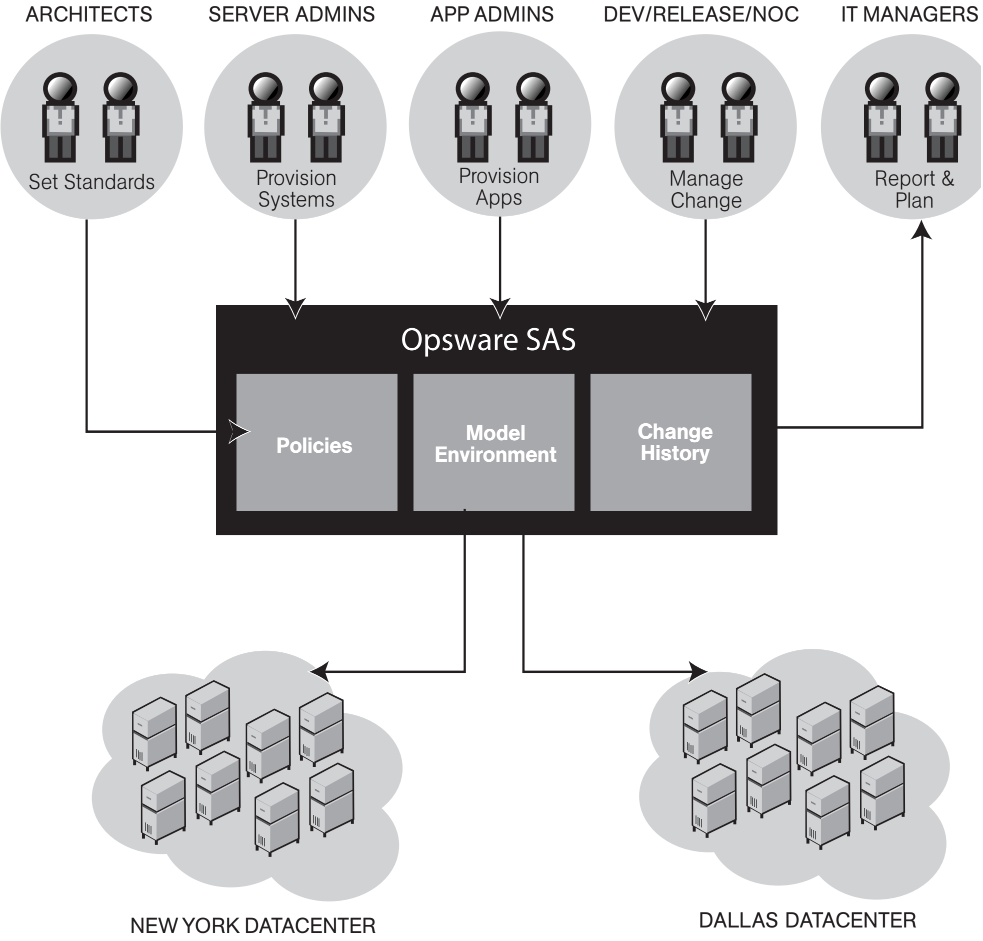

Types of SA Users

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

Table 1-1: Types of SA Users

SA USER	RESPONSIBILITIES
Data Center and Operations Personnel	After manually racking and stacking servers, manage customer facilities and boot bare-metal servers over the network or from an SA boot image.
System Administrators	Install operating systems and applications (for example, Solaris 5.7 or WebLogic 6.0 Web Server), upgrade servers, create operating system definitions, and set up software management policies.
Site Engineers and Customer Project Managers	Deploy custom code on servers.

In addition to the SA users listed above, this guide describes the following three types of users:

- **End Users** are responsible for all aspects of managing and provisioning the servers in an operational environment. In the SA documentation, these users are referred to as SA users or system administrators. These users log into the SAS Web Client and SA Client and use these interfaces to manage servers in their IT environment.
- **SA Administrators** are the users, with special training and information, who are responsible for installing and maintaining SA. In the SA documentation, these users are referred to as SA administrators. They use the Administration features in the SAS Web Client to manage SA and SA users (by adding user accounts and assigning permissions for different levels of operation and access), to add customers and facilities, and to change SA configurations. They monitor and diagnose the health of SA components. SA administrators need to understand how SA features operate to support users and SA.
- **Policy Setters** are the power users who are responsible for architecting what SA will do in the managed environment; for example, they determine which operating systems can be installed on your managed servers and how those operating systems will be configured during installation. Policy setters, for example, prepare specific features in SA by defining the Software Policies, preparing Operating System Definitions, and

acting as Patch Administrators to approve patches for installation in the operational environment.

SA Interfaces and Tools

Depending on the type of operation you need to perform with HP Server Automation, you select the appropriate user interface, as Figure 1-2 shows.

Figure 1-2: Interfaces in HP Server Automation

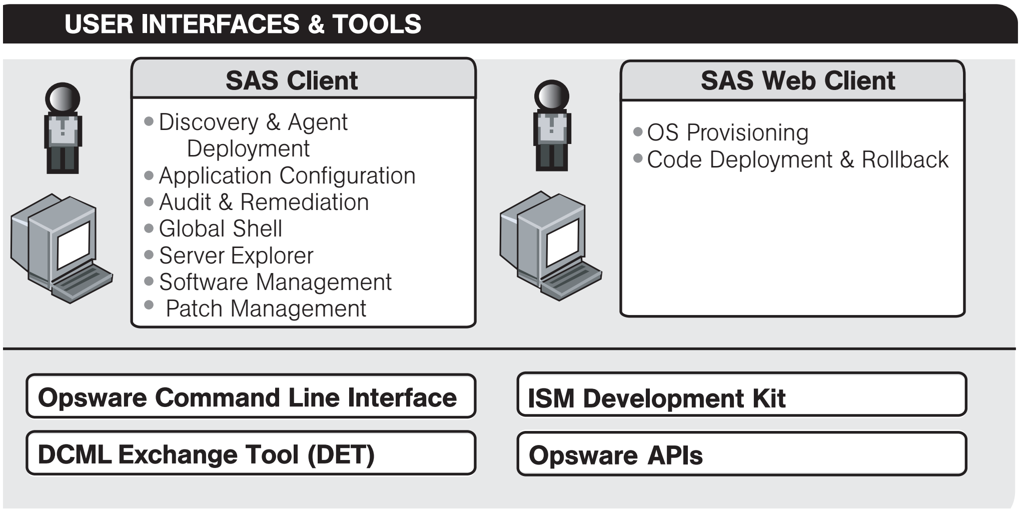

- **SAS Web Client**: The web-based user interface to HP Server Automation through which users can manage servers, provision applications and operation systems onto servers, run distributed scripts on servers, and deploy code and content to servers, among other things.
- **SA Client**: A Java Web-Start application that extends the SAS Web Client features and provides the following new features:
  - Discovery and Agent Deployment
  - Device Explorer
  - Virtualization Director
  - Server Automation Visualizer (SAV)
  - Audit and Remediation
  - Compliance View

  - Reports
  - Software Management
  - Script Execution
  - Patch Management for Windows
  - Patch Management for Unix
  - Application Configuration Management
  - Global Shell
  - NA Integration

- **Command Line Interface (OCLI)**: A command line interface that users can use to upload packages to the SA Software Repository (version 1), and perform many other HP Server Automation operations (version 2).
- **DCML Exchange Tool (DET)**: A utility that enables users to export almost all server management content from any core – standalone or multimaster mesh – and import it into any other core. HP Server Automation can also provide pre-packaged server management content appropriate for new installations that can be imported into a core after initial setup. See the *SA Content Utilities Guide* for information about using this utility.
- **ISM Development Kit**: A development kit that consists of command-line tools and libraries for creating, building, and uploading ISMs. An ISM is a set of files and directories that include application bits, installation scripts, and control scripts. See the *SA Content Utilities Guide* for information about using the ISM Development Kit.
- **SA APIs**: A set of APIs and a command-line interface (CLI) that facilitate the integration and extension of HP Server Automation. This platform allows other IT systems – such as existing monitoring, trouble ticketing, billing, and virtualization technology – to exchange information with HP Server Automation. This broadens the scope of how IT can use HP Server Automation to achieve operational goals.

SA Features

SA is made up of a set of features components that automate particular IT processes.

The features are designed to replace ad hoc, error-prone, manual processes. For example, by using the OS Provisioning feature, users can set standards for different types of servers and automatically provision the servers, saving time and ensuring that operating system builds are consistent. By using the Patch Management feature, users can establish polices about how patches are installed. HP Server Automation uniformly enforces those polices.

The following features are currently available as part of HP Server Automation:

- Operating System Provisioning
- Code Deployment & Rollback
- Configuration Tracking
- Script Execution
- Discovery and Agent Deployment
- Device Explorer
- Virtualization Director
- Server Automation Visualizer (SAV)
- Audit and Remediation
- Compliance View
- Reports
- Software Management
- Patch Management for Windows
- Patch Management for Unix
- OS Provisioning
- Application Configuration Management
- Global Shell
- NA Integration

All SA features support cross-platform environments and are designed to automate both new and existing data center environments. See the following figure.

Figure 1-3: HP Server Automation Features

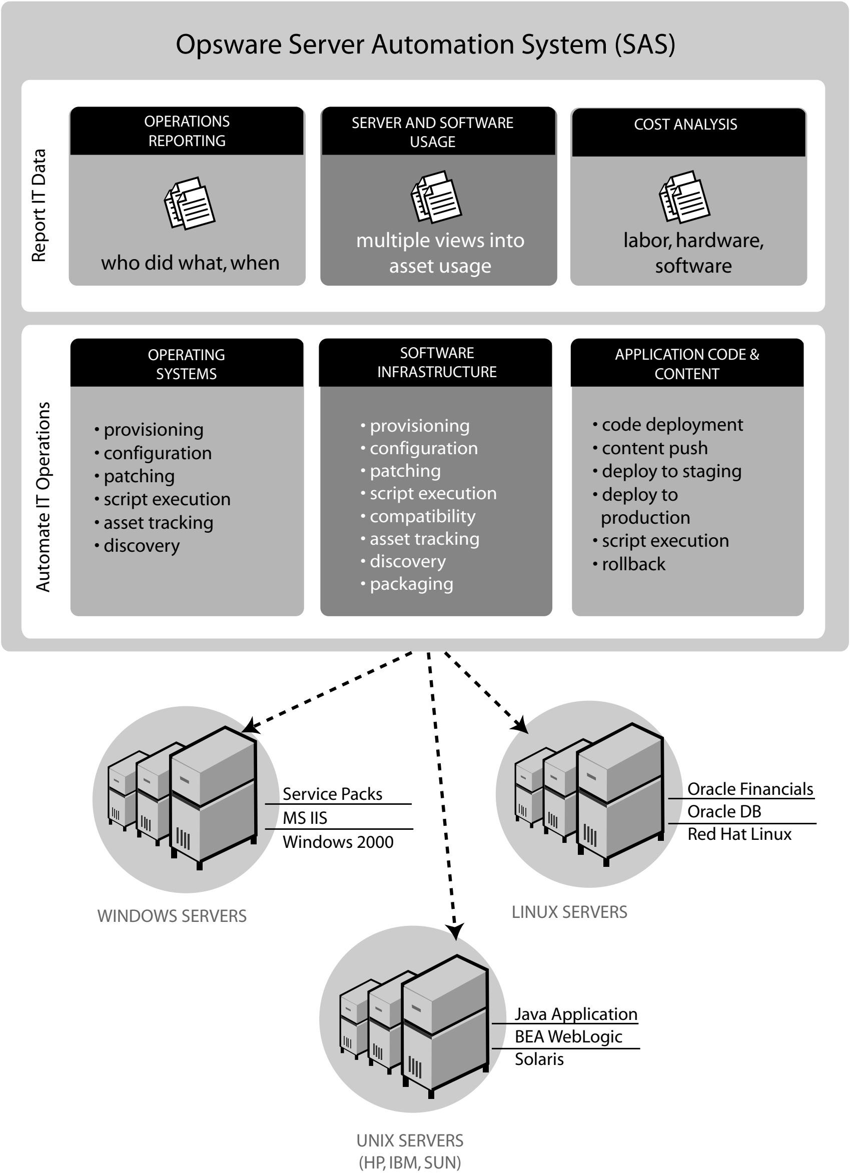

Operating System Provisioning

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

Benefits of the OS Provisioning feature include the following items:

- **Integration with the other features of SA**

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

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

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

- **Flexible architecture designed to work in many environments**

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

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

- Preparing the hardware for OS installation using an OS installation profile
- Creating OS sequences that define a server build policy, including application policies, patch policies, device groups, and remediation policies
- Installing a base operating system and default OS configuration using an OS sequence
- Applying the latest set of OS patches, the exact list depends on the applications running on the server

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

Code Deployment & Rollback

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

Specifically, CDR enables you to perform the following actions:

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

Configuration Tracking

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

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

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

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

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

Script Execution

The Script Execution feature enables you to share and run ad-hoc or saved scripts across an entire farm of SA-managed servers.

By executing scripts with SA instead of manually, administrators benefit by using the following features:

- Parallel script execution across many Unix and/or Windows servers, saving time and ensuring consistency.
- Role-based access control, ensuring only authorized administrators can execute scripts on hosts to which they have access.
- The ability to control access to scripts by storing them in private or in public libraries.
- The ability to see and download script output one server at a time or in a consolidated report, which captures output from all servers in a single place.
- The ability for scripts to be mass-customized. Administrators can access information in SA about the environment and the state of servers. This is critical to ensuring that the right scripts are executed on the right servers.
- A comprehensive audit trail that reports who, what, when, and where a particular script was executed.

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

- Using known system state and configuration information to customize script execution, users can tailor each script by referencing and accessing the rich store of information in SA, such as the customer or business that owns the server, whether the server is a staging or production server, which facility the server is located in, and custom name-value pairs.

- By sharing scripts without compromised security, users can share scripts with each other without compromising security because SA maintains strict controls on who can execute scripts on which servers and generates a comprehensive audit trail of script execution.

Discovery and Agent Deployment

The SA Discovery and Agent Deployment (ODAD) feature allows you to deploy Server Agents to a large number of servers in your facility and place them under SA management.

Using the ODAD features, you can perform the following tasks:

- Scan your network for servers.
- Select servers for SA Agent installation.
- Select a communication tool and provide user/password combinations.
- Choose agent installation options and deploy agents.

Device Explorer

The Device Explorer lets you view information about servers in your managed environment.

From the Server Explorer, you can perform the following tasks:

- Create a server snapshot, perform a server audit, audit application configurations, create a package, and open a remote terminal session on a remote server.
- Browse a server’s file system, registry, hardware inventory, software and patch lists, and services.
- Browse SA information such as properties, configurable applications, and even server history.

From the Groups Browser, you can perform the following tasks:

- Audit system information, take a server snapshot, and configure applications.
- View and access group members (servers and other groups).
- View group summary and history information.

Virtualization Director

The Virtualization Director feature enables you to provision and manage virtual servers for Solaris 10 local zones and VMware ESX 3 virtual machines (VMs). Using the SA Client, you can perform the following tasks:

- View both hypervisor and virtual servers and their relationships in the SA Client, so you can find out the hypervisors that are hosting your virtual machines and local zones.
- View virtual servers and their relationship in the HP Server Automation Visualizer (SAV).
- Provision VMware ESX and Solaris 10 hypervisors on bare metal servers.
- Provision VMware virtual machines (VM) using an OS sequence.
- Create, start, stop, modify, and remove Solaris local zones.
- Deploy agents on unmanaged virtual servers using the Agent Discovery and Deployment for VMware ESX VMs.
- Search for virtual servers in your data center using the Search tool.
- Create dynamic Device Groups based upon virtual server characteristics (zones or VMs).

Server Automation Visualizer (SAV)

The Server Automation Visualizer (SAV) feature is designed to help you optimally understand and manage the operational architecture and behavior of distributed business applications in your IT environment. Since these applications are complex collections of services that typically run across many servers, as well as network and storage devices, it can become increasingly 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.

SAV helps you see (visualize) this type of information through physical and logical drawings.

Audit and Remediation

The Audit and Remediation feature 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, perform server comparisons against a baseline, and create custom audit policies that define company or industry server configuration compliance standards, and which can be used inside of audits, snapshot specifications, and audit policies.

Using Audit and Remediation, you can perform the following tasks:

- Compare servers or snapshots to reference servers or snapshots
- Create audits for repeated use
- Create audit policies that define compliance and security standards for your organization
- Associate audits with individual servers or dynamic server groups
- Remediate problems at multiple levels, including files, directories, patches, registry keys, and packages

Compliance View

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

Reports

The 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 reports are presented in graphical and tabular format, and are actionable—where 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 (in .html and .xls formats) to facilitate use within your organization.

Software Management

The Software Management feature in HP Server Automation provides a powerful mechanism to model software by using software policies and to automate the process of deploying software and configuring applications on a server in a single step. In addition, the Software Management feature provides a structure to organize your software resources in folders and define security permissions around them. This feature allows you to verify the compliance status of a server and remediate non-compliant servers.

The Software Management feature in SA Client provides the following functions:

- Create an organizational structure for software
- Define security boundaries for folders
- Define a model-based approach to manage the IT environment in your organization
- Enable sharing of software resources among user groups
- Deploy and configure applications simultaneously
- Deploy multiple application instances on one server
- Establish a software deployment process
- Verify compliance status of servers to software policies
- Generate reports
- Comprehensively search for software resources and servers

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 SA Client user interface, you can identify and install patches that support security vulnerabilities for the Windows 2000, Windows 2003, and Windows NT4.0 operating systems. These patches include Service Packs, Update Rollups, and hotfixes.

Patch Management for Unix

The Patch Management for Unix feature enables you to identify, install, and remove Unix patches to maintain a high level of security across managed servers in your organization. With the SA Client, you can identify and install patches that support security vulnerabilities for the AIX, HP-UX, Linux, and Solaris operating systems.

OS Provisioning

OS Provisioning in the SA Client allows you to install an operating system, applications, and packages and packages on unprovisioned servers by creating OS Installation Profiles and OS Sequences. From the Devices list, you can view all unprovisioned servers in your facility and provision those servers by running an OS Sequence.

OS Sequences allow you to set up a complete server build (policy) that represents the ideal manner in which a particular OS should be installed, which includes the proper OS Installation Profile that should be used, as well as Application and Patch policies, the servers to install the OS on, and how these policies should be remediate either before or after the OS is installed.

Application Configuration Management

Application Configuration Management (ACM) allows you to create templates so you can modify and manage application configurations associated with server applications. ACM enables you to manage, update, and modify those configurations from a central location, ensuring that applications in your facility are accurately and consistently configured.

Using ACM, you can perform the following tasks:

- Manage configurations based on files and objects, such as Windows registry, IIS metabase, WebSphere, COM+, and more.
- Preview configuration changes before applying them.
- Edit and push configuration changes to individual servers or server groups.
- Use information in the SA data model to set configuration values.
- Manage configurations of any application by building configuration templates.
- 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.

Global Shell

The Global Shell feature enables you to manage servers by using a command-line interface. You can remotely perform the following tasks:

- Complete routine maintenance tasks on managed servers.
- Troubleshoot, identify, and remediate problems on managed servers.

Global Shell consists of a file system and a command-line interface to that file system for managing servers in SA. The file system is known as the SA Global File System (OGFS). All object types in the OGFS (such as servers, customers, and facilities) are represented as directory structures in this file system.

The Global Shell feature also manages user permissions for accessing the file system, Windows Registry, and Windows Services objects on managed servers.

NA Integration

The NA Integration feature enables you to closely examine detailed information about managed servers and the network devices connected to them so that you can determine how they are related and then, subsequently, coordinate and implement those changes. This feature supports an integrated approach to using NA and SA so that you can perform actions on device groups, such as combine event history, determine compliance, and identify duplex mismatches across servers and network devices in your environment.

HP Server Automation Terms and Concepts

This section discusses the following topics:

- Agent-Server Architecture of SA Technology
- Server Management in Multiple Facilities
- Multimaster Support

Agent-Server Architecture of SA Technology

The agent-server architecture of HP Server Automation enables server management. The server portion of HP Server Automation consists of multiple, integrated components, each with a unique purpose. Each server managed by HP Server Automation runs an intelligent agent (the Server Agent).

Figure 1-4: HP Server Automation *Agent-Server Architecture*

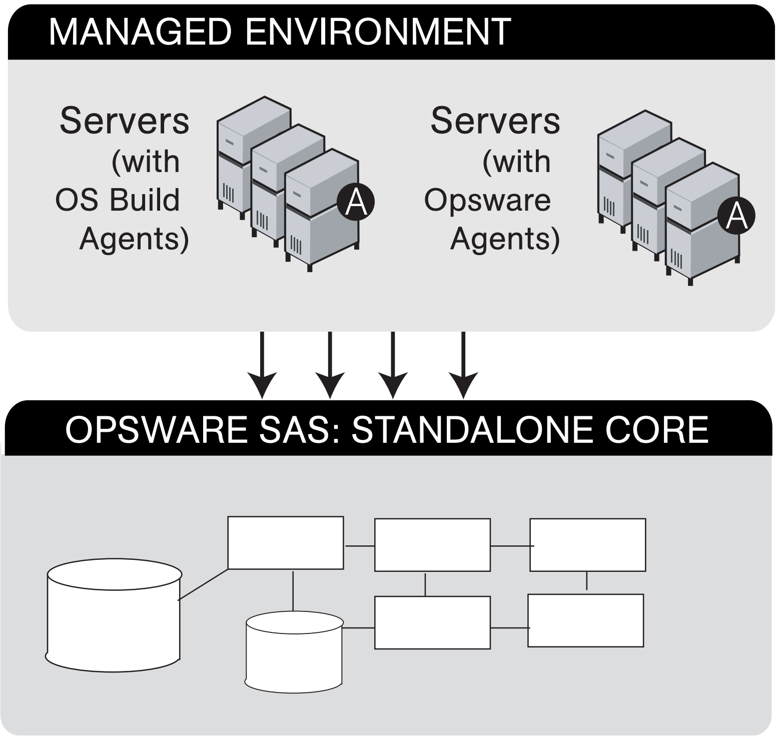
MANAGED ENVIRONMENT
Servers (with OS Build Agents)
A
Servers (with Opsware Agents)
A
OPSWARE SAS: STANDALONE CORE

See the *SA Administration Guide* for a detailed description of the components that make up the HP Server Automation server portion of the architecture.

The Server Agent is the agent of change on a server. Whenever HP Server Automation needs to make changes to servers, it does so by sending requests to the Server Agents. Depending on the request, the Server Agent on a server might use global HP Server Automation services in order to fulfill the request. For example, the Server Agent might often make requests to the Model Repository, the central database for all HP Server Automation components, and the Software Repository, the central repository for all software that HP Server Automation manages.

Some functions that the Server Agent supports are:

- Software installation and removal
- Configuration of software and hardware
- Periodically reporting server status
- Auditing of the server

A Server Agent is idle unless HP Server Automation is trying to perform some change on the server. In addition, each Server Agent periodically contacts the Data Access Engine and registers itself. The Data Access Engine is an XML-RPC interface to the model repository. The Data Access Engine sends this data to the Model Repository, which allows the Model Repository to keep track of server status, and know when particular servers are disconnected from or reconnected to the network.

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

See "Communication Test Troubleshooting" on page 457 in Appendix A for information about how to install an Server Agent on a server so that HP Server Automation can manage it.

Server Management in Multiple Facilities

The managed environment might span several facilities. A facility refers to the collection of servers that a single SA Model Repository manages, and the database that stores information about the managed environment. For example, one facility might be dedicated to an organization's Intranet, while another facility might be dedicated to the web services offered to the public. Your HP Server Automation can contain facilities (a full HP Server Automation is installed) and Satellite facilities. See Figure 1-5.

Figure 1-5: Server Management in Multiple Facilities

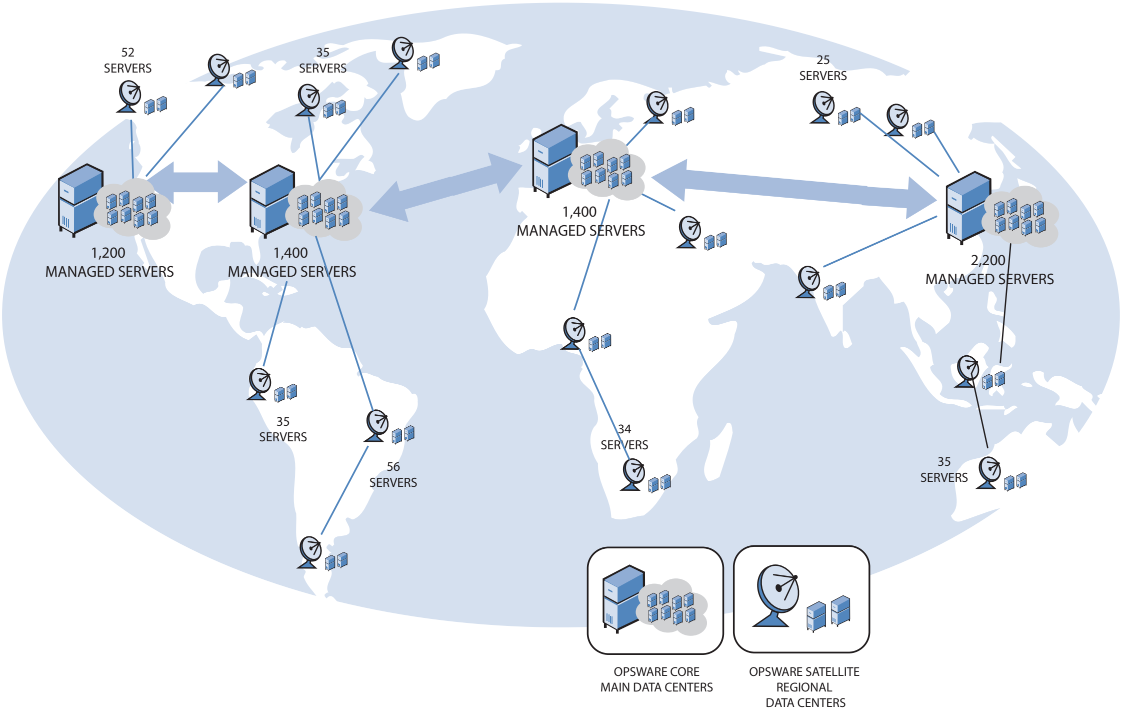

See the *SA Planning and Installation Guide* for more information about the types of installations that HP Server Automation supports.

Users can manage servers in any facility from an SAS Web Client in any facility. When a user updates data in a facility, the Model Repository for that facility is synchronized with the Model Repositories located in all remote facilities.

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

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

Multimaster Support

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

Through the Model Repository Multimaster Component, HP Server Automation provides the ability to easily rebuild server and application environments, provision additional capacity, distribute updates, and share software builds, templates and dependencies – across multiple facilities and from one user interface. See Figure 1-6.

Figure 1-6: Multimaster Support

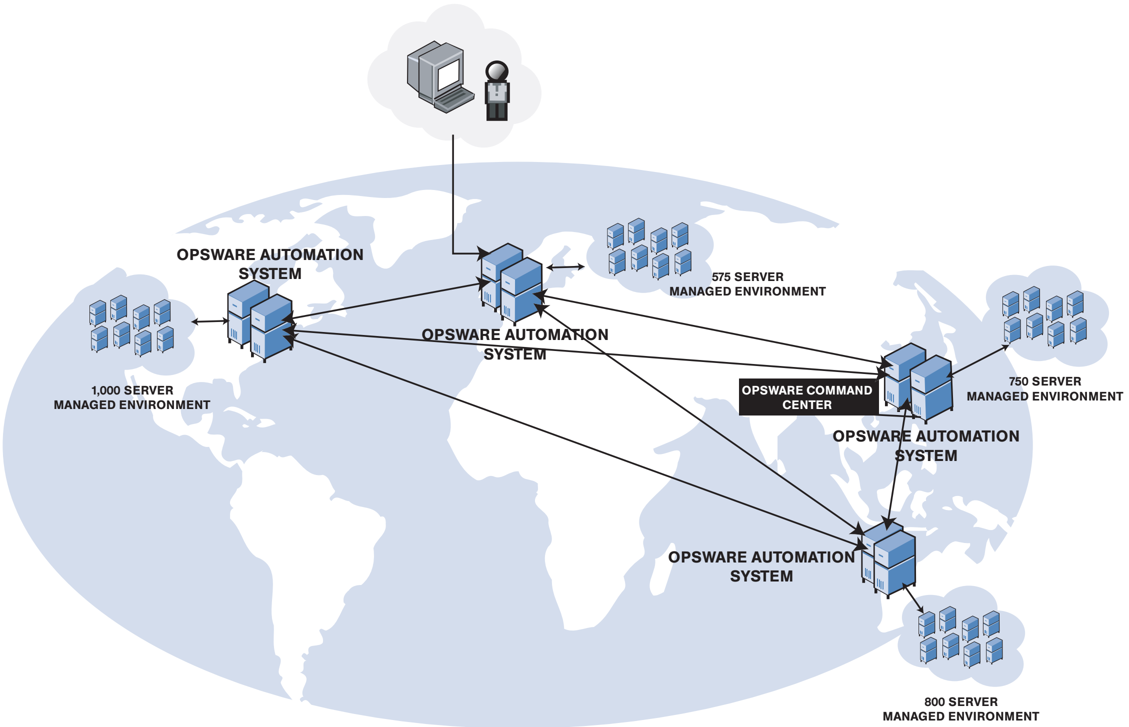
OPSWARE AUTOMATION SYSTEM
1,000 SERVER MANAGED ENVIRONMENT
OPSWARE AUTOMATION SYSTEM
575 SERVER MANAGED ENVIRONMENT
OPSWARE COMMAND CENTER
750 SERVER MANAGED ENVIRONMENT
OPSWARE AUTOMATION SYSTEM
OPSWARE AUTOMATION SYSTEM
800 SERVER MANAGED ENVIRONMENT

Chapter 2: Getting Started with the SAS Web Client

IN THIS CHAPTER

This section provides an overview of the SAS Web Client and discusses the following topics:

- Getting Started with the SAS Web Client
- Access to Features in the SAS Web Client
- User Interface

Getting Started with the SAS Web Client

The following section describes Supported Browsers for the SAS Web Client (a web application) and Browser Configuration Requirements.

Supported Browsers for the SAS Web Client

The following table lists the supported browsers for the SAS Web Client.

Table 2-1: Supported Browsers for the SAS Web Client

BROWSER	WINDOWS VISTA	WINDOWS 2000	WINDOWS 2003	WINDOWS XP
Microsoft Internet Explorer 5.5		X		
Microsoft Internet Explorer 6.0		X	X	X
Microsoft Internet Explorer 7.0	X	X	X	X

Table 2-1: Supported Browsers for the SAS Web Client (continued)

BROWSER	WINDOWS VISTA	WINDOWS 2000	WINDOWS 2003	WINDOWS XP
Mozilla 1.6		X	X	X
Firefox 1.0		X	X	X
Firefox 2.0	X	X	X	X

Browser Configuration Requirements

To use the SAS Web Client, your browser must be configured in the following manner:

- The browser must accept cookies and be able to use Java.
- The browser must support SSL and should provide 128-bit encryption (recommended).
- Using a pop-up blocker might prevent some functions from working correctly. Either disable the pop-up blocker completely or use the supported browser's native pop-up blocking function instead of a third-party product.

Access to Features in the SAS Web Client

This section explains how to configure your user profile to access features within the SA System. This section contains the following topics:

- System Management of SA User Information
- Best Practices for Selecting Passwords
- Updating User Profiles and Passwords

System Management of SA User Information

An SA administrator creates additional users who can use the SAS Web Client, and the SA administrator assigns them temporary passwords. Once added to the system, an SA user can update their personal information, password, and time zone and date display preferences by using the My Profile link.

An SA user cannot change their access permissions by using the My Profile link. If an SA user needs additional access permissions, they contact their SA administrator.

Best Practices for Selecting Passwords

The HP Server Automation enforces a security policy that allows only authorized users to log into the SAS Web Client. SA users are advised to select a password based on the following guidelines:

- Change your password frequently to ensure that your account information is secure.
- Select a password that is easy to remember so that you don't have to write it down.
- Use a mixture of upper and lower case letters, numbers, and punctuation in your password.
- Do not share your password.
- Do use a password that you can type quickly, without having to look at the keyboard.

Updating User Profiles and Passwords

As an SA user you can change your name, contact information, password, and preferences such as time zone and date display. You cannot change your access permissions. Contact your SA administrator to change your access permissions.

Perform the following steps to change your profile and password.

1. Log into the SAS Web Client. The SAS Web Client Home Page appears.
2. Click the My Profile link located at the top of the page. The My Profile Page appears.
3. To change your profile information, enter new information on the User Identification tab of the My Profile Page.
4. Click **Save**.
5. To change your password, click the Change Password Link. The Password Change page appears.
6. Enter your old password.
7. Enter your new password.
8. Confirm your new password in the Confirm Password field.
9. Click **Save**. A confirmation page appears indicating that your password was successfully changed.
10. Click **Okay**. The Profile page appears.

User Interface

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

- Requirements for Logging In
- Overview of the SAS Web Client User Interface
- My Profile
- Search
- My Servers
- Mouseover Icon Tooltips

Requirements for Logging In

In order to log in and access SAS Web Client features, your SA administrator must have created a user name and password for you, and assigned user permissions that control the features you can access, the actions you can perform, and the resources you can access.

Overview of the SAS Web Client User Interface

The SAS Web Client user interface consists of the five following sections:

- Home Page
- Tasks
- My Jobs
- My Customers
- Navigation Panel

Home Page

Figure 2-1 shows the Home page as it appears when you log in or when you click the Home link in the navigation bar.

Figure 2-1: SA System Home Page

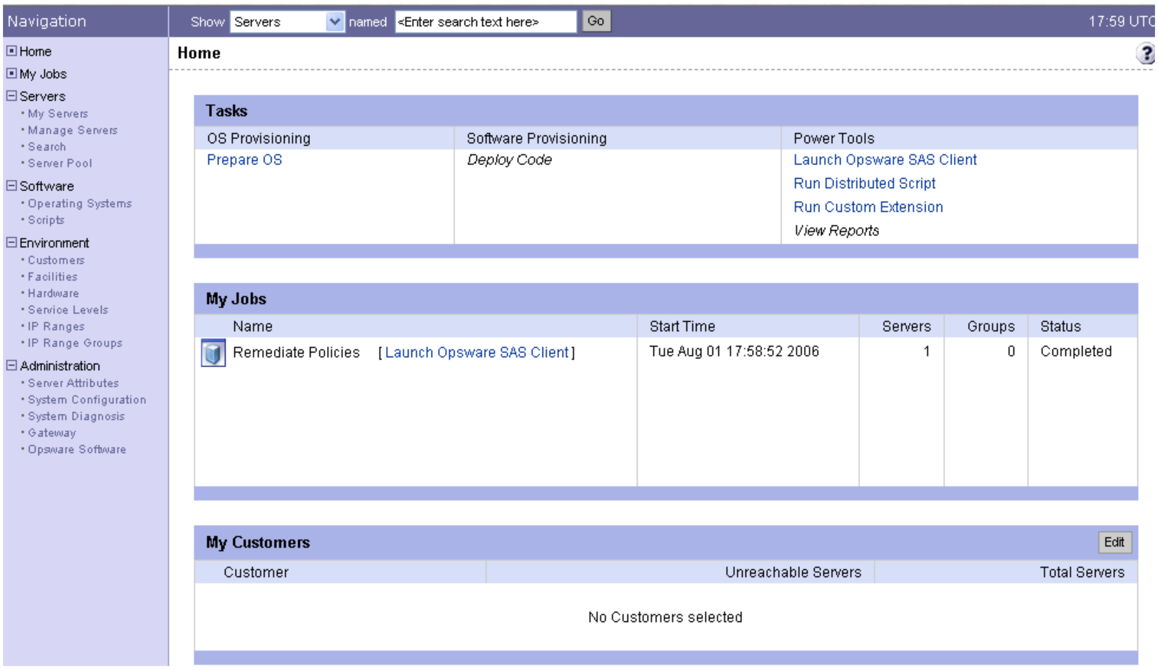

Home

Tasks

OS Provisioning	Software Provisioning	Power Tools
Prepare OS	*Deploy Code*	Launch Opsware SAS Client
		Run Distributed Script
		Run Custom Extension
		View Reports

My Jobs

Name	Start Time	Servers	Groups	Status
Remediate Policies [Launch Opsware SAS Client]	Tue Aug 01 17:58:52 2006	1	0	Completed

My Customers

Customer	Unreachable Servers	Total Servers
No Customers selected		

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

Tasks

The Tasks area of the Home page displays links to the wizards that you have permissions to access, a link to the Code Deployment page if you have that permission. If you do not have permissions to a task in this area, the task name still displays, but it is italicized and it is not an active link. Figure 2-2 shows the Tasks area with all permissions enabled.

Figure 2-2: Tasks Area of the Home Page

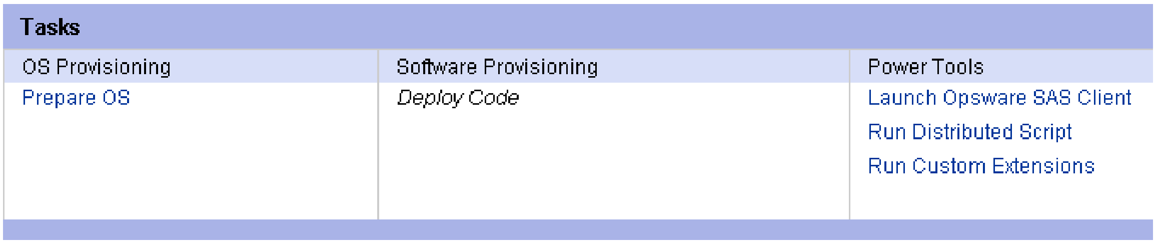

Tasks

OS Provisioning	Software Provisioning	Power Tools
Prepare OS	*Deploy Code*	Launch Opsware SAS Client Run Distributed Script Run Custom Extensions

My Jobs

The My Jobs area of the Home page is populated with details of the jobs that you have run, jobs that are currently in progress, or jobs that you have scheduled to run, including the name of the job, the start time, the number of servers affected by the job, and the status of the job. If there are more than six jobs, you can see the rest of them by clicking the See All link, which also shows the total number of jobs in parentheses, as Figure 2-3 shows.

If the job was run in the SA Client, then there will be a link next to the job that (when clicked) will launch the SA Client. You can view the more detailed information about the job in the SA Client.

Figure 2-3: My Jobs Area of the Home Page

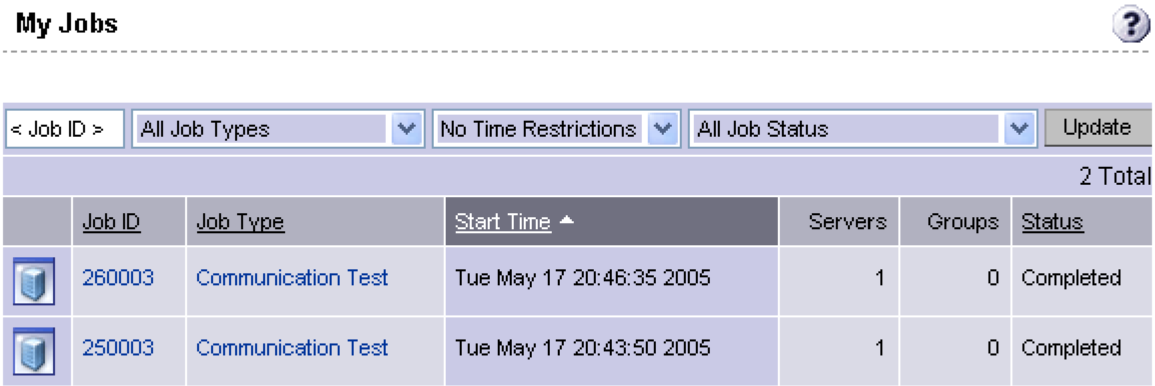

My Jobs

2 Total

	Job ID	Job Type	Start Time ▲	Servers	Groups	Status
	260003	Communication Test	Tue May 17 20:46:35 2005	1	0	Completed
	250003	Communication Test	Tue May 17 20:43:50 2005	1	0	Completed

My Customers

The My Customers area of the Home page is populated with customer information, including unreachable servers associated with a customer and the total number of servers associated with that customer. To select the customers to display in the My Customers area of the Home page, click **Edit** and select the check box next to the customer name. See Figure 2-4.

Figure 2-4: My Customers Area of the Home Page

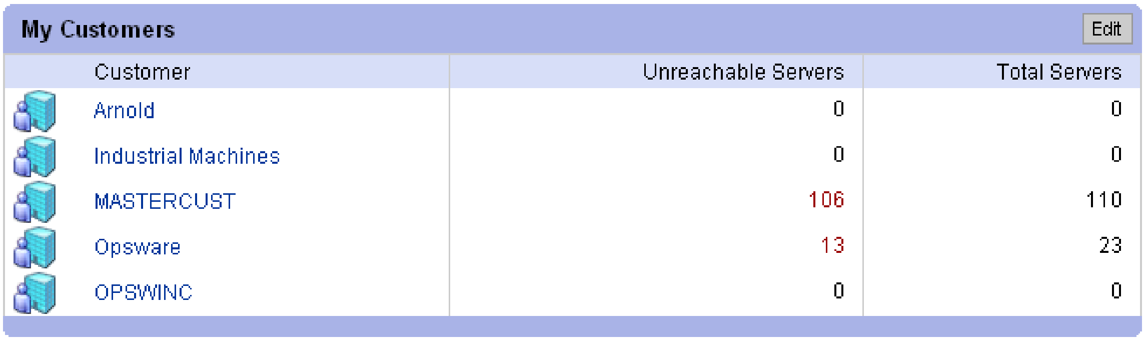

Customer	Unreachable Servers	Total Servers
Arnold	0	0
Industrial Machines	0	0
MASTERCUST	106	110
Opsware	13	23
OPSWINC	0	0

Navigation Panel

The navigation panel on the left side of the SAS Web Client shows all possible features, as shown in Figure 2-5. The features you can access and the actions you can perform depend upon your user profile, as defined by the SA administrator.

Figure 2-5: Navigation Panel, All Permissions View

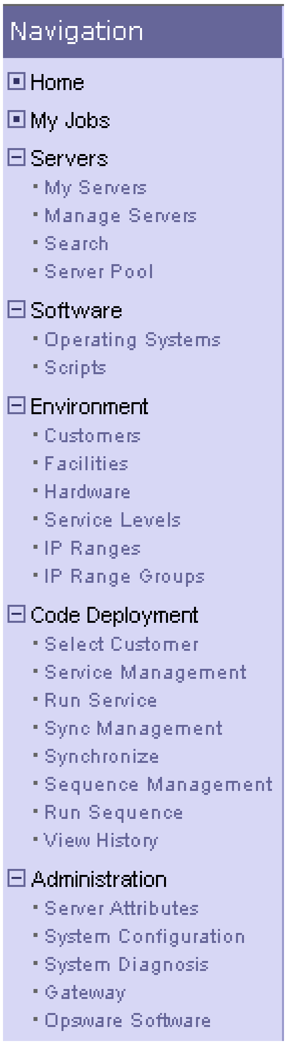

The items that appear in the navigation panel depend on the permissions the user has. Clicking an item on the navigation panel displays that feature in the main part of the Home page. For a user with all permissions, the following links appear:

Home: Displays the top level of the SAS Web Client. The Home page is described in this section of the guide. Wizards are documented in their respective functional areas of the system.

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

Servers: Expands to display these selections:

- **My Servers**: Use to add any server or server group to your own personal view of servers. My Servers provides an efficient way to manage servers when your operational environment contains hundreds or thousands of servers.
- **Manage Servers**: Use filters to display a list of servers and perform operations on them such as edit server values, assign, run scripts, and add servers to My Servers. See “Server Management in SAS Web Client” on page 275 in Chapter 8 for more information.
- **Search**: Find specific servers using default criteria or user-defined criteria. See “Server Search” on page 242 in Chapter 7 for more information.
- **Server Pool**: Use filters to display a list of servers, install operating systems on the servers, and delete the servers. See *SA User’s Guide: Application Automation* for more information.

Software: Expands to display these selections:

- **Operating Systems**: Prepare operating systems for installation and delete existing operating systems. See *SA User’s Guide: Application Automation* for more information.
- **Scripts**: Run scripts, upload scripts, and create new scripts. See “Script Execution” on page 493 in Chapter 8 for more information.

Environment: Expands to display these selections:

- **Customers**: Create new customers and edit or delete existing customers. See the *SA Administration Guide*.
- **Facilities**: Create new facilities, edit facility properties, and assign and edit custom attributes for facilities. See the *Opsware® SAS Planning and Installation Guide*.

- **Hardware**: A read-only view of servers categorized by the hardware manufacturer and model, and their related information. See "Hardware Information for Managed Servers" on page 271 in Chapter 7 for more information.
- **Service Levels:** Define service levels and custom attributes. See "Service Levels" on page 342 in Chapter 8 for more information.
- **IP Ranges**: Identify and create IP ranges and IP range types. See the *SA Policy Setter's Guide* for more information.
- **IP Range Groups**: Create IP Range Groups. See the *SA Policy Setter's Guide* for more information.

Code Deployment: Expands to display these selections:

- **Deployment Home**: The exact CDR links that you see in the Code Deployment area are based on the permissions that you have for the customer you want to work with.
- **Service Management**: Create, modify, and delete service definitions. Services define the location and commands to manipulate applications on hosts.
- **Run Service**: Perform service operations on one or more hosts, or request that a service operation be performed on your behalf. Service operations include starting or stopping applications, cutting over or rolling back code, and backing up or restoring code.
- **Sync Management**: Create, modify, and delete synchronization definitions. Synchronizations define the path for pushing code from a source host to one or more destination hosts.
- **Synchronize**: Perform a synchronization to one or more hosts, or request that a synchronization be performed on your behalf.
- **Sequence Management**: Create, modify, and delete sequence definitions. Sequences allow the grouping of service operations and synchronization operations to define higher level code deployment operations.
- **Run Sequence**: Perform a pre-defined sequence of service operations and synchronizations on one or more hosts, or request that a sequence be performed on your behalf.
- **View History**: Get information about previously run Code Deployment operations. See "Code Deployment and Rollback" on page 401 in Chapter 11 for more information.

Administration: Expands to display the following features. For more information about these features, see the *SA Administration Guide* and *SA Policy Setter's Guide*.

- **Users & Groups**: Administrators use this feature to create user groups, define permissions for those groups, create new administrators, and add users to groups.
- **Server Attributes**: Define and edit server use attributes, enable them for code deployment, and define deployment stages. Also define and edit server deployment stage attributes.
- **System Configuration**: Contains the configuration parameters that define how the SA System works in your environment. This selection is only used at the direction of SA.
- **System Diagnosis**: Runs a series of tests on SA components to make sure that they are functioning correctly.
- **Gateway**: Allows you to connect Satellites with this or other cores.
- **Opsware Software**: Provides a view of the properties, custom attributes, installation order, and history of the software attached to the SA in the system.

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

Navigation

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

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

- Clicking a hyperlinked name to display a page

  For example, if you select Software ➤ Scripts (assuming that you have the correct permissions), the page that appears shows all of the scripts that have been uploaded so far.

- Selecting a tab to display a page

For example, when you select Manage Servers and click one of the hyperlinked server names, the resulting page shows a row of tabs, like Figure 2-6 shows.

Figure 2-6: Example of Tabs

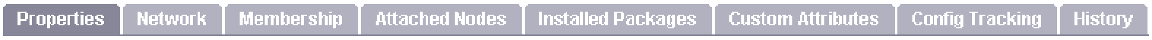

Each tab displays a page, each with its own buttons and functionality.

- Clicking a button to display a page

  For example, when you select the Custom Attributes tab, a page appears with several buttons: **Add**, **Delete**, and **Copy**, which are common to each page called by these tabs, and **Add Custom Attribute**, which is unique to this particular tab. You will find similar functionality on all tabbed pages in the SA System.

My Profile

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

Search

At the top of the Home page open the dialog box as Figure 2-7 shows.

Figure 2-7: Search Function

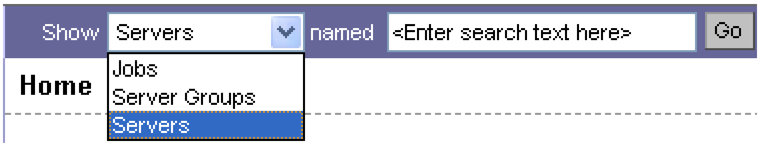

You can search for servers, jobs, and server groups, by making the selection from the drop-down list, and then entering an identifying string in the text box.

My Servers

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

Mouseover Icon Tooltips

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

Chapter 3: Getting Started with SA Client

IN THIS CHAPTER

This section discusses the following topics:

Overview of the SA Client

The SA Client is a powerful Java client for the HP Server Automation (SA) system. It provides the look-and-feel of a Microsoft Windows desktop application with the cross-platform flexibility of Java.

The SA Client provides the following features:

- Discovery and Agent Deployment
- Server Explorer
- Virtualization Director
- Audit and Remediation
- Service Automation Visualizer (SAV)
- Software Management
- Patch Management for Windows
- Patch Management for Unix

- Application Configuration Management
- Global Shell
- Integration with NA

For descriptions of all the SA Client features, see "SA Features" on page 39.

In order to visualize networking information with Network Automation (NA) inside of Service Automation Visualizer (SAV), you must have both a licensed version of NA integrated with your SA core, plus an additional license to run SAV showing NA data.

Additionally, in order to view storage devices and SAN information from the Storage Automation System (ASAS) inside of SAV, you must have both a licensed version of ASAS integrated with your SA core, plus an additional license to run SAV showing ASAS data.

SA Client and Server Automation Launcher

The Server Automation Launcher is a self-contained Java application that allows you to access the SA Client from any core in your mesh. You can use the Server Automation Launcher to log in to and download the latest version of the SA Client. If the SA Client has been upgraded on a specific core or on a core in a different mesh, you can choose which core you would like to use for downloading the SA Client.

The Server Automation Launcher also allows you to configure advanced settings, such as debug settings, locale settings, proxy server settings, and more.

This section contains the following topics:

- SA Client and Server Automation Launcher Requirements
- Installing the Server Automation Launcher
- Launching the SA Client
- Server Automation Launcher Advanced Options

SA Client and Server Automation Launcher Requirements

The SA Client is a Java application that installs and runs with its own Java Runtime Environment (JRE). The SA Client will not interfere with any other versions of JRE installed on your system. The JRE will not be used (and is not usable) by any other Java application on the target computer, and it will not set itself as the default JRE on the target computer.

The SA Client is supported on the following Microsoft operating systems:

- Windows Server 2003
- Windows 2000
- Windows XP
- Windows Vista

The minimum system requirements to run the SA Client are as follows:

- Minimum 1 GB of DRAM
- Minimum 100 MB of disk space each for SA Client and Server Automation Launcher
- If using the SA Client to connect to a core with a residential VoIP connection, a minimum 384 Kbps DSL connection is recommended.
- You must be logged in as a user with sufficient permissions to install software on the computer. (You do not need to be an administrator user to install the launcher.)

To run the SA Client, you must download and install the Server Automation Launcher (accessible from the SAS Web Client). In order to install the launcher, you must be a Windows user that is able to install applications on your system.

Installing the Server Automation Launcher

In order to run the SA Client, you need to download and install the SA Client Server Automation Launcher, which is a Java application that allows you to access the SA Client from any core in your mesh. When you install the SA Client launcher, it installs all of the necessary Java applications (Java Web Start and JRE) you need to run the SA Client.

If you plan to have multiple users install the Server Automation Launcher on the same computer, SA recommends that each user choose a unique path to install the application. For example, if one user has already installed the Server Automation Launcher at this location, `C:\SA`, then if another user logs in to the same computer and attempts to install to that same default location, they will get an error and not be able to install. If this occurs, choose a new location.

To install the Server Automation Launcher, perform the following steps:

1. Open a Web browser and enter the URL that points to the SAS Web Client.
2. On the SAS Web Client login page, click the Download Server Automation Launcher.
3. Download the Server Automation Launcher installation file and double-click to start the Server Automation Launcher installer
4. In the Welcome window, click **Next** to begin the Server Automation Launcher installation.
5. In the License Agreement window, select the "I accept the agreement" and then click **Next** to proceed with the installation.
6. In the Select Destination Directory window, accept the default installation directory, or click **Browse** to select a custom location. Click **Next**.
7. In the Select Start Menu Folder window, accept the default name and click **Next**.
8. In the Select Additional Tasks window, accept the default options or choose your own, and then click **Next** to install the Server Automation Launcher.
9. When the installation has completed, click **Finish** to exit.

Launching the SA Client

The launcher allows you to log in to an SA core using the SA Client. If your organization has installed and configured one or more multimaster meshes in your data center, you can choose to log into any SA core in your meshes. The SA Client gives you access to all devices in your data center as well as all the features you need to automate and manage your data center.

You also have the option of choosing which core you want to use to download the latest version of the SA Client, separate from the core you log in to. For example, when you log into a core and that core has a new version of the SA Client, the new version will automatically be downloaded when you log in. Using the launcher, you can choose one core to log in to and a separate core to download the latest client. This gives you the freedom to not have to download the SA Client every time you log into a different core.

The Server Automation Launcher allows you to log on to a SA 6.50 and above core. If you attempt to log into a pre-SA 6.50 core, you will get a 404 page not found Java Web Start error message and not be able to log on to the core.

If you are running the Server Automation Launcher on Windows 2000, you may see a missing DLL error message when you log on. This error will not affect the log on procedure. To fix this so the error message does not appear, install this Microsoft update: `http://support.microsoft.com/default.aspx?scid=kb;en-us;259403&Product=vc6`.

To launch the SA Client, perform the following steps:

1. Start the Server Automation Launcher from one of two locations:

   - On your desktop, double-click the HP Server Automation Client icon.

     Or

   - From the **Start** menu, select ➤ **All Programs** ➤ **HP Business Service Automation** ➤ **HP Server Automation**.

2. In the Log In to HP Server Automation window, enter your SA user name, password, and the SA core server you want to log in to, as shown in Figure 3-1.

Figure 3-1: Log In the SA Client Window

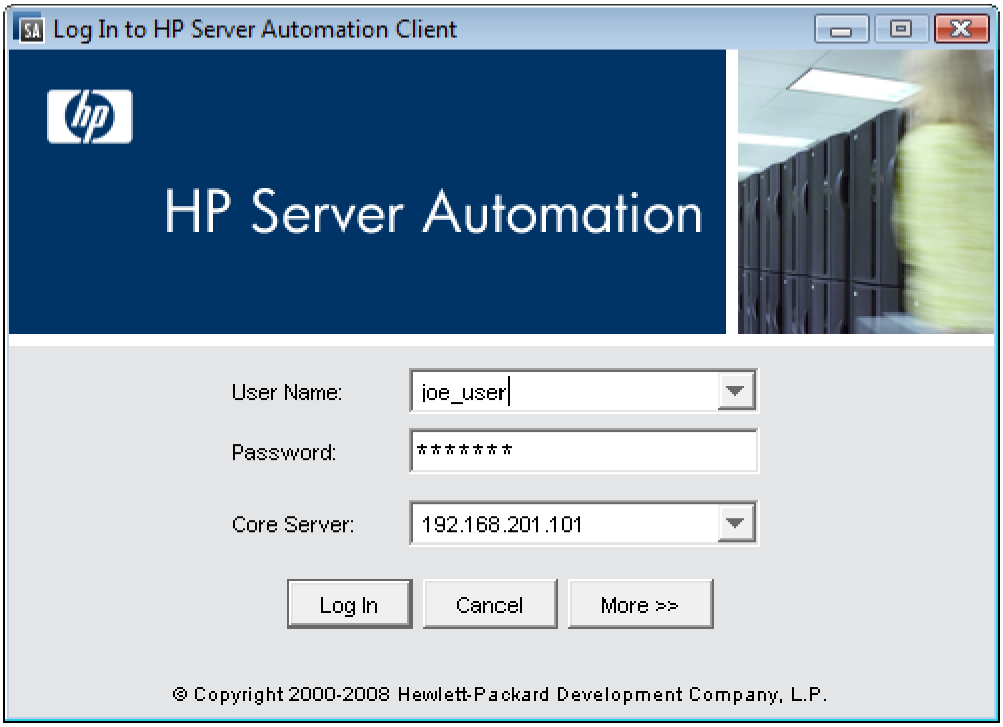

User Name: joe_user

Password: *******

Core Server: 192.168.201.101

© Copyright 2000-2008 Hewlett-Packard Development Company, L.P.

The user name is not case sensitive. If you have access to more than one core server in a mesh, you can enter the core server's IP address or name in the core server field. If you do not specify a port with the host:port notation, port 443 is used.

If this is the first time you are logging in to a specific core, the launcher will download the latest version of the SA Client when you log in. If you would like to differentiate between the core you log in to and the core from which you download the latest version of the SA Client, you can change those options by clicking **More** in the log in window and configuring your Client Host Server. For information on this and other advanced SA Client Server Automation Launcher options, see "Server Automation Launcher Advanced Options" on page 75.

3. Click **Log In**.

4. If you are asked to accept the certificate from the core server, click **Yes**. The SA Client appears.

Server Automation Launcher Advanced Options

You can configure the following advanced options for the SA Client:

- **Debug Settings**: Gives you control over the level of detail as well as the type of information included in SA Client log file.
- **Client Download Server**: Allows you to change the host server from which you want to download the SA Client.
- **Proxies**: Allows you to configure the SA Client proxy server settings
- **HP Server Automation Home**: Allows you to change the default location SA Client is downloaded and saved on your local computer, and to delete the SA Client's cache, and to change the location of SA Client log files.

To configure the Server Automation Launcher's advanced options, perform the following steps:

1. Start the Server Automation Launcher from one of two locations:

   - On your desktop, double-click the Server Automation Launcher icon

     Or

   - From the **Start** menu, select ➤ **All Programs** ➤ **HP Business Service Automation** ➤ **HP Server Automation**.

2. In the Log In to HP Server Automation window, you can set the following configuration:

   - **Username**: Enter your SA username.
   - **Password**: Enter your SA user password.
   - **Core Server**: Choose the SA core server to log in to.

3. Next, click **More**. In the expanded launcher window, you can now configure the following settings:

   - **Locale**: Choose a locale to match the localized version of the SA Client, either Japanese (ja), Korean (ko), for English (en). English is the default.
   - **Debug Settings**: Debugging options that are captured in the following log file: `C:\<user_home>\Application Data\Opsware\Deployment\log\*.log`

     - **Enable Debug Logging (Fine)**: Enables debugging and sends SA Client operations and errors to the log file.

   - **Enable Server Method Call Logging**: Adds server method calls to the log file.

- **Show Console**: Displays the Java Console window while the SA Client runs.

4 Click **Advanced Settings**.

5 In the Advanced Settings window, you can configure the following Server Automation Launcher options:

Client Download Server

You can configure the Server Automation Launcher so the default core you log in to is different from the core you use to access the latest version of the SA Client. This can be useful if you do not want to download a new version of the SA Client each time you log in to a different core running the same version of SA.

- **Use Core Server**: Select this option to use the default server specified in the Core Server field in the main Log In window to be the server from which you want to download the SA Client.
- **Use**: Enter a core server you want to use to download the SA Client.

Proxies

By default, the SA Client uses the proxy server settings configured for the default browser on your local system. For example, if your default browser has no proxy server settings configured, neither will the SA Client. You can change those proxy server settings here:

- **None**: Do not use a proxy server to connect to the SA Client.
- **Use Browser**: Use the proxy server settings specified in your default browser.
- **Manual**: Enter the proxy server hostname and port.
- **No Proxy Hosts**: If you want to add proxy server overrides, add them here, separated by commas. (This is only enabled when proxy server settings is set to Manual.)

HP Server Automation Home

- **Location**: The location where the SA Client is downloaded and saved on your local computer, along with all log files generated when the SA Client runs.

  Note that starting with SA 7.50, this location also controls where the SA Client data cache is stored.

The default home location is `<user_home>\Application Data\Opsware`, which is private to each user. If you choose to change this location, be aware that other users may have access to the new directory. You are responsible for setting the permissions on the new directory if you want to prevent unwanted access to your SA Client home.

- **Delete Application Cache**: Clicking this completely removes all downloaded copies of the SA Client. This ensures the launcher will download the latest SA Client from the core the next time the user logs in.
- **Delete Logs**: Delete all log files created by previous sessions of the SA Client. (All SA Client log files are located at: `<user_home>\Application Data\Opsware\Deployment\log\*.log`.)

6. When you are finished setting the options, click **OK** to save your settings.
7. Click **Log In** to log in to the SA Client.

SA Client User Interface

The SA Client user interface provides easy access to all of the SA Client features and functionality. The SA Client user interface has six main areas:

- Menus
- Navigation Pane
- Search
- Content Pane
- Details Pane
- Status Bar

Figure 3-2: SA Client User Interface

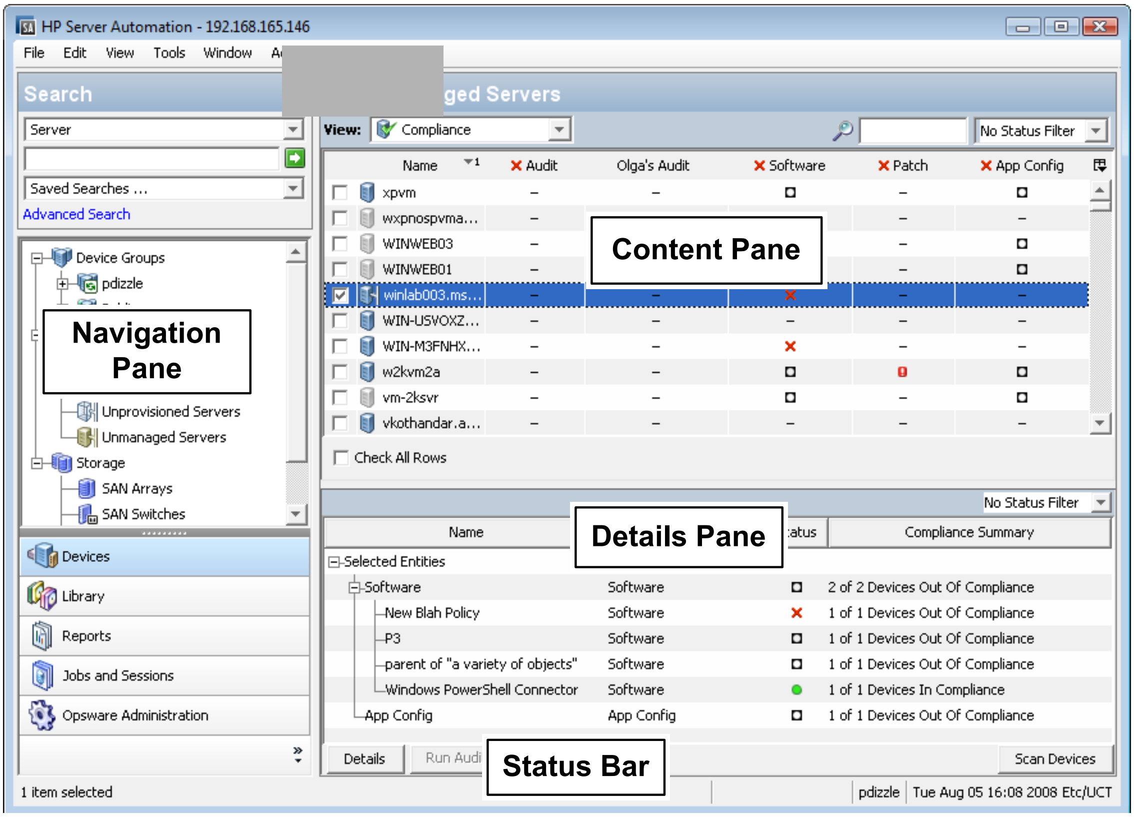

	Name	✗ Audit	Olga's Audit	✗ Software	✗ Patch	✗ App Config
□	xpvm	–	–	□	–	□
□	wxpnospvma...	–			–	–
□	WINWEB03	–			–	□
□	WINWEB01	–			–	□
☑	winlab003.ms...	–	–	✗	–	–
□	WIN-USVOXZ...	–	–	–	–	–
□	WIN-M3FNHX...	–	–	✗	–	–
□	w2kvm2a	–	–	□	!	□
□	vm-2ksvr	–	–	□	–	□
□	vkothandar.a...	–	–	–	–	–

Name		atus	Compliance Summary
⊟ Selected Entities			
⊟ Software	Software	□	2 of 2 Devices Out Of Compliance
New Blah Policy	Software	✗	1 of 1 Devices Out Of Compliance
P3	Software	□	1 of 1 Devices Out Of Compliance
parent of "a variety of objects"	Software	□	1 of 1 Devices Out Of Compliance
Windows PowerShell Connector	Software	●	1 of 1 Devices In Compliance
App Config	App Config	□	1 of 1 Devices Out Of Compliance

Menus

The SA Client includes the following menus:

- **File**: This enables you to open a new SA Client window, or close the current window, or exit all open SA Client windows.
- **Edit**: This enables you to cut, copy, paste, and delete text.
- **View**: This refreshes the current view and shows the latest information from the core that you are currently logged into (such as compliance test information for the compliance dashboard). You can also access SA Client features in the Navigation pane, such as Devices (groups of devices, managed and unmanaged servers), Reports (Compliance Dashboard, Reports) Software Library (application configuration, patch management), OS sequences and OS installation profiles, Jobs and Sessions (job logs and shell sessions), and Opsware Administration (patch settings and patch compliance rules). This also allows you to show or hide the Action pane and the Preview pane.
- **Tools**: This enables you to open a Global Shell session, open the Service Automation Visualizer, or access the SA Client options.
- **Actions**: Depending upon the feature that you have selected in the Navigation pane, this menu enables you to perform numerous functions related to all main SA Client features.
- **Window**: This enables you to access multiple instances of SA Client windows, if more than one window is open.
- **Help**: This menu provides help for the SA Client. Help F1 provides context-sensitive help relevant to the current feature window selected or opened (same as F1). The contents and index will open the SA Client help system to the main table of contents. (The About SA Client menu provides version and system information.)

Navigation Pane

To access SA Client features, select a feature in the Navigation pane, as shown in Figure 3-3. When you select a feature, its contents appear in the Content pane. You can access functions related to it through the Actions menu.

Figure 3-3: Navigation Pane

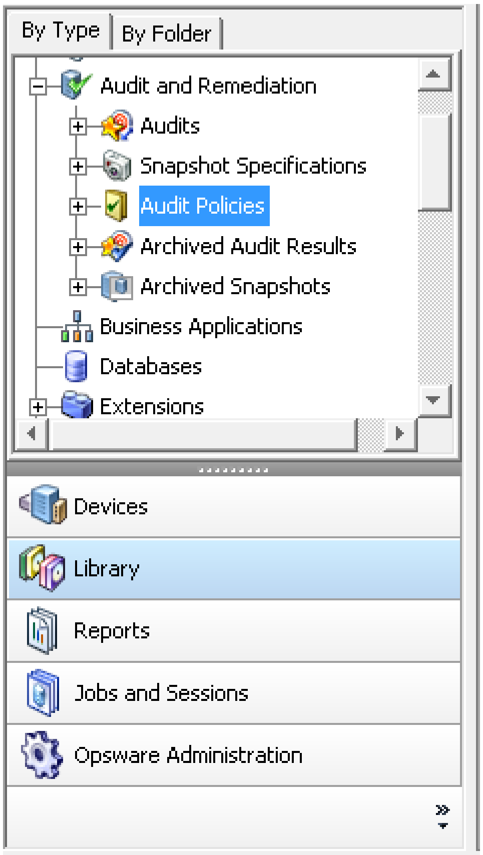

Search

The Search feature allows you to search for any information in HP Server Automation, such as SA Client Server, Device Group, Folder, Job, Software, Patch, Application Configuration, Software Policy, Patch Policy, Audit, and Snapshot Results. For more information on how to use the search tool, see “SA Client Search” on page 88.

Figure 3-4: Search in the SA Client

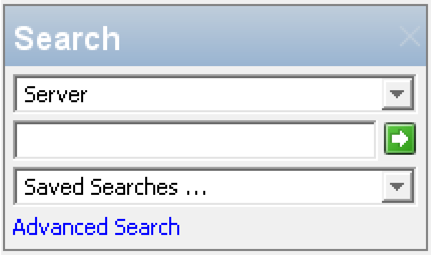

Content Pane

Depending on the selection in the Navigation pane (Devices, Library, Reports, Jobs and Sessions, Opsware Administration), the Content pane lists the following information:

- All managed servers and device groups, including unmanaged servers – both physical and virtual
- Agent deployment information
- Application Configurations and configuration templates
- Software Policies
- Audit and Remediation audits, audit policies, and snapshots
- Patches and patch policies
- OS installation profiles and OS sequences
- Packages
- Reports and the Compliance Dashboard
- Custom attributes
- Jobs that the user has run
- Access to the Global Shell sessions
- Patch configuration and patch compliance rules

Figure 3-5 is an example of the Content pane for managed servers. You can perform actions on features in the Content area using the Action Menu or Action Pane, or you can right-click to perform various actions or double-click to open.

Figure 3-5: Content Pane Showing Server History

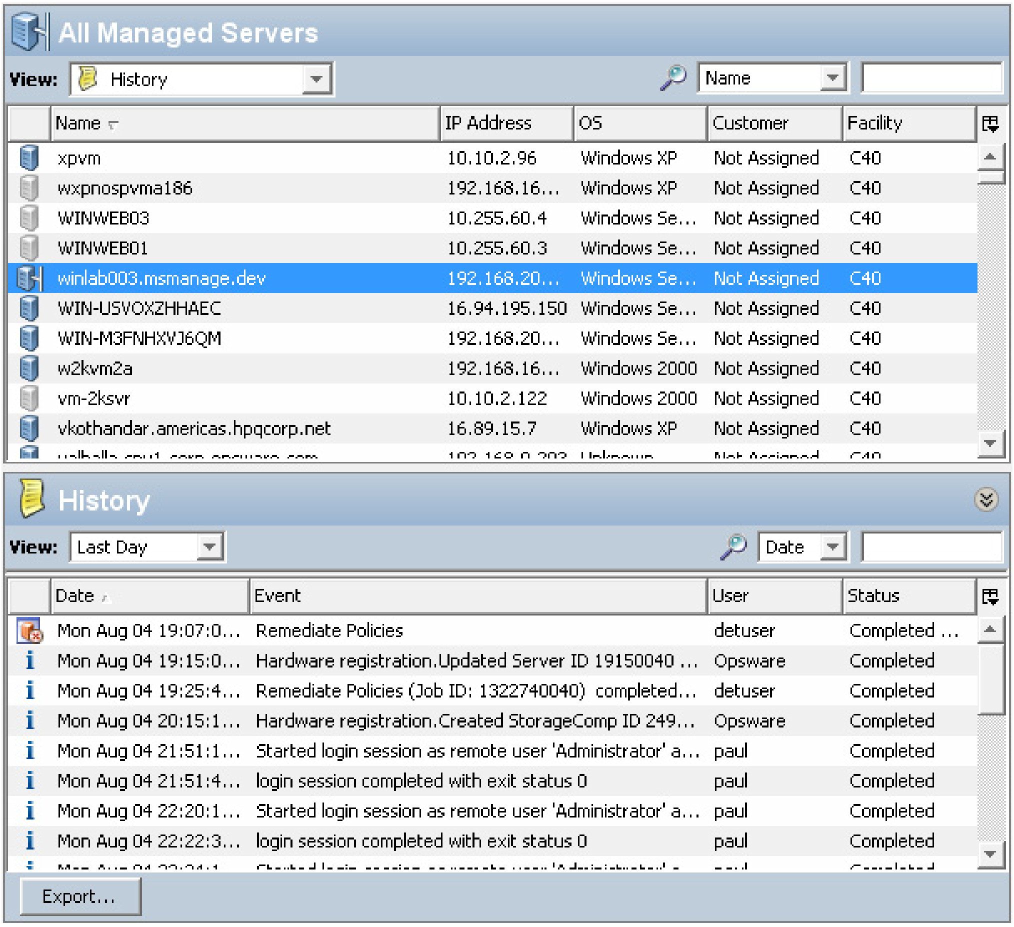

Name	IP Address	OS	Customer	Facility
xpvm	10.10.2.96	Windows XP	Not Assigned	C40
wxpnospvma186	192.168.16...	Windows XP	Not Assigned	C40
WINWEB03	10.255.60.4	Windows Se...	Not Assigned	C40
WINWEB01	10.255.60.3	Windows Se...	Not Assigned	C40
winlab003.msmanage.dev	192.168.20...	Windows Se...	Not Assigned	C40
WIN-USVOXZHHAEC	16.94.195.150	Windows Se...	Not Assigned	C40
WIN-M3FNHXVJ6QM	192.168.20...	Windows Se...	Not Assigned	C40
w2kvm2a	192.168.16...	Windows 2000	Not Assigned	C40
vm-2ksvr	10.10.2.122	Windows 2000	Not Assigned	C40
vkothandar.americas.hpqcorp.net	16.89.15.7	Windows XP	Not Assigned	C40
valhalla.cpu1.corp.opsware.com	192.168.0.203	Unknown	Not Assigned	C40

Date	Event	User	Status
Mon Aug 04 19:07:0...	Remediate Policies	detuser	Completed ...
Mon Aug 04 19:15:0...	Hardware registration.Updated Server ID 19150040 ...	Opsware	Completed
Mon Aug 04 19:25:4...	Remediate Policies (Job ID: 1322740040) completed...	detuser	Completed
Mon Aug 04 20:15:1...	Hardware registration.Created StorageComp ID 249...	Opsware	Completed
Mon Aug 04 21:51:1...	Started login session as remote user 'Administrator' a...	paul	Completed
Mon Aug 04 21:51:4...	login session completed with exit status 0	paul	Completed
Mon Aug 04 22:20:1...	Started login session as remote user 'Administrator' a...	paul	Completed
Mon Aug 04 22:22:3...	login session completed with exit status 0	paul	Completed
Mon Aug 04 22:24:1...	Started login session as remote user 'Administrator' a...	paul	Completed

Content Pane Tools

From inside the Content pane, you can perform the following actions:

- With the View drop-down list, you can change the view of a selected feature. For example, you can select a server from the Content pane, and then from the View drop-down list, choose Software Policies. This shows all software policies attached to the server, as shown in Figure 3-6.

Figure 3-6: View Drop-down List

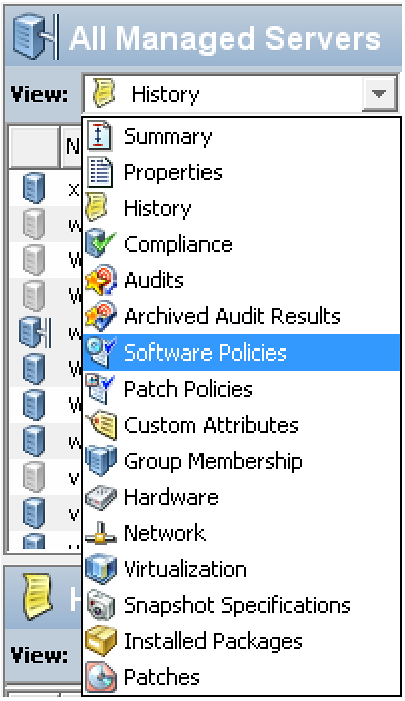

- In the column headings of the Content pane, you can sort data about a feature. For example, for a managed server, you can sort by Summary, Properties, IP address, OS, and so on.

- With the search tool, you can search the Content pane by feature or by attribute, as shown in Figure 3-7.

Figure 3-7: Search Tool

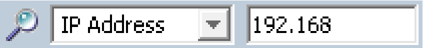

Details Pane

The Details pane allows you to preview information about servers, device groups, patches, and patch policies selected in the Content pane without having to open a new window.

You can use the Details pane to perform the following actions:

- Preview information about a server, device group, patch, or patch policy. To do so, select it in the Content pane.
- Filter the type of information you view in the Details pane. From the top of the content area, choose a view from the View drop-down list.
- Deactivate the Details pane. To do so, from the **View** menu, select **Details Pane ➤ Minimize**.

For example, if you are viewing Windows 2003 patches from the Library, you can select a patch in the Content pane and see information about the patch in the Details pane. This is shown in Figure 3-8.

Figure 3-8: Main Application Windows Showing Patch Properties Information in the Preview Pane

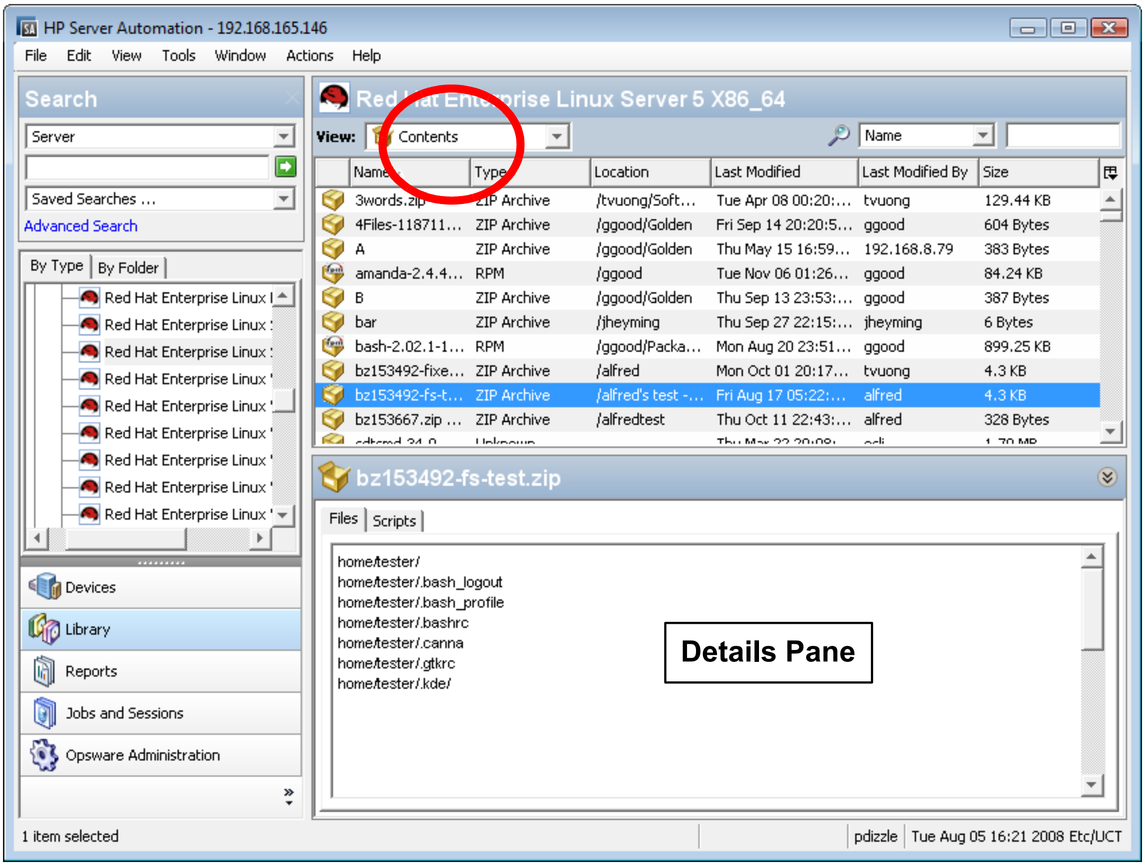

Name	Type	Location	Last Modified	Last Modified By	Size
3words.zip	ZIP Archive	/tvuong/Soft...	Tue Apr 08 00:20:...	tvuong	129.44 KB
4Files-118711...	ZIP Archive	/ggood/Golden	Fri Sep 14 20:20:5...	ggood	604 Bytes
A	ZIP Archive	/ggood/Golden	Thu May 15 16:59...	192.168.8.79	383 Bytes
amanda-2.4.4...	RPM	/ggood	Tue Nov 06 01:26...	ggood	84.24 KB
B	ZIP Archive	/ggood/Golden	Thu Sep 13 23:53:...	ggood	387 Bytes
bar	ZIP Archive	/jheyming	Thu Sep 27 22:15:...	jheyming	6 Bytes
bash-2.02.1-1...	RPM	/ggood/Packa...	Mon Aug 20 23:51...	ggood	899.25 KB
bz153492-fixe...	ZIP Archive	/alfred	Mon Oct 01 20:17...	tvuong	4.3 KB
bz153492-fs-t...	ZIP Archive	/alfred's test -...	Fri Aug 17 05:22:...	alfred	4.3 KB
bz153667.zip ...	ZIP Archive	/alfredtest	Thu Oct 11 22:43:...	alfred	328 Bytes

```
home/tester/
home/tester/.bash_logout
home/tester/.bash_profile
home/tester/.bashrc
home/tester/.canna
home/tester/.gtkrc
home/tester/.kde/
```

To view other types of information about the selected patch, from the View drop-down list, choose a view.

Details Pane Show Filter

Some features displayed in the Details pane allow you to further filter the feature. Using the Show drop-down list, you can choose different views of the feature.

For example, if you are viewing all of the servers that the patch policy is attached to, in the Details pane, you can filter either Servers with Policies Attached or Servers with Policies Not Attached, as shown in Figure 3-9.

Figure 3-9: Details Pane Show Drop-down List

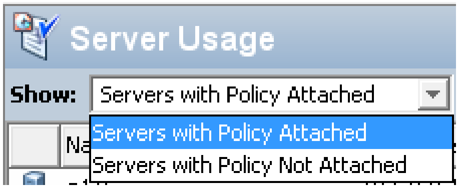

Server Usage
Show:
Servers with Policy Attached
Na
Servers with Policy Attached
Servers with Policy Not Attached

Status Bar

At the bottom of the SA Client Application window, the status bar provides the following information:

- Informational text (left hand side) about the selected object
- A progress bar that shows progress on retrieving information from the core
- The user ID
- The current HP Server Automation time

Figure 3-10: SA Client Status Bar

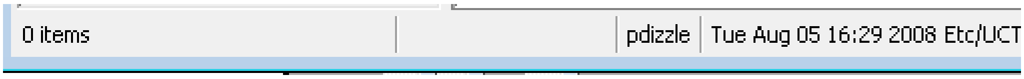

Sharing SA Client Objects with Drag and Drop

You can easily drag and drop servers and other SA Client objects out of the SA Client and into an email, a chat window, or a web browser. When you drag an object out of the SA Client, a URL is constructed that enables you to launch the object in the SA Client.

You can share such SA Client objects as servers, application configurations, audits, a Business Application, Patch Policies, OS Profiles, and more – basically, any device (server, storage, network, and so on) or any object from in the SA Library that is searchable.

You can also use the URL that is created during drag and drop as a link on a web page, which gives you easy access to those SA objects you are most interested in.

In order to drag and drop networking devices from the SA Client, you must have a licensed version of Network Automation (NA) integrated with your SA core. Additionally, in order to drag and drop storage devices from the SA Client, you must have a licensed version of ASAS integrated with your SA core,

If you have not purchased NA or ASAS but would like to, contact your SA sales representative.

To drag and drop a SA Client object, perform the following steps:

1. From inside the SA Client, select a server, network or storage device, device group, Patch Policy, a Business Application, an audit, or any other object.
2. Drag the selected object to one of the following locations:
   - Web browser
   - Chat window text entry box
   - Email
3. When you or another user clicks the link, you are asked to save the file. Once the file is downloaded and saved, click **Open**.
4. In the Log In to HP Server Automationwindow, enter your SA user name, password, and the SA core server you want to log in to and click **OK**.

   The user name is not case sensitive. If you have access to more than one core server in a mesh, you can enter the core server's IP address or name in the core server field. If you do not specify a port with the host:port notation, port 443 is used.

   If the SA Client is already opened, and the object you are opening belongs to the core you are logged into, the object appears in its own window without requiring you to log in.

SA Client Search

In the SA Client, you can search for any information about your operational environment that is available in SA using the SA Client Search feature. The Search feature enables you to search for Application Configuration, Application Configuration Template, Audit, Device Group, Folder, Job, Patch, Patch Policy, Server, Snapshot Result, Software, and Software Policy.

You can also search for storage Database, SAN Switch, Server, and Storage Systems objects if you have purchased ASAS. For more information, see the *ASAS User’s Guide*.

HP Application Storage Automation System (ASAS) is a separately licensed product that requires SA. To visualize or search for any storage data in ASAS, you must license a version of ASAS. If you have not purchased ASAS and would like to, contact your sales representative.

The SA Client Search feature enables you to perform the following actions:

- Perform a simple search by using keywords.
- Perform an advanced search by creating search queries.
- Save a search query.
- Perform actions on search results.
- Email search results.
- Print search results.
- Customize search results formatting.

You cannot search for Custom Fields on servers using the SA Client Search feature. To search for Custom Fields on servers, use Server Search in the SAS Web Client. See “Server Search” in *SA User’s Guide: Server Automation* for more information.

Performing a Simple Search

In the SA Client, a simple search enables you to search for items by entering a keyword in the search text field. When you enter a keyword or text, simple search carries out a “contains” operator search on the entered text.

The search operation returns only items on which you have at least read permissions. To perform an action on an item, you must have write permission on that item.

To carry out a simple search, perform the following steps:

1. From the Navigation pane, select Search.
2. To display Search in the Navigation pane, select **Search Pane** from the **View** menu.
3. From the drop-down list, select the item you want to search as shown in Figure 3-11.

Figure 3-11: Simple Search in the SA Client

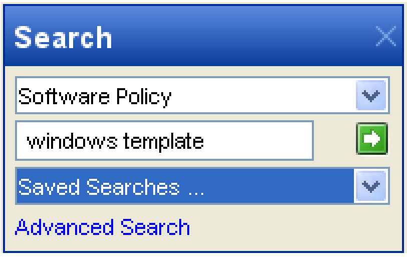

Search

Software Policy

windows template

Saved Searches ...

Advanced Search

4. Enter the search text in the text field. The SA Client search feature does not support a wildcard search and the search is not case sensitive.
5. Click ➡ to search. The search results appear in the Content pane.
6. (Optional) Click on any column heading to sort the search results. You can also change the order of the columns by dragging the column heading and dropping it into the desired location.
7. (Optional) Click **Save** to save your search query. The Save Search window appears. Enter the name of the search and click **Save**. The name of the saved search cannot be more than 64 characters. See “Creating a Device Group Using Search” on page 189 in Chapter 5 for information about saving a search query as a device group.
8. (Optional) Click **Export** to export search results. The Export Results window appears. Enter the location, the file name and the file type and then click **Export Results**.
9. (Optional) After you export the search results, you can email the search results by attaching the exported file to the email. You can also print the search results by printing the exported file from an application that supports .csv or .html files.

10 (Optional) To perform an action on the search results, select an item from the Content pane and then from the **Actions** menu, select the appropriate action.

Performing an Advanced Search

In the SA Client, an Advanced search enables you to create complex queries on search items. In a search query, you can specify up to five rules and combine each rule with a logical "And" or "Or" operator. You cannot use both the "And" and "Or" operator in a single search query. Each rule is a combination of an attribute/operator/value that enables you to search for a specific attribute value for the selected search item. Depending on the attribute that you select, the options for the operator and value are displayed. You can specify the attribute values by entering text or a numerical value in the text field, by selecting a value from the drop-down list, or by selecting multiple values from a list of values in the Select Values window.

The Select Values window appears when you need to specify multiple values for a rule containing an "equals" or "not equals" operator. In this window, you can select multiple values from the list or add values to the Available field.

The search operation returns only items on which you have at least read permissions. To perform an action on an item, you must have write permission on that item.

To carry out an advanced search, perform the following steps:

1 From the Navigation pane, select Search. To display Search in the Navigation pane, select the **Search Pane** from the **View** menu.

2 Click **Advanced Search**. The Advanced Search page appears in the Content pane. By default, the search item Server is selected in the first drop-down list and one search rule is added to the search.

3 From the first drop-down list, select the item you want to search.

Figure 3-12: SA Client Search

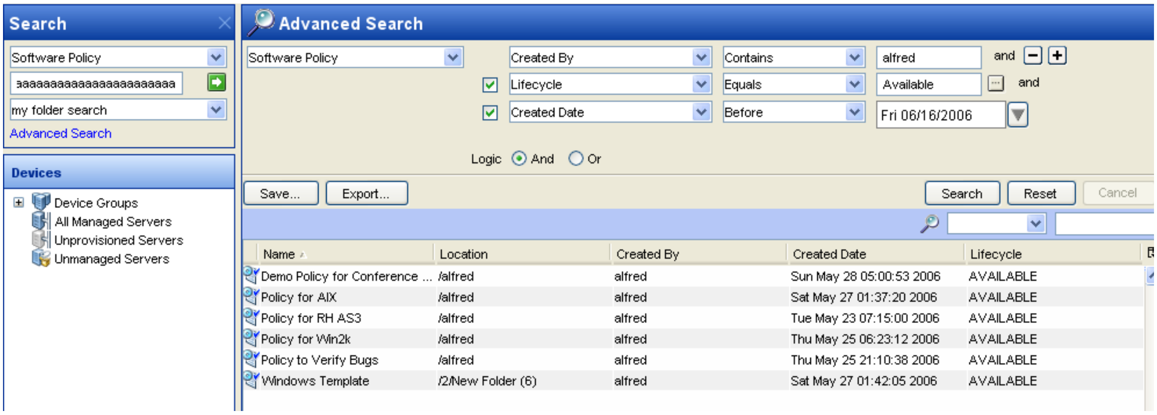

Software Policy

	Created By	Contains	alfred	and
☑	Lifecycle	Equals	Available	and
☑	Created Date	Before	Fri 06/16/2006	

Logic ● And ○ Or

Name	Location	Created By	Created Date	Lifecycle
Demo Policy for Conference ...	/alfred	alfred	Sun May 28 05:00:53 2006	AVAILABLE
Policy for AIX	/alfred	alfred	Sat May 27 01:37:20 2006	AVAILABLE
Policy for RH AS3	/alfred	alfred	Tue May 23 07:15:00 2006	AVAILABLE
Policy for Win2k	/alfred	alfred	Thu May 25 06:23:12 2006	AVAILABLE
Policy to Verify Bugs	/alfred	alfred	Thu May 25 21:10:38 2006	AVAILABLE
Windows Template	/2/New Folder (6)	alfred	Sat May 27 01:42:05 2006	AVAILABLE

4 Create a rule by selecting the attribute from the second drop-down list. Depending on the attribute that you select, options available for the operator and values for the rule will change.

5 Select the operator from the third drop-down list. The operator selected defines how the search text is treated.

6 Enter a value in the field or select a value from the drop-down list or click ⋯ to select multiple values from the **Select Values** window.

7 Click ⊞ to add additional rules and repeat steps 4 to 6. Click ⊟ to delete any rules.

8 Select the logic (And/Or) to be applied for every rule in the query.

9 Click **Search** to run the search query. The search results appear in the Content pane.

10 (Optional) Click **Reset** to clear the search query rules or click **Cancel** to cancel the search operation.

11 (Optional) Click on any column heading to sort the search results. You can also change the order of the columns by dragging the column heading and dropping it into the desired location.

12 (Optional) Click **Save** to save your search query. The Save Search window appears. Enter the name of the search query and click **Save**. The name of the saved search cannot be more than 64 characters. The saved search query appears in the Saved Searches. See "Creating a Device Group Using Search" on page 189 in Chapter 5 for information about saving a search query as a device group.

13 (Optional) Click **Export** to export search results. The Export Results window appears. Enter the location, file name and file type (.csv or .html) and then click **Export Results**.

14 (Optional) After you export the search results, email the search results by attaching the exported file to the email. Or, print the search results by printing the exported file from an application that supports .csv or .html.

15 (Optional) To perform an action on the search results, select an item from the Content pane and then from the **Actions** menu, select the appropriate action.

Running a Saved Search Query

To run a saved search query, perform the following steps:

1 From the Navigation pane, select Search.

2 From the Saved Searches drop-down list, select a search query. The query appears in the Content pane.

3 Click **Search** to run the query. The search results will appear in the Content pane.

Deleting a Saved Search

To delete a saved search query, perform the following steps:

1 From the Navigation pane, select Search.

2 From the Saved Searches drop-down list, select a saved search query. The saved search appears in the Content pane.

3 Click **Save**. The Saved Search window appears.

4 Select the saved search and click **Delete**. Click **Delete** on the Confirmation window to delete a saved search.

Accessing SA Client Options

You can configure the following options for the SA Client:

- **General Options**: This enables you to set options such as choosing the core you want to log into by default, how to handle caching, and so on.
- **Unmanaged Servers**: This enables you to set options for the Agent Discovery and Deployment feature.
- **Terminal and Shell**: This enables you to configure your Terminal (UNIX) and RDP (Windows) client for the Global Shell and Remote Terminal connections.
- **Patch Policies**: This enables you to specify that a confirmation message will display when you try to remove a patch policy or a patch policy exception from a managed server.
- **HP Network Automation**: This enables you to reset the name of the NA host that you log into, restore the previously saved (default) host name, and launch the NA login window.
- **HP Service Automation Visualizer**: This enable you to specify timeout values for launching SAV and the manner in which you want SAV to scan virtual server relationships.

Accessing SA Client Options

To set SA Client options, perform the following steps:

1. From the **Tools** menu, select **Options**.
2. From the left side of the Set Options window, choose an option.

General Options

The following general options enable you to select your default core:

Core Server Defaults

This option allows you to configure the core that you log into from the SA Client. Options include:

- **Host**: This enables you to choose the name or IP address of the SAS Web Client host that you log into by default.

- **Port**: This enables you to choose the port number of the SAS Web Client host that you log into by default; 443 is the default.

Cache

This option enables you to configure the caching of data displayed inside the SA Client. You can configure the following cache settings:

- **Check for updates every X minute(s)**: This enables you to enter a value for how many minutes will pass before the cache is refreshed.
- **Update Cache**: This enables you to check instantly for new information from the core.
- **Reload the Cache**: This enables you to immediately reload (refresh) the cache.

Progress Information

This option shows the progress of a job. When a job finishes, the Progress window closes.

Equals Operator Limit in Search and Reports

This sets limits on the number of available value selections in the Advanced Search and Reports interfaces. To prevent delays and excessive system load, the list of available values is not populated when the number of values exceeds this setting. Values are added by entering them in a text box.

Unmanaged Servers

These options allow you to control the operation of the ODAD. You can set the following ODAD options:

- Installer Options
- Protocols
- Advanced

For more information, see "Agent Management" on page 105.

Installer Options

From the Installer Options window, you can set the Installer options and control the installation of an Server Agent on a server. The Installer Options window enables you to perform the following actions:

- **Start the Server Agent after Installation**: This enables you to start the Server Agent after installing it on the server. By default, the Server Agent Installer does not start the Server Agent.

- **Ignore prerequisite check failures**: This enables you to ignore prerequisite check failures and forces Server Agent installation.
- **Set the server's time from the Opsware core**: This enables you to synchronize the time on the server in which the Server Agent is installed with the SA core.
- **Install Windows Installer (MSI) if required**: This enables you to install MSI along with the Server Agent. If MSI is already installed, this option has no effect.
- **Install Windows Management Instrumentation (WMI) if required**: This enables you to install WMI along with the Server Agent. If the WMI is already installed, this option has no effect.
- **Reboot Windows servers after agent installation**: This enables you to reboot Windows servers after the Server Agent installation is complete.
- **Install Red Hat Package Manager (RPM) on AIX and Solaris**: This enables you to install the RPM handler with the Server Agent. SA recommends that you always include this option when you install Server Agents on Solaris and AIX servers.
- **Reset Server Agent configuration, if present**: This enables you to replace the existing Server Agent configuration.
- **Delete gateway address list, if present**: This enables you to delete the SA Gateway address list, if present and no longer required.
- **Overwrite staged Server Agent installer**: This enables you to overwrite the existing Server Agent Installer.
- **Log Level**: This enables you to set the log level for log messages. With this option, you can specify levels for error, warning, info, and trace.

Protocols

The Protocol window allows you to specify the standard port to connect to the servers for deployment. The following protocols are used:

- **SSH**: This enables you to determine the standard port to connect to the servers for deployment using the SSH protocol.
- **Rlogin**: This enables you to determine the standard port to connect to the servers for deployment using the Rlogin protocol.
- **Telnet**: This enables you to determine the standard port to connect to the servers for deployment using the Telnet protocol.

- **NetBIOS**: This enables you to determine the standard port to connect to the servers for deployment using the NetBios protocol.
- **WTS**: This enables you to determine the standard port to connect to the servers for deployment using the WTS protocol.

Advanced

The Advanced Installer Options window allows you to set the following options:

- **Immediately do a full hardware registration**: This enables you to force the Server Agent Installer to report full hardware information to the core.
- **Immediately do software registration**: This enables you to force the Server Agent Installer to report full software information to the core.
- **Suppress Server Agent reachability check**: This enables you to disable this check during installation. By default, the installer triggers the core to check if the server is reachable.
- **Disallow anonymous SSL connections if Server Agent is dormant**: This enables you to configure the Server Agent so that browsers cannot connect without a valid certificate.
- **Force creation of new device record if conflict found**: This enables you to suppress this functionality. During registration, the Data Access Engine creates a new device record.
- **Fail if initial hardware registration fails (do not go dormant)**: This enables you to ensure that the Server Agent does not become dormant, if it fails to report hardware information.
- **Do not open Windows Firewall for core-agent communications:** By default, the Server Agent Installer will modify the Windows Firewall configuration on Windows XP and Windows 2003 (r2) servers to allow the SA core to contact the managed server on port 1002. If you select this option, the firewall configuration will not be modified. The server may not be manageable by SA in this case.
- **Remediate software policies**: This enables you to remediate the server against any software policies attached to the server.
- **Attach to software Policy ID**: This enables you to attach the server to the software policy ID.

- **Reconcile Type (deprecated)**: (Applicable only for upgrades) This enables you to reconcile the server against any nodes assigned to the server. The reconcile type can be Full or Add only.
- **Attach to template ID (deprecated)**: (Applicable only for upgrades) This enables you to assign the nodes contained in the template to the server.
- **Install directory (UNIX)**: This enables you to specify the directory where the Server Agent will be installed. By default, the Server Agent is installed in `/opt/opsware/agent` on UNIX, and `%ProgramFiles%\Opsware\agent` on Windows.

  Extra Installer options: This enables you to specify any other installer options. For example:

  `--log file <path>` allows you to specify the path to the installer log file. By default, the installer log files are placedin the `/tmp` on Unix or `%SYSTEMDRIVE%\WINDOWS\SYSTEM` on Windows.

  `--workdir <path>` allows you to specify the path to the working directory to use while the installation is in progress.
- **nmap parameters**: This option allows you to specify parameters used when scanning for unmanaged servers. If you find that the Discovery and Agent Deployment is unable to correctly locate and identify unmanaged servers due to the network firewall configuration, you can specify a different set of scan parameters. See the nmap documentation for more information.

Terminal and Shell

These settings define the command that the SA Client invokes on your PC to open a Global Shell or Remote Terminal session. (For instructions on using an `ssh` client instead of the SA Client, see “Opening a Global Shell Session” on page 388.)

Terminal Client

This setting specifies the terminal client that the SA Client uses for Remote Terminal sessions on Unix managed servers and for Global Shell sessions. The default value is:

```
cmd /c start /w cmd /c "telnet %h %p && echo %m && pause > nul"
```

The `telnet` program emulates a command-line terminal session. The `%h` represents the host and the `%p` is for the port. (See Table 3-1.)

If you change the Terminal Client setting from the default value, make sure that the command blocks until the terminal application terminates. The terminal application must not run in the background. If you specify `cmd /c start`, include the `/w` switch to make `cmd` block until the underlying command (such as `telnet`) completes.

You are not required to use telnet as the terminal application. For example, to use a PuTTY client, specify the following command:

```
"C:\\Program Files\\putty\\putty.exe" -telnet %h %p
```

RDP Client

This setting specifies the Remote Desktop Protocol (RDP) client that the SA Client uses for Remote Terminal sessions on Windows managed servers. The default value is the Microsoft Terminal Services Client:

```
mstsc "%r"
```

The SA Client supports the Windows XP version of the Remote Desktop Connection Software, which can be downloaded from the following URL:
http://www.microsoft.com/windowsxp/downloads/tools/rdclientdl.mspx

The specified terminal client must be installed on your PC. To verify the existence of the terminal client, click **Test**.

The command can include variables such as `%h` and `%p`. When the terminal client is launched, these variables are replaced with the values shown in Table 3-1. To override a replacement value, specify a constant instead of a variable. For example, you might specify 435 for the port instead of `%p`.

Table 3-1: Variables for the Terminal and RDP Client Options

VARIABLE	DESCRIPTION	REPLACEMENT VALUE
`%e`	The character encoding.	For Remote Terminal sessions, the encoding of the managed server. For Global Shell sessions, the value of the Encoding field.
`%h`	The host name that the `telnet` client connects to.	The value of the `localhost` of the managed server.

Table 3-1: Variables for the Terminal and RDP Client Options (continued)

VARIABLE	DESCRIPTION	REPLACEMENT VALUE
`%m`	A locale-specific message on how to close the window.	For English locales, click the Enter key to close this window.
`%p`	The port that the telnet client connects to.	A randomly chosen port.
`%r`	The name of the Remote Desktop (RDP) connection file. This variable is used only for the Microsoft Terminal Services Client (`mstsc`).	A temporary RDP file generated at runtime by the SA Client.
`%t`	The title displayed in the terminal window.	For Remote Terminal sessions, the name of the managed server. For Global Shell sessions, the string “Global Shell.”

Encoding

This setting sets the encoding for Global Shell and the Remote Terminal sessions. This option is the replacement value of the `%e` variable in the command specified by the Terminal Client field. The default value of the Encoding option is UTF-8.

Patch Policies

This option allows you to specify that a confirmation message will display when you try to remove a patch policy or a patch policy exception from a managed server.

HP Network Automation

This options allows you to reset the name of the NA host that you log into, restore the previously saved (default) NA host name, and test whether SA can communicate with NA by using the new host name.

- **Host**: This option specifies the name of a server that is acting as a proxy for the NA host. Only the format of the host name is verified.
- **Restore Default**: This option restores the previously saved NA host name.
- **Test**: This option opens the NA login window to verify whether the host name is valid.

HP Service Automation Visualizer

For SAV, you can specify the following options:

- Virtualization Settings
- Scan Time-Out Preference
- Discovery Settings
- Reset All Settings

Virtualization Settings

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

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

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

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

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

To change the virtualization settings, perform the following steps:

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

Scan Time-Out Preference

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

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

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

Discovery Settings

If servers are scanned and it is determined that they are dependent on external IP addresses, when this option is selected SAV attempts to determine which servers or network devices those IP addresses refer to.

Keep in mind that this could cause scan time to increase, depending on the numbers of servers you selected for the scan and how many remote dependencies are discovered.

For recurring background business application snapshots, this detection is always done and cannot be turned off.

Reset All Settings

Restores all SAV settings to their defaults.

Browsing Job Logs

A job is any major process run by the SAS Web Client or the SA Client, such as a communication test, a software installation, an Audit Servers, Create Package, Create Virtual Zone, Install Software, Push Configurations, Run Communication Test, Run OS Sequence, Scan Configuration Compliance, Uninstall Patch and so on.

The Job Logs window shows the details of all jobs run under the currently logged in user name and jobs that are currently in progress. It also displays jobs scheduled to run, including the name of the job, the start time, the number of servers affected by the job, and the status of the job.

If the job status is Pending Approval, then the job is blocked until it is approved by a process that is external to HP Server Automation. If jobs are blocked indefinitely, the SA Administrator should check the settings of the Approval Integration window or the configuration of the backend connector.

The display of a job's start time and finish time is determined by original user preferences set in the SAS Web Client, which may be different than those of the current user. If a user is working in multiple times zones, it is a good idea to make sure that these preferences are set to display the time zone in the date.

To see the details of a finished job, open a job. If the job is recurring (scheduled to be run), then opening the job will cause the Schedule Job window to display. You will only be able to modify the scheduled job if you created the job, or have Edit All Jobs permissions. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide* for more information.

In order to view a job in the SA Client, you must have permissions to run or execute the feature action. For example, if you wanted to view Application Configuration Push jobs in the SA Client, and you had the Manage Application Configurations permission set to Read, but not Write, you would not be able to see any Application Configuration Push jobs in the SA Client.

At the top of the Job Logs window, you can search by the following fields:

- **Job ID**: This enables you to enter the job ID.
- **Job Type**: This enables you to choose a type of job from this list.
- **Time Restrictions**: This enables you to limit the search for a job by a time restriction such as the last week, last two weeks, or last month.
- **Ticket ID**: This enables you to enter a ticket ID here if one was given for the job.
- **Job Status**: This enables you to choose a job status, such as job completed, completed with errors, cancelled, and so on.

- **User name**: This enables you to enter your user name to see only those jobs you have run, or enter nothing here to show all jobs. Click on the column header to sort the list. Only users with the View All Jobs or Edit All Jobs permissions will be able to view all jobs in the core. If you do not have these permissions, you will not be able to see all jobs.

To view a job, select it, right-click, and select **Open**.

To cancel a non-recurring job, from the Job Logs window, right-click the job and select **End Job**. You cannot cancel a job while it is running (that is, the job status is In Progress).

To cancel a recurring (scheduled) job, from the Recurring Schedules window, right-click the job and select **Delete Schedule**.

Recurring Schedules

The recurring schedules window shows all jobs that are scheduled to run on a recurring basis. You can choose to view all recurring jobs, or filter the list of recurring jobs by the following methods:

- **Any ID**: This enables you to enter the job ID.
- **Any Type**: This enables you to choose a type of job from this list.
- **Any Ticket**: This enables you to enter a ticket ID here if one was given for the job.
- **Any User**: This enables you to enter your user name to see only those jobs you have scheduled to run on a recurring basis, or enter nothing here to show all recurring jobs. Click on the column header to sort the list. Only users with View All Jobs or Edit All Jobs permissions will be able to view all jobs in the core. If you do not have these permissions, you will not be able to see all jobs.

To view a recurring job schedule, select it, right-click, and select **Open**.

To delete a recurring job, select it, right-click, and select **Delete Schedule**.

Chapter 4: Agent Management

IN THIS CHAPTER

This section discusses the following topics:

- Agent on Managed Servers
- Discovery and Agent Deployment
- Agent Installation Using ODAD
- Viewing Unmanaged Server Information
- Agent Reachability Communication Tests

Agent on Managed Servers

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

- Overview of the Server Agent on Managed Servers
- Security for Agents Running on Managed Servers
- Agent Functionality on Managed Servers
- Server Information that the Agent Tracks

Overview of the Server Agent on Managed Servers

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

Every twelve hours, the Server Agent submits a minimal hardware information, such as changes in IP address, node name and so on, for the managed server on which it is running. Every 7 days, the Server Agent submits a full hardware inventory for the managed server on which it is running. Hardware registration also occurs during Server

Agent installation or software installation. Every 24 hours, the Agent submits software inventory and information for the managed server to the core. It reports the software found on the server and the core determines any differences. See Figure 4-1.

Figure 4-1: Server Management Tasks Performed by Agent

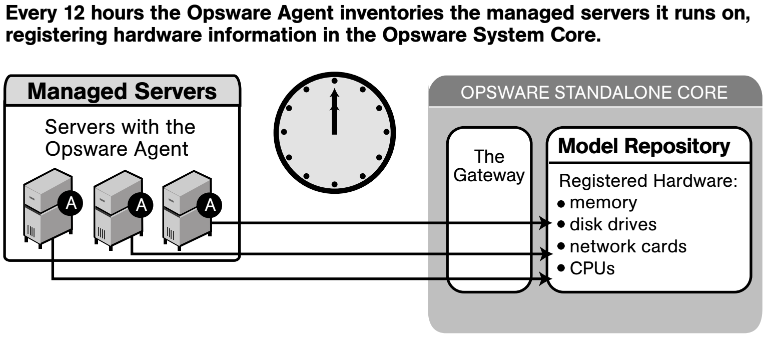

Every 12 hours the Opsware Agent inventories the managed servers it runs on, registering hardware information in the Opsware System Core.

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

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

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

If the Agent experiences an error in reporting, or has not reported within 24 hours, you can run a Communication Test to troubleshoot the problem.

If you modify the server hardware, it could take up to for the change to appear in the SAS Web Client user interface, depending on the time that the Agent for that server contacted the core.

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

In some cases, not all of a server's hardware information is reported. For example, if the Agent was installed with its default settings, not all hardware information is reported to the SAS Web Client until an hour after the agent is installed. There also might be a problem retrieving certain hardware information, such as a disk failure, that could prevent some hardware information from being reported. In these cases, the server's property page lists unreported information as not set.

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

Security for Agents Running on Managed Servers

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

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

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

To further safeguard the SSL/TLS-based communication channel, the two Agents participating in the code deployment also need to have a common shared secret provided by the Command Engine. Before one Agent can begin a file transfer to another

Agent, the two agents must verify that they have a common shared secret provided on a per-session basis by the Command Engine. This safeguard prevents unauthorized users from copying files from one managed server to another.

Agent Functionality on Managed Servers

The Agent is designed such that:

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

SA runs with administrator privileges (root on UNIX servers and Local System on Windows servers), because it performs tasks that require administrator privileges, such as installing patches and rebooting servers.

The core performs client authentication and, checks to see if the presenting certificate belongs to that particular server. SA does this by comparing the certificate to the server's IP address that is generated when the Agent is initially installed. If the certificate is not valid or the originating IP address does not match the IP address stored in the Model Repository, authentication fails and the Agent cannot continue communication with SA.

If an unauthorized user were able to log on to a managed server with administrator privileges and compromise a server's security, the user would have only limited access to the following information in the Model Repository:

- The server's own hardware inventory (already available to someone logged on with administrator privileges)
- The server's own software inventory (already available to someone logged on with administrator privileges)
- The custom attribute information

Server Information that the Agent Tracks

For each managed server, the Agent reports software, networking, and hardware information, as shown in Figure 4-2 and Figure 4-3. By communicating with the core and reporting the installed hardware and software for the server, SA determines what software should be installed on a server.

Figure 4-2: Manager Servers: Properties Page: Server Information

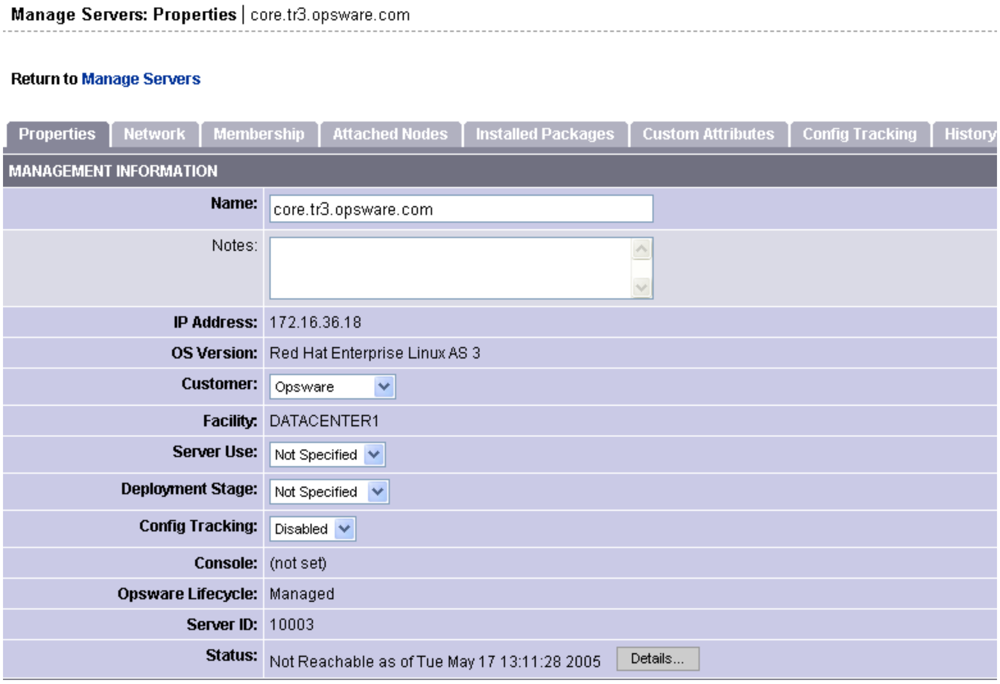

Manage Servers: Properties | core.tr3.opsware.com

Return to **Manage Servers**

MANAGEMENT INFORMATION

Name:	core.tr3.opsware.com
Notes:	
IP Address:	172.16.36.18
OS Version:	Red Hat Enterprise Linux AS 3
Customer:	Opsware
Facility:	DATACENTER1
Server Use:	Not Specified
Deployment Stage:	Not Specified
Config Tracking:	Disabled
Console:	(not set)
Opsware Lifecycle:	Managed
Server ID:	10003
Status:	Not Reachable as of Tue May 17 13:11:28 2005

Figure 4-3: Server Properties-Hardware and Additional Information

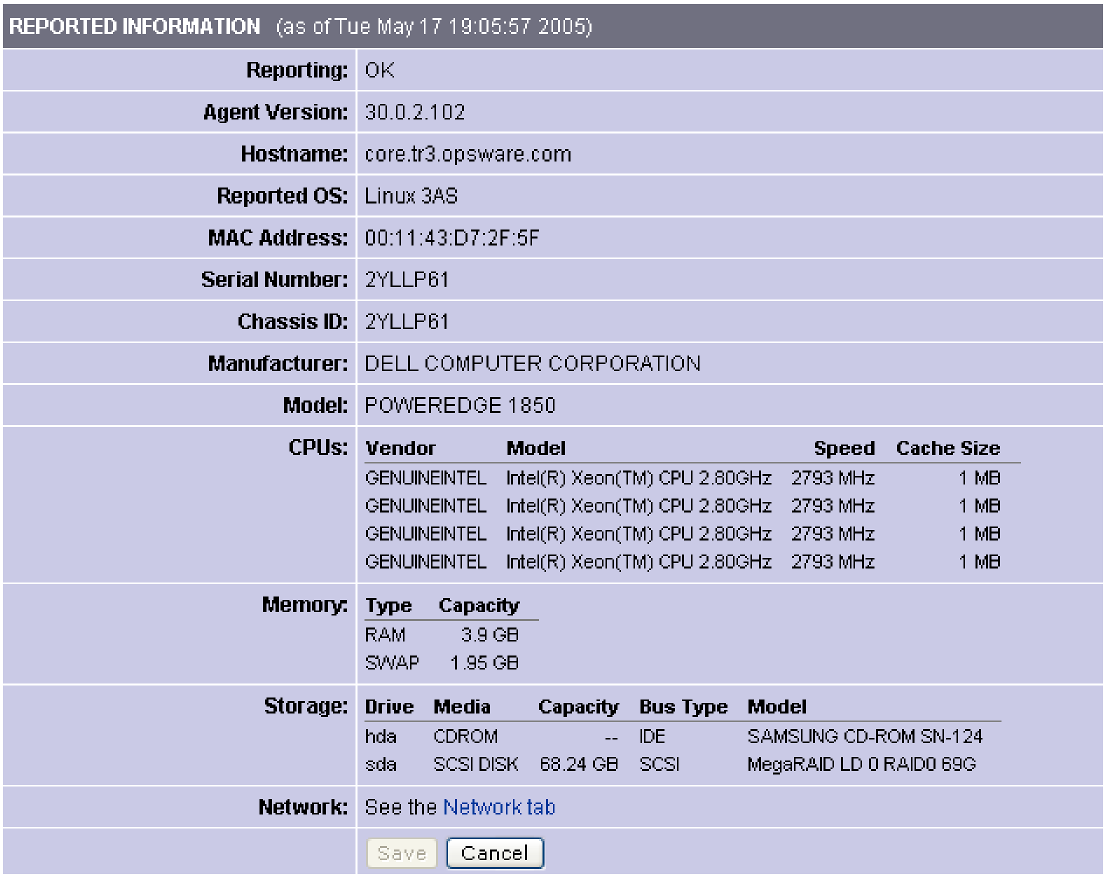

REPORTED INFORMATION (as of Tue May 17 19:05:57 2005)

Reporting:	OK
Agent Version:	30.0.2.102
Hostname:	core.tr3.opsware.com
Reported OS:	Linux 3AS
MAC Address:	00:11:43:D7:2F:5F
Serial Number:	2YLLP61
Chassis ID:	2YLLP61
Manufacturer:	DELL COMPUTER CORPORATION
Model:	POWEREDGE 1850

CPUs:

Vendor	Model	Speed	Cache Size
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2793 MHz	1 MB
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2793 MHz	1 MB
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2793 MHz	1 MB
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2793 MHz	1 MB

Memory:

Type	Capacity
RAM	3.9 GB
SWAP	1.95 GB

Storage:

Drive	Media	Capacity	Bus Type	Model
hda	CDROM	--	IDE	SAMSUNG CD-ROM SN-124
sda	SCSI DISK	68.24 GB	SCSI	MegaRAID LD 0 RAID0 69G

Network: See the Network tab

Software Information

The software that should be installed is recorded in the Model Repository. To access that information, click the Install List or Installed Packages tabs. This shows the software that should be installed on the server or all software that is installed on the server.

To display the list of what software packages should be installed on that server, select the Install List tab from the Manage Servers: Server Properties page.

To display the list of software that is reportedly installed on the server, select the Installed Packages tab from the Managed Servers: Server Properties page.

Partially-installed Solaris packages do not show up in the Installed Packages list, even though the package was partially installed.

Hardware Information

SA tracks hardware information in a variety of manners. Table 4-1 shows how the Agent obtains the server and hardware information about each managed server.

Table 4-1: Hardware Information that the Agent Reports for Servers

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
Name	The user-configurable name for the server. By default, SA uses the configured host name of the server until a user edits it.	**Windows**: Uses the fully qualified DNS name of the server. **Linux, Solaris, AIX, HP-UX**: Uses the current host name of the server that the `hostname` command returns.
Reported OS	The version number of the server's operating system.	**Windows**: Uses the Windows version number as reported by the operating system. This information includes the major version number, the minor version number, the Windows build number, and the Service Pack level. **Linux, Solaris, AIX, HP-UX**: Uses the operating system version that the `uname` command returns.
OS Version	The OS version specified for the OS definition.	Specified by the user who prepared the OS with the Prepare Operating System Wizard. See the *SA Policy Setter's Guide* for more information.
Serial Number	The serial number of the system. SA attempts to report a chassis ID, if possible.	**Windows, Linux**: Obtained from the system BIOS. **Solaris, AIX, HP-UX**: Obtained from the system ROM.
Manufacturer	The manufacturer of the server if available.	**Windows, Linux**: Obtained from the system BIOS. **Solaris, AIX, HP**: Obtained from the system ROM.
Model	The model of the server if available.	**Windows, Linux**: Obtained from the system BIOS. **Solaris, AIX**: Obtained from the system ROM. **HP-UX**: Output of model command (which is read from the system ROM).

Table 4-1: Hardware Information that the Agent Reports for Servers (continued)

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
Memory	The amount of physical RAM and the total amount of virtual memory paging space configured.	**Windows**: Uses the Windows 2000 API `GlobalMemoryStatus().` **Linux:** Obtained from information in the file `/proc/meminfo`. **Solaris**: Obtained from the `sysconf` and `swapctl` APIs. **AIX**: Uses the `lsattr` command for memory information and the `lsps` command for paging space. **HP-UX**: Uses the `pstat` system call.
Processors	Information about each of the processors in the system.	**Windows**: If WMI is available, iterates over all instances of `Win32_Processor`. If WMI is not available, parses the registry key `HARDWARE\DESCRIPTION\System \CentralProcessor`. There is one sub-key for each processor. **Linux**: Obtained from information in the file `/proc/meminfo`. **Solaris, HP-UX**: Uses system APIs to enumerate the processors in the system. **AIX**: Uses the `lscfg` command.
Storage	Information about each installed disk drive or RAID array.	**All Platforms**: Uses system APIs to discover and probe disk drives and RAID arrays.
Server ID	The internal ID that SA uses to identify the server.	In most cases, the server ID is the same as the MID.

Table 4-1: Hardware Information that the Agent Reports for Servers (continued)

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
MID	The MID (Machine ID) is a unique number that SA assigns when the server first registers. The server stores the MID and reports it each time the server registers.	**Windows**: The MID is stored in the file `%ProgramFiles%\Common Files\\cogbot\mid` if present. **Linux, Solaris, AIX, HP-UX**: The MID is stored in the file `/etc/opt//agent/mid`

In addition to hardware and software reporting, the Agent reports networking information. See "Network Configuration" on page 280 in Chapter 7 for descriptions of the networking information reported and how you can modify that information using the Network tab in the SAS Web Client.

Starting an Agent

To start an Agent, log onto the managed server and enter the following command.

Unix:

```
/etc/init.d/-agent start
```

Windows:

```
net start agent
```

To stop an agent, enter the same command, specifying `stop` instead of `start`.

Discovery and Agent Deployment

The Discovery and Agent Deployment (ODAD) feature helps to deploy Agents to a large number of servers through the SAS Client. This feature helps to identify servers on which to install an Agent, specify the deployment actions to be carried out on each server, select the login protocols to connect to each server, specify the Agent Installer options for installing an Agent, and generate reports on Agent installation status. See Figure 4-4.

Figure 4-4: Agent Deployment Process

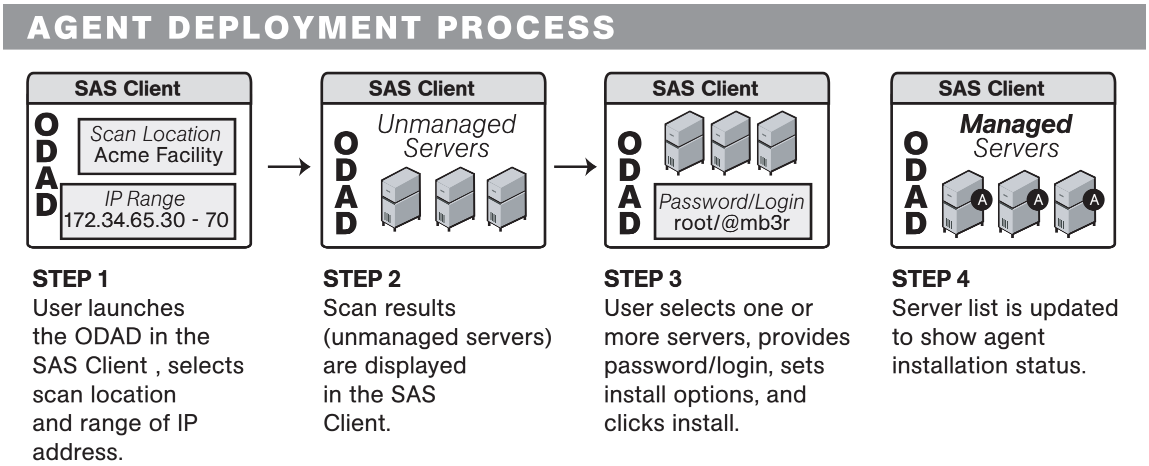

STEP 1
User launches the ODAD in the SAS Client , selects scan location and range of IP address.

STEP 2
Scan results (unmanaged servers) are displayed in the SAS Client.

STEP 3
User selects one or more servers, provides password/login, sets install options, and clicks install.

STEP 4
Server list is updated to show agent installation status.

This overview section contains the following topics:

- Discovering Servers for Installing an Agent
- Setting Deployment Actions for Each Server
- Specifying Login Settings
- Agent Installer Options
- Reports on Server Status

See "Installing Agents Using ODAD" on page 120 in this chapter for information about how to install an Agent.

Discovering Servers for Installing an Agent

Using the Discovery and Agent Deployment feature, you can select a location to scan for servers. After selecting a location, you can specify the IP addresses or IP address ranges to perform a network scan to identify servers in which to install an Agent. Instead of performing a network scan, you can also import a file containing a list of IP addresses or IP address ranges. When the scan is complete, a list of scanned servers is shown.

For each server, this feature determines the status of the server, its IP address, its host name, detected operating system, and open ports used to connect to the server.

Setting Deployment Actions for Each Server

Once you have identified the servers, you can select the servers and perform the following deployment actions:

- Verify installation prerequisites.

  When you select this action, verification checks are performed to ensure that the Agent is successfully installed on the server. The verification checks include:

  - Checking for sufficient disk space for Agent installation on the server
  - Verifying that no other application is using port 1002
  - Verifying if ports to the Gateway are accessible

- Verify prerequisites, and copy the Agent Installer to servers.

  When you select this action, verification checks are performed to ensure that the Agent is successfully installed on the server and the Agent Installer is copied to the server.

- Verify prerequisites, copy installer, and install Agent.

  When you select this action, verification checks are performed to ensure that the Agent is successfully installed on the server, the Agent Installer is copied to the server, and the Agent is installed on the server.

Specifying Login Settings

After selecting the deployment action, you can select the network protocols to connect to the server and specify the user name and password to login to each server. The Agent needs administrator-level privileges (root on Unix servers and administrator on Windows servers) to manage a server. Therefore, Agent installation is performed as root on Unix operating systems and as administrator on Windows operating systems. ODAD tries to log into each of the selected servers with the specified user name and password and performs the specified deployment actions.

Agent Installer Options

The Discovery and Agent Deployment feature allows you to specify the installer options listed in Table 4-2.

Table 4-2: Agent Installer Options

OPTION	DESCRIPTION
Start the Agent after Installation	Starts the Agent after installing it on the server. By default, the Agent Installer does not start the Agent.
Ignore prerequisite check failures	Ignores the prerequisite check failures and forces Agent installation.
Set the server's time from the core	Synchronizes the time on the server in which the Agent is installed to the core.
Install Windows Installer (MSI) if required	Installs MSI along with the Agent. If MSI is already installed, this option has no effect.
Install Windows Management Instrumentation (WMI) if required	Installs WMI along with the Agent. If WMI is already installed, this option has no effect.
Reboot Windows servers after agent installation	Reboots windows servers after Agent installation is complete.
Install Red Hat Package Manager (RPM) on AIX and Solaris	Installs the RPM handler with the Agent. Always include this option when you install Agents on Solaris and AIX servers.
Reset Agent configuration, if present	Replaces the existing Agent configuration.
Delete gateway address list, if present	Deletes the Gateway address list if present and is no longer required.
Overwrite staged Agent installer	Overwrites the existing Agent Installer.
Log Level	Sets the log level for log messages. With this option, you can specify the log levels of error, warning, info, and trace.

Table 4-2: Agent Installer Options (continued)

OPTION	DESCRIPTION
Use Gateways	Specifies the Gateways used during Agent installation.
Immediately do a full hardware registration	Forces the Agent Installer to report full hardware information to the core.
Immediately do software registration	Forces the Agent Installer to report full software information to the core.
Suppress Agent reachability check	By default, the installer triggers the core to check if the server is reachable. This option disables this check during installation.
Disallow anonymous SSL connections if Agent is dormant	Configures the Agent so that browsers cannot connect without a valid certificate.
Force creation of new device record if conflict found	During registration, the Data Access Engine creates a new device record. This option suppresses this functionality.
Fail if initial hardware registration fails (do not go dormant)	Does not allow the Agent to become dormant, if it fails to report hardware information.
Do not open Windows Firewall for core-agent communications	By default, the Agent Installer will modify the Windows Firewall configuration on Windows XP and Windows 2003 (r2) servers to allow the core to contact the managed server on port 1002. If you select this option, the firewall configuration will not be modified. The server may not be manageable by in this case.
Remediate software policies	Remediates the server against any software policies attached to the server.
Attach to software Policy ID	Attaches server to the software policy *<ID>*.

Table 4-2: Agent Installer Options (continued)

OPTION	DESCRIPTION
Extra Installer options	Allows you to specify any other installer options. For example: `--log file <path>` allows you to specify the path to the installer log file. By default, the installer log files are placedin the `/tmp` on Unix or `%SYSTEMDRIVE%\WINDOWS\SYSTEM` on Windows. `--workdir <path>` allows you to specify the path to the working directory to use while the installation is in progress.
Install directory (UNIX)	Allows you to specify the directory where the Agent will be installed. By default, the Agent is installed in `/opt//agent` on UNIX, and `%ProgramFiles%\\agent` on Windows.

Reports on Server Status

After the deployment action is complete, the SA Client displays the results and updates the status icons for the servers as shown in Table 4-3.

Table 4-3: Server Status

ICONS	SERVER STATUS
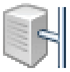	The server is unmanaged.
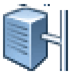	The server is managed by .

Table 4-3: Server Status (continued)

ICONS	SERVER STATUS
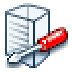	The server failed prerequisite checks.
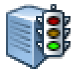	The server passed prerequisite checks.
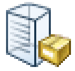	The server passed prerequisite checks and the Agent Installer was copied to the server.
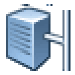	The Agent was successfully deployed.
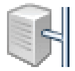	The Agent was not successfully deployed.

A server is considered to be managed by when the ODAD determines that the Agent is listening for TCP connections on port 1002.

For a failed deployment action, you can view the errors on each server. You can also log into the server from this feature and correct the errors.

Using this feature, you can create the following reports:

- All the servers in the current network scan
- Selected servers in the current network scan
- All the servers in the current network scan with successful deployments
- Servers in the current network scan with failed deployments

You can save and export the reports to a CSV, HTML, or text format file.

Agent Installation Using ODAD

The Discovery and Agent Deployment (ODAD) feature helps you to deploy Agents to a large number of servers through the SA Client. This section contains the following topics:

- Prerequisite Setup for Discovery and Agent Deployment
- Permissions Required for Discovery and Agent Deployment
- Installing Agents Using ODAD

Prerequisite Setup for Discovery and Agent Deployment

Before you can install Agents on Windows servers using ODAD, you must install the Windows Agent Deployment Helper. For instructions, see “Installing the Windows Agent Deployment Helper” in the *SA Planning and Installation Guide*.

Permissions Required for Discovery and Agent Deployment

To use the Discovery and Agent Deployment feature you must have certain permissions. See *SA Administration Guide* for more information about the permissions required to use the ODAD feature.

To obtain the required permissions for scanning and deploying Agents, contact your administrator.

Installing Agents Using ODAD

To install Agents using ODAD, launch the SA Client from the Opsware SAS Web Client. See “SA Client and Server Automation Launcher Requirements” on page 71 in Chapter 3 for more information.

Perform the following steps to install an Agent:

1. Log into the SA Client. The SA Client Page appears.
2. From the navigation pane, select Devices and then select Unmanaged Servers.
3. Select a location to scan for servers from the Scan in drop-down list.

4 Select Supply IP Address Ranges from the drop-down list to specify a set of IP address ranges to scan. Enter the IP address range in the From and To field. Click the plus (+) button (next to the To field) to add another IP address range. You can add a maximum of five IP address ranges. Click the minus (-) button to delete the IP address range field.

Or

Select Explicit list of IPs from the drop down list to specify the list of IP addresses to scan, separated by spaces (commas not supported). Click the ellipsis (...) button as shown in Table 4-5 to display a text editor that allows you to load and scan a file with IP ranges to scan.

Figure 4-5: Loading a File with IP Addresses

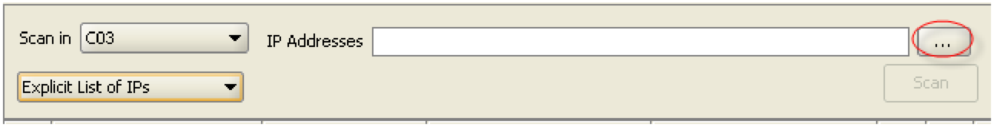

Scan in C03

Explicit List of IPs

5 Click **Scan** to scan for servers. When the scan is complete, the list of scanned servers is shown. For each server, ODAD determines the status of the server, its IP address, its host name, the detected operating system and open ports that can be used to connect to the server.

If you have a firewall enabled, ODAD may not be able to accurately detect a managed server’s actual installed operating system. Some firewalls can interfere with the methods ODAD uses to detect the operating system. ODAD must be able to access at minimum one open port and one closed port to gather the information needed to determine the operating system. If you find that ODAD has not identified any operating system or has misidentified the operating system, you may need to configure your firewall to allow network packets from the SA Core.

See Figure 4-6. The actual operating system of the server can be only determined if ODAD is able to successfully log into the server.

Figure 4-6: Scan Results

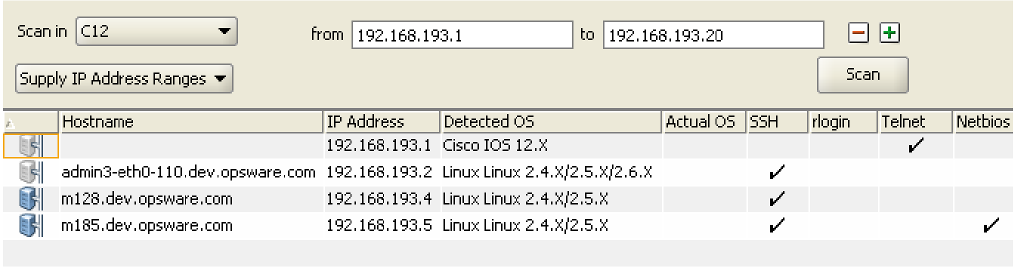

Scan in C12 from 192.168.193.1 to 192.168.193.20

Supply IP Address Ranges

	Hostname	IP Address	Detected OS	Actual OS	SSH	rlogin	Telnet	Netbios
		192.168.193.1	Cisco IOS 12.X				✓	
	admin3-eth0-110.dev.opsware.com	192.168.193.2	Linux Linux 2.4.X/2.5.X/2.6.X		✓			
	m128.dev.opsware.com	192.168.193.4	Linux Linux 2.4.X/2.5.X		✓			
	m185.dev.opsware.com	192.168.193.5	Linux Linux 2.4.X/2.5.X		✓			✓

6 Click on any of the column headings to sort the server list by that column. If you want to hide managed servers, select **Hide -managed Servers** from the **View** menu.

7 Select servers on which you want to deploy the Agent. The SA Client supports hot keys to make multiple selections.

8 From the **Actions** menu, select **Deploy Agent**. The Deploy Agent dialog box appears as shown in Figure 4-7.

Figure 4-7: Deploy Agents

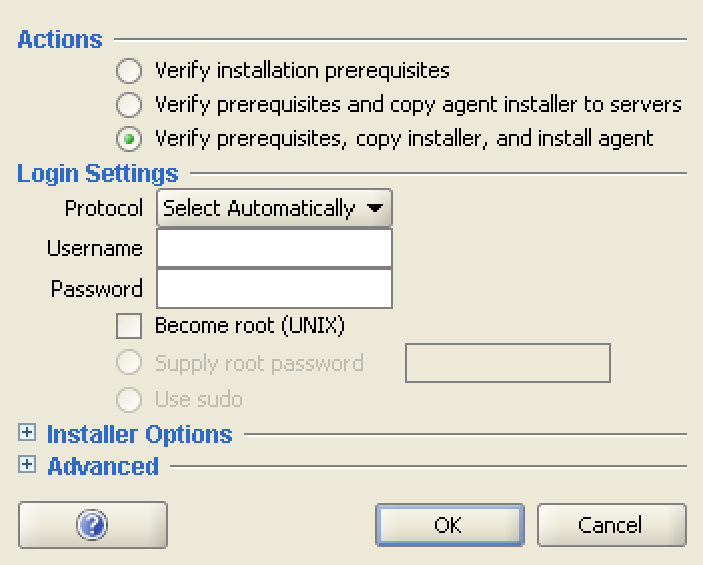

Actions

○ Verify installation prerequisites

○ Verify prerequisites and copy agent installer to servers

● Verify prerequisites, copy installer, and install agent

Login Settings

Protocol: Select Automatically

Username:

Password:

□ Become root (UNIX)

○ Supply root password

○ Use sudo

⊞ Installer Options

⊞ Advanced

9 Select one of the following deployment actions:

- Verify installation prerequisites.
- Verify prerequisites and copy agent installer to servers.
- Verify prerequisites, copy installer, and install agent.

See “Setting Deployment Actions for Each Server” on page 115 in this chapter for more information.

10 Select a network protocol to log in and connect to the server from the drop-down list.

Or

Choose Select Automatically to allow ODAD to select an appropriate protocol for each server.

11 Enter the user name to log into the server. For Windows, log in as administrator. For Unix, log in as root.

12 If logging in as root is not allowed, select the Become root (UNIX) checkbox. Select “Supply root password” and enter the password or select Use sudo**,** if sudo access is enabled for that account.

If you log in using sudo, the sudoers configuration file (usually `/etc/sudoers`) must allow your account to run any command with root privileges. This is typically accomplished by using the "ALL" alias in the sudoers file. If only certain commands are permitted, then deploying agents using ODAD will fail.

See “Specifying Login Settings” on page 115 in this chapter for more information.

13 If you are unable to deploy the agent to a UNIX server by logging in as root, the system you are deploying to may be configured to disallow direct root logins. In such cases ODAD allows you to log in as a non-root user and then escalates your privileges via either the "su" command or the "sudo" command.

Perform the following steps to deploy agents as a non-root user:

1. Enter the unprivileged user name in the Username field to log into the server.
2. Enter the unprivileged password in the Password field.
3. Select the Become root (UNIX) checkbox. Select “Supply root password” and enter the password or select “Use sudo”. If you choose to use sudo, the unprivileged account must be able to run any command as root.

14 Specify the Agent Installer options to control the way the Agent is installed on a server. See “Agent Installer Options” on page 116 in this chapter for more information.

15 Click **OK** to deploy Agents to selected servers.

16 After the deployment action is completed, the SA Client displays the results and updates the status icons for the servers. You can view information on an unmanaged server and generate reports on Agent Installation status. See “Reports on Agent Installation” on page 127 in this chapter for more information.

Viewing Unmanaged Server Information

After the deployment action is completed, you can review the results and generate reports. You can view the summary and history information for an unmanaged server. For a failed deployment action, you can view the errors on each server. You can also log into the server from ODAD and correct the errors.

This section discusses the following topics:

- Summary Information for an Unmanaged Server
- History Information for an Unmanaged Server
- Remote Terminal Sessions on Unmanaged Servers
- Reports on Agent Installation

Summary Information for an Unmanaged Server

The Unmanaged Server Summary browser shows the following information about the unmanaged server:

- **Host Name**: The host name of the unmanaged server, if defined in the Domain Name System (DNS).
- **IP Address**: The IP address of the unmanaged server.
- **Detected OS**: The operating system detected on the server after performing the network scan. The listed operating system is a *best guess* made by comparing the Managed Server’s response to a network probe of a list of known operating system *fingerprints*. The listed operating system may not always be accurate, for example when a firewall exists between the SA Core and the Managed Server and does not allow or alters the network probe.

- **Actual OS**: The actual operating system detected on the server after the deployment action. This entry will be blank until ODAD can successfully log on to the managed Server and identify the operating system type and version.
- **MAC Address**: The Media Access Control (MAC) address, which is the network interface card's unique hardware number of the unmanaged server. The MAC is used as the server's physical address on the network. The MAC address is only detected if the server is on the same physical network as the Gateway.
- **NIC Vendor**: The vendor name for the Network Interface Card (NIC) driver. The NIC vendor is only detected if the server is on the same physical network as the Gateway.
- **Open Ports**: The discovered open ports on an unmanaged server. ODAD does not perform a comprehensive search for open ports. There may be open ports not listed.
- **Number of Deployment Attempts**: The number of deployment attempts on an unmanaged server.
- **Last Deployment Attempt Date/Time)**: The Date/Time of the last deployment attempt.
- **Last Deployment Attempt Message**: The message stating the possible cause for the failure of the last deployment attempt.

Viewing Summary Information of an Unmanaged Server

Perform the following steps to view the summary of an unmanaged server.

1. Log into the SA Client. The SA Client Page appears.
2. From the navigation pane, select Devices and then select Unmanaged Servers.
3. From the Unmanaged Servers page, select the unmanaged server. From the **Actions** menu, select **Open**.

   The Unmanaged Server Browser page appears.

4 From the left Navigation pane, click Summary to view the summary of the unmanaged server. See Figure 4-8.

Figure 4-8: Unmanaged Server Summary Page

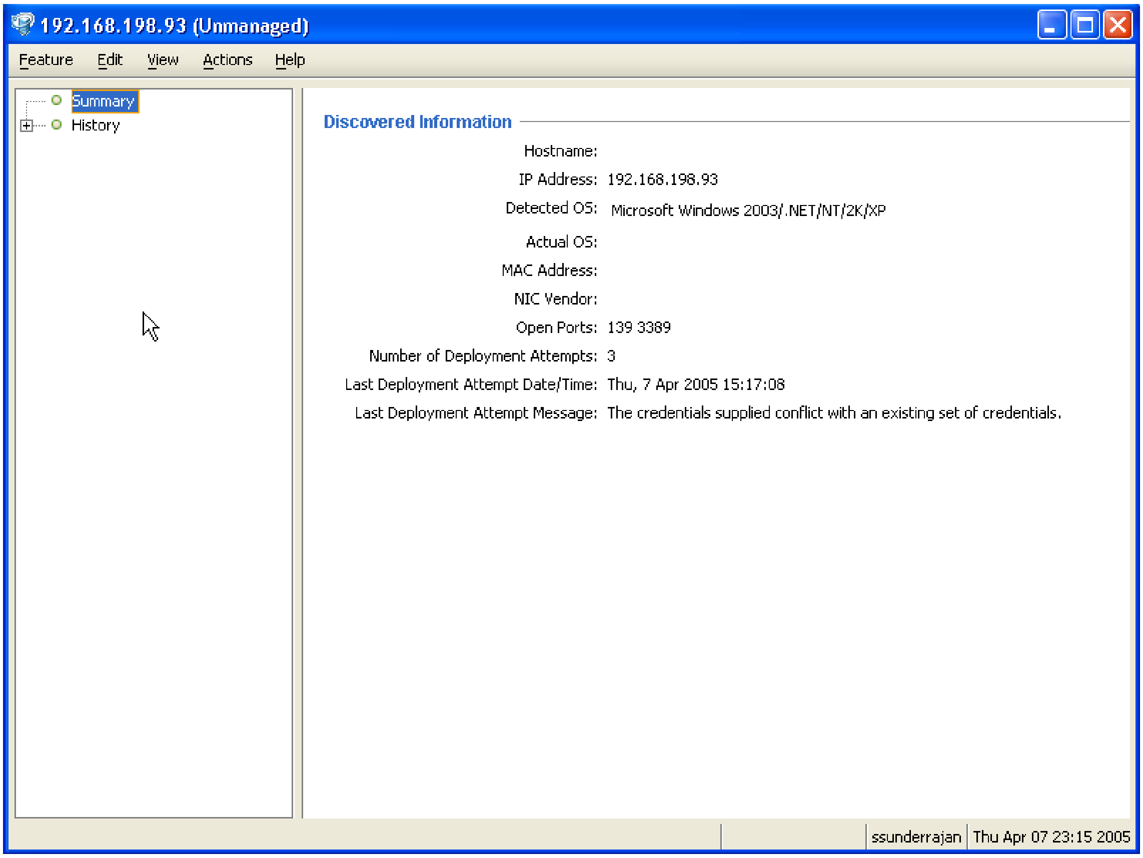

Hostname:
IP Address: 192.168.198.93
Detected OS: Microsoft Windows 2003/.NET/NT/2K/XP
Actual OS:
MAC Address:
NIC Vendor:
Open Ports: 139 3389
Number of Deployment Attempts: 3
Last Deployment Attempt Date/Time: Thu, 7 Apr 2005 15:17:08
Last Deployment Attempt Message: The credentials supplied conflict with an existing set of credentials.

History Information for an Unmanaged Server

The Unmanaged Server History Browser shows a history of all actions executed on unmanaged servers such as:

- A summary of the actions performed on the unmanaged server
- Details of all the actions performed on the unmanaged server
- Log information

The history is read-only.

Viewing History Information of an Unmanaged Server

Perform the following steps to view the history of an unmanaged server:

1. Log into the SA Client. The SA Client Page appears.
2. From the navigation pane, select Devices and then select Unmanaged Servers.
3. From the Unmanaged Servers page, select the unmanaged server. From the **Actions** menu, select **Open**.

   The Unmanaged Server Browser page appears.
4. From the left Navigation pane, click History to view the history of all actions executed on unmanaged servers.

Remote Terminal Sessions on Unmanaged Servers

Using the ODAD feature, you can open terminal sessions on an unmanaged server. You can log into the server using the appropriate protocol to correct the errors on the server or perform any other operations.

You can use the SA Client Remote Terminal and Shell preferences to configure the Terminal (UNIX) and RDP (Windows) clients used to launch terminal sessions. See “Overview of the Global Shell” on page 377 in Chapter 10 for more information.

Opening Remote Terminal Sessions on an Unmanaged Servers

Perform the following steps to open remote terminal sessions on a server:

1. Log into the SA Client. The SA Client Page appears.
2. From the navigation pane, select Devices and then select Unmanaged Servers.
3. From the Unmanaged Servers page, select the unmanaged server. From the **Actions** menu, select the appropriate log in with protocol.

Reports on Agent Installation

Using the Discovery and Agent Deployment feature, you can create the following reports:

- All the servers in the current network scan
- Selected servers in the current network scan
- Servers in the current network scan with successful deployments
- Servers in the current network scan with failed deployments

You can also save and export the reports to a CSV, HTML, or text file format.

Creating Reports on Agent Installation

Perform the following steps to create reports:

1. Log into the SA Client. The SA Client Page appears.
2. From the navigation pane, select Devices and then select Unmanaged Servers.
3. From the Unmanaged Servers page, select the unmanaged server. From the **Actions** menu, select **Export to** the desired report format. The Save Report dialog box appears.
4. From the drop-down list, select the type of report as shown in Figure 4-9.

Figure 4-9: Generating Reports

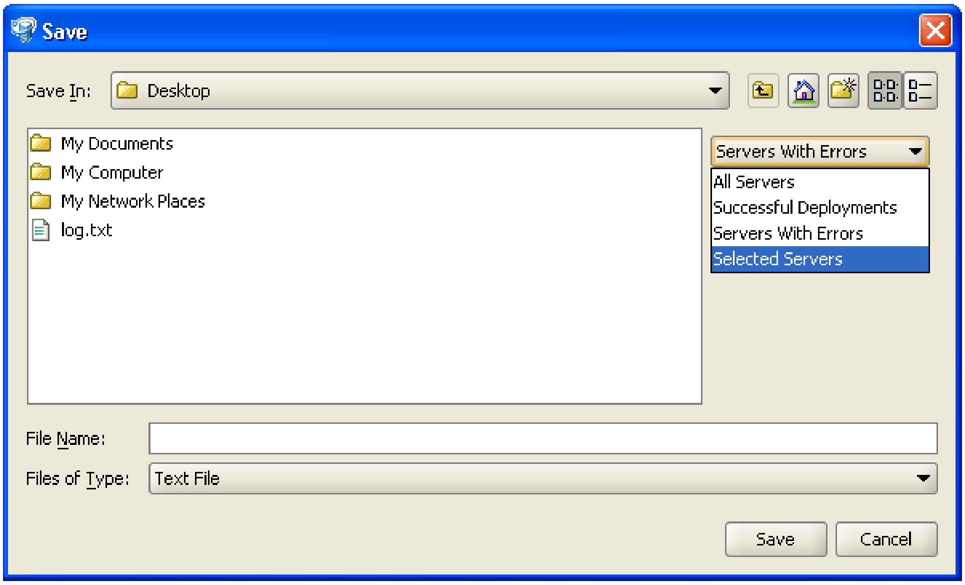

5. Enter the location and file name to save the report.

Example Report

The following example shows a report of servers with failed deployments.

```
Servers which the Agent could not be deployed to:

Hostname : Unknown

IP Address : 192.168.198.93
```

```
Detected Operating System : Microsoft Windows 2003 Server or XP
SP2

Open Ports : 139 3389

MAC Address : Unknown

NIC Vendor : Unknown

# of Deployment Attempts : 2

Last Deployment Message : The credentials supplied conflict with
an existing set of credentials.

Last Deployment Attempt Date/Time : "Wed, 6 Apr 2005 16:18:46"

Failed Phase : Check

Return Code : 2013

Suggested Resolution : The Agent Deployment Helper was unable to
log in to the unmanaged server.

An incorrect login name or password was specified. Try a
different login name and/or password.
```

Agent Reachability Communication Tests

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

- Overview of Communication Tests
- Communication Tests and Unreachable Agents
- Communication Test Types
- Communication Test Errors
- Additional Information on Communication Tests
- Running a Communication Test on an Individual Server
- Running a Communication Test on Multiple Servers
- Viewing Servers by Communication Status
- Searching for Unreachable Servers

- Viewing My Jobs Communication Tests
- Exporting the Unreachable Server Status List to CSV

Overview of Communication Tests

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

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

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

You have the ability to check Agent reachability for individual servers, selected servers, or all servers under the management of SA. Each time that you run a Communication Test, this test is saved in the My Jobs panel, which allows you to view a history of all the Communication Tests that you have run. You can even export the current reachability status of all managed servers to a CSV file.

Communication Tests and Unreachable Agents

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

Communication Test Types

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

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

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

See “Agent Functionality on Managed Servers” on page 108 in this chapter for information about the Agent and its relationship to managed servers.

Communication Test Errors

After you run a Communication Test, three icons indicate the success or failure of agent reachability as Table 4-4 shows.

Table 4-4: Agent Unreachability Status Icons

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

Table 4-5 describes each type of Communication Test and all possible errors for each test.

Table 4-5: Communication Test Types with Possible Results

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

Table 4-5: Communication Test Types with Possible Results (continued)

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

Table 4-5: Communication Test Types with Possible Results (continued)

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

Table 4-5: Communication Test Types with Possible Results (continued)

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

See "Communication Test Troubleshooting" in *User's Guide: Server Automation* for information about how to troubleshoot Communication Test Errors.

Additional Information on Communication Tests

In the case that the Communication Test cannot be performed on a server, you see an error named Unexpected Error with a small plus (+) button next to it. Click the plus (+) button to see the Additional Information window, which provides traceback information regarding the error. You can send this information to Customer Support to solve a problem of this nature. See Figure 4-10.

Figure 4-10: Additional Information on an Unexpected Error

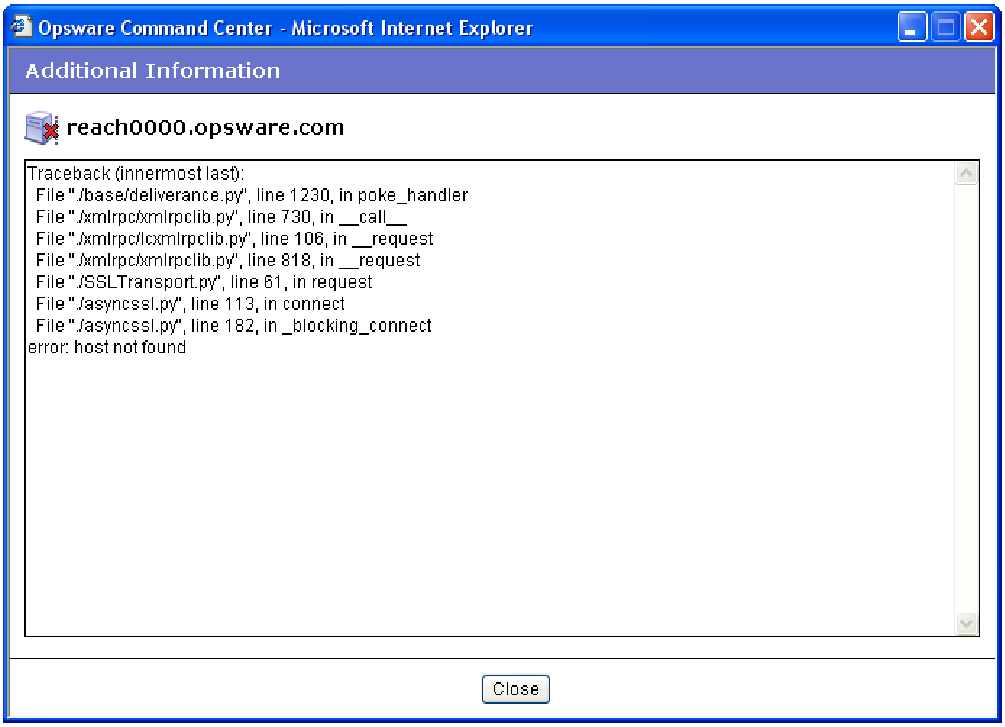

```
Traceback (innermost last):
  File "./base/deliverance.py", line 1230, in poke_handler
  File "./xmlrpc/xmlrpclib.py", line 730, in __call__
  File "./xmlrpc/lcxmlrpclib.py", line 106, in __request
  File "./xmlrpc/xmlrpclib.py", line 818, in __request
  File "./SSLTransport.py", line 61, in request
  File "./asyncssl.py", line 113, in connect
  File "./asyncssl.py", line 182, in _blocking_connect
error: host not found
```

Running a Communication Test on an Individual Server

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

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

2. From the Manage Servers list, click the display name of the server that you want to perform a Communication Test on.

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

3. To see the results of the last Communication Test for this server, click **Details**. The Communication Test window for the server appears, as Figure 4-11 shows.

Figure 4-11: Communication Test Results on an Individual Server

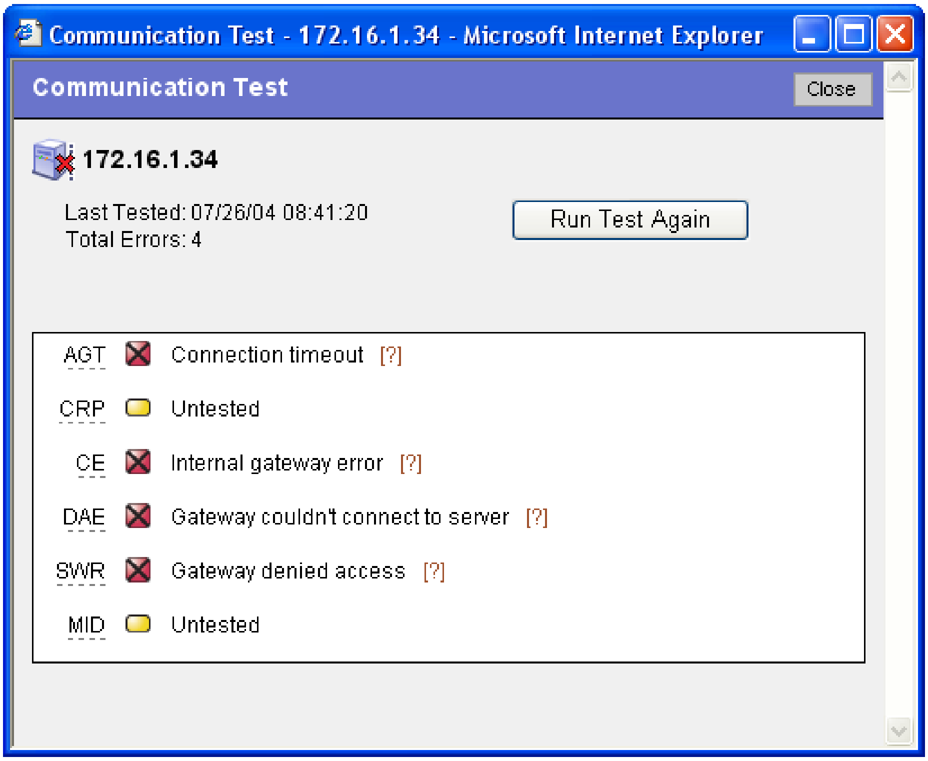
Communication Test - 172.16.1.34 - Microsoft Internet Explorer
Communication Test
Close
Run Test Again

172.16.1.34

Last Tested: 07/26/04 08:41:20
Total Errors: 4

AGT	☒	Connection timeout [?]
CRP	▭	Untested
CE	☒	Internal gateway error [?]
DAE	☒	Gateway couldn't connect to server [?]
SWR	☒	Gateway denied access [?]
MID	▭	Untested

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

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

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

Running a Communication Test on Multiple Servers

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

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

2 In the Manage Servers Summary View page, select the servers for which you want a Communication Test.

3 Choose **Tasks ➤ Run ➤ Communication Test**. The Communication Test window opens and the test is initiated. The top of the test window shows the status report for the Communication Test, which indicates the time of the test, how many servers in the test were reachable, which servers were not reachable, and a progress bar. This summary information is shown in Figure 4-12.

Figure 4-12: Communication Test Summary

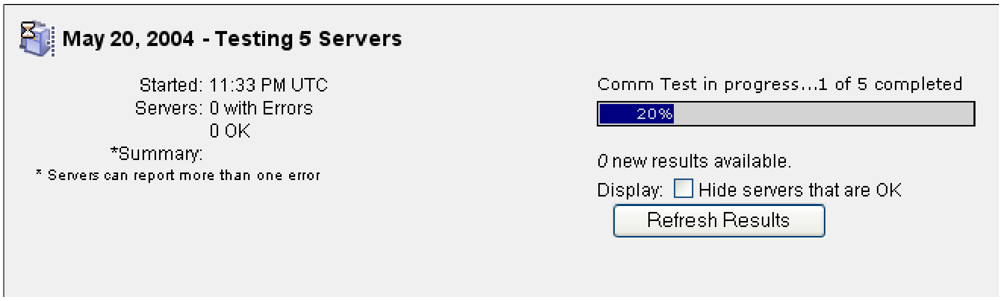

May 20, 2004 - Testing 5 Servers

Started: 11:33 PM UTC
Servers: 0 with Errors
0 OK
*Summary:
* Servers can report more than one error

Comm Test in progress...1 of 5 completed

20%

0 new results available.

Display: □ Hide servers that are OK

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

Below the summary section is a list of all Agents that were not reachable and their details, as Figure 4-13 shows.

Figure 4-13: Communication Test Results on an Unreachable Agent

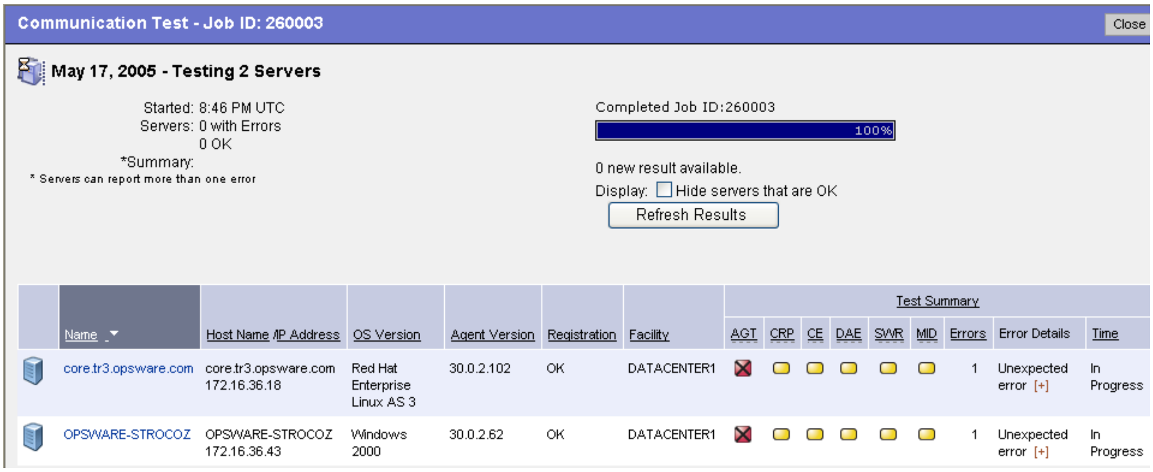

Communication Test - Job ID: 260003

May 17, 2005 - Testing 2 Servers

Started: 8:46 PM UTC
Servers: 0 with Errors
0 OK
*Summary:
* Servers can report more than one error

Completed Job ID:260003
100%
0 new result available.
Display: □ Hide servers that are OK

							Test Summary								
	Name	Host Name /IP Address	OS Version	Agent Version	Registration	Facility	AGT	CRP	CE	DAE	SWR	MID	Errors	Error Details	Time
	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	Red Hat Enterprise Linux AS 3	30.0.2.102	OK	DATACENTER1	☒	●	●	●	●	●	1	Unexpected error [+]	In Progress
	OPSWARE-STROCOZ	OPSWARE-STROCOZ 172.16.36.43	Windows 2000	30.0.2.62	OK	DATACENTER1	☒	●	●	●	●	●	1	Unexpected error [+]	In Progress

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

Viewing Servers by Communication Status

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

1. From the navigation panel, click Servers ➤ Manage Servers.
2. From the **View** menu, choose **Communication**. A list of the most recent Communication Test that was run on all managed servers appears. All servers listed in this view are listed as unreachable or reachable since the last regularly scheduled Communication Test was run.

3. To sort the information in this view, click any of the column headings. For example, you might want to view the servers by number of errors. Click the Error column title by name, by OS version, and so on.

4. To export this view to the CSV file format, choose **Export to CSV** from the **Resource** menu.

5. To run a new Communication Test for some or all of the servers in this view, select the servers that you want to run a Communication Test on and then choose **Run ➤ Communication Test** from the **Tasks** menu.

Searching for Unreachable Servers

Another way you can discover unmanaged servers (those with unreachable agents) is to search for all servers that have a status of unreachable. See “Using the Search Feature” on page 171 in *User’s Guide: Server Automation* for more information.

Perform the following steps to search for unreachable agents:

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

2. From the Server Search page, choose the Agent Reporting attribute from the first list, as Figure 4-14 shows.

Figure 4-14: Agent Reporting Server Search Attribute

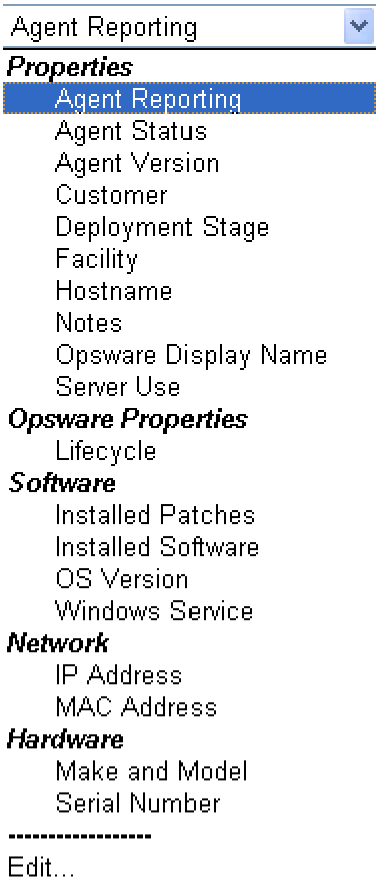

More criteria displays for the search parameters.

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

4. Click **Search**. Wait a few moments for the results (the speed of the search results depends on how many managed servers are being searched). If any of your managed servers are not reachable, these servers will be listed in the search results.

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

Viewing My Jobs Communication Tests

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

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

Exporting the Unreachable Server Status List to CSV

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

1. From the navigation panel, click Servers ➤ Manage Servers.
2. From the **View** menu, choose **Communication**. You see a list of all servers that are in an unreachable or reachable state.

To export a list of these servers to the CSV file format, select the check box next to each server that you want to include in the report, then from the **Resource** menu choose **Export to CSV**.

Chapter 5: Exploring Servers and Groups in the SA Client

IN THIS CHAPTER

This section discusses the following topics:

- Exploring Servers in the SA Client
- Server Status
- Ways to Use the Device Explorer
- Device Explorer Interface
- Basic Server Management Tasks
- Device Groups
- Device Group Explorer

Exploring Servers in the SA Client

The SA Client allows you to view a list of all your servers in your data center, which can exist in various states of SA management. All your servers can be accessed from the Devices pane in the main SA Client interface, as shown in Figure 5-1.

Figure 5-1: Servers in the Device Pane

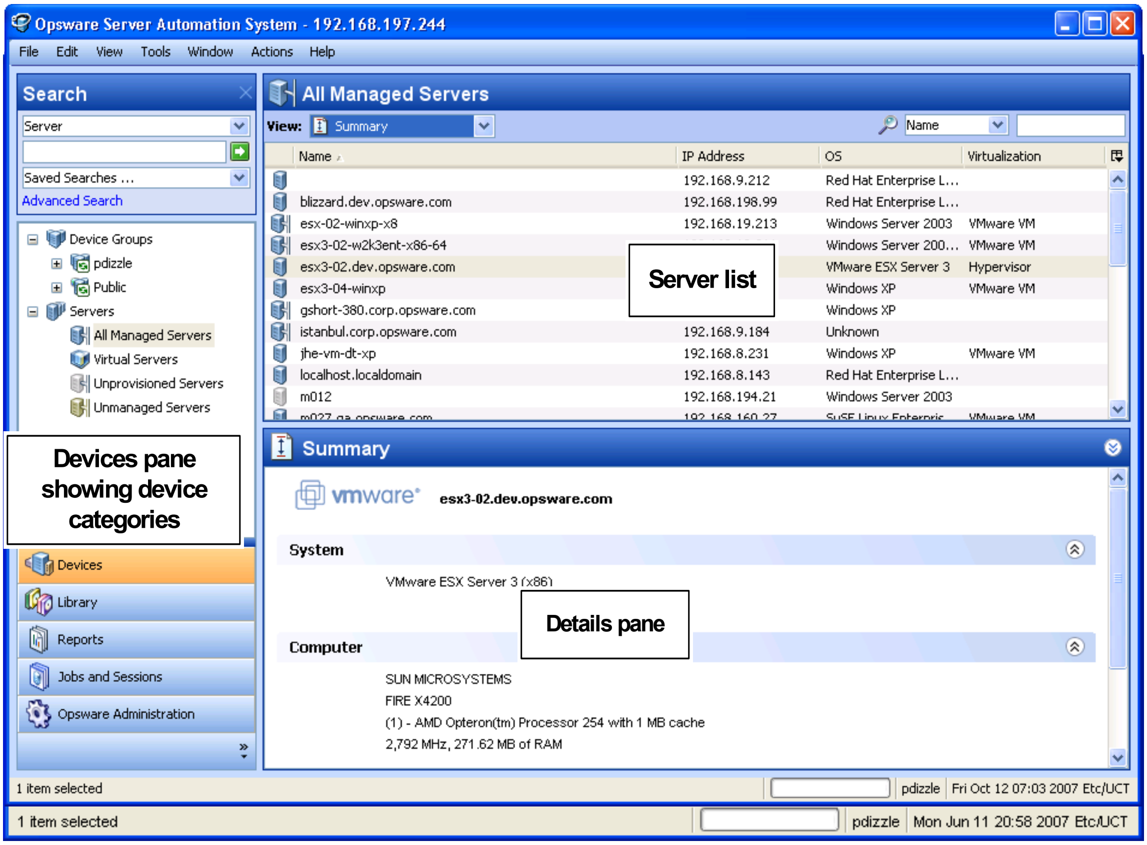

Name	IP Address	OS	Virtualization
	192.168.9.212	Red Hat Enterprise L...	
blizzard.dev.opsware.com	192.168.198.99	Red Hat Enterprise L...	
esx-02-winxp-x8	192.168.19.213	Windows Server 2003	VMware VM
esx3-02-w2k3ent-x86-64		Windows Server 200...	VMware VM
esx3-02.dev.opsware.com		VMware ESX Server 3	Hypervisor
esx3-04-winxp		Windows XP	VMware VM
gshort-380.corp.opsware.com		Windows XP	
istanbul.corp.opsware.com	192.168.9.184	Unknown	
jhe-vm-dt-xp	192.168.8.231	Windows XP	VMware VM
localhost.localdomain	192.168.8.143	Red Hat Enterprise L...	
m012	192.168.194.21	Windows Server 2003	
m027.qa.opsware.com	192.168.160.27	SuSE Linux Enterpris	VMware VM

esx3-02.dev.opsware.com

System

VMware ESX Server 3 (x86)

Computer

SUN MICROSYSTEMS
FIRE X4200
(1) - AMD Opteron(tm) Processor 254 with 1 MB cache
2,792 MHz, 271.62 MB of RAM

1 item selected | pdizzle | Fri Oct 12 07:03 2007 Etc/UCT

1 item selected | pdizzle | Mon Jun 11 20:58 2007 Etc/UCT

In order to visualize networking information with Network Automation (NA) inside of the SA Client, you must have both a licensed version of NA integrated with your SA core, plus an additional license for NA.

Additionally, in order to view storage devices and SAN information from the Storage Automation System (ASAS) inside of the SA Client, you must have both a licensed version of ASAS integrated with your SA core, plus an additional license to run SA Client showing ASAS data. If you have not purchased NA or ASAS, but would like to, contact your SA sales representative.

Server Status

In the SA Client, you can determine the status of a server by the type of icon used next to the server in the list, as defined in Table 5-1.

Table 5-1: Server Status Icons in SA

SERVER ICON	DESCRIPTION
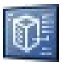	**Planned** Indicates that a device record has been created for the server, but an OS Build Agent has not yet been installed on it. Servers in this stage cannot be provisioned until the OS Build Agent is installed. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
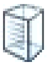	**Unprovisioned – Unreachable** Indicates a server that has been registered with the core via the OS Build Agent, but has not reported as ready for provisioning recently. This may be due to networking problems between the server and the SA core or the server having been disconnected or powered off. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
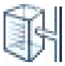	**Unprovisioned – Reachable** Indicates a server that has been registered with the core via the OS Build Agent and is available to have a target OS installed on it. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.

Table 5-1: Server Status Icons in SA (continued)

SERVER ICON	DESCRIPTION
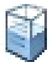	**Provisioning – Unreachable** Indicates a server on which the OS Provisioning feature was in the process of installing the target OS, but for some reason stopped because the server is unable to communicate with the SA core. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
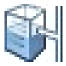	**Provisioning – Reachable** Indicates a server on which the OS Provisioning feature is in the process of installing the target OS. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
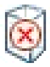	**Provisioning Failed – Unreachable** Indicates an available server on which an error occurred while the OS Provisioning Subsystem was installing a target OS, and that the server is not able to communicate with the SA core. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
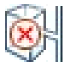	**Provisioning Failed – Reachable** Indicates an available server on which an error occurred while the OS Provisioning Subsystem was installing a target OS. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.

Table 5-1: Server Status Icons in SA (continued)

SERVER ICON	DESCRIPTION
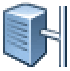	**Managed – Reachable** Indicates a server has an An Server Agent is running on it and that it is able to communicate with the SA core. In the SA Client, appears in the All Managed Servers list and Virtual Servers list. In the SAS Web Client, appears in the My Jobs panel in the home page, in the list in the My Jobs page, in the Manage Servers list, and in the server lists in the SA wizards.
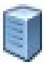	**Managed – Unreachable** Indicates a managed server cannot communicate with the SA core (it is Not Reachable). If you want to discover reasons why the managed server is unreachable, you can run a Communication Test. See “Agent Reachability Communication Tests” on page 129 in Chapter 4 for more information. In the SA Client, appears in the All Managed Servers list and Virtual Servers list. In the SAS Web Client, appears in the My Jobs panel in the home page, in the list in the My Jobs page, in the Manage Servers list, and in the server lists in the SA wizards.

Table 5-1: Server Status Icons in SA (continued)

SERVER ICON	DESCRIPTION
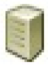	**Unmanaged** Indicates the server is unmanaged, which means it does not have an Server Agent installed on it. For unmanaged virtual servers, this means that someone created a VMware Virtual Machine (VM) or Solaris local zone outside of SA and thus does not have an Server Agent installed on it. For virtual servers, this state could also mean that the VMware VM is has not yet been provisioned, or that your user belongs to a group that does not have permissions to perform operations on this virtual server. For more information on installing an Server Agent on an unmanaged server, see Chapter 4, “Agent Management” on page 105 of this guide.
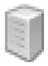	**Deactivated** Indicates a server that was deactivated in HP Server Automation so that it is currently not managed and is no longer reachable. Appears in the Manage Servers list and in the server lists in the SA wizards (however, it is not selectable in the wizards).
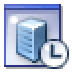	**Scheduled** Indicates a server that is scheduled for an operation (install software, uninstall software, and so forth). In the SA Client, appears in the Job Logs list. In the SAS Web Client, appears in the My Jobs panel in the home page and in the list in the My Jobs page.
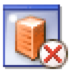	**Error** Indicates a managed server on which an error occurred while HP Server Automation was installing or uninstalling software. Appears in the My Jobs panel in the home page and in the list in the My Jobs page.

Table 5-1: Server Status Icons in SA (continued)

SERVER ICON	DESCRIPTION
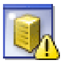	**Warning** Indicates a managed server on which a warning occurred while HP Server Automation was installing or uninstalling software. Appears in the My Jobs panel in the home page and in the list in the My Jobs page.
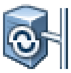	**Application Configuration Out of Sync** Indicates a managed server on which the configuration file on the server is out of sync with the Application Configuration Template (SA model). Appears only in the Application Configuration feature and the server list in the SA Client.
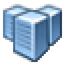	**Static Device Group** Indicates a static server group. The same states that apply to single servers apply to groups. See “Device Groups” on page 183 for more information about the different types of server groups.
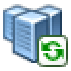	**Dynamic Device Group** Indicates a dynamic server group. The same states that apply to single servers apply to groups.
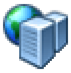	**Public Static Device Group** Indicates a public and static server group. The same states that apply to single servers apply to groups.
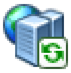	**Public Dynamic Device Group** Indicates a public and dynamic server group. The same states that apply to single servers apply to groups.

Ways to Use the Device Explorer

The HP Server Automation Client Device Explorer allows you to browse and manage servers, devices (such as network or storage), and groups of servers in your environment.

Using the Device Explorer you can perform the following actions on individual servers:

- Browse basic device system information, such as device, operating system, memory, Server Agent version, and more.
- View device compliance information and view the details of any policies attached to the server, such as all audits, patch policies, software policies, as well as any application configurations attached to the server.
- Run audits of the server, remediate any software policies attached to the server, and push application configurations on to a server
- Browse live and up to date information about a server’s file system, registry, hardware inventory, hardware, ethernet and SAN connections, installed software and patch lists, runtime state, user and user group membership, services, snapshots, and more
- View server group membership
- View virtual servers, either hypervisors or VMware VMs and Solaris local zones.
- Add and delete custom attributes.

Network and Storage Devices in the Device Explorer

For more information on the types of information you can view in the Device Explorer for network and storage devices, see the following topics:

- “Device Information in SA” on page 358 for network devices
- “Storage (ASAS-only)” on page 167
- The *Opsware® ASAS User’s Guide* for storage devices

Device Explorer Interface

The Device Explorer consists of two main sections: the Views pane and the Content pane. The Views pane lists server objects from the managed server, and the Content pane displays content for each of the server's objects. When you select a server object in the Views pane, its corresponding content appears in the Content pane. See Figure 5-1.

Figure 5-1: Device Explorer Interface

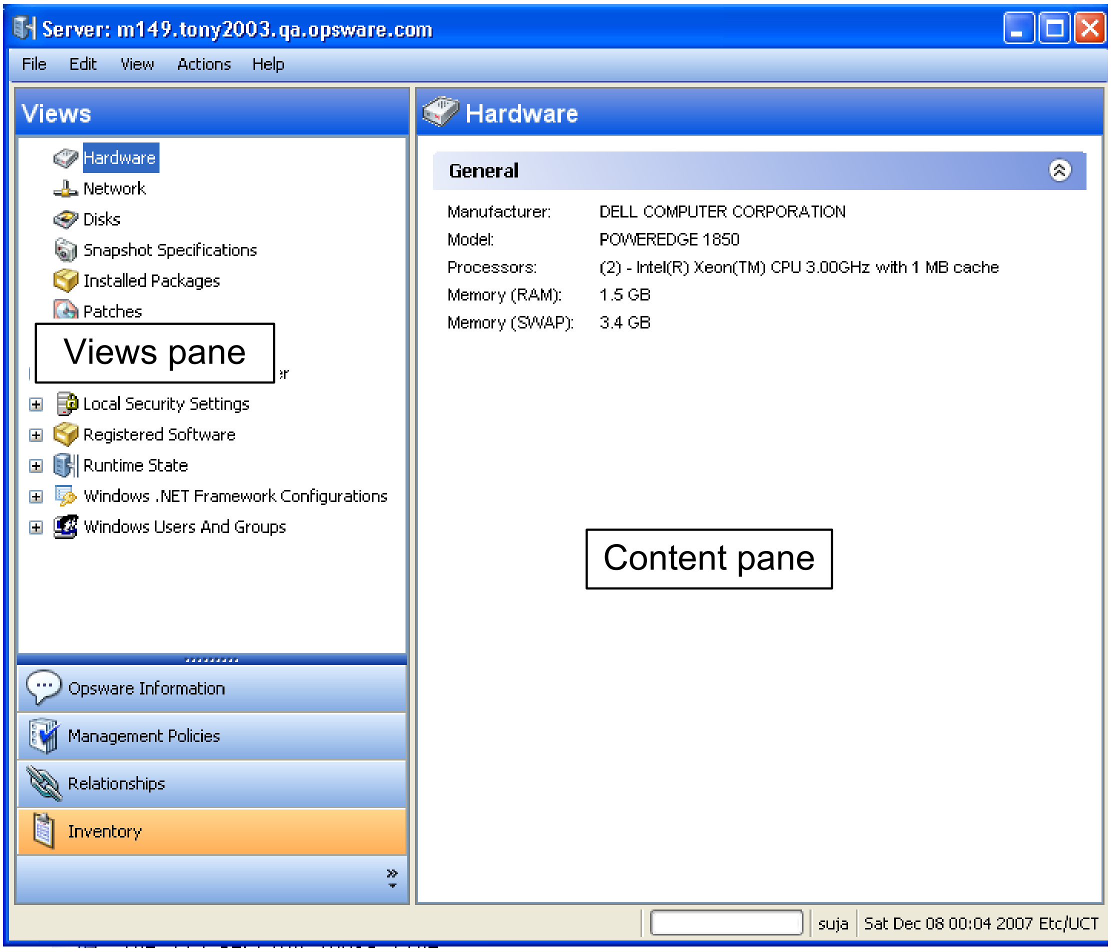

Manufacturer: DELL COMPUTER CORPORATION
Model: POWEREDGE 1850
Processors: (2) - Intel(R) Xeon(TM) CPU 3.00GHz with 1 MB cache
Memory (RAM): 1.5 GB
Memory (SWAP): 3.4 GB

Device Explorer Views Tabs

The Device Explorer Views tabs organize four different types of information about your device:

- **Opsware Information**: Shows general property and system information, such as computer manufacturer, hardware type, system, processor and memory, OS version, Server Agent version and status, SA customer assignment, history of changes to the server, and more.
- **Management Policies**: Displays a roll up of all compliance policies attached to the server, as well as compliance for individual compliance policies, such as audits, software and patch policies, application configurations, and any custom user-create policies. It also shows any custom attributes created on the server.
- **Relationships - Group Membership**: Shows all groups that the selected server is a member of and allows you to modify group membership (if your user has sufficient permissions).
- **Inventory**: Displays a list of live server configuration objects captured directly from the server, such as registered hardware, network connections, snapshots taken of the server, installed packages, patches, and software, runtime information about processes running on a server, local security settings, users and groups memberships, and more.

Accessing the Device Explorer

To access the Device Explorer, perform the following steps:

1. Launch the SA Client and then from the Navigation pane, select **Devices ➤ All Managed Servers**.
2. A list of servers will display in the Content pane.

   If the list of servers is long, use the search tool 🔍 to locate a server (upper right corner) by name, IP address, OS, customer, facility, or description. If you search by user name, the text entry is case insensitive.

   You can also sort the list by clicking a column heading, such as name, IP address, OS, customer, and so on. To reverse sort, click the column heading a second time.
3. Open a server from the Content pane. This opens the Device Explorer. From the **Actions** menu, you can perform many types of operations, such as:

- Open in Service Automation Visualizer (SAV) if your core is licensed to run SAV
- Run a script on the server
- Create or run an audit or snapshot of the server
- Scan software, application configuration, or patch compliance
- Add to a device group
- Export patch information to a .csv file
- And more

Action menu items change according to the server object selected.

For example, if you select the Configured Applications object from the server object tree, then from the **Actions** menu, you can add, remove, or open an application configuration, create a package, and so on.

Opening a Remote Terminal

You can open a remote terminal for any managed server, but not a group of servers. To do so, perform the following steps:

1 Launch the SA Client. From the Navigation pane, select **Devices ➤ All Managed Servers**.

2 Select a managed server and open it.

1. In the Device Explorer window, from the **Actions** menu, select **Launch Remote Terminal**.
2. Log in to the remote terminal.

See “Opening a Remote Terminal” on page 399 in Chapter 10 for more information.

Opsware Information

The Device Explorer – Opsware Information allows you to review the following information:

- Summary
- Properties
- Server History

Summary

The Summary view in the Device Explorer lists the following information:

- **System**: This displays operating system information.
- **Computer**: This displays server manufacturer, hardware, and system details.
- **Opsware Agent**: This displays communication status, the time when the server was last registered, the number of applications and patches registered with HP Server Automation, and so on.

Properties

The Properties view in the Device Explorer lists the following information for the server that you are viewing:

- Management Information
- Reported Information

Management Information

- **Name**: This displays the name of the managed server.
- **Notes**: This displays any notes listed.
- **IP Address**: This displays the IP address of the managed server.
- **OS Version**: This displays the operating system (platform) that the managed server is running on.
- **Customer**: This displays an account within HP Server Automation that has access to designated resources, such as servers and software.
- **Facility**: This displays the location of the server. Users can manage servers in any facility from an SAS Web Client.

- **Realm (link speed)**: This displays the minimum bandwidth limit between the Server Agent and the core (if the agent is going through gateways).
- **Server Use**: This displays how an organization is using the managed server; for example, a server could be a staging server, a production server, a development server, and so on.
- **Deployment Stage**: This displays the stages of deployment for a server; for example, a server could be live or offline.
- **Opsware Lifecycle**: This displays the server's stage in the managed server lifecycle; for example, unprovisioned, available, managed, or deactivated.
- **Server ID**: This displays the internal ID that SA uses to identify the server.
- **Status**: This displays whether or not the server is reachable and thus managed by SA. "OK" means that the server (its Server Agent) is reachable; unreachable means that there is a communication problem and SA cannot communicate with the server.

Reported Information

- **Reporting**: This displays information about the ability of the server's agent to communicate with the core. Statuses include Has not reported, OK, Registration in progress, and Reporting error.
- **Agent Version**: This displays the version number of the agent.
- **Hostname**: This displays the host name of the managed server.
- **Reported OS**: This displays the operating system (platform) that the managed server is running on.
- **MAC Address**: This displays the Media Access Control (MAC) address. This is the network interface card's unique hardware number. The MAC address is used as the server's physical address on the network.
- **Serial Number**: This displays the serial number of the system. HP Server Automation attempts to report a chassis ID if possible.
- **Chassis ID**: This displays a unique hardware-based identifier that the Server Agent discovers, typically derived from some property of the server's chassis. As a common source for this ID, HP Server Automation uses an interface's MAC address or the host ID on Solaris servers, or the serial number for one of the interfaces.
- **Encoding**: This displays the character encoding of the managed server, such as Shift_JIS (Japanese) or Windows 1252 (Western).

From this window, you can also open a remote terminal on the selected server.

Server History

The Server History view shows changes made to the selected server. For example, it displays who modified a server, what change was made, when it was modified, and so on. Server History specifically shows when a user has performed one of the following actions:

- Added the server to a group
- Removed the server from a group
- Reassigned the server from one group to another
- Login sessions
- Jobs that were run on the server, such as snapshots, audits, patch and software policy remediations
- And more

Entries are generated when actions are performed for managed servers in the SAS Web Client. The History is read-only. Double-click an entry to see more detailed information, such as:

- **Date**: The date when the last change occurred.
- **Device Name**: Name of the server or device where the change was made.
- **User**: The user who made the change.
- **Details**: A description of the change.

Use the View drop-down list to sort the server history list according to a range of time, such as last week, the last two months, and so on.

Management Policies

The Device Explorer – Management Policies allows you to review the following information:

- Compliance
- Audits
- Patch Policies - Windows Only

- Software Policies
- Configured Applications
- Custom Attributes

Compliance

The Compliance view of the Device Explorer displays overall compliance levels – a roll up of all compliance policies attached to the server – and compliance for individual compliance policies, such as audits, software and patch policies, application configurations, and any custom user-created policies.

You can select and expand a compliance category and view all tests in each category. For each test you can view policy details and remediate any tests that are out of compliance. Each compliance category contains an expandable list that contains all the policies of this type. The top-level node shows the rollup compliance status of all policies in this category. If you expand the list, you can see each individual policy and its compliance status as well, for a detailed breakdown of compliance for the server.

You can also sort the list to show all policies, or filter it to show only the policies for one compliance categories, such as all Audit policies and their compliance statuses. You can also sort by status filter, such as, show all compliance tests that are compliant, non-compliance, are currently scanning, and so on.

For more information on server and device group compliance, see "Server Compliance" on page 277.

Compliance Categories

Compliance categories a server include:

- **Audit**: A roll up compliance status of all scheduled audit that target this server appears by on the top node of the Audit category in the Details pane. This category displays the overall compliance status of all recurring audits that run on this server. To see the individual audits that use this server as a target, expand the Audit list, which shows each audit's compliance status in the Status column.
- **Software**: A roll up compliance status of all software policies attached to the server appears the top node of the Software category in the Details pane. Software compliance indicates whether or not all software policies attached to the selected server are compliant with the actual server configuration.

A software policy can include 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. To see the individual software policies attached to this server, expand the Software list, which shows each software policy's compliance status in the Status column.

- **App Config**: A roll up compliance status of all application configurations attached to the server appears the top node of the App Config category in the Details pane. An Application Configuration (App Config) policy defines how specific application configurations files should be configured 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 out of compliance. To see the individual application configurations attached to this server, expand the App Config list, which shows each application configuration s compliance status in the Status column.

  Patch (Windows only): A roll up compliance status of all patch policies attached to the server appears the top node of the Patch category in the Details pane. Patch compliance determines whether all patches in a patch policy and a patch policy exception were installed successfully. To test patch compliance, servers are scanned to determine whether they conform to their attached policies and exceptions, based on compliance status 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 out of compliance. To see the individual patch policies attached to this server, expand the Patch list, which shows each patch policy's compliance status in the Status column.

Audits

Audits view shows a list of all audits associated with the server, where the selected server is either the source or the target of an audit.

Use the following information for the selected server:

- **Audit - Server is Target**: Shows all audits where the selected server is the target of an audit.
- **Audit - Server is Source**: Shows all audits where the selected server is used as the source of an audit.

Archived Audit Results

This list displays all audits results associated with this server that have been deliberately archived by a user. In some cases, audits that run regularly can accumulate many audit results. In the main Audit and Remediation feature, you can select to archive audit results if you want to save them for later viewing. All audit results for the selected server are displayed here.

For more information, see "Audit and Remediation" on page 153.

Patch Policies - Windows Only

This window displays all patch policies associated with the selected Windows managed server (or server group).

Show Options

The Show drop-down list displays the following patch policy information:

- **Policies Attached to the Server**: This displays all policies attached to the Windows server, or policies attached to a server group to which the selected Windows managed server belongs.
- **Policies Not Attached to the Server**: This displays a list of all patch policies relevant to the selected server that are not attached to the server.

Software Policies

The Software Policies view displays all software policies associated with the selected server (or group of servers).

Show Options

You can use the Show drop-down list to filter the following types of software policies:

- **Policies Attached Servers**: This displays all policies attached to the server, or policies attached to a group of servers to which the selected managed server belongs.
- **Non-Compliant Policies**: This displays the software policies attached to the selected server which are non-compliant with the actual server configuration.
- **Policies which Require Scan**: This displays the software policies for which the software compliance information is not yet calculated.
- **Policies Currently Scanning**: This displays the software policies for which software compliance information is being calculated.

From the **Actions** menu, you can perform actions such as attaching a policy, detaching a policy, remediating a server, and scanning software compliance.

Configured Applications

The Installed Configurations tab allows you to browse and edit all Application Configurations attached to the managed server.

Each Application Configuration displays the following information:

- Configured Applications (Top Level)
- Application Instances (Folder Containing Children of Configured Applications)

Configured Applications (Top Level)

To view an application configuration, expand the Configured Application folders. The top level of the Configured Application hierarchy lists all application instances being managed by the Application Configuration Management System.

For each installed configuration, the following information appears:

- **Check Mark**: A check mark next to the name indicates that the Application Configuration (all the template files) has been checked as working.
- **Name**: The name of the application configuration.
- **Instance Name**: The name of the application configuration instance.
- **State**: If you add or delete more configurations to this server, this column displays added or removed until you click **Save Changes**. To add another application configuration to this server, from the **Actions** menu, choose **Add Configuration**.
- **Version**: The version number of the Application Configuration.
- **# Files**: The number of files the Application Configuration contains.
- **Last Modified**: The date when the Application Configuration was last modified.
- **Modified By**: The user who modified the Application Configuration.
- **Description**: The description from the Application Configuration's properties window.

Application Instances (Folder Containing Children of Configured Applications)

Underneath the Configured Application icon is a list of applications, and each application can contain one or more instances. Configuration values can be set at the two following levels:

- At the application level (folder), for all instances of the application.
- At the instance level, for individual application instances.

For each application instance, the following fields appear:

- **Template**: If the Application Configuration has more than one configuration template, you can select it from this list.
- **Filename**: The name of the configuration template file.
- **Encoding**: Choose a character encoding for the source configuration file that the Application Configuration will be managing.
- **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 to show a read-only view of all inherited values in the configuration template. The Source column will show where each key value pair in the template is inherited from in the inheritance hierarchy.

Value Set Editor

The Contents pane (right side of window) allows you to edit the key-value pairs for the configuration template:

- **Name**: The value set element name from the configuration file. Elements that are required appear in bold font.
- **Setting**: This allows you to either enter a literal value or choose an attribute from the Server's settings, such as the 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 or custom attribute for the value, click the ellipsis (...) button to access the Set Value dialog box.
- **Value**: The actual value that will be applied to the configuration file on the server.
- **Source**: This indicates the source of the value (where it is inherited from). The value can be applied at the server instance level or inherited from its ancestors (in descending order). See the following examples:

- Application configuration defaults ➤ Customer/Facility ➤ server group ➤ server group application defaults ➤ application instance

Or

- Application configuration defaults ➤ Customer/Facility ➤ server ➤ server application defaults ➤ application instance

These values can be edited and applied to the configuration file on the server by clicking **Push**. To preview the changes before you apply them, click **Preview**.

To select an application configuration template, select the drop-down list located in the upper right corner.

Run Script

The **Run Script** button executes the Application Configuration’s Data-manipulation script. If the application configuration contains no data-manipulation script, the **Run Script** button is disabled.

Preview

This launches a file differencing window that allows you to compare the contents between the Application Configuration and the actual configuration file on the managed server.

Push

This applies any changes made to the Application Configuration and also saves them.

Schedule

Launches the Schedule Job window so you can schedule the Application Configuration push to run at a later date and time or on a recurring basis.

Save Changes

This saves all changes made to the Application Configuration. Note that this does not apply changes made to the Application Configuration on the selected server.

Cancel Changes

This cancels any modifications without saving.

For more information, see “Application Configuration Management” on page 561.

Custom Attributes

This window displays the custom attributes set to a server or group of servers. You can add, edit or remove custom attributes from this window.

Custom attributes can be one of the following two types:

- Inherited from another source, such as a customer, a software policy, group of servers, ISM control, and so on.
- Attached directly to the server.

To override inherited custom attribute values, click the inherited arrow icon once, and enter a new value in the value field. Press ENTER. The inherited arrow changes to indicate that the custom attribute value has been overridden.

Relationships - Group Membership

The Relationships view of the Device Explorer shows all groups of which this server is a member. Members of a group can include servers, network devices (for NA-enabled cores), storage devices and assets (for ASAS-enabled cores), and other device groups.

You can select a group, right-click and select **Open** to view its contents. You can also select to Add to Group to select other devices to add to the group.

If the list of groups is long, use the search tool to locate a server (upper right corner) by name, description, access, type, modified by, and so on. If you search by device group name, the text entry is case insensitive.

Inventory

The Inventory tab shows gives you access to the following server objects discovered on the selected server:

- Hardware
- Network
- Storage (ASAS-only)
- Disks
- Virtualization
- Snapshot Specifications
- Installed Packages

- Patches
- Files
- Windows COM+ Objects
- Windows IIS Metabase
- Windows Registry
- Services
- Windows .Net Framework Configuration
- Internet Information Server
- Local Security Settings
- Registered Software
- Runtime State
- Users and Groups

Hardware

The Hardware view lists all the reported hardware on the selected managed server. This includes the following information:

- **Processors**: This lists the processor information for all processors on the managed server.
- **Memory**: This lists the total amount and the types of memory on the managed server.
- **Storage**: This lists all storage devices on the managed server.
- **Network Interfaces**: This lists the network interfaces on the managed server, including Ethernet cards (NICs) and any Fibre Channel Adapters (for ASAS-enabled SA cores), ports – including the switch ports that a port is connected to as well as any zones.

Network

The Network view shows all network connections (or SAN, for ASAS-enabled cores), providing such details as IO address, subnet mask, MAC address, duplex, and more for all network interfaces.

Storage (ASAS-only)

If your core is ASAS-enabled, the Inventory tab includes a Storage view that provides information about SAN, NAS and DAS (Direct Attached Storage) assets related to the selected server.

This view also provides a summary of the storage (if applicable) consumed by and allocated to the selected server, as well as insight into Database storage (if applicable) consumed by the selected server.

This Storage view provides details about the following storage assets:

- **File Systems**: Shows a list of local and remote file systems (SAN-based storage). It provides information such as mount location, type of file system, storage capacity and free space. From the View drop-down list, you can select four different views into File Systems:
  - **Properties**: Displays information like mount location, description for mount point; mount point, type of file system, storage capacity and free space.
  - **Volumes**: Displays a list of volumes based on the selected file system.
  - **Disks**: Displays a list of disks on which the selected file system is dependent.
  - **Connectivity**: Displays supply chain information for the selected file system in a tree format.
- **Volumes**: Shows a list of volumes consumed by the selected server. These volumes could be local or remote (SAN volumes). From the View drop-down list, you can select four different views into Volumes:
  - **Properties**: Displays the name, type, service type, status and storage capacity for the selected volume.
  - **Composition**: Displays *upstream* and *downstream* storage dependencies. Upstream storage dependency means that the storage asset depends on the selected volume; downstream storage dependency means that the selected volume is dependent upon other storage assets.
  - **Disks**: Displays the list of disks on which the selected volume is dependent.
  - **Access Path**: Displays data which is mostly interesting if the selected volume is a remote SAN volume. This sub-view provides LUN Mapping information for the remote SAN Volume – the target storage array, target storage array port, target storage volume, LUN number and initiator port.
  - **Connectivity**: Displays the supply chain information for a volume in a tree format.

This view is useful in the context of remote SAN volumes.

- **Unmounted Volumes**: Shows a list of volumes that are available to the selected server but that are not used by the server. These volumes are typically remote SAN volumes which are mapped to the selected server but the server is not using them. From the View drop-down list, you can select four different views into Unmounted Volumes:
  - **Properties**: Displays name, type (raid type), service type, status, storage capacity and target device for the selected Unmounted volume.
  - **Composition**: Displays *downstream* storage dependencies, which means that the unmounted volumes are dependent upon the selected storage asset.
  - **Disks**: Displays the list of disks on which the selected unmounted volume is dependent.
  - **Access Path**: Displays data which is mostly interesting if the selected unmounted volume is a remote SAN volume. This sub-view provides LUN Mapping information for the remote SAN Volume – the target storage array, target storage array port, target storage volume, LUN number and initiator port.
  - **Connectivity**: Displays the supply chain information for a unmounted volume in a tree format. This view is useful in the context of remote SAN volumes.
- **Disks**: Shows information about local and remote (SAN-based) disks. Detailed information includes name, manufacturer, model, device (server), storage capacity, status and if its spare or not. From the View drop-down list, you can select two different views into Disks Volumes:
  - **Volumes**: Displays list of volumes based on the selected disk.
  - **File Systems**: Displays a list of file systems based on the selected disk.
- **Manager Software**: Shows information about Volume Manager software and MultiPath software on the selected server, including vendor, version and details about logical volumes managed like name, type (RAID Type), service type, status, storage capacity and number of paths.

Disks

The Disks views provides local disk information for the managed server, including such information as disks names, manufacturer, the device that contains the disk, model number of the disk, its capacity, its status (identifies the disk health, such as OK, ONLINE, Disable, Not Ready, Error, READONLY), and whether the disk is used as a spare (Yes) or not used as a spare (No).

Virtualization

The Virtualization view displays information about virtual servers – both hypervisors and virtual guest servers. For example, for a hypervisor server, the Virtualization view displays all virtual servers being hosted on the server. For virtual servers, the Virtualization view displays all virtual server information as well as the hypervisor server name that is hosting it.

Solaris 10 – Device Explorer

You can use the Device Explorer to view a Solaris 10 global zone hypervisor and local zone server information. A global zone or local zone as seen through the Device Explorer looks nearly the same as a regular physical server, except that it has an extra property named “Virtualization,” which provides the following information:

- **Hypervisors**: When you view a Solaris 10 global zone hypervisor in the Device Explorer, the Virtualization view shows all hosted local zones. From here, you can also stop and start the local zones.
- **Solaris Local Zones**: When you view a Solaris local zone in the Device Explorer, the Virtualization view indicates the name of its hypervisor, its reserved CPU shares, and virtual hardware information.

VMware ESX 3 – Device Explorer

You can use the Device Explorer to view a VMware ESX VM or a hypervisor’s server information. A VM as seen through the Device Explorer looks nearly the same as a regular physical server, except that it has an extra property named “Virtualization,” which provides the following information:

- **Hypervisors**: When you view a VMware ESX hypervisor in the Device Explorer, you can see all of its hosted VMs.
- **VMware ESX VMs**: When you view a VMware ESX VM in the Device Explorer, the Virtualization view shows you the VM’s virtual server properties, its virtual network configuration, and its data store configuration.

For more information, see Chapter 6, “Virtualization Director” on page 209 of this guide.

Snapshot Specifications

The Snapshot Specifications view shows a list of all Snapshot Specifications where the selected server is listed as a target. To view the results of one of the Specifications, select it in the Contents pane (right side). When you select a Snapshot Specification, a list of all snapshot results appear in the pane below (for all servers that are targets of the Snapshot Specification).

To open a Snapshot Specification, select one, right-click, and select **Open** (or double-click it). To view the results of a snapshot, select one from the lower pane, right-click, and select **Open** (or, double-click it).

For more information, see "Snapshots" on page 255.

Installed Packages

The Installed Packages view enables you to view any installed packages on the selected managed server that are managed by the SA system. For each package, you can view name, type, size, last modified, and description. To sort the list by these categories, click the title of each column. See the *SA Policy Setter's Guide* on how to create a package.

For information about using the Device Explorer to see packages that exist on the manager server but are not yet managed by SA, see "Patches" on page 170.

Patches

The Patches view shows all patches related to the selected server that are recognized and registered by the SA system, you can use the Show drop-down list to filter the following types of patch information:

- **Patches Installed**: This option displays all patches that have been installed on the server.
- **Patches Recommended By Vendor**: This option displays all application and operating system patches that have been recommended by Microsoft (MBSA 2.0.1) for the selected server. If multiple patches have the same QNumber, Patch Management detects the application files that are already installed on a managed server and, subsequently, recommends the correct patch to install.
- **Patches with Policies or Exceptions**: This option displays patches in policies attached to the selected server, or patches that have always install exceptions, *and* have one of the following conditions:
  - The patches are not currently installed and are recommended by the vendor.

  - The patches are currently installed.

- **Patches Needed**: This option displays all patches that should be installed on the selected server but are not. These include patches that are in policies attached to that server, or patches that have always install exceptions, *and* are recommended by the vendor.
- **Patches with Exceptions**: This option displays all patches that have exceptions (such as always install or never install) a*nd* have one of the following conditions:
  - The patches are not currently installed and are recommended by the vendor.
  - The patches are currently installed.
- **All Patches**: This option displays all patches that are associated with the operating system of the server.

See the *SA User's Guide: Application Automation* for more information on Unix Patch Management and Windows Patch Management.

Files

The Files view enables you to browse the file system of a managed server. The File System has two main sections (similar to the Windows file system explorer): the server's directories and the contents of the selected directory.

The left side navigation panel of the Device Explorer shows all the directories of the selected server, and the right side of the Device Explorer lists the contents of the selected directory.

For each file, the SA Client lists the file's name, size, type, and date modified. To sort the files by any of these categories, click on the top of the column.

Depending upon your user permissions, you might not have access to a particular server's file system. In such a case, you cannot select and view the server's file system in the Device Explorer. If you have access to a server's file system, then you will see user names, such as Administrator, root, and Local System. These are user names used to access that server's file system as the selected user.

Viewing File Contents

To view file contents, perform the following steps:

1. Launch the SA Client. From the Navigation pane, select **Devices ➤ All Managed Servers**.
2. A list of servers will display in the Content pane. Select a server and open it. This opens the Device Explorer.
3. From the left side of the Device Explorer, select a File System object.
4. You are prompted to select a user name to log into the computer, such as Administrator, LocalSystem, or root. Select a user.
5. To view the contents of disk drives or folders, expand the icon. Select a directory.
6. From the **Actions** menu, select **View Contents**. The file content view pane appears at the bottom of the window.
7. To change the character encoding, select an item from the Encoding drop-down list.

Ways to Copy Files

You can copy files from a server to another directory on the same server, to a directory on another SA-managed server, or to your local computer (where the SA Client is running). You can also copy a file from your local computer to a directory on the managed server. A few restrictions apply when copying files on a managed server's file system using the Device Explorer:

- You cannot copy folders/directories.
- You can only copy to servers that you have permissions to write to and to view.
- You can only copy one file at a time.
- You cannot undo a deletion – once you delete a file, it's gone.

Copying Files Between Managed Servers

To copy files between managed servers, perform the following steps:

1. Launch the SA Client. From the Navigation pane, select the **Devices ➤ All Managed Servers**.
2. To launch the Device Explorer, open a server from the server list.
3. From the left side of the Device Explorer, select a File System object. To view the contents of disk drives or folders, expand the icon.
4. Navigate to the directory that contains the file that you want to copy and select it.
5. From the **Actions** menu, select **Copy To**.

6 In the Copy To dialog box, select from the following locations in the drop-down list:

- Managed Servers (other managed servers in the core)
- This Server
- Local File System

7 Navigate to the desired directory.

8 Click **Select**.

Copying Files from Your Computer to a Managed Server

To copy files from your computer to a managed server, perform the following steps:

1 Launch the SA Client. From the Navigation pane, select **Devices ➤ All Managed Servers**.

2 To launch the Device Explorer, open a server from the server list.

3 From the left side of the Device Explorer, select a File System object.

4 Navigate to the target directory where you want to copy the file.

5 Use your system's file system explorer to select the file that you want to copy, then drag the file to the desired location in the Device Explorer.

Deleting Files

Once you delete a file, it cannot be recovered. (However, before you delete, you are prompted with a confirmation dialog box.)

To delete a file, perform the following steps:

1 Launch the SA Client. From the Navigation pane, select the **Devices ➤ All Managed Servers**.

2 To launch the Device Explorer, open a server from the server list.

3 From the left side of the Device Explorer, select a File System object.

4 Select a file to delete from the Content pane.

5 From the **Actions** menu, select **Delete**.

6 Click **Yes** in the confirmation dialog box.

Renaming Files

To rename a file, perform the following steps:

1. Launch the SA Client. From the Navigation pane, select the **Devices ➤ All Managed Servers**.
2. To launch the Device Explorer, open a server from the server list.
3. From the left side of the Device Explorer, select the File System object. To view the contents of a folder, expand the folder.
4. Select the file that you want to rename, and from the **Actions** menu, select **Rename**.
5. Enter a new name for the file, then press ENTER. Pressing the ESC key on your keyboard will cancel the rename operation.

Creating a Configuration Template from a File

For any file on a managed server, you can create an configuration template.

To create an configuration template from a file, select the file. From the **Actions** Menu, select **Create Configuration Template**. See *SA User's Guide: Application Automation* for more information.

Creating a Package from a File

For any file on a managed server, you can create an installable software package. For each package, you can specify the customer assignment, the reboot requirements, and the pre/post install and pre/post uninstall scripts.

To create a package from a file on the managed server file system, select the file. From the **Actions** menu, select **Create Package**. See *SA Policy Setter's Guide* on how to create a package.

Windows COM+ Objects

This window displays a read-only view of all the COM+ objects on the selected managed server. In the Server Explorer window, the Views pane displays the following two folders for browsing COM+ objects:

- **All Objects**: This is a flat list of all the COM+ objects on the managed server.
- **Component Categories**: This contains an alphabetical list of all COM component categories.

To view the contents of a COM+ object, perform the following steps:

1. Select the All Objects or the Component Categories folder. In the Content pane, expand the folder until you reach an object.
2. To view the contents of a COM+ object, from the **Actions** menu, select **View Contents**. The contents will then display.
3. If the content of the COM+ object uses a different encoding, choose the appropriate encoding type from the Encoding drop-down list.

You must have specific user permissions to view Windows COM+ objects. If you are unable to access the Windows Registry, contact your SA administrator.

Windows IIS Metabase

This window displays a read-only view of the IIS Metabase on the selected Windows managed server. You can use this window to browse the IIS Metabase much like one of the metabase browsing tools such as metaedit or the IIS Metabase Browser.

The left side of the Metabase window displays the hierarchical layout of the metabase tree. Selecting an item in the tree on the left shows the data items associated with the selected key in the right hand view. Clicking the (+) symbol to the left of a key item will expand the item's child keys.

To view Windows IIS Metabase items, select the top-level Windows Metabase icon in the Server Explorer, right-click, and select a user. Your user must have permissions to view Windows Metabase items. If your user is unable to access the Windows Registry, contact your SA Administrator.

Windows Registry

This window displays a read-only view of the Windows registry on the selected Windows managed server. You can navigate to this registry much like the regedit tool on the Windows operating system.

Folders on the left side of the window represent keys in the registry. Clicking a folder on the left displays entries in a key in the right window.

To view Windows Registry items, select the top-level Windows Registry icon in the Server Explorer and select a user from the menu. Your user must have proper permissions to view Windows Registry keys. If your user is unable to access the Windows Registry for the selected managed server, contact your SA Administrator.

The HKEY_CLASSES_ROOT might have thousands of entries and can take time to load.

Services

The Services window shows you a list of all running services on the selected managed server. Depending on the installed operating system, you can perform different operations on the services:

- For Windows services, you can start, stop, pause, resume, and restart a service. You can also set the service to start manually, to start automatically when the system is rebooted, or to be disabled altogether.
- For Linux servers (supported by Red Hat and SuSE versions), you can perform any action that a particular service supports. Supported actions may vary from service to service, for example, start, stop, restart, condrestart, or status. You can also specify the run levels that you want a service to run under.

To perform an operation on a service, select the service and right-click.

Windows .Net Framework Configuration

The Windows .Framework Configuration window allows you to view real time information about Assembly Cache and Configured Assembly List for a Windows server.

For each Assembly Cache you can view information such as Assembly Name, Version, Locale, Public Key Token, Cache file (GAC or ZAP), Processor Architecture, Custom, and File name.

For every Configured Assembly List, you can view information such as Assembly Name, Public key token, Codebases, Binding policy, File name, File data.

To view the information for the Windows .Framework Configuration server object, you will need the appropriate permissions. See the *SA Administration Guide* for information about the permissions required for the server object.

When you are accessing the Windows .Framework Configuration for the first time, it may take a few minutes to load the server object. Subsequent usage of the server object will be significantly faster.

In the SA Client you can manage Windows .Framework Configuration information by adding it to a software policy or audit.

- Using the Software Management feature, you can deploy the Windows .Framework Configuration on a managed server. See the *SA User's Guide: Application Automation* for information about the Software Management feature.
- Using the Audit and Remediation feature you can specify the audit rules for assembly cache and configured assemble list and then remediate any differences found between the target server and the audit rule. See the *SA User's Guide: Application Automation* for information about Audit and Remediation.

Internet Information Server

The Internet Information Service (IIS) window allows you to view the real time information about IIS for a Windows server. For every Windows server you can view the information such as Server name, Server type, Server state, Log file path, and Document file path.

To view the information for the Internet Information Service server object, you will need the appropriate permissions. See the *SA Administration Guide* for information about the permissions required for the server object.

When you access the Internet Information Server for the first time, it may take a few minutes to load the server object. Subsequent usage of the server object will be significantly faster.

Using the Audit and Remediation feature you can specify audit rules for the Internet Information Service to compare Internet Information Service configurations against a baseline server, or user-defined values, or a server snapshot. See the *SA User's Guide: Application Automation* for information about Audit and Remediation.

Local Security Settings

The Local Security Settings window allows you to view the real time information about security settings for a Windows managed server. For every Windows server you can view security settings such as Password policy, Audit policy, User rights, and Security options.

To view the information for the Local Security Settings server object, you will need the appropriate permissions. See the *SA Administration Guide* for information about the permissions required for the server object.

When you access the Local Security Settings for the first time, it may take a few minutes minutes to load the server object. Subsequent usage of the server object will be significantly faster.

In the SA Client you can manage the Local Security Settings information by adding it to a software policy or audit.

- Using the Software Management feature, you can deploy the Local Security Settings on a managed server. See the *SA User's Guide: Application Automation* for information about the Software Management feature.
- Using the Audit and Remediation feature you can specify the audit rules for the Local Security Settings and then remediate any differences found between the target server and the audit rule. See the *SA User's Guide: Application Automation* for information about Audit and Remediation.

Registered Software

The Registered Software view allows you to view real time information of all the packages and patches installed on a managed server. For each package or patch you can view information such as Name, Version, Release, Unit Type, and Installed Unit.

To view the information for Registered Software, you will need the appropriate permissions. See the *SA Administration Guide* for information about the permissions required for the server object.

When you access Registered Software for the first time, it may take a few minutes minutes to load the server object. Subsequent usage of the server object will be significantly faster.

Using the Audit and Remediation feature you can specify the audit rules for Registered Software to compare package and patch configurations against a baseline server, or user-defined values, or a server snapshot. See the *SA User's Guide: Application Automation* for information about Audit and Remediation.

Runtime State

The Runtime State window allows you to view real time information about the run time data for a managed server. It provides information about the DNS servers, Routes, and Processes for every managed server.

To view the information for the Runtime State server object, you will need the appropriate permissions. See the *SA Administration Guide* for information about the permissions required for the server object.

When you access the Local Security Settings for the first time, it may take a few minutes minutes to load the server object. Subsequent usage of the server object will be significantly faster.

Using the Audit and Remediation feature you can specify the audit rules for the Runtime State to compare Runtime state configurations against a baseline server, or user-defined values, or a server snapshot See the *SA User's Guide: Application Automation* for information about Audit and Remediation.

Users and Groups

The Windows Users and Group window allows you to browse and manage users and groups information on a Windows server. The Unix Users and Group window allows you to browse and manage users and groups information on a Unix server.

For every SA User you can view information such as Name, Description, Country code, Home directory, Password, Number of Logons, and Last logoff and logon time.

For every SA User Group you can view information such as Group Name and Description.

To view the information for the User’s and Groups server object, you will need the appropriate permissions. See the *SA Administration Guide* for information about the permissions required for the server object.

When you access the Users and Groups for the first time, it may take a few minutes minutes to load the server object. Subsequent usage of the server object will be significantly faster.

In the SA Client you can manage the Users and Group information by adding it to a software policy or audit.

- Using the Software Management feature, you can deploy the Users and Groups on a managed server. See the *SA User’s Guide: Application Automation* for information about the Software Management feature.
- Using the Audit and Remediation feature you can specify the audit rules for the Users and Groups and then remediate any differences found between the target server and the audit rule. See the *SA User’s Guide: Application Automation* for information about Audit and Remediation.

Basic Server Management Tasks

You can perform the following basic server management tasks in the SA Client:

- Refreshing Server Status
- Deactivating a Server
- Rebooting a Server
- Opening a Remote Terminal

Refreshing Server Status

Refresh a server to see if anything has changed on the server since you last looked at it in the SA Client. Refreshing a server retrieves the latest server information from the model repository and displays it. Refreshing a server’s status is good idea from time to time to make sure you are looking at current data on the server.

For more information on server statuses, see “Ways to Use the Device Explorer” on page 152.

To refresh a server's status, perform the following steps:

1. Launch the SA Client. From the Navigation pane, select **Devices**.
2. Select one of the server categories, such as Device Groups, All Managed Servers, Unprovisioned Servers, or Virtual Servers. (You cannot refresh server status of unmanaged servers or a device group.)
3. Select a single server or multiple servers (Shift + select).
4. Right-click and select **Refresh Server Status**.

Deactivating a Server

Deactivating a server removes the server from management by SA. You might want to deactivate a server, for example, if you are moving the server to a warehouse for storage. Or, you might want to deactivate a server when you need to rebuild it from scratch, without using the OS Provisioning feature.

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

After you deactivate a server, you can reactivate it by re-installing an Server Agent with the Server Agent Installer and the --clean command line option. See Agent Reachability Communication Tests for more information.

Deactivating a server accomplishes the following tasks:

- Removes custom attributes from the server.
- Deletes any configuration tracking policies from the server that are associated with backups.
- Sets the server life cycle value to Deactivated.

To deactivate a server, perform the following steps:

1. Launch the SA Client. From the Navigation pane, select the **Devices ➤ All Managed Servers**.
2. Select a server, right-click, and select **Deactivate Server**.

Rebooting a Server

You can reboot a single server or a group of server immediately, or schedule the reboot for a later time. If you choose to reboot a group of servers, all servers contained in the group will be rebooted.

In order to be able to reboot a server, your user needs to belong to a user group that has the Reboot Server permission. For more information, contact your SA Administrator.

If you are rebooting a hypervisor server that is hosting other virtual servers, then all virtual servers being hosted by that hypervisor will be shut down as well. Whether or not the hosted virtual servers get rebooted depends upon the individual virtual server's configuration.

To reboot a server, perform the following steps:

1. Launch the SA Client. From the Navigation pane, select **Devices ➤ All Managed Servers** or **➤ Device Groups**.
2. Select a server or group of servers, right-click, and select **Reboot Server**.
3. In the Reboot Server window, step one lists the server or servers you have selected to reboot. Click **Next**.
4. In the Scheduling page, choose if you want to reboot the server or group of servers immediately, or at some later time and date. To run the reboot at a later time, select Analyze and Run Task At, and then choose a day and time.
5. Click **Next**.
6. In the Notifications page, by default your user will not have a notification email sent when the reboot finishes, whether or not the reboot job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
7. (Optional) You can specific if you want the email to be sent upon success of the reboot job (✓) and/or failure of the reboot job (✗).
8. (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when SA Professional Services has integrated SA with your change control systems. It should be left blank otherwise.
9. Click **Next**.
10. In the Summary View page, click **Start Job** to reboot the selected server or group of servers. When the job has run, click **View Results** to view the results of the reboot job.

11 In the Job Status page, you can see the progress of the job if you ran the job immediately. If the job is scheduled to run, you can close the window, and to view the job details, from the left side of the SA Client, select **Jobs and Sessions ➤ Job Logs**.

SA Tasks that Reboot a Server

There are a few other SA tasks that will initiate a server reboot, depending on the options set in the task:

- Installing or uninstalling a patch on Windows or UNIX. See “Setting Reboot Options for Remediation” on page 379 (Windows) or “Setting Reboot Options for a Unix Patch Installation” on page 445 for more information.
- Remediating a Patch policy. See “Remediating Patch Policies” on page 377 for more information.
- Installing a package and remediating a software policy. See “Specifying Options for Remediation” on page 477 for more information.

Opening a Remote Terminal

You can open a remote terminal for any managed server, but not a group of servers. To do so, perform the following steps:

1 Launch the SA Client. From the Navigation pane, select **Devices ➤ All Managed Servers**.

2 Select a managed server and open it.

   1. In the Device Explorer window, from the **Actions** menu, select **Launch Remote Terminal**.
   2. Log in to the remote terminal.

Device Groups

This section contains information about managing device groups in the SA Client. See “Device Groups” on page 183 in Chapter 5 for information about managing groups in the SAS Web Client. For information about storage device groups in the ASAS Client, see the *ASAS User’s Guide*.

The Device Groups feature provides a useful way for gathering servers into collections or organize groups of servers. Grouping servers enables you to perform the same action such as installing patches, remediating servers on all of the servers simultaneously, instead of performing the action on individual servers, one at a time.

Device groups can consist of individual servers as well as other device groups.

In SA, some of the recommended ways of grouping servers include:

- Grouping servers by OS version
- Grouping servers by customer
- Grouping servers by facility
- Grouping servers by deployment stage
- Grouping servers by use
- Group servers by virtual technology
- Grouping servers by operational boundaries, for example, grouping together all servers that require identical application configuration
- Grouping servers to control access, for example, creating device groups that are associated with a specific user group

A device group acts as a container for a collection of servers. Operations that can be performed on a server can be performed on a device group. As a result, a device group enables you to perform an action on multiple servers and avoid repeating the same operation on each individual server. When any of the operations, such as installing software or patch and configuring application configurations, are performed on a device group, they are actually performed on the servers within the group, and not on the group itself.

Characteristics of Device Groups

Device groups have the following characteristics:

- Individual servers can be included in many device groups or not included in any device groups.
- Adding servers to a device group does not remove those servers from the list of all servers that appear in the All Managed Servers list in the SA Client.
- Device groups can contain servers and other device groups.

- Device groups are hierarchical (they can be nested) with the following caveats:
  - Private and public groups cannot be mixed in a hierarchy, but static and dynamic groups can.

    See “Types of Device Groups” on page 185 in this chapter for more information about private and public groups.
  - The rules for a dynamic group are not inherited from a parent dynamic group to a dynamic subgroup.

    See “Dynamic Device Groups” on page 186 in this chapter for more information about the characteristics of dynamic groups.
  - Groups do not inherit custom attributes from their parents.
  - When you run an operation on a parent device group that contains child subgroups, the operation also applies to all the servers in the child groups below the current parent device group. For example, remediating a parent device group remediates all servers in the subgroups, but the child subgroups do not inherit the software policy attached to the parent device group.
  - When an Application Configuration operation within the SA Client is applied to groups, and those groups contain subgroups, the operation does not apply to all the servers in the subgroups. It only applies to the group upon which the operation was directly applied.

Types of Device Groups

There are two types of device groups, private groups and public groups, and each can be either static or dynamic.

Public and Private Device Groups

Public Device Groups

Public device groups can be created, edited, or deleted by any user who has Manage Public Device Groups permissions. Public groups are visible to all users, and can be used by any SA user. Only users with Manage Public Device Group permission can add members to a static public device group or change the rules that govern the dynamic public device group. Public device groups can also be used for modeling.

Accessing servers in a public device groups also depends on the Device Group permission in the SAS Web Client. See *SA Administration Guide* for more information about setting device group permissions.

Public Device Group Modeling

With SA modeling, the desired state of a server is defined and then applied to servers. In the case of public device groups (static and dynamic), you can define a model consisting of application configurations, patch policies, software policies, and custom attributes, which will be applied to all servers in the group. The modeling information is attached to the group, but not to any subgroups.

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

Private Device Groups

Private device groups can be created by any user belonging to a user group that has access to the Manage Servers list. Only the user who creates the private device groups can see and manage them. These groups are not visible to the other SA users. Private groups behave the same way as public groups, with the exception that modeling is not available for private groups. See “Public Device Groups” on page 185 in this chapter for more information.

Static and Dynamic Device Groups

Static Device Groups

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

Dynamic Device Groups

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

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

HP Server Automation calculates server group membership each time any of the following actions occur:

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

Additionally, HP Server Automation automatically recalculates dynamic group membership every hour.

Ways to Create Device Groups

In the SA Client, you can create a device group from the Manage Servers list or by performing a server search, and saving the resulting list of servers or the rules as a group.

You can create the following types of device groups:

- Creating a Static Device Group
- Creating a Dynamic Device Group
- Creating a Device Group Using Search

Creating a Static Device Group

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

A static group without servers added can be converted to a dynamic group. Conversely, a static group that has servers cannot be converted to a dynamic group.

To create a public static device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

To create a static device group, perform the following steps:

1. From the navigation pane, select **Devices ➤ Device Groups**. The list of device groups appears in the Content pane.

2. To create a public static device group, select public group and then from the **Actions** menu, select **New ➤ Static Group**.

   Or

   To create a private static device group, select private group and then from the **Actions** menu, select **New ➤ Static Group**.

   The name of the device group that you just created is New Device Group (n), where n is a number based on the number of new device groups already in existence.

3. Enter the name of the device group in the Content pane

4. Press the Enter key to save the device group.

To create a device group in a specific location, navigate to the desired location in the device group hierarchy and then select **New** from the **Actions** menu.

After you create a static device group, you can add servers to that group. See "Adding a Server to a Static Device Group" on page 192 in this chapter for more information.

Creating a Dynamic Device Group

Dynamic device groups are rule-based, and the servers in dynamic groups will be added or removed automatically based on the rules defined by you.

When a dynamic device group is converted to a static device group, all the servers remain in the static device group, but the rules used to define the server membership are lost.

To create a public dynamic device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

To create a dynamic group, perform the following steps:

1. From the Navigation pane, select **Devices ➤ Device Groups**. The list of device groups appears in the Content pane.

2. To create a public dynamic device group, select public group and then from the **Actions** menu, select **New ➤Dynamic Group**.

Or

To create a private dynamic device group, select private group and then from the **Actions** menu, select **New ➤ Dynamic Group**.

The name of the device group that you just created is New Device Group (n), where n is a number based on the number of new device groups already in existence.

3. Enter the name of the device group in the Content pane.
4. Press the Enter key to save the device group.

To create a device group in a specific location, navigate to the desired location in the device group hierarchy and then select **New** from the **Actions** menu.

After you create a dynamic device group, you can define the rules for the group. See "Adding a Server to a Dynamic Device Group" on page 194 in this chapter for more information.

When a server matches the dynamic server group criteria, the server will appears quickly in the server group browser, under the server membership tab. However, there is a delay for the server to appear until the cache is reloaded in the following locations:

– From Devices ➤ Device groups ➤ Public, then clicking on the group ➤ Preview pane (Members)

– From double-clicking on a device group

From selecting the Device Group Explorer window ➤ Summary tab (members)

Creating a Device Group Using Search

To create a public device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

To create a device group using search, perform the following steps:

1. From the Navigation pane, select **Search**.
2. Click **Advanced Search**. The Advanced Search page appears in the Content pane.

3. From the first drop-down list, select Server.
4. Create a rule by selecting the attribute from the second drop-down list. Depending on the attribute that you select, options available for the operator and values for the rule will change.
5. Select the operator from the third drop-down list. The operator selected defines how the search text is treated.
6. Enter a value in the field or select a value from the drop-down list or click to select multiple values from the Select Values window.
7. Click ⊞ to add additional rules and repeat steps 4 to 6. Click ⊟ to delete any rules.
8. Select the logic (And/Or) to be applied for every rule in the query.
9. Click **Search** to run the search query. The search results appear in the Content pane.

10 Click **Save** to save your search query. The Save Search window appears as shown in the figure below. In the Search Type drop-down list, select Dynamic Device Group or Static Device Group. Specify the location and then enter the name of the device group in the Save As text box and click **Save**. The name of the saved search cannot be more than 64 characters.

Figure 5-2: The Save Search Window in the SA Client

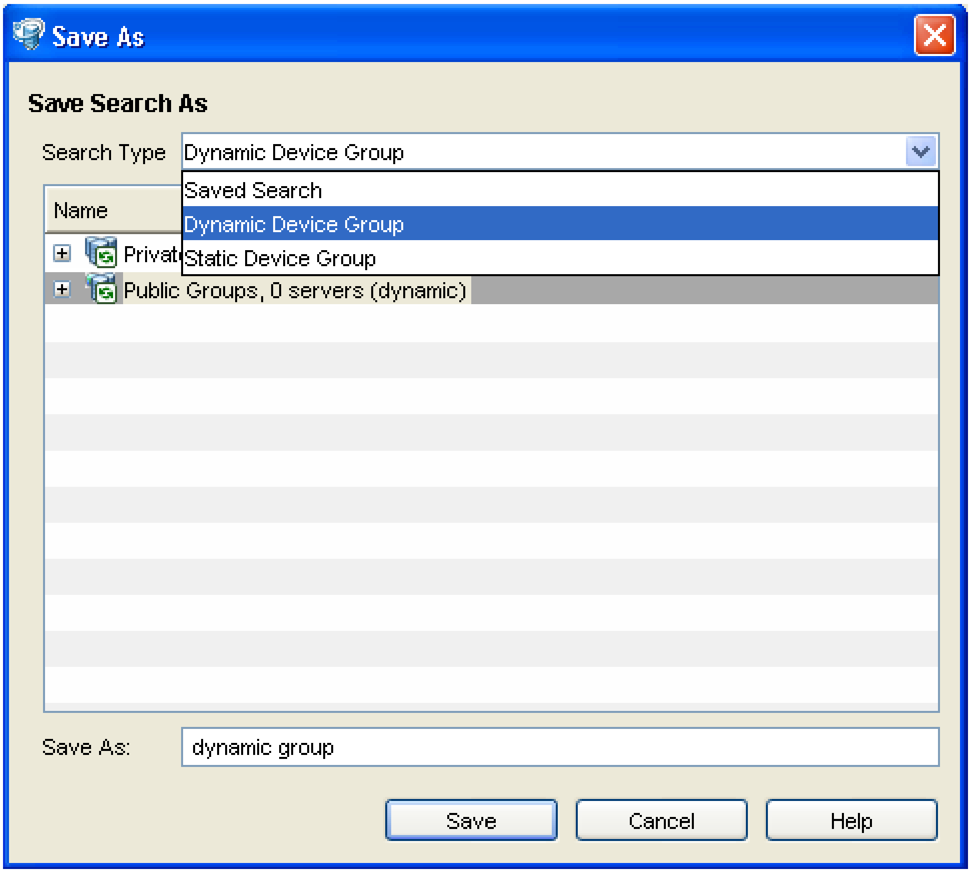

When you save the device group as a dynamic group, if the rules change or the servers in the environment change, servers will be added to or removed from the device group automatically. When you save the device group as a static group, the servers will be added to the device group, but all rules will be lost.

If the search query does not return any results, then you can only save the device group as a dynamic device group.

Adding a Server to a Static Device Group

To add a server to a public static device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

Method I

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

1. From the Navigation pane, select **Devices ➤ Device Group**. The device groups appear in the Content pane.
2. Select a static device group and from the **Actions** menu, select **Open**. The Device Group Explorer appears.
3. From the Views pane, select Server Membership.

4 From the **Actions** menu, select **Add Members**. The Add Members to Static Group window appears.

Figure 5-3: The Add Members to Static Group Window in the SA Client

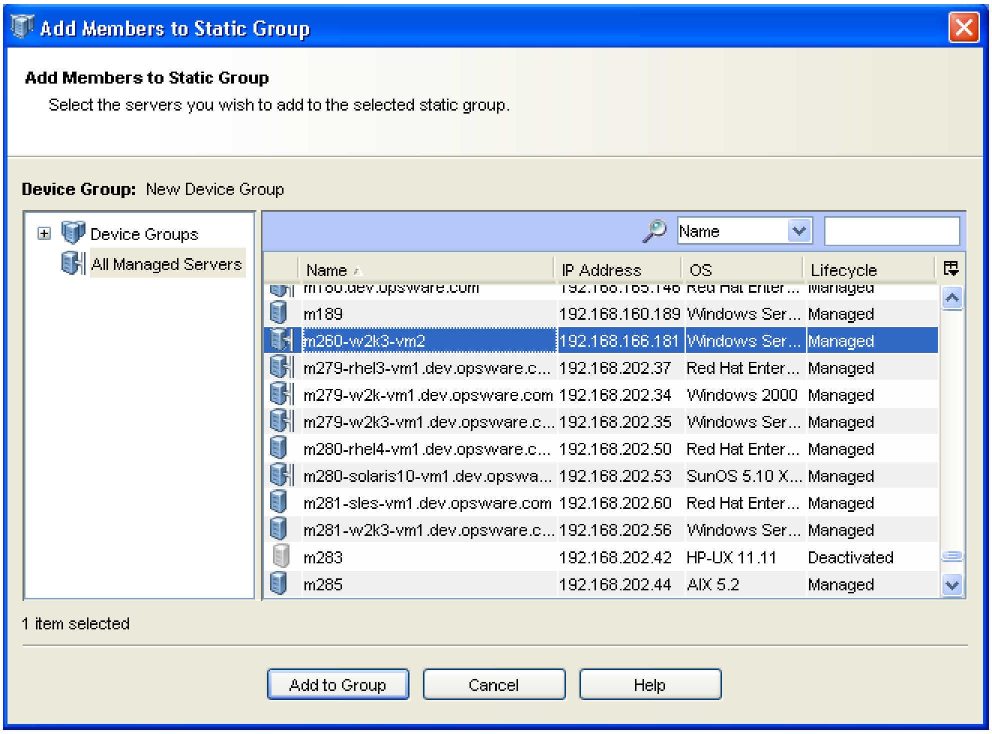

Add Members to Static Group

Select the servers you wish to add to the selected static group.

Device Group: New Device Group

Name	IP Address	OS	Lifecycle
m180.dev.opsware.com	192.168.165.146	Red Hat Enter...	Managed
m189	192.168.160.189	Windows Ser...	Managed
m260-w2k3-vm2	192.168.166.181	Windows Ser...	Managed
m279-rhel3-vm1.dev.opsware.c...	192.168.202.37	Red Hat Enter...	Managed
m279-w2k-vm1.dev.opsware.com	192.168.202.34	Windows 2000	Managed
m279-w2k3-vm1.dev.opsware.c...	192.168.202.35	Windows Ser...	Managed
m280-rhel4-vm1.dev.opsware.c...	192.168.202.50	Red Hat Enter...	Managed
m280-solaris10-vm1.dev.opswa...	192.168.202.53	SunOS 5.10 X...	Managed
m281-sles-vm1.dev.opsware.com	192.168.202.60	Red Hat Enter...	Managed
m281-w2k3-vm1.dev.opsware.c...	192.168.202.56	Windows Ser...	Managed
m283	192.168.202.42	HP-UX 11.11	Deactivated
m285	192.168.202.44	AIX 5.2	Managed

1 item selected

5 Select the servers or device groups to add to the static device group.

6 Click **Add to Group**. The selected severs or device groups appear in the Device Group Browser.

7 From the **File** menu, select **Save** to save the device group.

Method II

Perform the following steps to add a server to a static device group:

1 From the Navigation pane, select **Devices ➤ All Managed Servers**. The list of managed servers appears in the Content pane.

2 From the Content pane, select the servers and then from the **Actions** menu, select **Add to Device Group**. The Add to Group window appears.

3 Select the static device group to add the servers and then select Add to Group. The selected servers are added to the static device group.

Adding a Server to a Dynamic Device Group

Servers are added to dynamic device groups automatically, based on the rules created for the group. You can change the membership of a dynamic group by adding, deleting, or updating the dynamic group rules.

To add a server to a dynamic public device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

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

1. From the Navigation pane, select **Devices** ➤ **Device Groups**. The list of device groups appear in the Content pane.
2. From the Content pane, select the device group and then from the **Actions** menu, select **Open**. The Device Group Explorer appears.

3 From the Views pane, select Server Membership. In the Content pane, you can specify the rules for the dynamic device group. See Figure 5-4.

Figure 5-4: Defining Rules for a Dynamic Device Group

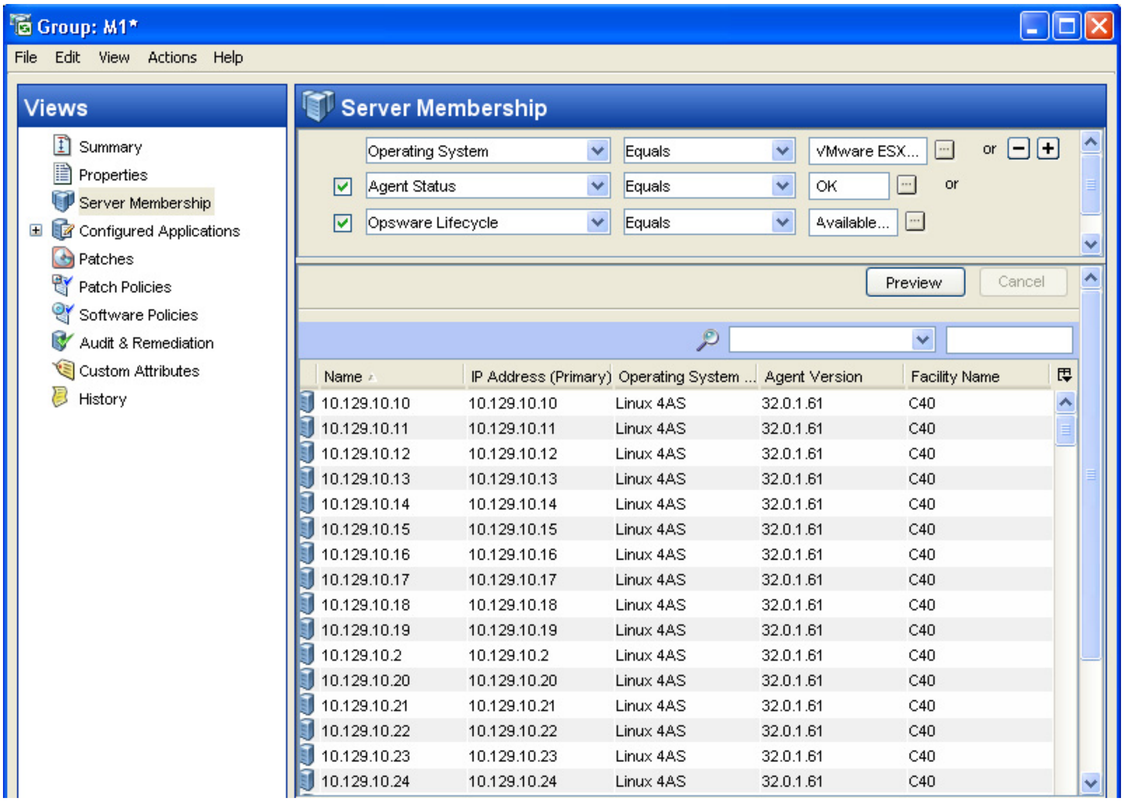

	Operating System	Equals	VMware ESX...	or
☑	Agent Status	Equals	OK	or
☑	Opsware Lifecycle	Equals	Available...	

Name	IP Address (Primary)	Operating System ...	Agent Version	Facility Name
10.129.10.10	10.129.10.10	Linux 4AS	32.0.1.61	C40
10.129.10.11	10.129.10.11	Linux 4AS	32.0.1.61	C40
10.129.10.12	10.129.10.12	Linux 4AS	32.0.1.61	C40
10.129.10.13	10.129.10.13	Linux 4AS	32.0.1.61	C40
10.129.10.14	10.129.10.14	Linux 4AS	32.0.1.61	C40
10.129.10.15	10.129.10.15	Linux 4AS	32.0.1.61	C40
10.129.10.16	10.129.10.16	Linux 4AS	32.0.1.61	C40
10.129.10.17	10.129.10.17	Linux 4AS	32.0.1.61	C40
10.129.10.18	10.129.10.18	Linux 4AS	32.0.1.61	C40
10.129.10.19	10.129.10.19	Linux 4AS	32.0.1.61	C40
10.129.10.2	10.129.10.2	Linux 4AS	32.0.1.61	C40
10.129.10.20	10.129.10.20	Linux 4AS	32.0.1.61	C40
10.129.10.21	10.129.10.21	Linux 4AS	32.0.1.61	C40
10.129.10.22	10.129.10.22	Linux 4AS	32.0.1.61	C40
10.129.10.23	10.129.10.23	Linux 4AS	32.0.1.61	C40
10.129.10.24	10.129.10.24	Linux 4AS	32.0.1.61	C40

4 Create a rule by selecting the attribute from the first drop-down list. Depending on the attribute that you select, options available for the operator and values for the rule will change.

5 Select the operator from the second drop-down list.

6 Enter a value in the field, or select a value from the drop-down list, or click ⋯ to select multiple values from the Select Values window.

7 Click ⊞ to add additional rules and repeat steps 4 to 6. Click ⊟ to delete any rules.

8 Select the logic (And/Or) to be applied for every rule in the query.

9 Click **Preview** to view the servers, which are members of the device group.

10 From the **File** menu, click save to **Save** the rules.

Removing Servers from a Static Device Groups

In a static device group, servers have to be removed manually. Removing a server only removes the server from the device group and not from the Manage Server list. Servers can belong to more than one device group, so if you want to remove a server from each device group that it belongs to, then you must locate and remove each instance of the server from all the device groups.

To remove servers from a public device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

Method I

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

1 From the Navigation pane, select **Devices ➤ Device Groups**. The list of device groups appears in the Content pane.

2 Select a static device group and from the **Actions** menu, select **Open**. The Device Group Explorer appears.

3 From the Views pane, select Server Membership. The list of servers and device groups in the device group appear in the Content pane.

4 Select the server or device group. From the **Actions** menu, select **Remove Members**. The server or device group selected is removed from the static device group.

Method II

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

1 From the Navigation pane, select **Devices ➤ Device Groups**. The Device groups appear in the Content pane.

2 Select the device group and double-click to display its members.

3 Select the server displayed in the Content pane. From the **Actions** menu, select **Remove Members**. The server selected is removed from the static device group.

Removing Servers from a Dynamic Group

Servers are added or removed from dynamic device groups automatically, based on the rules created for the group. To remove servers, you can update the rules. These rules apply only to the group being created or modified, and not to any of the subgroups.

To remove servers from a public dynamic device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

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

1. From the Navigation pane, select **Devices ➤ Device Groups**. The list of device groups appear in the Content pane.
2. From the Content pane, select the device group and then from the **Actions** menu, select **Open**. The Device Group Explorer appears.
3. From the Views pane, select Server Membership. In the Content pane, you can specify the rules for the dynamic device group.
4. Update a rule by selecting the attribute from the first drop-down list. Depending on the attribute that you select, options available for the operator and values for the rule will change.
5. Select the operator from the second drop-down list.
6. Enter a value in the field or select a value from the drop-down list or click … to select multiple values from the Select Values window.
7. Click ⊞ to update additional rules and repeat steps 4 to 6. Click ⊟ to delete any rules.
8. Select the logic (And/Or) to be applied for every rule in the query.
9. Click **Preview** to view the servers which are members of the device group.
10. From the **File** menu, click **Save** to save the rules.

Moving a Device Group

Moving a device group moves the location of the device group to the desired location. In the SA Client, you can move device groups from a private group to a public group. You cannot move device groups from a public group to a private group. If you move a device group containing sub groups, the sub groups are also moved to the desired location.

To move a public device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

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

1. From the Navigation pane, select **Devices ➤ Device Groups**. The list of device groups appears in the Content pane.

2 From the Content pane, select the device group and then from the **Actions** menu, select **Move Group**. The Move Device Group window appears as shown below.

Figure 5-5: The Move Device Group Window in the SA Client

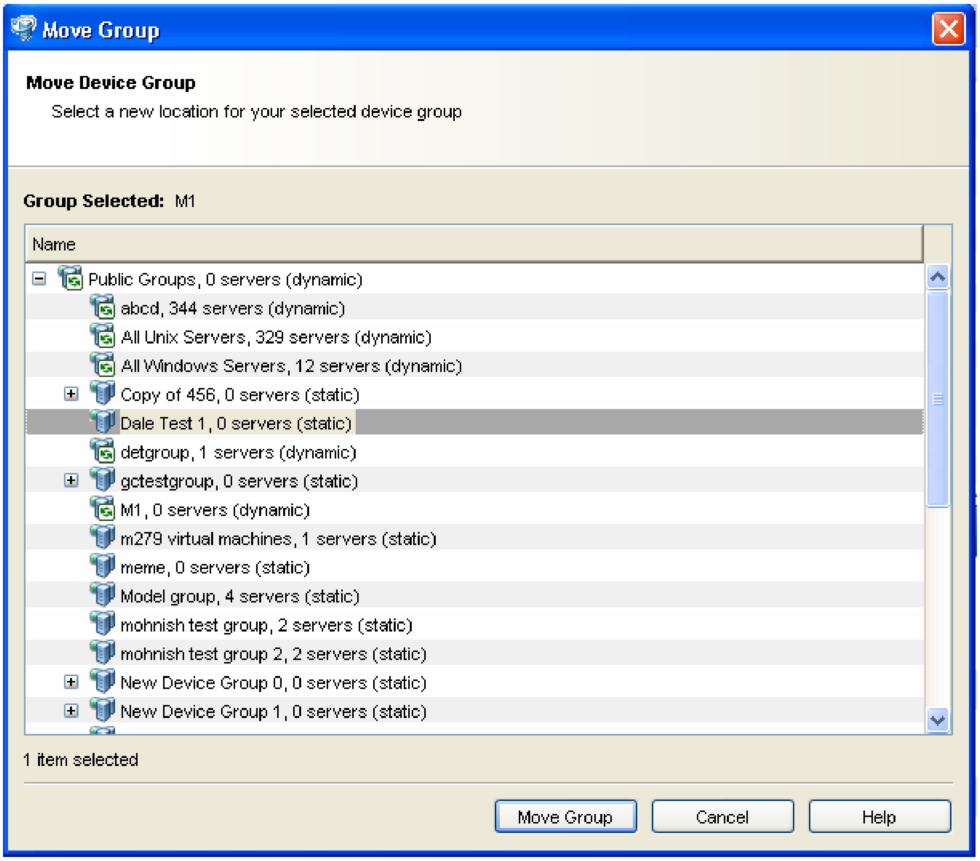

Move Group

Move Device Group

Select a new location for your selected device group

Group Selected: M1

Name

- Public Groups, 0 servers (dynamic)
  - abcd, 344 servers (dynamic)
  - All Unix Servers, 329 servers (dynamic)
  - All Windows Servers, 12 servers (dynamic)
  - Copy of 456, 0 servers (static)
  - Dale Test 1 , 0 servers (static)
  - detgroup, 1 servers (dynamic)
  - gctestgroup, 0 servers (static)
  - M1 , 0 servers (dynamic)
  - m279 virtual machines, 1 servers (static)
  - meme, 0 servers (static)
  - Model group, 4 servers (static)
  - mohnish test group, 2 servers (static)
  - mohnish test group 2, 2 servers (static)
  - New Device Group 0, 0 servers (static)
  - New Device Group 1 , 0 servers (static)

1 item selected

3 Select the desired location to move the device group.

4 Click **Move to Group**. The device groups are moved to their new location.

Duplicating a Device Group

When you duplicate a device group, the servers are copied to a new group and they remain in the original group. In the SA Client, you can only select one device group at a time to duplicate.

To duplicate a public device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

To duplicate an existing device group, perform the following steps:

1. From the Navigation pane, select **Devices ➤ Device Groups**. The list of device groups appears in the Content pane.
2. From the Content pane, select a device group and then from the **Actions** menu, select **Duplicate Group**. A copy of the device group is created.

Deleting a Device Group

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

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

- Software policies or patch polices are attached to the group or a subgroup of the group.
- Access control boundaries are attached to the groups or a subgroup of the group.
- The device group contains servers and other device groups.

To delete a device group, you must have Manage Public Device Group permissions. To obtain these permissions, contact your SA Administrator.

To delete a server group, perform the following steps:

1. From the Navigation pane, select **Devices ➤ Device Groups**. The list of device groups appears in the Content pane.
2. From the Content pane, select the device group and then from the **Actions** menu, select **Delete Group**.
3. Click **Delete** on the confirmation dialog to delete the device group.

Device Group Explorer

The Device Group Explorer allows you to view and manage the properties of a device group in the SA Client. From the Device Groups Explorer, you can perform the following actions:

- View the properties and members of a device group.
- Change a dynamic device group to a static device group.
- Add or remove members from a device group.
- View rollup compliance information for group members targeted by compliance policies such as Audit, Software, App Config, and Patch.
- Add application configurations to a device group.
- View and manage patches, patch policies and software policies associated with the servers in the group.
- Create an audit.
- Take a snapshot.
- View and create custom attributes.
- View server history.

See “Device Group Explorer” on page 203 in this chapter for more information.

To access a Device Group Explorer, perform the following steps:

1 From the Navigation pane, select **Devices ➤ Device Groups**. The device groups appear in the Content pane.

2 Select a device group and from the **Actions** menu select **Open**. The Device Group browser appears as shown below.

Figure 5-6: The Device Group Explorer Window

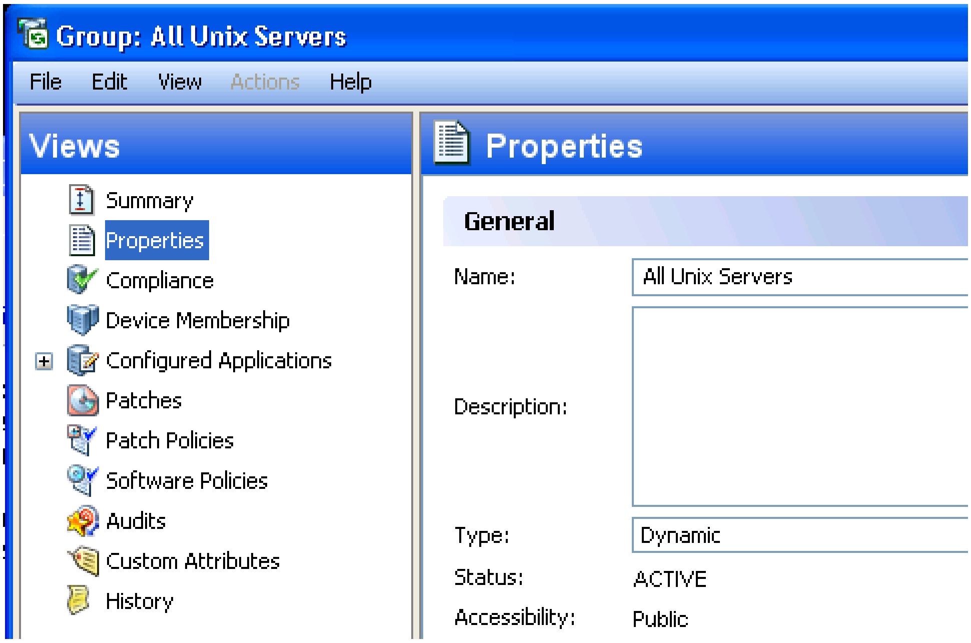

Name: All Unix Servers

Description:

Type: Dynamic

Status: ACTIVE

Accessibility: Public

3 To view the properties or perform an action on the device group, select one of the following views:

- **Summary**: This window allows you to view the type of device group and the members associated with the device group.
- **Properties**: This window allows you to view the properties of the device group and edit the Name, Description, and the Type of the device group.
- **Compliance**:
- **Device Membership**: This window allows you to view the members associated with the device group and also manage members in the device group.
- **Configured Applications**: This window allows you to view and edit application configurations attached to a device group, and add an application configuration to the device group.
- **Patches**: This window allows you to view all the patches associated with the device group, add patches to a patch policy, and install patches on the device group.

- **Patch Policies**: This window allows you to view all the patch policies associated with the device group, and remediate the servers in the device group with the patch policies.
- **Software Policies**: This window allows you to view all the software policies associated with the device group and remediate the servers in the device group with the software policies.
- **Audit and Remediation:** This window allows you to create and run audits and snapshots.
- **Custom Attributes**: This window displays the custom attributes set to a device group and allows you to add, edit or remove custom attributes from a device group.
- **History**: This window allows you to view the changes made to a device group. For each action on the group (but not group members), the history displays the date, the action occurred, and the user who performed the action (if the group is public).

Device Group Explorer

The Device Groups Explorer provides access to all groups of servers and other devices in the core. From the Device Groups Browser, you can perform the following actions:

- Browse groups of servers and access servers inside of groups:
- View patches, patch policies and software policies associated with the servers in the group.
- Create an audit.
- Take a snapshot.
- View and create custom attributes.
- View server history.

Accessing the Device Groups Explorer

To access the Device Groups Explorer, perform the following steps:

1. Launch the SA Client and select **Devices ➤ Device Groups**.
2. Select a device group and open it. (You can expand a group to find sub groups.)
3. If the list of device groups is long, sort the list by clicking a column name, such as name, IP address, OS, customer, or facility.

4 For each device group, you can also perform an audit, take a snapshot, or configure applications.

Device Group Explorer Interface

The Device Group Explorer allows you to review the following information about a group of servers:

- Summary
- Properties
- Server Membership
- Configured Applications for Device Groups
- Patches for Device Groups
- Audit and Remediation
- Custom Attributes for Device Groups
- History Properties for Device Groups

Summary

The Summary view lists the following information:

- **Properties**: This displays if the device group is Private, Public, Dynamic or Static.
- **Members**:This displays the total number of members in the device group.

Properties

The Properties view lists the following property information such as name, type, status, accessibility for the group of servers that you are browsing.

Server Membership

From inside each device group, you can view all members – managed servers and other groups of servers and other devices – that belong to the group. For each server that belongs to the group, the system displays its name, IP address, OS, customer, facility, and any description.

Configured Applications for Device Groups

If the device group is public, then you can add an Application Configuration to the group. The Application Configuration applies to all servers and groups in this group.

- The Installed Configurations tab allows you to browse and edit all Application Configurations attached to the device group.
- The Backup Configurations tab provides a history of all changes made to the selected application configuration template, and allows you to revert to a previous version of the configuration.

See *SA User's Guide: Application Automation* for more information.

Patches for Device Groups

This window displays all patches associated with the selected server group.

Show Options

You can use the Show drop-down list to filter the following types of patch information displayed in the Device Groups Explorer.

- **Patches with Exceptions (Windows Only)**: This option displays all patches that have exceptions for Windows servers (such as always install or never install) a*nd* have one of the following conditions:
  - The patches are not currently installed and are recommended by the vendor.
  - The patches are currently installed.
- **All Patches**: This displays all patches that are associated with the operating system of a server.

Patch Contents

This section displays the following patch contents information.

- **Icon**: This displays a dimmed patch icon when the patch has not yet been uploaded to the Software Library.
- **Name**: This indicates the QNumber of a patch that is a hotfix or an update rollup. Service pack patches do not have a QNumber.
- **Compliance**: This shows one of the following three levels of patch compliance, as defined by a patch administrator:
  - **Non-compliant** (red): This indicates that the patch is installed on the server, but that it is not in the policy, or the patch is not installed on the server but is in the policy.

  - **Partial** (yellow): This indicates that the policy and exception do not agree, and that the exception does not have data in the Reason field.
  - **Compliant** (green): This indicates any 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 has a never install exception because it is not recommended by the vendor.

  In the Preview pane, move the cursor over the icon or text in the Compliance column to view patch compliance information about a server.

- **Type**: This indicates the type of patch, such as Windows Hotfix or Windows Update Rollup.
- **Bulletin**: (Optional) This indicates the Microsoft Security Bulletin ID number for this patch.
- **Severity**: (Optional) This displays one of following Microsoft severity ratings for this patch:
  - **Critical**: This indicates a patch whose exploitation could allow the propagation of an internet worm, without user action.
  - **Important**: This indicates 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**: This indicates a patch whose exploitability is mitigated to a significant degree by factors, such as default configuration, auditing, or difficulty of exploitation.
  - **Low**: This indicates a patch whose exploitation is extremely difficult, or whose impact is minimal.
- **Release Date**: This shows the date that Microsoft released this patch.
- **Exception**: This indicates the type of patch policy exception set for the selected server.
- **Installed**: This indicates if the patch is installed on the selected server.

- **Recommended**: A check mark indicates that this patch was recommended by the vendor (MBSA 2.0) during the last software registration.
- **Description**: This shows a description of the server.

Patch Policies for Device Groups

This window displays all patch policies associated with the selected device group. You can use the Show drop-down list to filter the type of patch policies to display in the Server Explorer.

Show Options

This section displays the following patch information:

- **Policies Attached to Device Group**: This displays all policies attached to the device group, or policies attached to a server group to which the selected managed server belongs.
- **Policies Not Attached to the Server**: This displays a list of all patch policies relevant to the selected server group that are not attached to the group.

Patch Contents

This section displays the following patch content information:

- **Name**: This displays the name of the patch policy.
- **OS**: This displays the operating system associated with the patch policy.
- **Description**: This shows a description of the patch policy.

Software Policies for Device Groups

This Software Policies view displays all software policies associated with the selected server (or group of servers). You can perform actions such as attaching a policy, detaching a policy, remediating a server, and scanning software compliance from the **Actions** menu.

Audit and Remediation

This window allows you to create and run audits and snapshots.

Custom Attributes for Device Groups

This window displays the custom attributes set to a server or device group. You also can add, edit or remove custom attributes from this window.

Custom attributes can be one of the following two types:

- Inherited ⬇ from another source, such as a customer, a software policy, group of servers, ISM control, and so on.
- Attached directly to the server.

To override inherited custom attribute values, click the inherited arrow icon ⬇ once, and enter a new value in the value field. Press ENTER. The inherited arrow changes ➡ to indicate that the custom attribute value has been overridden.

History Properties for Device Groups

The History view shows changes made to the selected device group. Entries are generated when actions are performed on a device group in the SAS Web Client. The History is read-only. Each entry shows the following information:

- **Date**: The date when the last change occurred.
- **Event**: A description of the change.
- **User**: The user who made the change.

Use the View drop-down list to sort the device group history list according to a range of time, such as last week, the last two months, and so on.

Chapter 6: Virtualization Director

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Virtualization Director
- Viewing Virtual Servers
- Solaris Zone Management
- VMware Virtual Machine Management
- Shared Steps for Virtual Server Management

Overview of Virtualization Director

SA's Virtualization Director automates virtual server management. You can perform the same operations on virtual servers using the Virtualization Director as on physical servers, including audit, remediation, application configuration, software management, and patching. This integrated management solution enables you to use the same policies on virtual servers that you currently use on physical servers. You can also automate virtual server creation and other administrative tasks.

Supported Virtualization Platforms

In this release, SA supports the following virtualization platforms:

- Solaris 10 Zones
- VMware ESX Server 3

Virtualization Director Features

You can perform the following tasks using Virtualization Director:

- In a single operation, create a virtual server, provision the operating system, apply patches, and install application software (VMware ESX only).
- Discover, audit, remediate, and report on virtual servers.

- View the physical-to-virtual relationships between hypervisors and virtual servers. You can easily find out which virtual servers are hosted by a given hypervisor.
- Map dependencies for the entire application environment in the HP Server Automation Visualizer (SAV), including virtual and physical servers, software, network devices, and storage devices.
- Provision VMware ESX and Solaris 10 hypervisors on bare metal servers.
- On Solaris zones, perform the following operations: create, modify, remove, start, and stop.
- On VMware virtual machines, perform the following tasks: create, modify, remove, power on, power off, suspend, and reset.
- Find virtual servers in the SA Client by searching on virtualization properties.
- Create dynamic device groups which group servers based on virtual server properties. Device groups enable you to easily perform operations on multiple servers.
- Control access to virtual and physical servers with a common authentication and authorization structure.

Like all features in SA, Virtualization Director is governed by permissions. To obtain these permissions, contact your SA administrator. For information on permissions, see the *SA Administration Guide*.

Virtualization Terminology

Due to different implementations of virtualization by different vendors, terminology can vary with platform and operating system.

The following terms are non-platform specific:

- **Hypervisor** - A hypervisor is a virtualization platform that allows multiple, heterogeneous operating systems to run on a single host computer at the same time. A Type 1 hypervisor (or bare-metal architecture) is software that runs directly on the hardware platform. Any guest operating systems runs at the second level above the hardware. VMware ESX and Solaris 10 Zones are type 1 hypervisors. A Type 2 hypervisor (or hosted architecture) is software that runs within the operating system environment. Any guest operating systems run at the third level above the hardware.

- **Virtual Server** - The logical server (or virtual machine) that runs on a hypervisor. To an application, the virtual server appears to be a physical server. Each virtual server can run its own operating system, and each server can be independently rebooted. Examples of virtual servers are VMware virtual machines and Solaris zones.
- **Physical Server** - An operating system running on a hardware device. In this chapter, the term physical server refers to a non-virtual server.

Table 6-1 maps the terms hypervisor and virtual server to the values displayed in the SA Client and to corresponding terms in the vendor's platform documentation.

Table 6-1: Virtualization Terminology

TERM THAT IS NOT PLATFORM SPECIFIC	VIRTUALIZATION FIELD IN SA CLIENT	OS FIELD IN SA CLIENT	RELATED TERMS IN VENDOR DOCUMENTATION
Hypervisor	Hypervisor	SunOS 5.10 *	Solaris 10, global zone
	Hypervisor	VmWare ESX Server	ESX Server, computing server, host
Virtual Server	Solaris Zone	SunOS 5.10**	zone, non-global zone **
	VMware VM	(varies)	virtual machine, guest operating system

* Strictly speaking, SunOS is the name of the operating system, and Solaris is the operating system (SunOS) plus a graphical user environment (Solaris). Some system utilities such as `uname` display the string SunOS for the operating system, but most people refer to the operating system as Solaris.

** A previous version of this chapter referred to a non-global zone as a local zone.

OS Provisioning and Agent Installation for Hypervisors

Before you can create or manage a virtual server with SA, you must install an SA Server Agent on the hypervisor that hosts the virtual server.

To provision an operating system or install an Agent on a hypervisor, you follow the same procedures as with other physical servers. As Figure 6-1 shows, you can provision an operating system using the OS Provisioning feature, that automatically installs an Agent along with the operating system. Alternatively, you can install an Agent using the Discovery and Agent Deployment (ODAD) utility on a pre-existing hypervisor. After an Agent has been installed, the hypervisor will be listed in both the All Managed Servers and Virtual Servers lists of the SA Client.

Figure 6-1: Hypervisor OS Provisioning and Agent Installation

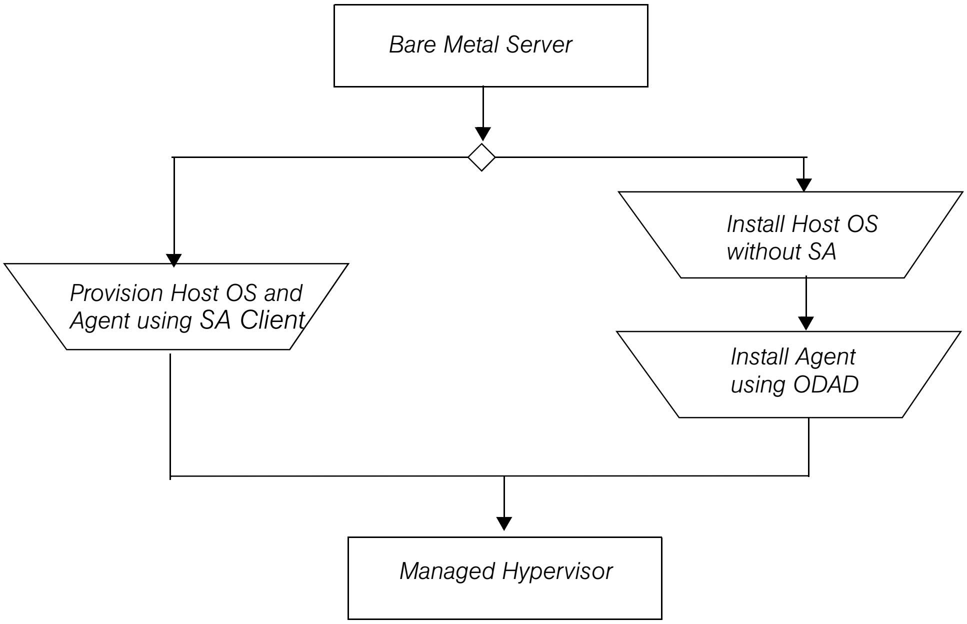
Bare Metal Server
Install Host OS without SA
Provision Host OS and Agent using SA Client
Install Agent using ODAD
Managed Hypervisor

To find out more about the tasks shown in Figure 6-1, see the documentation referenced in Table 6-2.

Table 6-2: Documentation for Hypervisor OS Provisioning and Agent Installation

TASK IN FIGURE 6-1	WHERE TO GO FOR MORE INFORMATION
Install a Host OS without using SA	OS vendor documentation

Table 6-2: Documentation for Hypervisor OS Provisioning and Agent Installation (continued)

TASK IN FIGURE 6-1	WHERE TO GO FOR MORE INFORMATION
Install an Agent using ODAD	"Installing an Agent on an Unmanaged Zone" on page 221 "Installing an Agent on an Unmanaged Virtual Machine" on page 227
Provision the Host OS and Agent using the SA Client	"Operating System Provisioning" chapter of the *SA User's Guide: Application Automation*

OS Provisioning and Agent Installation for Solaris Zones

Figure 6-3 shows the process for creating a non-global zone and bringing it under management by SA. You can create a zone either with the SA Client or Solaris CLI. When the SA Client creates a zone, it also installs an Agent on the zone. After the Agent is installed, you can manage the zone using SA.

Figure 6-2: Solaris Zone OS Provisioning and Agent Installation

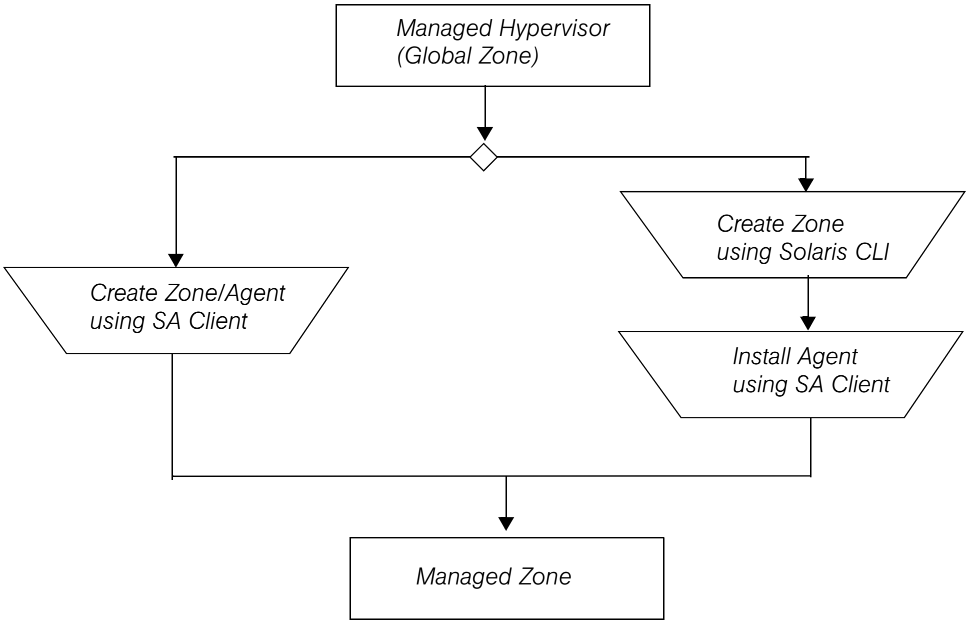
Managed Hypervisor (Global Zone)
Create Zone/Agent using SA Client
Create Zone using Solaris CLI
Install Agent using SA Client
Managed Zone

To find out more about the tasks shown in Figure 6-2, see the documentation referenced in Table 6-3.

Table 6-3: Documentation for Zone OS Provisioning and Agent Installation

TASK IN FIGURE 6-2	WHERE TO GO FOR MORE INFORMATION
Create a Zone using the SA Client	"Creating a Zone" on page 219
Create a Zone using the Solaris CLI	Solaris 10 documentation
Install an Agent using the SA Client	"Installing an Agent on an Unmanaged Zone" on page 221

OS Provisioning and Agent Installation for Virtual Machines

Like a physical server, a VMware virtual server must have an Server Agent installed in order to be managed by SA. Figure 6-3 shows the two different approaches you can take to install an Agent on a virtual machine. In the figure, the path on the left uses SA for the entire process. When you create a virtual machine using the SA Client, you can specify an OS sequence and provision the OS immediately, or you can provision the OS at a later time. As with a physical server, provisioning an OS using the SA Client on a virtual server automatically installs an Agent. The path on the right shows the process when the virtual machine is created and provisioned without using SA. This situation typically occurs when you need to bring existing virtual machines under SA management.

Figure 6-3: VMware Virtual Machine OS Provisioning and Agent Installation

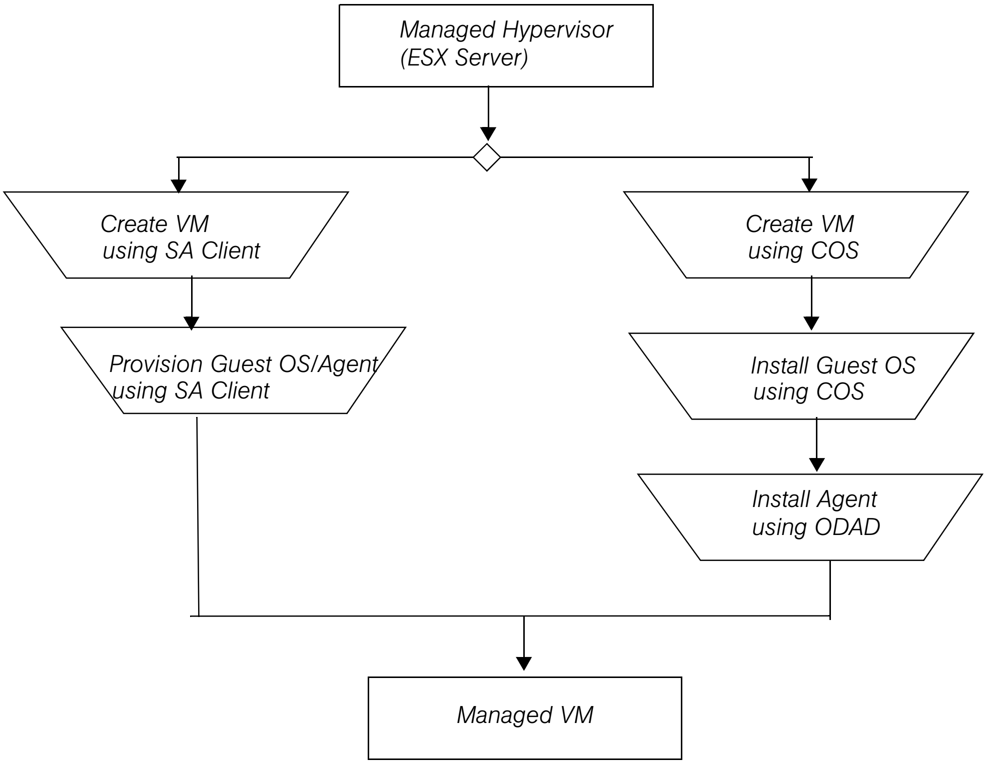

To find out more about the tasks shown in Figure 6-3, see the documentation referenced in Table 6-4.

Table 6-4: Documentation for VM OS Provisioning and Agent Installation

TASK IN FIGURE 6-3	WHERE TO GO FOR MORE INFORMATION
Create a VM using COS	VMware documentation
Create a VM using the SA Client	“Creating a Virtual Machine” on page 223
Install an Agent using ODAD	“Agent Management” section of the *User’s Guide: Server Automation*
Install a Guest OS using COS	VMware documentation

Table 6-4: Documentation for VM OS Provisioning and Agent Installation (continued)

TASK IN FIGURE 6-3	WHERE TO GO FOR MORE INFORMATION
Provision a Guest OS using the SA Client	"Operating System Provisioning" chapter of the *SA User's Guide: Application Automation*

Viewing Virtual Servers

The tasks described in this section apply to both VMware virtual machines and Solaris zones.

Searching for Servers by Virtualization Properties

The Search feature of the SA Client enables you to find servers and other objects according to certain criteria. You can search for virtual servers or hypervisors by the following properties:

- Virtualization type, such as Hypervisor, Solaris Zone, or VMware VM.
- Solaris zone properties, such as Zone CPU Shares and Zone Device Dirs.
- VMware virtual machine properties, such as VM Data Store Name and VM Virtual Processors.
- Properties common to both physical and virtual servers, such as Operating System and Software Policy.

For step-by-step instructions, see "SA Client Search" on page 88.

Viewing Physical-to-Virtual Relationships

To view hypervisors and virtual servers in the SA Client, perform the following steps:

1. From the Navigation pane, select Devices.
2. Expand Servers.
3. To view the hypervisors and virtual servers as a tree structure, select Virtual Servers.

   The Content pane displays the hypervisors. To see which virtual servers belong to a hypervisor, expand the hypervisor in the tree. The Virtualization view of a hypervisor also lists its virtual servers.

4. To identify the hypervisor of a virtual server displayed in the All Managed Servers list, select the virtual server, select the Virtualization view, and note the value of the Hosted By field.

Viewing Configuration Properties of Virtual Servers

In the SA Client, perform the following steps:

1. From the Navigation pane, select **Devices**.
2. Expand Servers.
3. Select the virtual server from either the Virtual Servers or All Managed Servers list.

   In the All Managed Servers list, to see if a server is a hypervisor or a virtual server, choose Virtualization from the column selector.
4. For the View, select **Virtualization**.

   The Content pane displays the configuration properties of the virtual server.

Refreshing Virtual Server Information

Because virtual servers can be created and managed outside of SA, the SA Client view of virtual servers must be refreshed periodically. SA automatically refreshes the virtual server information every 24 hours. To refresh the virtual server information immediately, perform the following steps:

1. From the Navigation pane, select **Devices**.
2. Expand Servers.
3. In the Virtual Servers list, right-click the hypervisor whose information you want to refresh.
4. Depending on the virtualization platform, select one of the following menu items:
   - **Solaris Zones ➤ Refresh Zones**
   - **VMware Virtual Machines ➤ Refresh Virtual Machines**

Visualizing Virtual Servers

This section briefly describes virtual servers and hypervisors in SAV. For more information on SAV, see the Server Automation Visualizer chapter in the *SA User's Guide: Application Automation*.

SAV enables you to manage the operational architecture and behavior of distributed applications in your IT environment. The SAV Server Map displays the relationships between hypervisors and virtual servers. For VMware ESX hypervisors, you can also view vSwitches alongside virtual machines, in addition to the connections between the virtual machines and vSwitches' port groups. The buttons on the Virtualization tool bar in SAV enable you to start, stop, restart, and pause virtual servers.

To view virtual servers in SAV, perform the following steps:

1. In the SA Client, from the **Navigation** pane, select **Devices**.
2. Expand **Servers**.
3. From the **Virtual Servers** list, select the hypervisor.
4. From the **Actions** menu, select **Open with ➤ Server Automation Visualizer**

   After scanning is complete, the SAV window appears.
5. In the SAV window, select the **Server Map** tab.

Reporting on Virtual Servers

The SA Client reporting feature provides the following virtualization reports:

- **All Virtual Servers**: Lists all managed hypervisors and virtual servers.
- **Virtual Servers by Virtualization Technology**: Displays a pie chart with slices for each virtualization technology (vendor), as well as a list that summarizes the number of servers for each technology.
- **Virtual Servers by Hypervisor**: Lists properties of individual virtual servers, grouped by hypervisor.
- **Resource Allocation by Hypervisor:** Shows resource allocation (such as CPU shares or memory) by hypervisor.

For instructions on running reports, see the *SA User’s Guide: Application Automation*.

Solaris Zone Management

This section provides step-by-step instructions for managing non-global zones with the SA Client. The read-only tasks for non-global zones are in “Viewing Virtual Servers” on page 216. To perform any of these tasks, the global zone must be reachable by SA.

Creating a Zone

The SA Client enables you to create a fully operable, non-global zone on a Solaris 10 global zone (hypervisor).

Before a zone can be created with the SA Client for the first time, the global zone must have its hardware registered with SA. This hardware registration only needs to occur once for a global zone. Creating a zone also installs an Agent on the zone, bringing the zone under management by SA. After Agent installation, the zone appears in the All Managed Servers and the Virtual Servers lists in the SA Client.

To create a non-global zone, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.

2. Expand Servers and select Virtual Servers.

3. In the Content pane, right-click the global zone that will host the new zone, and select **Solaris Zones ➤ Create Zone**.

4. In the Zone Definition Method step of the Create Virtual Zone window, select one of the following methods:

   - **Filling Out a Data Form**: With this method, you will enter the zone configuration parameters in the fields of the Zone Definition step, which follows the current window in this task.

   - **Entering a Zone Creation Command Script**: If you select this method, in the Zone Definition step, you will enter or upload a command script that contains the zone creation and configuration commands. This command script has the same syntax as the *command_file* specified by the following Solaris command:

     ```
     zonecfg -z zone_name -f command_file
     ```

     SA does not validate the contents of the zone creation command script, so ensure that your script works. For the syntax of the commands allowed in the script, see the Solaris 10 documentation on `zonecfg`.

5. Click **Next** to proceed to the Zone Definition step.

6. In the Zone Definition step of the Zone Creation window, fill out the following fields:

   - **Server**: Enter the SA server name for the new virtual server. (Beneath the server name, you can see the global zone (hypervisor) that will be hosting this new zone.)

- **Zone Name**: Enter a name for the new zone. You might want to give the zone a name that indicates its purpose, such as the name of the application that will be running in the zone.
- **Locale**: Enter the language code for the zone.
- **Terminal Type**: Specify the terminal client that the SA Client uses for Remote Terminal sessions on the zone.
- **Auto Reboot**: Select to automatically reboot the zone if the physical server running the global zone is rebooted.
- **Time Zone**: Select a time zone for the zone.
- **Root Password**: Enter the root password for logging into the zone.
- **Confirm Password**: Enter the root password again.

7 If you selected the Entering a Zone Creation Command Script option, in the Zone Configuration Commands section, enter (or paste) the contents of the zone creation command script. Or, click **Import File** to import the zone creation command script.

8 If you selected the Filling Out a Data Form option, enter data in the following fields:

- **CPU Shares Reservation**: Enter an integer that allocates the shares of the CPU resource.

For the CPU Shares Reservation to take effect, the Fair Share Scheduler (FSS) must be enabled on the Solaris 10 hypervisor server on which you are creating the new zone. By default, the FSS is not enabled.

- **IP Address**: Click **Add** to enter the IP address and the name of the virtual network interface. A zone that requires network connectivity must have one or more dedicated IP addresses. These addresses are associated with the physical network interface of the global zone.
- **Device**: (Optional) Click **Add** to enter a path name to a device, for example, `/dev/pts*`, on the global zone hypervisor. This action gives the new zone access to a device on the global zone.

- **Inherited Package Directory**: (Optional) Click **Add Sparse Root Directory** to enter a default package directory inherited by the new zone from the global zone. Or, click **Add** to enter a path name to the package directory you want the zone to inherit. These package directories are read-only.
- **Mount Directory**: (Optional) Click **Add** to mount a file system type and path for the zone. This action grants the zone access to a physical disk or file system of the non-global zone. A ufs file system creates a file system mount within the zone; the lofs type is a loopback file system mount to the global zone.

9 Click **Next**.

10 Continue through the remaining steps of the Create Zone window, as explained in the following sections:

- "Scheduling Step for Virtual Servers" on page 229
- "Notifications Step for Virtual Servers" on page 229
- "Job Status Step for Virtual Servers" on page 230

Installing an Agent on an Unmanaged Zone

To bring a zone under management by SA, you must install an Agent on it. Before installing an Agent on a zone, verify the following requirements:

- The zone is running.
- All services on the zone are up and running. If they are not up and running, perhaps the zone does not have the necessary sysconfig information, that is provided by a user when the zone is booted for the first time. See the Sun Solaris 10 documentation on zone configuration for more information.

You can install an Agent on an unmanaged Solaris zone in two ways:

- With ODAD, as described in the "Agent Management" chapter of the *User's Guide: Server Automation*.
- Through the Solaris zone management interface.

To install an Agent on an unmanaged zone through the zone management interface, perform the following steps:

1 In the SA Client, from the Navigation pane, select the Devices tab.

2 Expand Servers and select Virtual Servers.

3. In the Content pane, expand the global zone that hosts the zone you want to modify.
4. Right-click the non-global zone and select **Solaris Zones ➤ Install Agent**.

Modifying a Zone

You can modify a subset of the zone parameters that were defined when you created the zone. For example, if an HTTP server running on a zone is sluggish because of increased usage, you can increase the CPU Share Reservation parameter of the zone. After you modify a zone with the SA Client, the zone is automatically re-started.

To modify the parameters of a non-global zone, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, expand the global zone that hosts the zone you want to modify.
4. Right-click the non-global zone and select **Solaris Zones ➤ Modify Zone**.

   The Modify Virtual Zone window appears. In the Zone Definition step of the Modify Virtual window, edit the parameters you want to change.
5. Click **Next**.
6. Continue through the remaining steps of the Modify Zone window, as explained in the following sections:

- "Scheduling Step for Virtual Servers" on page 229.
- "Notifications Step for Virtual Servers" on page 229.
- "Job Status Step for Virtual Servers" on page 230

Starting or Stopping a Zone

To start or stop a zone, perform the following steps:

1. From the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, expand the hypervisor that hosts the zone.
4. Right-click the zone and select **Solaris Zones ➤ Start Zone** or **Stop Zone**.

Or:

5. In the Virtual Servers list, open the hypervisor. In the Virtualization view, select the zone and click either **Start** or **Stop**.

If the zone is not running, the state displayed is **Installed**.

Removing a Zone

To remove a non-global zone, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, expand the global zone that hosts the zone you want to remove.
4. Right-click the non-global zone and select **Solaris Zones ➤ Remove Zone**.

VMware Virtual Machine Management

This section provides step-by-step instructions for managing VMware virtual machines with the SA Client. The read-only tasks for virtual machines are in “Viewing Virtual Servers” on page 216.

Creating a Virtual Machine

With the SA Client, you can create a virtual machine on a VMware ESX server (hypervisor). If you specify an OS Sequence during this task, SA also provisions an OS and installs an SA Agent on the virtual machine. After Agent installation, the virtual machine appears in the All Managed Servers and the Virtual Servers lists in the SA Client.

To create a VMware virtual machine, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, right-click the VMware ESX Server that will host the virtual machine, and select **VMware Virtual Machines ➤ Create Virtual Machine**.
4. In the Virtual Machine Definition step of the Create Virtual Machine window, enter data in the following fields:

- **Name**: The SA server name of the new virtual machine. Also, the name of the new virtual machine as it will appear within the ESX Server.
- **Memory Size**: The amount of memory for the virtual machine.
- **Virtual Processors**: The number of CPUs for the virtual machine.
- **Guest OS**: Select the operating system (OS) of the virtual machine. This field is just a label. Specifying this field does not provision the OS on the virtual machine.
- **OS Sequence**: (Optional) To provision the OS when the virtual machine is created, select an OS sequence. If it has a software policy, the OS sequence installs the software in the policy on the virtual machine immediately after the OS is provisioned.

5 Under Network Configuration, click **Add** to create a virtual network adapter (NIC). To connect the virtual network adapter when the virtual machine powers on, select "Connect at power on."

6 Under Data Store Configuration, click **Add** to create a virtual disk. To select a data store from the hypervisor, double-click the entry under Name. You can also change the default values for Size and Units by double-clicking those entries.

7 Click **Next**.

8 Continue through the remaining steps of the Create Virtual Machine window, as explained in the following sections:

- "Scheduling Step for Virtual Servers" on page 229.
- "Notifications Step for Virtual Servers" on page 229.
- "Job Status Step for Virtual Servers" on page 230

If you are installing Solaris 10 as a Guest OS on a VMWare ESX server, the following configurations are strongly recommended:

Solaris 10 32-bit

Guest OS: SOlaris 10 32 bit
SCSI controller: LSI logic
Network Adapter: flexible

Solaris 10 64-bit

Guest OS: Solaris 10 64 bit
SCSI controller: LSI logic
Network Adapter: E1000

Installing Software During Virtual Machine Creation

With the SA Client, you can perform the following three operations individually or as a single task:

- Create the virtual machine.
- Provision an OS on the virtual machine.
- Install software on the virtual machine.

To bundle these operations as a single task, perform the following steps:

1. Decide which OS will be provisioned on the virtual machine.
2. Locate or create the corresponding OS sequence.

   For instructions on creating OS sequences, see the *SA Policy Setter's Guide*.
3. In the OS sequence, be sure to select Enable Remediation.
4. Decide which software will be installed on the virtual machine.
5. Create a software policy.

   For instructions, see the *SA Policy Setter's Guide*.
6. Add the software to be installed to the software policy.
7. Launch the Create Virtual Machine task.

   For instructions, see "Creating a Virtual Machine" on page 223.
8. In the Virtual Machine Definition step of the Create Virtual Machine task, specify the OS sequence.

   When the virtual machine is created, the OS is provisioned and the software policy is remediated on the virtual machine. The remediation operation installs the contents of the software policy on the virtual machine.

Opening VMware VI Web Access From the SA Client

VMware Virtual Infrastructure Web Access is a browser-based application for managing virtual machines. (The short name for this application is VI Web Access.) Before running VI Web Access, you must install a browser plug-in, as described the VMWare, Inc. document, *Virtual Infrastructure Web Access Administrators Guide.*

To open VI Web Access from the SA Client, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, right-click the VMware ESX server and select **VMware Virtual Machines ➤ Open Web Access**.

Opening a Virtual Machine Console From the SA Client

VI Web Access enables you to open a browser-based console for a VMware virtual machine. With the console, you can monitor the progress of OS provisioning on a virtual machine. Within the console, press Ctrl + Alt to transfer control of your mouse and keyboard back to your computer.

To open VI Web Access console from the SA Client, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, right-click the virtual machine and select **VMware Virtual Machines ➤ Open Console**.

Provisioning an OS on a Virtual Machine

An unprovisioned virtual machine does not have an OS. To provision an OS on an existing VMware virtual machine, perform the following steps:

1. In the VMware console, boot a new "bare metal" virtual machine with either PXE boot or a boot image file. (You can open the console from VMware VI Web Access or VI Client.) In the SA Client, the virtual machine appears in the Unprovisioned Servers list.
2. Provision the OS on the virtual machine as if it were a physical server. For instructions, see the "Operating System Provisioning" chapter of the *SA User's Guide: Application Automation*.

During the OS provisioning process, SA installs an Agent on the virtual machine. After Agent installation, the virtual machine appears in the All Managed Servers list.

Installing an Agent on an Unmanaged Virtual Machine

Your environment will have unmanaged VMware virtual machines if they have been created and OS provisioned with tools other than SA. To bring a virtual machine under management by SA, you must install an Agent on it.

You can install an Agent on an unmanaged virtual machine with command-line tools or with ODAD. For more information on ODAD, see the “Agent Management” chapter of the *User’s Guide: Server Automation*.

Modifying a Virtual Machine

To modify the attributes of a VMware virtual machine, perform the following steps:

1. If you want to modify the following attributes, power off the virtual machine:

    - Memory Size
    - Virtual Processors
    - Network Configuration
    - Data Store Configuration

    For instructions, see “Powering a Virtual Machine On or Off” on page 228.

2. In the SA Client, from the Navigation pane, select the Devices tab.
3. Expand Servers and select Virtual Servers.
4. In the Content pane, expand the VMware ESX Server that hosts the virtual machine you want to modify.
5. Right-click the virtual machine and select **VMware Virtual Machines ➤ Modify Virtual Machine**.

    The Modify Virtual Machine window appears. In the Virtual Machine Definition step of this window, change the attributes.

6. Click **Next**.
7. Continue through the remaining steps of the Modify Virtual Machine window, as explained in the following sections:

    - “Scheduling Step for Virtual Servers” on page 229.

- “Notifications Step for Virtual Servers” on page 229.
- “Job Status Step for Virtual Servers” on page 230

Powering a Virtual Machine On or Off

To power a VMware virtual machine on or off, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, expand the VMware ESX Server that hosts the virtual machine you want to power on.
4. Right-click the virtual machine and select **VMware Virtual Machines**.
5. Select **Power On Virtual Machine** or **Power Off Virtual Machine**.

Suspending a Virtual Machine

The suspending action pauses the virtual machine activity. To suspend a VMware virtual machine, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, expand the VMware ESX Server that hosts the virtual machine you want to suspend.
4. Right-click the virtual machine and select **VMware Virtual Machines.**
5. Select **Suspend Virtual Machine**.

Resetting a Virtual Machine

The resetting action stops the virtual machine and reboots it. To reset a VMware virtual machine, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, expand the VMware ESX Server that hosts the virtual machine you want to reset.
4. Right-click the virtual machine and select **VMware Virtual Machines ➤ Reset Virtual Machine**.

Removing a Virtual Machine

If you remove a virtual machine, its entry is deleted from the Model Repository. To remove a VMware virtual machine, perform the following steps:

1. In the SA Client, from the Navigation pane, select the Devices tab.
2. Expand Servers and select Virtual Servers.
3. In the Content pane, expand the VMware ESX Server that hosts the virtual machine you want to remove.
4. Right-click the virtual machine and select **VMware Virtual Machines ➤ Remove Virtual Machine**.

Shared Steps for Virtual Server Management

This section describes the shared steps of tasks that launch SA jobs. These tasks include “Creating a Zone” on page 219 and “Creating a Virtual Machine” on page 223.

Scheduling Step for Virtual Servers

In the Scheduling step, you decide whether the job that performs the task runs immediately or at a later time.

To schedule the task, perform the following steps:

1. Click **Next** until you reach the Scheduling step.
2. Select one of the following options:
   - **Run Task Immediately**: This option allows you perform the task when you click **Start Job** in the Job Status step.
   - **Run Task At**: This option allows you to specify the date and time for the task.

Notifications Step for Virtual Servers

In this step, you can set email notifications to alert users on the success or failure of a job. You can also associate a Ticket ID with the job. These settings are optional.

To set email notifications, perform the following steps:

1. Click **Next** until you reach the Notifications step.
2. Click **Add Notifier**.

3. Enter the addresses in the Notification Email Address field.
4. To send email to the address if the job succeeds, select the ✓ icon. To send email if the job fails, select the ✗ icon.
5. Enter an ID to be associated with this job in the Ticket ID field.

Job Status Step for Virtual Servers

In the Job Status step, you can start the job, view the summary information of the job progress, and see the individual status of each action required for the job.

To start the job, perform the following steps:

1. Click **Next** until you reach the Job Status step.
2. Click **Start Job**.

   If you selected Run Task Immediately in the Scheduling step, the job begins now. If you scheduled the job for a later time, the job will run later. The job's progress information appears in this window.
3. To view the details of each action performed by the job, select a row in the table.
4. Click **End Job** to stop the job from running or click **Close** to exit this window.

You can also view jobs in the Jobs Log window of the SA Client. See the *User's Guide: Server Automation* for information about Job Logs.

Chapter 7: Server Tracking in the SAS Web Client

IN THIS CHAPTER

This section discusses the following topics:

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

Server Tracking

This section discusses the following topics:

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

Ways to Locate, List, and Display Servers

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

- By searching when you know the name, host name, or IP address of the server you want to provision or manage.
- By viewing the Manage Servers list and Server Pool list when you want to see a complete list of all your servers. You can refine the lists by using filters.

- By viewing servers sorted by hardware category. (Click Environment ➤ Hardware in the navigation panel.) The Servers tab in the Hardware pages shows each manufacturer and model that you have running in the operational environment. See "Hardware Information for Managed Servers" on page 271 in this chapter for more information.
- You can also browse managed servers and server groups using the SA Client.

Tracked Server Properties

Every server that SA manages has the following properties:

- IP addresses, host name, and the server ID
- All software that is installed on the server. Select the Installed Packages tab from the Manage Servers: Server Properties page to display the list of software that is reportedly installed on the server.

  HP Server Automation is able to determine what software is installed on a server because the Server Agent communicates with the SA core and reports the installed hardware and software for the server.

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

See "Server Information that the Agent Tracks" on page 109 in Chapter 4 for information about a complete list of all the hardware and software information that HP Server Automation tracks for managed servers.

Supported Operating Systems for Managed Servers

This section lists the supported operating systems for agents and the SA Client.

SA Server Agents

The following table lists the supported operating systems for SA server agents, which run on the servers managed by SA.

Table 7-5: SA Server Agent Supported Operating Systems

SUPPORTED OPERATING SYSTEMS FOR SA SERVER AGENT	VERSIONS	ARCHITECTURE
AIX	AIX 4.3 AIX 5.1 AIX 5.2 AIX 5.3	POWER POWER POWER POWER
HP-UX	HP-UX 10.20 HP-UX 11.00 HP-UX 11.11 HP-UX 11.23 (11i v2) HP-UX 11.31 (11i v3)	PA-RISC PA-RISC PA-RISC PA-RISC and Itanium PA-RISC and Itanium
Sun Solaris	Solaris 6 Solaris 7 Solaris 8 Solaris 9 Solaris 10 (Update 1, Update 2, Update 3, Update 4 and Update 5)	Sun SPARC Sun SPARC Sun SPARC Sun SPARC Sun SPARC, 64 bit x86, 32 bit x86 and Niagara
Fujitsu Solaris	Solaris 8 Solaris 9	Fujitsu SPARC Fujitsu SPARC
Windows	Windows NT 4.0 Windows 2000 Server Family Windows Server 2003 Windows XP Professional Windows Server 2008	32 bit x86 32 bit x86 32 bit x86 and 64 bit x86 32 bit x86 32 bit x86 and 64 bit x86

Table 7-5: SA Server Agent Supported Operating Systems (continued)

SUPPORTED OPERATING SYSTEMS FOR SA SERVER AGENT	VERSIONS	ARCHITECTURE
Red Hat Linux	Red Hat Enterprise Linux 2.1 AS	32 bit x86
	Red Hat Enterprise Linux 2.1 ES	32 bit x86
	Red Hat Enterprise Linux 2.1 WS	32 bit x86
	Red Hat Enterprise Linux 3 AS	32 bit x86 and 64 bit x86 and Itanium
	Red Hat Enterprise Linux 3 ES	32 bit x86 and 64 bit x86 and Itanium
	Red Hat Enterprise Linux 3 WS	32 bit x86 and 64 bit x86 and Itanium
	Red Hat Enterprise Linux 4 AS	32 bit x86 and 64 bit x86
	Red Hat Enterprise Linux 4 ES	32 bit x86 and 64 bit x86
	Red Hat Enterprise Linux 4 WS	32 bit x86 and 64 bit x86
	Red Hat Enterprise Linux Server 5	32 bit x86 and 64 bit x86
	Red Hat Enterprise Linux Desktop 5	32 bit x86 and 64 bit x86
SUSE Linux	SUSE Linux Enterprise Server 8	32 bit x86
	SUSE Linux Standard Server 8	32 bit x86
	SUSE Linux Enterprise Server 9	32 bit x86 and 64 bit x86
	SUSE Linux Enterprise Server 10	32 bit x86 and 64 bit x86
VMware	ESX Server 3.0	32 bit x86 and 64 bit x86
	ESX Server 3.0.1	32 bit x86 and 64 bit x86
	ESX Server 3.0.2	32 bit x86 and 64 bit x86
	ESX Server 3.5	32 bit x86 and 64 bit x86

On Red Hat Enterprise Linux 4 AS SA does not support SELinux (Security Enhanced Linux). By default, SELinux is enabled on Red Hat 4 AS. You must disable the SELinux feature on Red Hat 4 AS for the SA agent to function correctly. SA supports SELinux (Security Enhanced Linux) on Enterprise Linux 5.

Required Patches for Agent Installation

Table 7-6: Required Patches and Packages for Agent Installation

SERVER OPERATING SYSTEM	REQUIRED PATCHES
AIX 4.3 AIX 5.1	APAR IY39444 APAR IY39429 NOTE: If AIX 4.3.3.388, 4.4.4.89, or 5.1.0.3 is installed, the Agent Installer displays an error message that indicates the correct APAR to install on the server.
HP-UX (10.20, 11.00, 11.11/11i)	For HP-UX 10.20, PHCO_21018 Additionally, SW-DIST should be upgraded to the HP recommended patch level. You should continue to upgrade this package when HP recommends new versions.
Linux AS 3.0 Linux WS 3.0 Linux ES 3.0	Red Hat Enterprise Linux 3 Update 3
Solaris 10, 9, 8, 7, and 6	SUNWadmc SUNWcsl SUNWcslr (If available, depending on version) SUNWcsu SUNWesu SUNWlibms SUNlibmsr (If available, depending on version) SUNWswmt It is strongly recommended not to remove packages from the SUNWCreq minimal required install cluster, since many packages are interdependent and operation beyond that of basic SAS functionality may be affected.
Windows 2000	Service Pack 4
Windows NT 4.0	Service Pack 6a

SA Client

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

Table 7-7: SA Client Supported Operating Systems

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

Server Lists

This section discusses the following topics:

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

Types of Server Lists

The SAS Web Client displays lists for two types of servers, as Figure 7-1 shows.

Figure 7-1: Servers Section in the Navigation Panel

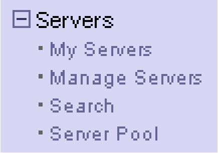

Server Pool: Servers in the Server Pool have registered their presence with HP Server Automation but do not have the target OS installed. An OS Build Agent is running on each server so that they can communicate with HP Server Automation.

See *SA User's Guide: Application Automation* on how to use the Server Pool when you install the target OS on a server.

Manage Servers: The Manage Servers list contains servers on which HP Server Automation can perform management tasks, because Server Agents are installed on them. However, HP Server Automation might not have provisioned all software running on the servers.

You begin the OS provisioning process by reviewing the servers in the Server Pool list. From the Server Pool, you can install a target OS by selecting a server and clicking **Install OS**.

Server Pool

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

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

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

  If the server is in an unreachable state (that is, if the server icon has a red "x" on it), you can run a Communication Test to help you troubleshoot why that server is unreachable. See "Agent Reachability Communication Tests" on page 129 in Chapter 4 for more information.
- The life cycle value, such as whether the server is available to have a target OS installed on it
- The facility in which the server is located
- The customer association
- Additional hardware information (Click the server name to open a window that displays specific hardware information.)

Manage Servers List

The Manage Servers list contains servers on which HP Server Automation can perform management tasks because Server Agents are installed on them. When an existing operational server has an Server Agent installed successfully, it appears in the Manage Servers list and the server icon indicates that it is fully manageable, as Figure 7-2 shows.

Figure 7-2: The Manage Servers List in the SAS Web Client

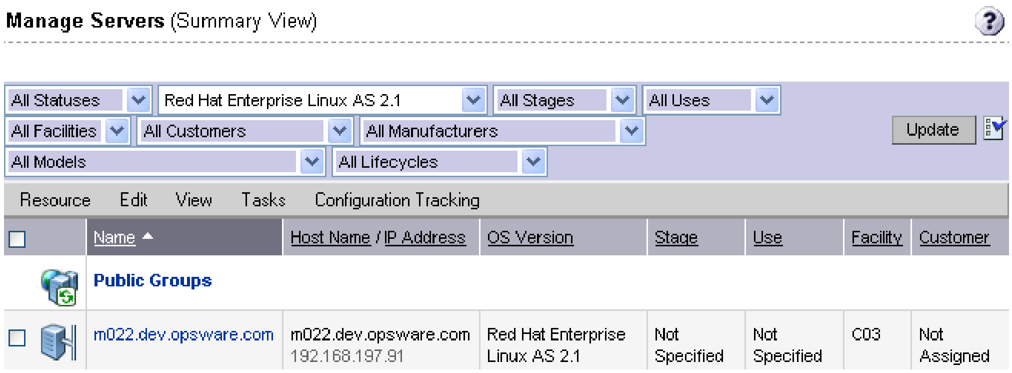

Manage Servers (Summary View)

All Statuses | Red Hat Enterprise Linux AS 2.1 | All Stages | All Uses
All Facilities | All Customers | All Manufacturers
All Models | All Lifecycles

□	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Facility	Customer
	Public Groups						
□	m022.dev.opsware.com	m022.dev.opsware.com 192.168.197.91	Red Hat Enterprise Linux AS 2.1	Not Specified	Not Specified	C03	Not Assigned

See "Agent on Managed Servers" on page 105 in Chapter 4 for more information. By default, servers in the Manage Servers list are sorted by the Name column. However, you can re-sort the list based on any of the column headings. For example, you can click the Hostname / IP Address column heading to re-sort the list by host name or IP address.

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

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

- The name of the server

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

- The host name of the server determined by the Server Agent
- The IP address configured for the server, which users can edit by using the network configuration feature in the
- The reported OS, which is obtained by the Server Agent that is running on the server
- The stage of the server, which specifies the stages of deployment for servers
- The server's use
- The facility in which the server is located

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

Filters on the Manage Servers List

The Manage Servers list displays the following filters that you can specify to qualify the servers that the SAS Web Client displays, as Figure 7-3 shows.

Figure 7-3: Filters in the Manage Servers List

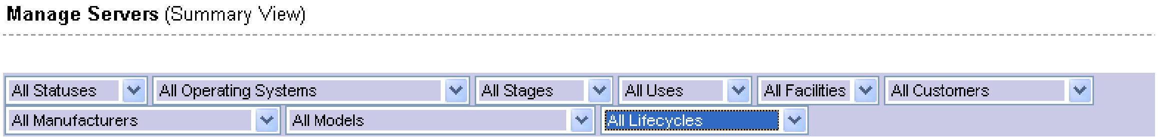

- **Status**: Specifies the ability of HP Server Automation to manage servers. HP Server Automation automatically detects the status of servers; a server's status is OK or Not Reachable.
- **OS**: Specifies the operating system on the server, which is obtained by the Server Agent that is running on the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

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

You can change the filters displayed in the Manage Servers page. To change the filters you want to be displayed on the Manage Servers page, click on the icon as shown in Figure 7-5 and specify the filters from the **Edit Filters** Menu.

Figure 7-4: Edit Icon

Figure 7-5 shows the filters that are in the Server Pool list.

Figure 7-5: Filters in the Server Pool List

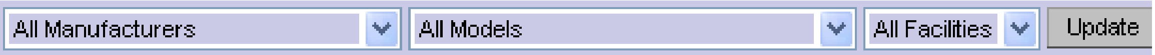

- **Manufacturers**: Specifies the manufacturer for the server as reported by the OS Build Agent running on the server.
- **Models**: Specifies the model of the server as reported by the OS Build Agent running on the server.
- **Facilities**: Specifies the location of the server. Users can manage servers in any facility from an SAS Web Client in any facility.

My Servers

This section contains the following topics:

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

Overview of My Servers

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

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

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

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

Adding Servers to My Servers

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

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

   Or

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

   See “Using the Search Feature” on page 244 in this chapter for more information. See “Server Searching by IP Address” on page 260 in this chapter for more information.

2 Select the servers that you want to add to My Servers.

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

4 Click **Close** to close the window.

5 Next, select the My Servers link at the top of the page. You see the selected servers added to My Servers, as Figure 7-6 shows.

Figure 7-6: Servers in My Servers

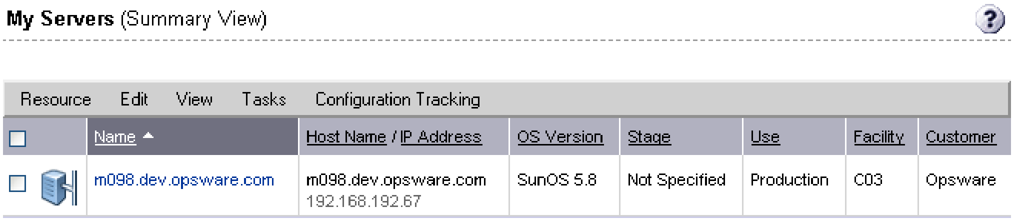

My Servers (Summary View)

□	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Facility	Customer
□	m098.dev.opsware.com	m098.dev.opsware.com 192.168.192.67	SunOS 5.8	Not Specified	Production	C03	Opsware

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

Removing Servers from My Servers

Perform the following steps to remove servers from My Servers:

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

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

Server Search

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

- Searching for a Server By Using the Search Box
- Ways to Use Search
- Using the Search Feature
- Rules for Server Search and for Creating Dynamic Groups
- Line Break Workaround for Server Search
- Conditions for Searching with Multiple Rules
- Server Searching by IP Address
- Example: Server Search
- Searching for a Server Group

Searching for a Server By Using the Search Box

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

1. On the Home page, click the down arrow in the top navigation panel to open the Search box, as Figure 7-7 shows.

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

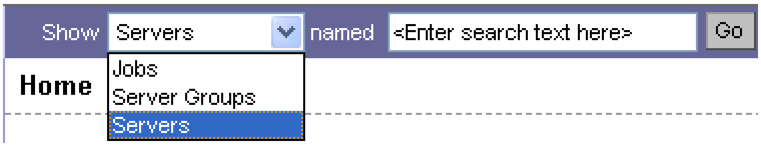

2. Verify that the Servers option is selected in the list.

3. Type the server’s IP address, host name, or name in the Search box and then click **Go**.

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

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

```
192.168.68.6
host02.coredev-va1.sample.com
192.168.*.19
host1*.xyz.samplecompany.com
```

   The resulting page contains one or more servers, depending on the type of search query that you specified. If no servers are found, the SAS Web Client displays a message that indicates that no servers were found that matched your query.

   See “Using the Search Feature” on page 244 in this chapter for information about how to formulate complex, multiple rule search queries.

Ways to Use Search

By using the SAS Web Client, you can perform searches in the following ways:

- From SA wizards

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

  What you can search for in the SA wizards is context sensitive to the type of operation that you are performing.

- From the navigation panel, click Servers ➤ Search.

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

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

Using the Search Feature

In the SA wizards, you can browse for servers, operating systems, or use the Search feature to search for these items.

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

1. In an SA wizard, select the Search tab. The following Search page appears. See Figure 7-8.

Figure 7-8: Search Tab in the Select Servers Step of an SA Wizard

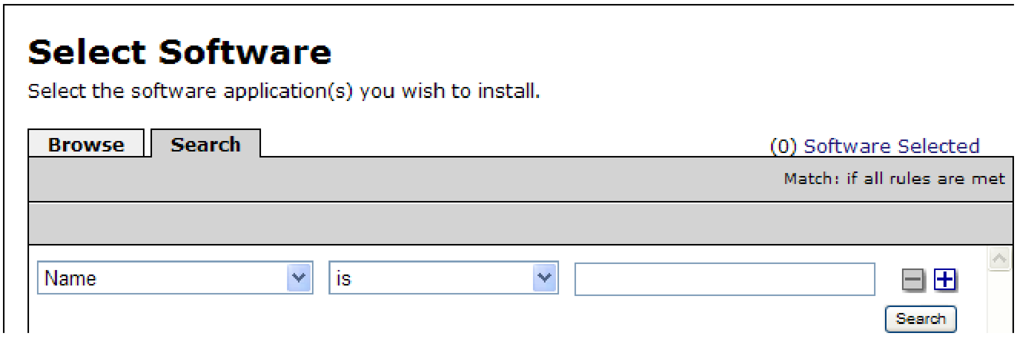

Select Software

Select the software application(s) you wish to install.

You can also use the Search tab at other steps in the wizards to search for operating systems, patches, applications, and templates.

Or

From the navigation panel, click Servers ➤ Search. The Search page appears, as Figure 7-9 shows.

Figure 7-9: Search Page

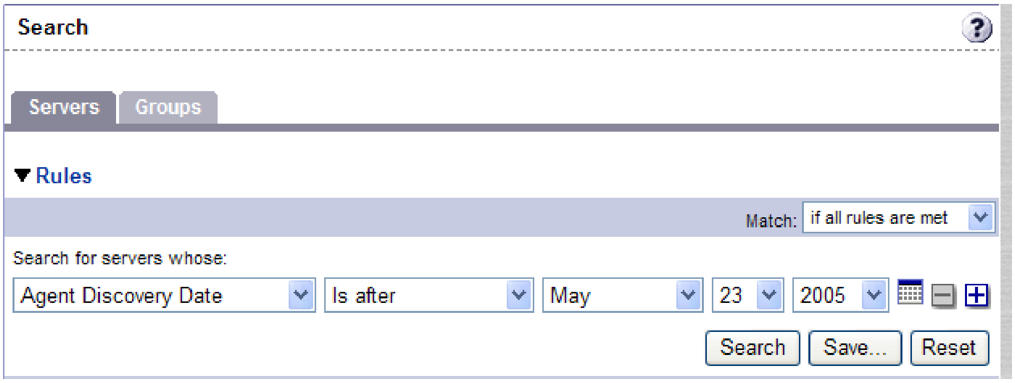

Match: if all rules are met

Search for servers whose:

Agent Discovery Date | Is after | May | 23 | 2005

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

2 Specify the rule that you want to search for by selecting it from the first list, as Figure 7-10 shows.

Figure 7-10: Search Rules List in the Search Page

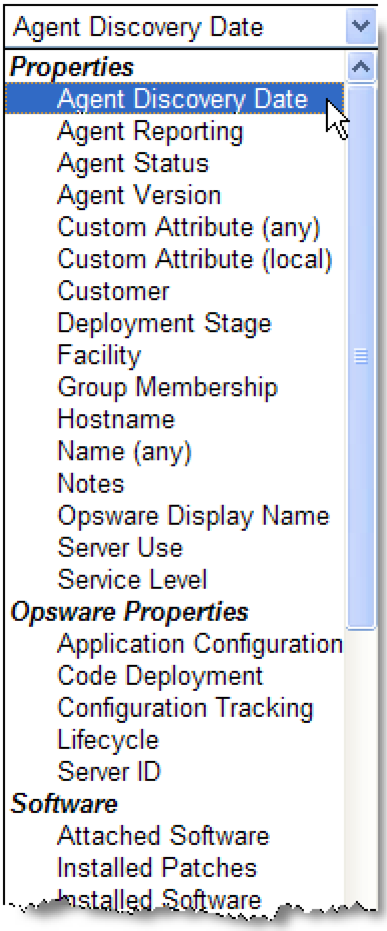

Agent Discovery Date

Properties
- Agent Discovery Date
- Agent Reporting
- Agent Status
- Agent Version
- Custom Attribute (any)
- Custom Attribute (local)
- Customer
- Deployment Stage
- Facility
- Group Membership
- Hostname
- Name (any)
- Notes
- Opsware Display Name
- Server Use
- Service Level

Opsware Properties
- Application Configuration
- Code Deployment
- Configuration Tracking
- Lifecycle
- Server ID

Software
- Attached Software
- Installed Patches
- Installed Software

Depending on the rule that you select, a popup window might appear in the page. For example, if you selected Deployment Stage, a popup window showing the stages appears in the page as Figure 7-11 shows:

Figure 7-11: Search Popup Window with Values

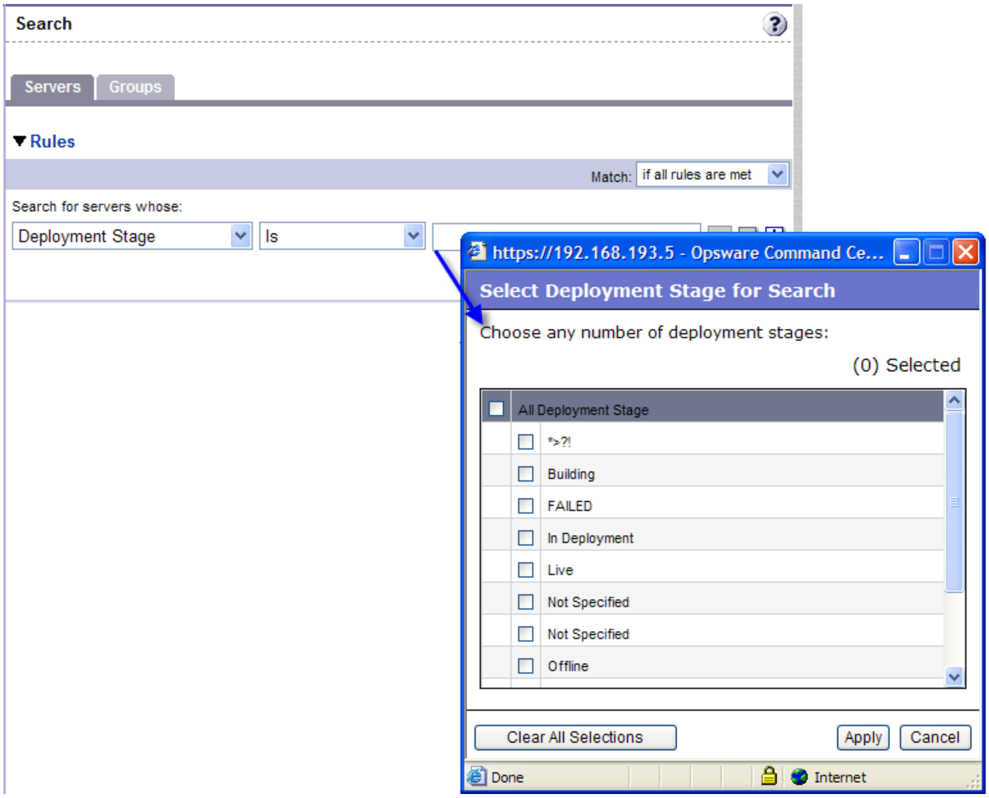

Match: if all rules are met

Search for servers whose: Deployment Stage Is

Choose any number of deployment stages:

(0) Selected

You cannot search in Notes that contain line breaks. See “Line Break Workaround for Server Search” on page 258 in this chapter for more information and a workaround.

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

3 In the second list, specify how you want HP Server Automation to search by selecting a value. The operator selected defines how the search text is treated. Negative operators might not be available in all cases.

See “Rules for Server Search and for Creating Dynamic Groups” on page 249 in this chapter for more information about the operator for each search rule.

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

5 (Optional) To add additional rules, click the plus (+) button as Figure 7-12 shows and repeat Steps 2 through 4.

Figure 7-12: Multiple Rules in a Search

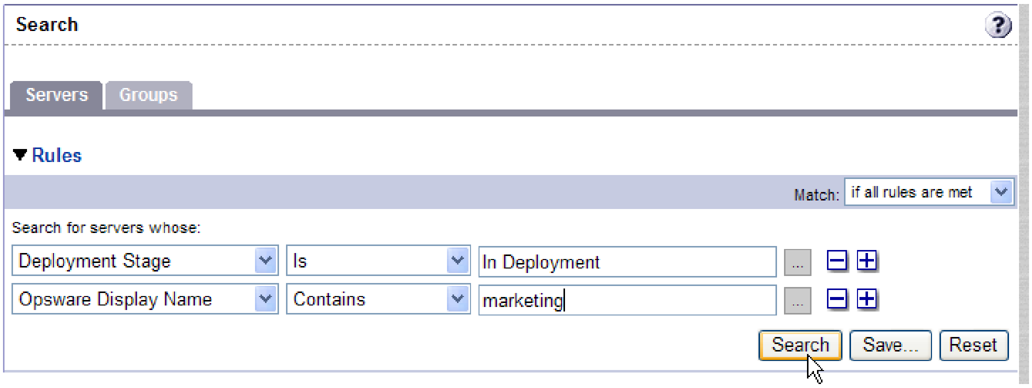

Search

Servers | Groups

Rules

Match: if all rules are met

Search for servers whose:

Deployment Stage | Is | In Deployment

Opsware Display Name | Contains | marketing

6 If you specified multiple rules for the search, select whether you want search results only if all rules are met or if any of the rules are met, as Figure 7-13 shows.

Figure 7-13: Operator Controlling Search Results

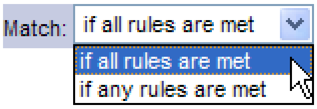

Match: if all rules are met

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

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

Figure 7-14: Displayed Search Results

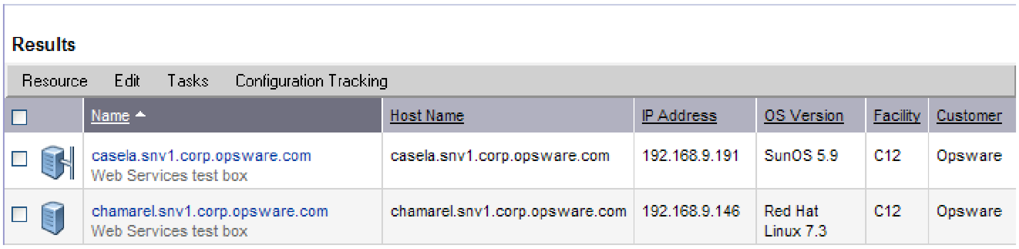

□	Name	Host Name	IP Address	OS Version	Facility	Customer
□	casela.snv1.corp.opsware.com Web Services test box	casela.snv1.corp.opsware.com	192.168.9.191	SunOS 5.9	C12	Opsware
□	chamarel.snv1.corp.opsware.com Web Services test box	chamarel.snv1.corp.opsware.com	192.168.9.146	Red Hat Linux 7.3	C12	Opsware

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

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

Rules for Server Search and for Creating Dynamic Groups

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

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

Table 7-8: Rules for Server Search and for Creating Dynamic Groups

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
PROPERTIES		
Agent Discovery Date: The date that the Server Agent was installed.	• Is after • Is before	Drop-down lists with the day, month, and year
	• Is within the last	User-entered text
	• Is today	N/A

Table 7-8: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Agent Reporting: Whether the Server Agent is reporting to HP Server Automation.	• Is • Is not	• Has not reported • OK • Registration in progress • Reporting error
Agent Status: Whether the Server Agent is reachable by HP Server Automation.	• Is • Is not	• Not reachable • OK
Agent Version: The version of the Server Agent – such as 14.2.3b.	• Contains • Does not contain • Is • Is not	User-entered text
Custom Attribute (any): The name of a custom attribute that is associated with the server through attachment or inheritance.	• Contains • Is	User-entered text
Custom Attribute (local): The name of a custom attribute that is locally attached to the server,	• Contains • Is	User-entered text
Customer: The customer or account that the server is associated with.	• Is • Is not	Popup window of customers

Table 7-8: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Deployment Stage: The stage the server performs within a lifecycle environment.	• Is • Is not	• In Deployment • Live • Not Specified • Offline The values that appear in this list are customizable; in addition to the values above, values specific to your environment might appear.
Facility: The collection of servers managed by an HP Server Automation installation.	• Is • Is not	Popup window of facilities When HP Server Automation is running multimaster mode, the list can contain many facilities.
Group Membership: Whether the server belongs to a group.	• Is	Popup window of groups
Host name: The host name of the server – such as m004.company.com.	• Contains • Does not contain • Is • Is not	User-entered text
Name (any): This enables searching for any name or IP address associated with a server.	• Contains • Does not contain • Is • Is not	User-entered text

Table 7-8: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Notes: The contents of the Notes field from the Properties tab for a server.	• Contains • Does not contain • Is • Is not	User-entered text
Opsware Display Name: The user-configurable name for the server.	• Contains • Does not contain • Is • Is not	User-entered text By default, HP Server Automation uses the configured host name of the server until a user edits it.
Server Use: How the server is being used – such as Development, Staging, Production.	• Is • Is not	• Development • Not Specified • Production • Staging You can customize values that appear in this list. In addition to the values above, values specific to your environment might appear.

Table 7-8: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Service Level: The user-defined category that can be used as an organizational tool.	• Is • Is attached here or below	Popup window of service levels Servers can be associated with multiple service levels.
OPSWARE PROPERTIES		
Application Configuration: Whether the server uses the Application Configuration feature.	• Is not used • Is used	N/A
Code Deployment: Whether the server uses the Code Deployment feature.	• Is not used • Is used	N/A
Configuration Tracking: Whether the Configuration Tracking feature is monitoring or backing up specific files or configurations on a server.	• Is off • Is on	N/A
Lifecycle: The server states that are part of bringing a server into HP Server Automation.	• Is • Is not	• Available • Build Failed • Deactivated • Installing OS • Managed

Table 7-8: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Server ID: The internal ID HP Server Automation uses to identify the server.	• Is • Is not	User-entered text In most cases, the Server ID is the same as the MID.
SOFTWARE		
Attached Software: The software that is assigned or modeled through the SA installation process.	• Is • Is attached here or below	Popup window of software
Attached Software Policies: The software policies that is assigned or modeled through the SA installation process.	• is	Popup window of software policies
Installed Patches: Whether a patch is installed on the server.	• Contains • Does not contain • Is • Is not	User-entered text
Installed Software: The package reported installed on the server.	• Contains • Does not contain • Is • Is not	User-entered text A package does not have to be installed by HP Server Automation to be reported as installed on a server.
OS Version: The OS version defined by OS definitions in the OS Provisioning feature.	• Is • Is not	Popup window of operating systems

Table 7-8: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Reported OS: For Windows – the version reported by the OS, for Unix – the version returned by the uname command.	• Contains • Does not contain • Is • Is not	User-entered text
Windows Service: The names of the Windows services that are reported to HP Server Automation.	• Contains • Does not contain • Is • Is not	User-entered text
NETWORK		
DNS Search Domains: The domains configured to be searched in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
DNS Servers: The IP addresses of the DNS servers configured in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
Default Gateway: The IP address of the default router.	• Contains • Does not contain • Is • Is not	User-entered text

Table 7-8: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
IP Address: Any Internet Protocol address for the server.	• Contains • Does not contain • Is • Is not	User-entered text
MAC Address: Any Media Access Control address, which is the network interface card's unique hardware number.	• Contains • Does not contain • Is • Is not	User-entered text
WINS Servers: The Windows Internet Naming Servers configured in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
HARDWARE		
CPU Make and Model: The vendor name and CPU model for the server – such as GENUINEINTEL Intel(R) Pentium(R) 4 CPU 2.60GHz.	• Is • Is not	Popup window of CPU makes and models
CPU Speed: The Central Processing Unit speed in gigahertz [GHz].	• Does not equal • Equals • Is greater than • Is less than	User-entered text A 600 Mhz machine should be entered as `0.6`.
Make and Model: The vendor name and server model for the server – such as Compaq - DL360.	• Is • Is not	Popup window of server makes and models

Table 7-8: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Number of CPUs: The number of CPUs on the server.	• Does not equal • Equals • Is greater than • Is less than	User-entered text
RAM: The amount of RAM on the server in megabytes [MB].	• Does not equal • Equals • Is greater than • Is less than	User-entered text To enter 1 Gigabyte, type `1024`.
Serial Number: The serial number of the server.	• Contains • Does not contain • Is • Is not	User-entered text
Storage Make and Model: The vendor name and storage model for the server – such as WDC - WD800BB-75DKA0.	• Is • Is not	Popup window of storage makes and models
CUSTOM FIELDS		
A **Numeric** field	• Does not equal • Equals • Is greater than • Is less than	User-entered text
A **String** field	• Contains • Does not contain • Is • Is not	User-entered text

Table 7-8: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
A **URI** field	• Contains • Does not contain • Is • Is not	User-entered text
A **File** field	• Contains • Does not contain • Is • Is not	User-entered text
A **Date** field	• Is after • Is before	Drop-down lists with the day, month, and year
	• Is within the last	User-entered text
	• Is today	N/A

Line Break Workaround for Server Search

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

```
line1 <line break>
```

```
line2
```

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

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

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

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

To work around this limitation, include an asterisk (*) where the line break occurs, as Figure 7-15 shows. For example:

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

Figure 7-15: Line Break Workaround in the Server Search Feature

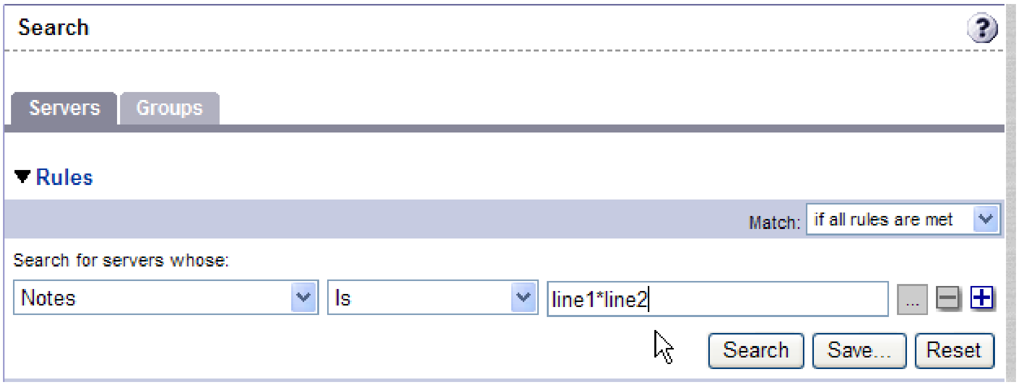

Match: if all rules are met

Search for servers whose:

Notes | Is | line1*line2

Conditions for Searching with Multiple Rules

When you perform a search with multiple rules, the following conditions apply to the way HP Server Automation provides search results:

- When it evaluates rules, the Search feature considers each rule individually, finding servers (or operating systems, patches, and so forth) that match the individual rule, and then the results of each rule are combined.
- You must select at least one rule, and it must have a value. Default filter rules count as rules with a value.
- Empty rules are ignored in searches. Users do not have to manually remove them for the search to proceed.
- You cannot search for empty values. For example, you cannot search for all servers where the Notes field is empty.

At certain steps in the SA wizards, the Search feature provides default values based on previous selections in the wizard.

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

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

Server Searching by IP Address

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

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

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

See "Communication Between Managed Servers and HP Server Automation" on page 277 in Chapter 8 for information about how HP Server Automation handles servers that are affected by static NAT.

Example: Server Search

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

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

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

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

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

- Installed Package contains apache
- Installed Package contains java

- If any rules are met is selected

Searching for a Server Group

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

   Select the Groups tab. The rules for server group search appear as Figure 7-16 shows. By default, one search rule is added to the search.

Figure 7-16: Rules for Server Group Search

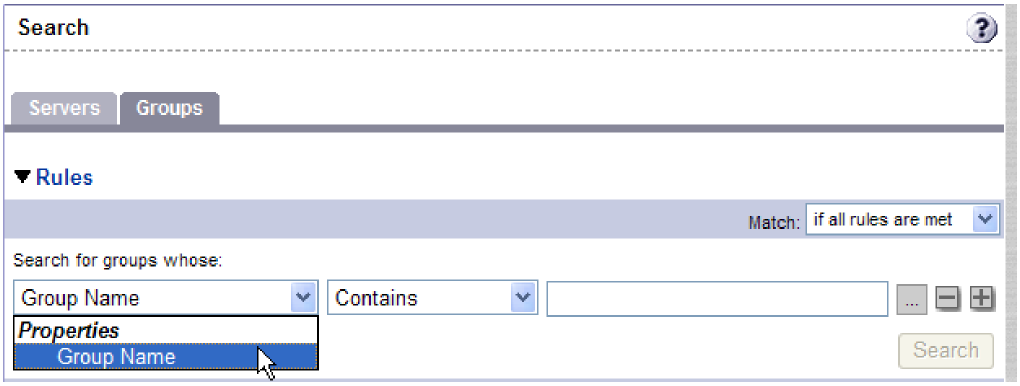

Search

Servers | Groups

▼ Rules

Match: if all rules are met

Search for groups whose:

Group Name | Contains

2. In the second list, specify how you want HP Server Automation to search by selecting either “Contains” or “Is”. The operator selected defines how the search text is treated.

3. Enter the text that you want to search for in the text box. The search text that you enter can include an asterisk (*) wildcard character. (The search text is case sensitive. **Search** is not enabled until you enter text in the text box.)

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

4 Click **Search**. The search results appear as Figure 7-17 shows.

Figure 7-17: Server Group Search Results

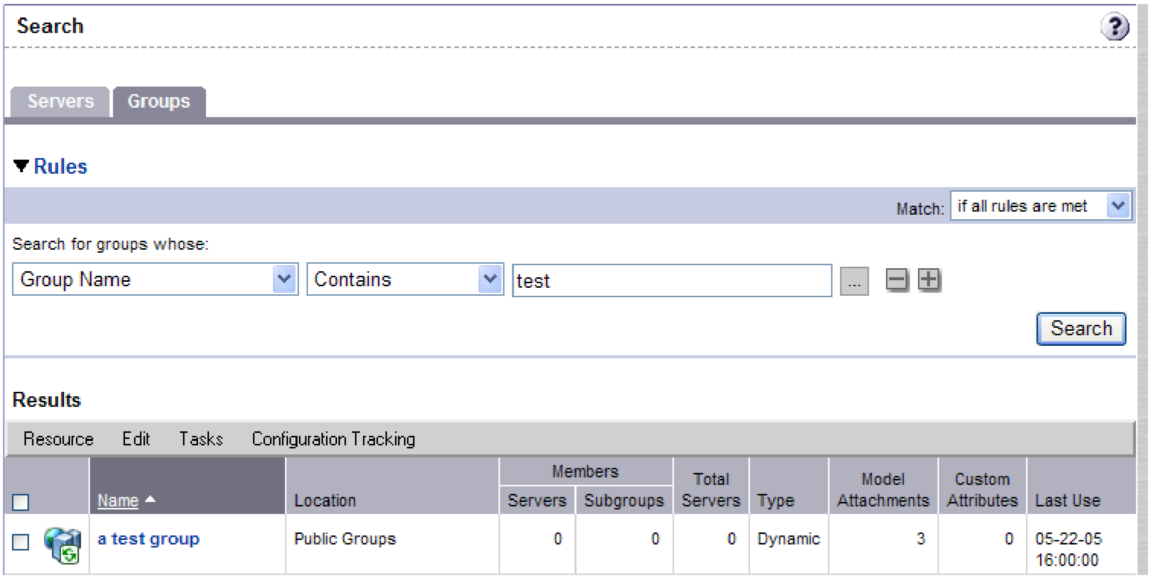

Match: if all rules are met

Search for groups whose: Group Name | Contains | test

Results

□		Name	Location	Members		Total	Type	Model	Custom	Last Use
				Servers	Subgroups	Servers		Attachments	Attributes	
□		a test group	Public Groups	0	0	0	Dynamic	3	0	05-22-05 16:00:00

Server Identification

This section provides information on server identification within HP Server Automation and contains the following topics:

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

Overview of Server Identification

HP Server Automation uses the following IDs to track managed servers:

- **MID**: Machine ID. The unique identifier that HP Server Automation uses to identify the server. The MID is usually equal to the server ID.

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

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

See “Example: HP Server Automation Swaps a Server’s Hard Disk” on page 264 in this chapter for information about swapping hard disks.

- **Server ID**: The primary key in the Model Repository (database) that represents a given server. The Server ID is used internally in HP Server Automation. Generally, users do not need this value for servers to manage them in HP Server Automation.
- **MAC Address**: Media Access Control address, which is the network interface card’s unique hardware number. The MAC is used as the server’s physical address on the network.
- **Chassis ID**: A unique hardware-based identifier that the Server Agent discovers, typically derived from some property of the server's chassis. As a common source for this ID, HP Server Automation uses an interface's MAC address or the host ID on Solaris servers, or the serial number for one of the interfaces.

Ways that Servers are Identified by HP Server Automation

Servers in the Manage Servers list are identified in the following ways when they register their hardware and software with HP Server Automation:

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

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

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

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

  To determine the value in the MAC Address column, HP Server Automation uses the hardware address by which the server contacted the SA Build Manager (a component of the OS Provisioning feature).

Example: HP Server Automation Swaps a Server's Hard Disk

The following steps show how HP Server Automation handles swapping a hard disk for a server:

1. A system administrator swaps the hard disk of Server A (MID `1230001`, chassis ID `AB:08`) with the hard disk of Server B (MID `98730001`, chassis ID `XY:96`).
2. The Server Agent on Server A registers its hardware with HP Server Automation. The MID for Server A equals `1230001` and the chassis ID equals `XY:96`.
3. HP Server Automation locates Server A by using the MID.
4. HP Server Automation updates the data it has for Server A in the Model Repository. It sets the chassis ID equal to `XY:96`.
5. The Server Agent on Server B registers its hardware with HP Server Automation. The MID for Server B equals `98730001` and the chassis ID equals `AB:08`.
6. HP Server Automation locates Server B by using the MID.
7. HP Server Automation updates the data it has for Server B in the Model Repository. It sets the chassis ID equal to `AB:08`.

Customer Accounts in HP Server Automation

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

HP Server Automation accommodates these requirements with customer accounts created by your SA administrator.

When your SA administrator creates a customer in HP Server Automation, a value for that customer is automatically added to the customer filter in the Manage Servers list, as Figure 7-18 shows.

Figure 7-18: Customer Filter in the Manage Servers List

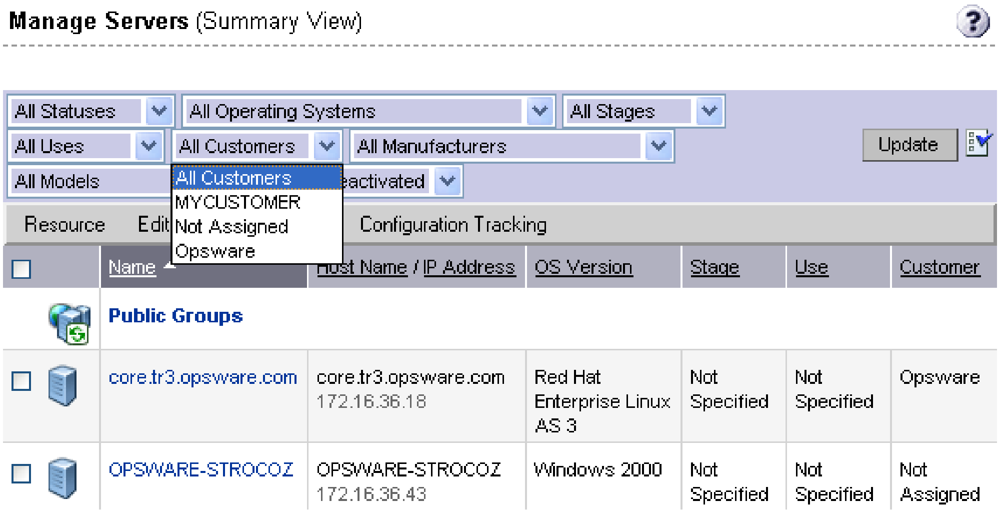

Manage Servers (Summary View)

□	Name	Host Name / IP Address	OS Version	Stage	Use	Customer
	Public Groups					
□	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	Red Hat Enterprise Linux AS 3	Not Specified	Not Specified	Opsware
□	OPSWARE-STROCOZ	OPSWARE-STROCOZ 172.16.36.43	Windows 2000	Not Specified	Not Specified	Not Assigned

By using customer accounts in the SAS Web Client, you can segregate servers that belong to different business units. By segregating servers, you can have separate accounting for each customer or different levels of security for different customers. You might want to segregate the servers based on the department or business unit.

By default, HP Server Automation is shipped with the following two customers:

- **Customer Independent**: A global customer in HP Server Automation. Resources that are associated with “Customer Independent” can be installed on any managed server, no matter what customer it is associated with.
- **Not Assigned**: The servers are not associated with a customer. You can install resources that are Customer Independent on Not Assigned servers. However, you cannot install or use any resources associated with a customer on a server that is not assigned to a customer.

When you install an Server Agent in a server, the server is associated with the Not Assigned customer if IP ranges were not created to automatically associate HP Server Automation managed servers with customers. See Figure 7-19.

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

Figure 7-19: Customers List Under Environment in the SAS Web Client

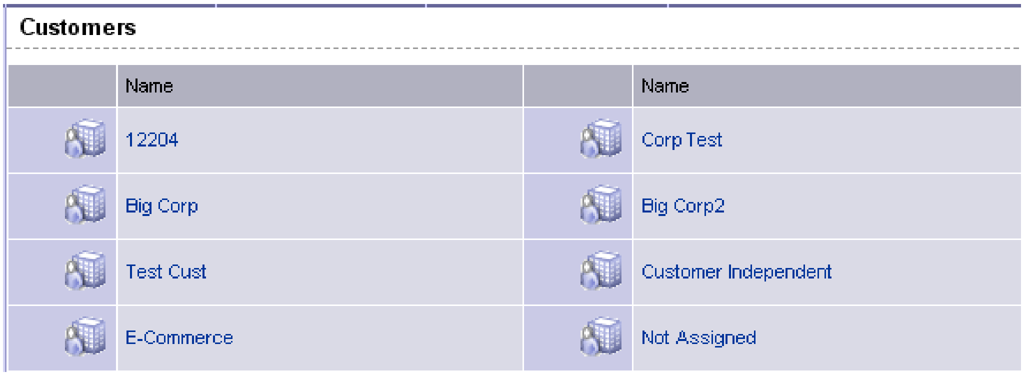

Customers

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

Associating Servers with Customers

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

- Install Server Agents on the servers.

  See "Agent Reachability Communication Tests" on page 129 in Chapter 4 for more information.

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

  See *SA User's Guide: Application Automation* for more information.

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

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

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

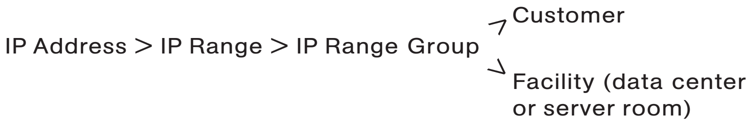

See the *SA Policy Setter's Guide* for more information.

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

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

See "Editing the Properties of a Server" on page 323 in Chapter 8 for information about how to change the customer association for a server.

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

See "Communication Between Managed Servers and HP Server Automation" on page 277 in Chapter 8 for information about how HP Server Automation uses management IP addresses for servers.

However, a server always belongs to the same facility (data center or server room) as its primary IP address. HP Server Automation enforces the relationship between server and facility at hardware registration. See Figure 7-20.

Figure 7-20: Primary IP Addresses in HP Server Automation

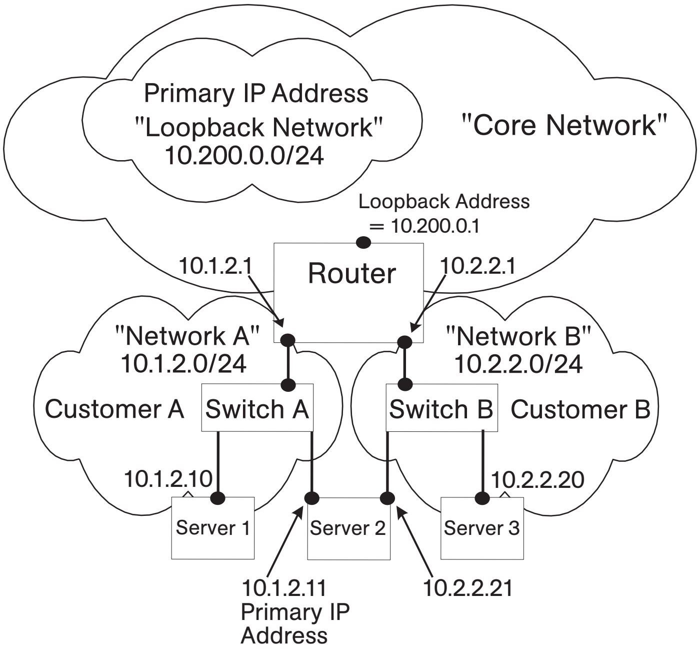

In this illustration, the following conditions apply:

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

Server Histories and Reports

This section provides information on server histories and reporting with HP Server Automation and contains the following topics:

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

Overview of Server History

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

Each History entry contains three pieces of information, as Table 7-9 shows.

Table 7-9: Description of the Entries in the Server History Tab

HISTORY ENTRY	DESCRIPTION
Event Description	Description of the operation performed, for example: `Run ISM Control: ad-hoc script (Job ID: 26190040) executed successfully.`
Modified By	The name of the SA user who made the change.
Date Modified	The date and time the change was made, for example: `Mon Aug 06 18:14:41 GMT+00:00 2001)`

Time Stamps used in History

Data is maintained for servers in HP Server Automation for the following periods of time:

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

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

If longer periods of time are required, Hewlett Packard recommends regular backups to enable offline storage of HP Server Automation data.

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

Viewing Server History

Perform the following steps to view the server history:

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

   Or

   Search for the server whose history you want to view.

   See "Using the Search Feature" on page 244 in this chapter for more information. See "Server Searching by IP Address" on page 260 in this chapter for more information.

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

3. Select the History tab.

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

Time Stamp for Server Operations

HP Server Automation maintains a comprehensive audit trail of the software that the SA users install, configure, and remove from a server. By using the SAS Web Client, you can view the history of the changes made to a server. Entries are generated when actions are performed for managed servers in the SAS Web Client. The history is read-only. See "Server Histories and Reports" on page 269 in this chapter for more information.

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

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

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

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

Additionally, Hewlett Packard recommends that after installing an Server Agent on the server (so that it becomes managed by HP Server Automation), you should synchronize the system clock on the server with the system clocks of the servers running the SA core components.

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

Hardware Information for Managed Servers

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

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

See "Server Information that the Agent Tracks" on page 109 in Chapter 4 for more information.

Viewing Managed Server Hardware Information

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

To view hardware of managed servers:

1. From the navigation panel, click Environment ➤ Hardware. You see the top level of the Hardware category in the managed environment, as shown in Figure 7-21.

Figure 7-21: Top Level Hardware

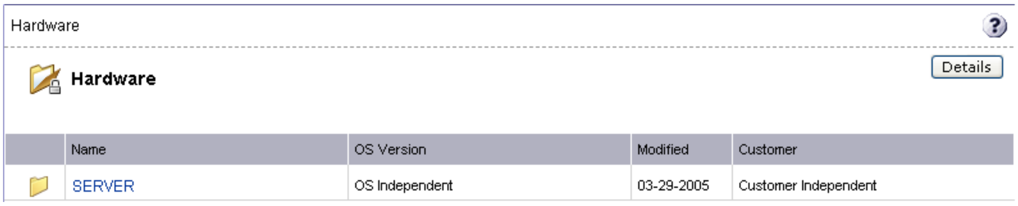

Hardware

Hardware

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

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

3. Drill down to the type of server you want to look at. For example, you might want to look at all Dell POWEREDGE 650s, as shown in Figure 7-22.

Figure 7-22: Hardware Home Page for Dell POWEREDGE 650s

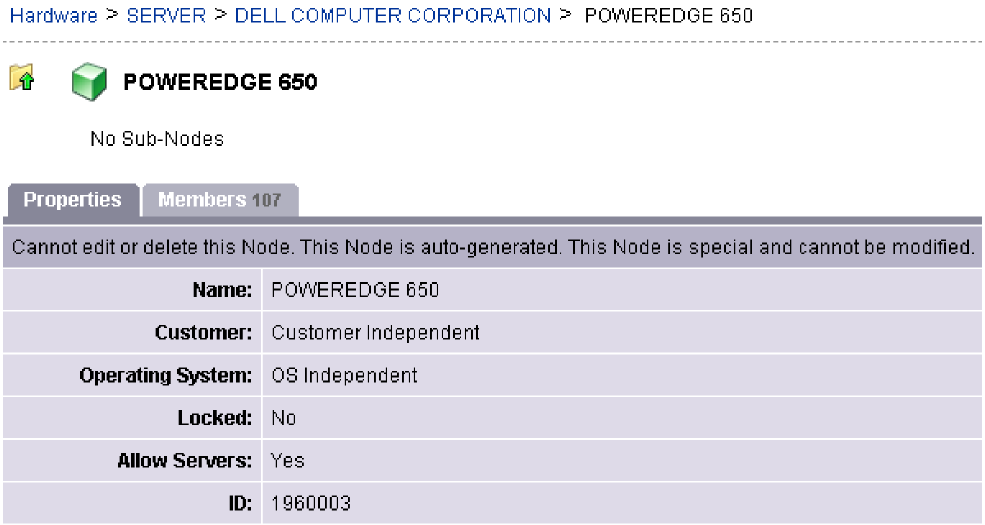

Hardware > SERVER > DELL COMPUTER CORPORATION > POWEREDGE 650

POWEREDGE 650

No Sub-Nodes

Cannot edit or delete this Node. This Node is auto-generated. This Node is special and cannot be modified.

Name:	POWEREDGE 650
Customer:	Customer Independent
Operating System:	OS Independent
Locked:	No
Allow Servers:	Yes
ID:	1960003

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

5. To view specific hardware information, from the **View** menu, choose **Hardware**. You now see more detailed information about all the Dell POWEREDGE 650s computers being managed by HP Server Automation.

Chapter 8: Server Management in SAS Web Client

IN THIS CHAPTER

This section discusses the following topics:

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

Overview of Server Management

HP Server Automation manages servers in an operational environment in the following ways:

- Provisioning servers with Microsoft Windows, Red Hat Linux, and Sun Solaris operating systems by using vendor-provided operating system bootstrapping technologies. Additionally, HP Server Automation integrates with AIX NIM and HP-UX Ignite installation technologies to provide a uniform method for OS provisioning across a heterogeneous environment.

See *SA User's Guide: Application Automation* for information about how HP Server Automation provisions Microsoft Windows, Red Hat Linux, Sun Solaris, and VMware ESX operating systems on managed servers.

See the *SA Administration Guide* for information about how HP Server Automation integrates with AIX NIM and HP-UX Ignite installation technologies.

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

  See "Configuration Tracking" on page 429 in Chapter 12 for more information.

- Automating the management and execution of server scripts.

  See See "Script Execution" on page 493 in Chapter 8 for more information.

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

Communication Between Managed Servers and HP Server Automation

This section provides information about communication between managed servers and HP Server Automation and contains the following topics:

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

Network Address Translation (NAT) for Managed Servers

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

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

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

Key Terms of SA Managed Server Communication

To understand how HP Server Automation communicates with managed servers, you must understand these three terms:

- **Management IP**: The IP address that HP Server Automation uses to communicate with the Server Agent on the server.

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

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

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

  During synchronizations, the Code Deployment & Rollback feature uses the primary IP address to communicate with the server. See "Code Deployment and Static NAT" on page 280 in this chapter for more information.

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

HP Server Automation does not support managed servers that have IPv6 addresses.

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

Locating the Management IP Address of a Managed Server

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

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

The Network tab shows (and allows the user to set) the management interface for the server by selecting it from a drop-down list, as Figure 8-1 shows.

Figure 8-1: Management IP Address Information in the Network Tab

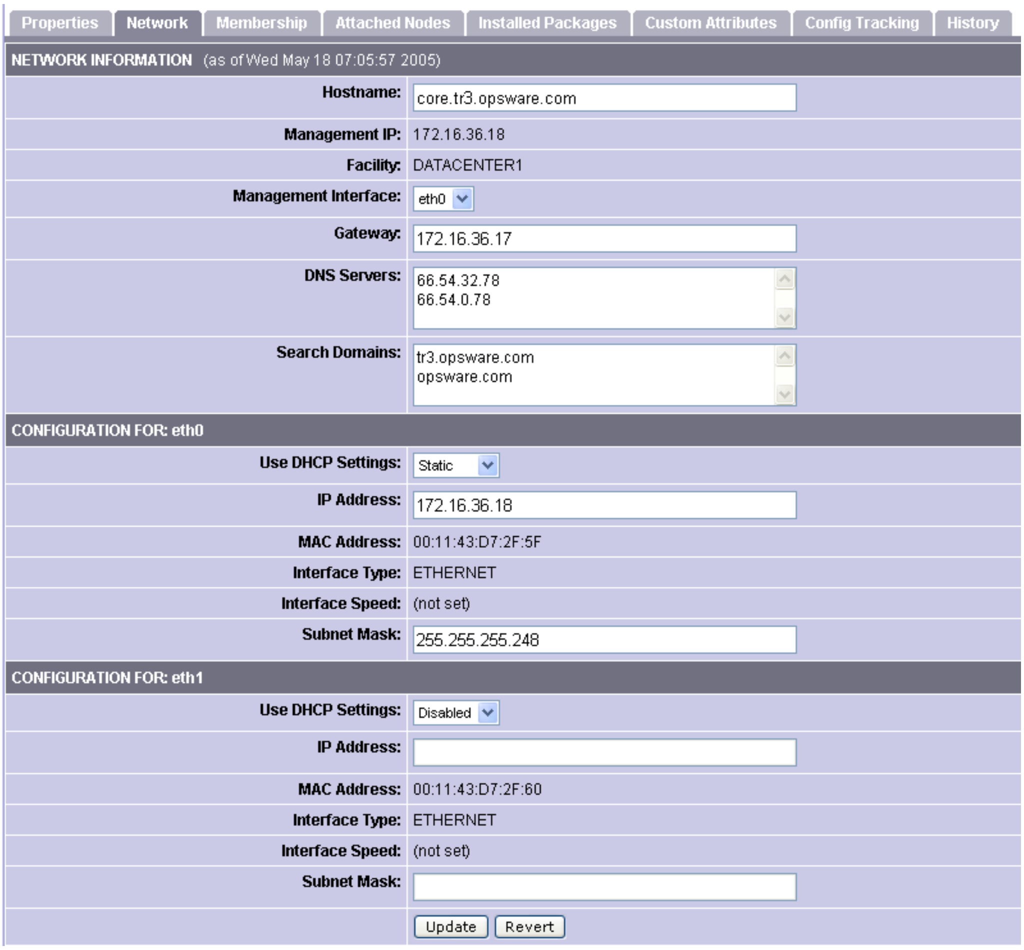

NETWORK INFORMATION (as of Wed May 18 07:05:57 2005)

Hostname:	core.tr3.opsware.com
Management IP:	172.16.36.18
Facility:	DATACENTER1
Management Interface:	eth0
Gateway:	172.16.36.17
DNS Servers:	66.54.32.78 66.54.0.78
Search Domains:	tr3.opsware.com opsware.com

CONFIGURATION FOR: eth0

Use DHCP Settings:	Static
IP Address:	172.16.36.18
MAC Address:	00:11:43:D7:2F:5F
Interface Type:	ETHERNET
Interface Speed:	(not set)
Subnet Mask:	255.255.255.248

CONFIGURATION FOR: eth1

Use DHCP Settings:	Disabled
IP Address:	
MAC Address:	00:11:43:D7:2F:60
Interface Type:	ETHERNET
Interface Speed:	(not set)
Subnet Mask:	

Figure 8-2 shows the Hardware view in the Manage Servers list, which displays the management interface for the server in the Network Info column. (To access the Hardware view, choose **Hardware** from the **View** menu.)

Figure 8-2: Hardware Tab in the Manage Servers List

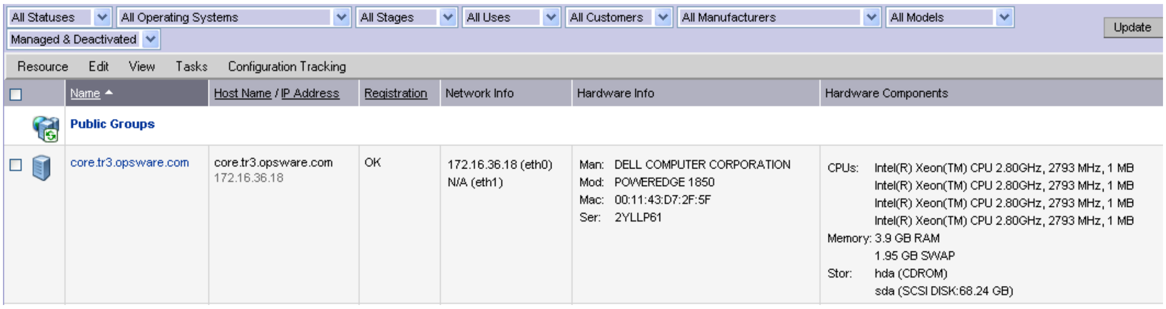

☐	Name ▲	Host Name / IP Address	Registration	Network Info	Hardware Info	Hardware Components
	Public Groups					
☐	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	OK	172.16.36.18 (eth0) N/A (eth1)	Man: DELL COMPUTER CORPORATION Mod: POWEREDGE 1850 Mac: 00:11:43:D7:2F:5F Ser: 2YLLP61	CPUs: Intel(R)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

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

Code Deployment and Static NAT

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

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

See “Code Deployment and Rollback” on page 401 in Chapter 11 for information about how to use CDR.

Setting the Primary IP Address of a Server

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

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

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

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

   Or

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

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

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

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

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

5. Click **Update**.

See "Using the Search Feature" on page 244 in Chapter 7 for more information. See "Server Searching by IP Address" on page 260 in Chapter 7 for more information.

NAT Table Mapping and Managed Servers

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

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

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

- If an IP address mapped to a particular server is mapped to a different server, both servers become unmanageable and might be marked Not Reachable on the Manage

Servers: Status page. This problem is resolved when one of the two servers reports its IP address during hardware registration. The other server remains unmanageable until the server registers with the Server Agent. Both servers eventually become manageable.

Server Groups

This section provides information on server groups within HP Server Automation and discusses the following topics:

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

Overview of Server Groups

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

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

The My Servers feature can also be used to gather servers and server groups, but it has different functionality than the Server Groups feature. You can add individual servers, server groups, and nested groups that you access frequently to My Servers. See “My Servers” on page 240 in Chapter 7 for information about using this feature.

Uses for Server Groups

Some recommended uses for server groups include:

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

Permissions Required for Working with Server Groups

Users must have the permissions shown in Table 8-1 in order to perform specific tasks related to server groups. Only administrators can set permissions.

Table 8-1: Permissions Required for Working with Server Groups

NAME OF PERMISSION	WHERE SELECTED	ENABLES YOU TO
Manage Servers		Create, edit, and delete private server groups, both static and dynamic.
Manage Public Server Groups	The Manage Servers Permissions section on the Other tab in Users and Groups	Create, edit, and delete public server groups, both static and dynamic.
Model Public Server Groups	The Manage Servers Permissions section on the Other tab in Users and Groups	Model public server groups. See "Public Group Modeling" on page 286 in this chapter for more information.

Table 8-1: Permissions Required for Working with Server Groups (continued)

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

Characteristics of Server Groups

When using server groups, groups have the following characteristics:

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

    See "Types of Server Groups" on page 285 in this chapter for more information about private and public groups.
  - The rules for a dynamic group are not inherited from a parent dynamic group to a dynamic subgroup.

    See "Dynamic Groups" on page 286 in this chapter for more information about the characteristics of dynamic groups.
  - Groups do not inherit modeling data from their parents, including custom attributes.
  - When you run an operation on a group that contains nested groups, the operation also applies to all the servers in the nested groups below the current group.

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

Types of Server Groups

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

Private Groups

If you belong to a user group that has access to the Manage Servers list, you can create groups that you alone can see and work with. Only you see your private groups. Other SA users cannot see them. Private groups behave the same way as public groups, with the exception that modeling is not available for private groups. See “Public Group Modeling” on page 286 in Chapter 8 for information about how the SA model affects groups.

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

Public Groups

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

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

Only public groups can be used for modeling.

Static Groups

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

Dynamic Groups

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

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

HP Server Automation calculates server group membership each time any of the following actions occur:

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

Additionally, HP Server Automation automatically recalculates dynamic group membership every hour.

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

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

Public Group Modeling

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

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

Manage Servers Display for Public and Private Groups

As Figure 8-3 shows, server groups appear at the top of the Manage Server list. Public Groups appear as a link at the top of the list. Private groups, if any have been created, appear next, followed by the list of individual servers.

Figure 8-3: Manage Servers Displaying Public and Private Groups — Servers and Group View

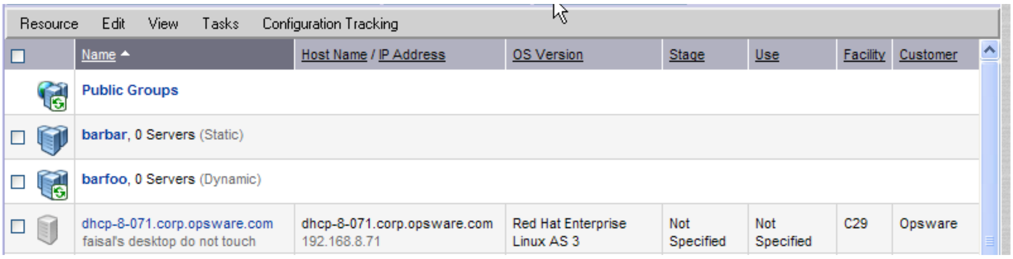

□		Name ▲	Host Name / IP Address	OS Version	Stage	Use	Facility	Customer
		Public Groups						
□		barbar, 0 Servers (Static)						
□		barfoo, 0 Servers (Dynamic)						
□		dhcp-8-071.corp.opsware.com faisal's desktop do not touch	dhcp-8-071.corp.opsware.com 192.168.8.71	Red Hat Enterprise Linux AS 3	Not Specified	Not Specified	C29	Opsware

Manage Servers List - Groups Only

When you select **View ➤ Groups Only** from the menu, information about the number of servers in a group, the number of groups in a group, and whether the group is static or dynamic appears in the list of groups, as Figure 8-4 shows.

Figure 8-4: Manage Servers List — Group Only View

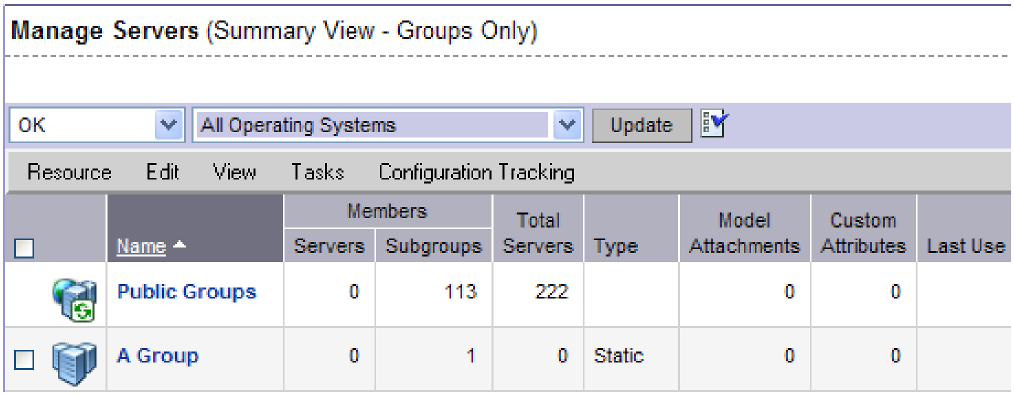
Manage Servers (Summary View - Groups Only)

□		Name ▲	Members		Total Servers	Type	Model Attachments	Custom Attributes	Last Use
			Servers	Subgroups					
		Public Groups	0	113	222		0	0	
□		A Group	0	1	0	Static	0	0	

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

Operations on Server Groups

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

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

- Run
  - Script
  - Custom Extension
  - Communication Test
- Customer
- Usage
- SA Client operations
  - Configure Application
  - Audit Application Configurations
  - Perform Server Audit
  - Create Server Snapshot
  - Install Patch
  - Install software

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

- Run Script
- Run Custom Extension

Server Groups Tabs

When you view group properties, the information for the group appears in the page, as Figure 8-5 shows.

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

- By clicking the This Group link at the top of each server group list.

- By selecting a group in the list and choosing **Resource ➤ Properties** from the menu.

Figure 8-5: Tabs Available for Server Groups

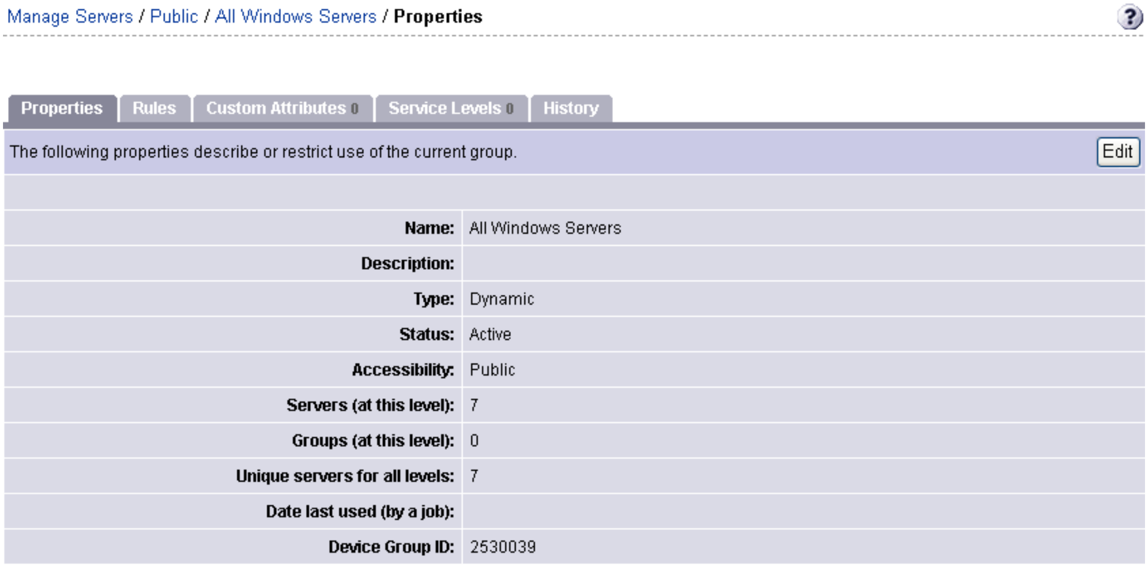

Manage Servers / Public / All Windows Servers / **Properties**
?
Properties
Rules
Custom Attributes 0
Service Levels 0
History
Edit

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

Name:	All Windows Servers
Description:	
Type:	Dynamic
Status:	Active
Accessibility:	Public
Servers (at this level):	7
Groups (at this level):	0
Unique servers for all levels:	7
Date last used (by a job):	
Device Group ID:	2530039

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

To perform any of these actions, you must have the correct HP Server Automation permissions. See "Permissions Required for Working with Server Groups" on page 283 in this chapter for more information.

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

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

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

See "Creating Groups by Using Search" on page 295 in this chapter for more information about specifying rules for a dynamic group.

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

  See "Custom Attributes for Servers" on page 338 in this chapter for more information about how custom attributes affect SA-managed servers.

- **Service Levels**: This allows you to attach service levels to a group.

  See "Service Levels" on page 342 in this chapter for more information.

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

Ways to Create Server Groups

Server Groups can be created by:

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

Creating Static Groups by Using the New Group Option

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

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

To create static server groups, perform the following steps:

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

2. From the **Resource** menu, choose **New Group**. The New Private Static Group dialog appears, as shown in Figure 8-6.

Figure 8-6: Default New Private Static Group Dialog

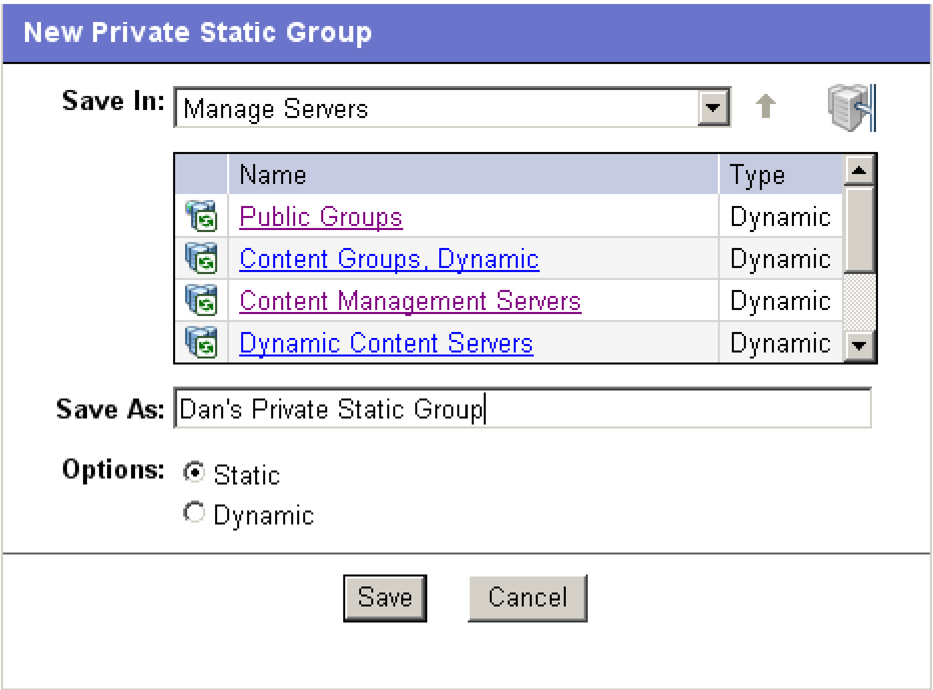

New Private Static Group

Save In: Manage Servers

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

Save As: Dan's Private Static Group

Options: ● Static ○ Dynamic

3. To create a Private Static group in the top level of Manage Servers, perform the following steps:

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

4. To create a Private Static group below another group, perform the following steps:

    1. Navigate to the group below which you want to create a new group.
    2. Enter the name of the group in the Save As text box.
    3. Click **Save**.

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

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

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

Creating Static Groups by Using the Copy to Group Dialog

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

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

2 Navigate to the servers and groups that you want to copy and click the check boxes next to the servers and groups you want to copy.

3 From the **Edit** menu, choose **Copy to Group**. The Copy to Group dialog box appears.

4 Navigate to the place in the group hierarchy where you want to copy the servers and group then by clicking the group links displayed in the Name and Type field.

5 In the Options field, select either of the following options:

- **Maintain Hierarchy**: Select this option to copy any server groups exactly as they are.
- **Expand to a Flat List**: Select this option to copy only the servers within the group to the new group.

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

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

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

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

7. Enter the name of your newly-created group.
8. Click **OK**. The name dialog closes, and the Copy to Group dialog reappears.
9. Click **OK** on the Copy to Group dialog.

Creating Dynamic Groups by Using the New Group Option

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

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

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

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

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears.
2. From the **Resource** menu, choose **New Group**. The New Private Static Group dialog appears.
3. To create a private dynamic group, navigate to the place in the group hierarchy where you want to create the group by clicking the group links displayed in the Name and Type field.

   Or

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

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

4. Enter the name of the group in the Save As text box.
5. In the Options field, select Dynamic.
6. Click **Create Rules**. The Manage Servers Properties page appears, with the Rules tab selected, as Figure 8-7 shows.

Figure 8-7: Manage Servers Properties Page with Rules Tab Selected

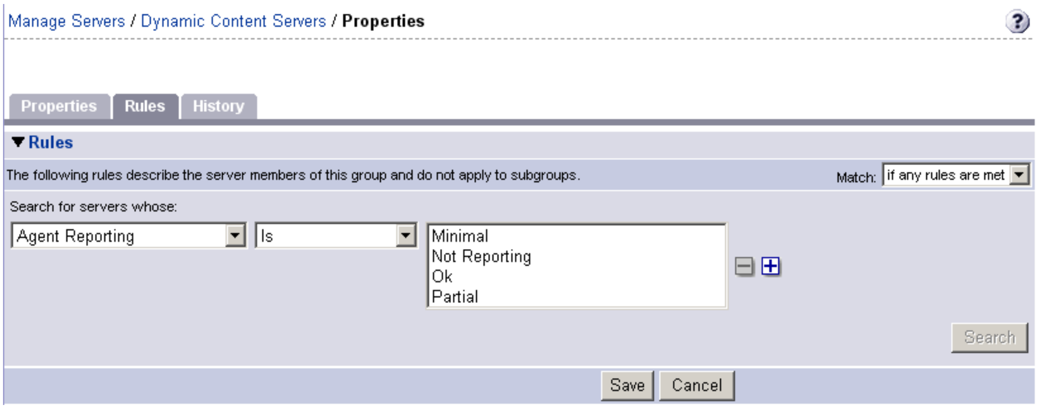

Manage Servers / Dynamic Content Servers / **Properties**

The following rules describe the server members of this group and do not apply to subgroups.

Match: if any rules are met

Search for servers whose:

Agent Reporting | Is | Minimal, Not Reporting, Ok, Partial

7. Select the criteria that apply to the servers you would like the group to include.

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

8. From the Match drop down list select either "If all rules are met" or "If any rules are met."

9 (Optional) Click **Search** to apply the rules to a server search to validate the results.

10 Click **Save** to save the rules that apply to your group.

Creating Groups by Using Search

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

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

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

2 Use the criteria to create the rules used for a server search and to identify servers for dynamic groups. The options in the user interface for specifying dynamic group rules and for using Search are the same. See “Criteria for Search and Dynamic Group Rules” on page 295 in this chapter for more information.

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

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

Or

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

The New Dynamic Group dialog appears.

Criteria for Search and Dynamic Group Rules

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

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

Table 8-2: Rules for Server Search and for Creating Dynamic Groups

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
PROPERTIES		
Agent Discovery Date: The date that the Server Agent was installed.	• Is after • Is before	Drop-down lists with the day, month, and year
	• Is within the last	User-entered text
	• Is today	N/A
Agent Reporting: Whether the Server Agent is reporting to HP Server Automation.	• Is • Is not	• Has not reported • OK • Registration in progress • Reporting error
Agent Status: Whether the Server Agent is reachable by HP Server Automation.	• Is • Is not	• Not reachable • OK
Agent Version: The version of the Server Agent – such as 14.2.3b.	• Contains • Does not contain • Is • Is not	User-entered text
Custom Attribute (any): The name of a custom attribute that is associated with the server through attachment or inheritance.	• Contains • Is	User-entered text
Custom Attribute (local): The name of a custom attribute that is locally attached to the server,	• Contains • Is	User-entered text

Table 8-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Customer: The customer or account that the server is associated with.	• Is • Is not	Popup window of customers
Deployment Stage: The stage the server performs within a lifecycle environment.	• Is • Is not	• In Deployment • Live • Not Specified • Offline The values that appear in this list are customizable; in addition to the values above, values specific to your environment might appear.
Facility: The collection of servers managed by an HP Server Automation installation.	• Is • Is not	Popup window of facilities When HP Server Automation is running multimaster mode, the list can contain many facilities.
Group Membership: Whether the server belongs to a group.	• Is	Popup window of groups
Host name: The host name of the server – such as m004.company.com.	• Contains • Does not contain • Is • Is not	User-entered text

Table 8-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Name (any): This enables searching for any name or IP address associated with a server.	• Contains • Does not contain • Is • Is not	User-entered text
Notes: The contents of the Notes field from the Properties tab for a server.	• Contains • Does not contain • Is • Is not	User-entered text
Opsware Display Name: The user-configurable name for the server.	• Contains • Does not contain • Is • Is not	User-entered text By default, HP Server Automation uses the configured host name of the server until a user edits it.
Server Use: How the server is being used – such as Development, Staging, Production.	• Is • Is not	• Development • Not Specified • Production • Staging You can customize values that appear in this list. In addition to the values above, values specific to your environment might appear.

Table 8-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Service Level: The user-defined category that can be used as an organizational tool.	• Is • Is attached here or below	Popup window of service levels Servers can be associated with multiple service levels.
OPSWARE PROPERTIES		
Application Configuration: Whether the server uses the Application Configuration feature.	• Is not used • Is used	N/A
Code Deployment: Whether the server uses the Code Deployment feature.	• Is not used • Is used	N/A
Configuration Tracking: Whether the Configuration Tracking feature is monitoring or backing up specific files or configurations on a server.	• Is off • Is on	N/A
Lifecycle: The server states that are part of bringing a server into HP Server Automation.	• Is • Is not	• Available • Build Failed • Deactivated • Installing OS • Managed

Table 8-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Server ID: The internal ID HP Server Automation uses to identify the server.	• Is • Is not	User-entered text In most cases, the Server ID is the same as the MID.
SOFTWARE		
Attached Software: The software that is assigned or modeled through the SA remediate operation – the installation process.	• Is • Is attached here or below	Popup window of software
Attached Software Policies: The software policies that is assigned or modeled through the SA installation process.	• is	Popup window of software policies
Installed Patches: Whether a patch is installed on the server.	• Contains • Does not contain • Is • Is not	User-entered text
Installed Software: The package reported installed on the server.	• Contains • Does not contain • Is • Is not	User-entered text A package does not have to be installed by HP Server Automation to be reported as installed on a server.
OS Version: The OS version defined by OS definitions in the OS Provisioning feature.	• Is • Is not	Popup window of operating systems

Table 8-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Reported OS: For Windows – the version reported by the OS, for Unix – the version returned by the uname command.	• Contains • Does not contain • Is • Is not	User-entered text
Windows Service: The names of the Windows services that are reported to HP Server Automation.	• Contains • Does not contain • Is • Is not	User-entered text
NETWORK		
DNS Search Domains: The domains configured to be searched in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
DNS Servers: The IP addresses of the DNS servers configured in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
Default Gateway: The IP address of the default router.	• Contains • Does not contain • Is • Is not	User-entered text

Table 8-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
IP Address: Any Internet Protocol address for the server.	• Contains • Does not contain • Is • Is not	User-entered text
MAC Address: Any Media Access Control address, which is the network interface card's unique hardware number.	• Contains • Does not contain • Is • Is not	User-entered text
WINS Servers: The Windows Internet Naming Servers configured in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
HARDWARE		
CPU Make and Model: The vendor name and CPU model for the server – such as GENUINEINTEL Intel(R) Pentium(R) 4 CPU 2.60GHz.	• Is • Is not	Popup window of CPU makes and models
CPU Speed: The Central Processing Unit speed in gigahertz [GHz].	• Does not equal • Equals • Is greater than • Is less than	User-entered text A 600 Mhz machine should be entered as `0.6`.
Make and Model: The vendor name and server model for the server – such as Compaq - DL360.	• Is • Is not	Popup window of server makes and models

Table 8-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Number of CPUs: The number of CPUs on the server.	• Does not equal • Equals • Is greater than • Is less than	User-entered text
RAM: The amount of RAM on the server in megabytes [MB].	• Does not equal • Equals • Is greater than • Is less than	User-entered text To enter 1 Gigabyte, type `1024`.
Serial Number: The serial number of the server.	• Contains • Does not contain • Is • Is not	User-entered text
Storage Make and Model: The vendor name and storage model for the server – such as WDC - WD800BB-75DKA0.	• Is • Is not	Popup window of storage makes and models
CUSTOM FIELDS		
A **Numeric** field	• Does not equal • Equals • Is greater than • Is less than	User-entered text
A **String** field	• Contains • Does not contain • Is • Is not	User-entered text

Table 8-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
A **URI** field	• Contains • Does not contain • Is • Is not	User-entered text
A **File** field	• Contains • Does not contain • Is • Is not	User-entered text
A **Date** field	• Is after • Is before	Drop-down lists with the day, month, and year
	• Is within the last	User-entered text
	• Is today	N/A

Adding a Server to Static Groups, Both Public and Private

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

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

2. Navigate to the static group that you want to add servers and server groups to.

3. Click the check box next to the group and select **Edit ➤ Add Servers**. The Select Servers and Groups to Add to [group name] window opens. The window is populated with the same servers and groups visible on the Manage Servers page. You can use the Status, OS, and Customer filters to change the servers and groups that appear on the list.

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

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

Adding a Server to Dynamic Groups, Both Public and Private

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

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

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

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

1. Click the check box next to the name of the server group to which you want to add a server group, and select **Edit ➤ Copy to Group**. Then follow the steps for creating a new group, either static or dynamic. See “Ways to Create Server Groups” on page 290 in this chapter for more information.

   Or

   Navigate to the server group that you want to add a server group to.
2. Click the check box next to “This Group” and select **Resource ➤ New Group**. Then follow the steps for creating a new group, either static or dynamic. See “Ways to Create Server Groups” on page 290 in this chapter for more information.

Removing Servers from Static Groups, Both Public and Private

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

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

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

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

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

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

Removing Servers from Dynamic Groups, Both Public and Private

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

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears.
2. Navigate to the Properties page of the dynamic group that you want to add rules to, by drilling down into the group and clicking the "This Group" link.
3. Select the Rules tab.
4. Click **Edit**.

5 Add one or more lines of criteria to the existing rules by following the procedure described in the “Creating Dynamic Server Groups by Using the New Group Option” topic, step 3 on page 295.

Moving Servers from One Static Group to Another

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

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

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

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

2 Navigate to the group containing the servers you want to move and select those servers.

3 From the **Edit** menu, choose **Move**. The Move [group name] Group dialog box appears.

4 Navigate to the place in the group hierarchy where you want to create the group by clicking the group links displayed in the Name and Type field.

5 Click **Move**.

Duplicating Server Groups

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

1 Navigate to the group that you want to duplicate.

2 Click the check box next to the group name.

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

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

Figure 8-8: Duplicate Group Dialog Box

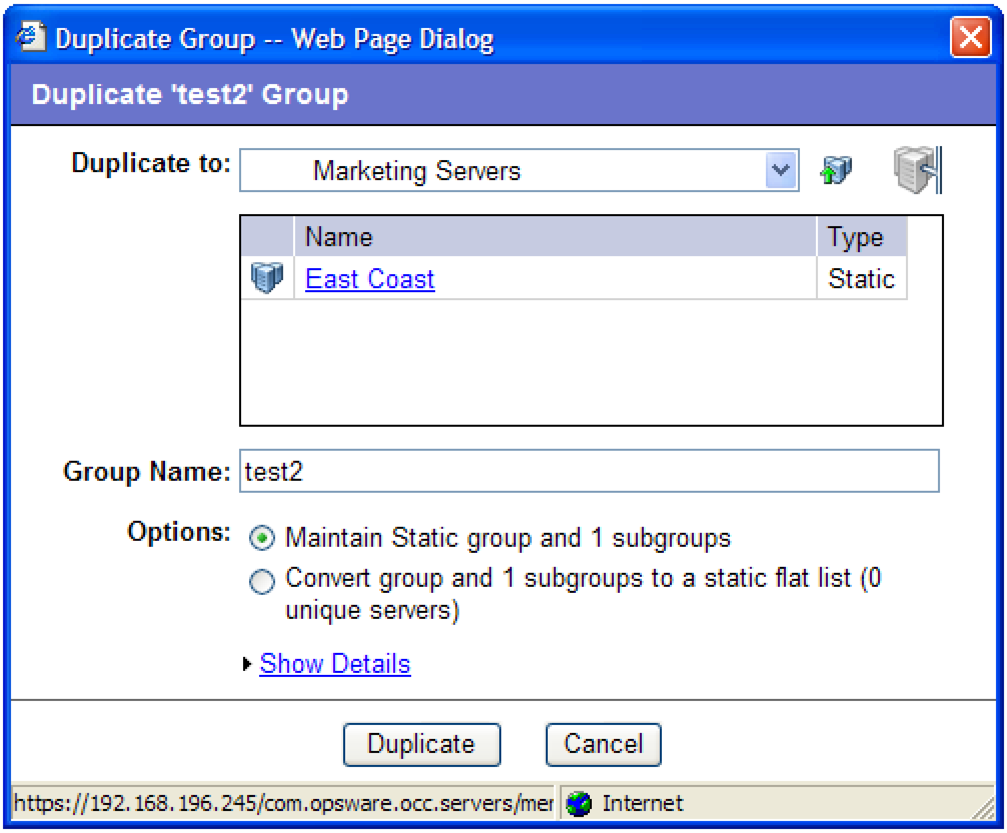

Duplicate Group -- Web Page Dialog

Duplicate 'test2' Group

Duplicate to: Marketing Servers

	Name	Type
	East Coast	Static

Group Name: test2

Options:
● Maintain Static group and 1 subgroups
○ Convert group and 1 subgroups to a static flat list (0 unique servers)

https://192.168.196.245/com.opsware.occ.servers/mer

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

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

5 Enter the name of the new group.

6. The values on the Options field will vary depending on whether the group duplicated is static or dynamic.

7. If the group is static, select one of the following options:

   - **Maintain Static group and <number> subgroups**: Selecting this option copies the group and its hierarchy as is.

   - **Convert group and <number> subgroups to a static flat list**: Selecting this option copies the group and flattens the hierarchy.

8. If the group is dynamic, select one of the following options:

   - **Maintain rules & group hierarchy**: Selecting this option copies the group and leaves it as a dynamic group, it also copies any subgroups.

   - **Convert to static group, maintain hierarchy**: Selecting this option copies the group, but turns it into a static group with the current servers defined by the rules and also copies any subgroups.

   - **Convert group and <number> subgroups to a static flat list**: Selecting this option copies the group, and turns it into a static group and flattens any hierarchy.

9. Click **Duplicate**. The newly-duplicated group appears in the selected destination.

Rules for Deleting Server Groups

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

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

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

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

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

- Software is attached to the group or a subgroup of the group.

- Access control boundaries are attached to the groups or a subgroup of the group.

- Servers and groups are selected together for deletion.

Deleting a Server Group

To delete a server group, perform the following steps:

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

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

Server Life Cycle

This section provides information on the server life cycle within HP Server Automation and contains the following topics:

- OS Provisioning and the Server Life Cycle
- Server Properties
- Server Management Tasks Related to the Server Life Cycle
- Changing the Use and Stage Values for Servers
- Editing the Properties of a Server
- Tasks Associated with Deactivating a Server
- Deactivating a Server
- Deleting a Server from HP Server Automation
- Server Management Jobs

OS Provisioning and the Server Life Cycle

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

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

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

During OS provisioning, servers progress through the HP Server Automation life cycle state changes:

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

Table 8-3 describes the HP Server Automation server life cycle values.

Table 8-3: HP Server Automation Life Cycle Values for Servers

MANAGED SERVER LIFE CYCLE VALUE	DESCRIPTION
Server Pool Values	
Planned	Indicates that a device record has been created for the server, but an SA build agent has not yet been installed. Servers in this stage cannot be provisioned until an OS build agent is installed. For more information on this unprovisioned lifecycle value, see your SA Administrator.
Available	Indicates a server on which the OS Build Agent was installed and is running, but the target OS has not been installed on the server. The OS Build Agent is a small agent that can run in the memory of the bare-metal server. See *SA User’s Guide: Application Automation* for more information about Operating System Provisioning.
Installing OS	Indicates that a user is installing the target OS on the server. The server stays in the Server Pool list until the installation process finishes successfully. Then the server moves to the Manage Servers list. See *SA User’s Guide: Application Automation* for more information about installing OS.
Build Failed	Indicates a server on which the OS Build Agent was installed and is running, but the installation of the target OS failed. The server remains in the Server Pool list with this status for 7 days before HP Server Automation deletes the entry. See *SA User’s Guide: Application Automation* for more information about Operating System Provisioning.

Table 8-3: HP Server Automation Life Cycle Values for Servers (continued)

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

Table 8-4 describes server states in SA in regard to the server life cycle.

Server Properties

Table 8-4: Server Status Icons in SA

SERVER ICON	DESCRIPTION
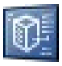	**Planned** Indicates that a device record has been created for the server, but an SA OS Build Agent has not yet been installed on it. Servers in this stage cannot be provisioned until the OS Build Agent is installed. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
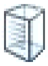	**Unprovisioned – Unreachable** Indicates a server that has been registered with the core via the SA OS Build Agent, but has not reported as ready for provisioning recently. This may be due to networking problems between the server and the SA core or the server having been disconnected or powered off. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
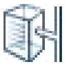	**Unprovisioned – Reachable** Indicates a server that has been registered with the core via the SA OS Build Agent and is available to have a target OS installed on it. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.

Table 8-4: Server Status Icons in SA (continued)

SERVER ICON	DESCRIPTION
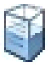	**Provisioning – Unreachable** Indicates a server on which the OS Provisioning feature was in the process of installing the target OS, but for some reason stopped because the server is unable to communicate with the SA core. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
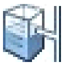	**Provisioning – Reachable** Indicates a server on which the OS Provisioning feature is in the process of installing the target OS. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
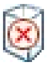	**Provisioning Failed – Unreachable** Indicates an available server on which an error occurred while the OS Provisioning Subsystem was installing a target OS, and that the server is not able to communicate with the SA core. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.
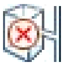	**Provisioning Failed – Reachable** Indicates an available server on which an error occurred while the OS Provisioning Subsystem was installing a target OS. In the SA Client, appears in the Unprovisioned Server list. In the SAS Web Client, appears in the My Jobs panel home page, in the My Jobs page, and in the Server Pool list.

Table 8-4: Server Status Icons in SA (continued)

SERVER ICON	DESCRIPTION
🖥	**Managed – Reachable** Indicates a server has an An Server Agent is running on it and that it is able to communicate with the SA core. In the SA Client, appears in the All Managed Servers list and Virtual Servers list. In the SAS Web Client, appears in the My Jobs panel in the home page, in the list in the My Jobs page, in the Manage Servers list, and in the server lists in the SA wizards.
🖥	**Managed – Unreachable** Indicates a managed server cannot communicate with the SA core (it is Not Reachable). If you want to discover reasons why the managed server is unreachable, you can run a Communication Test. See "Agent Reachability Communication Tests" on page 129 in Chapter 4 for more information. In the SA Client, appears in the All Managed Servers list and Virtual Servers list. In the SAS Web Client, appears in the My Jobs panel in the home page, in the list in the My Jobs page, in the Manage Servers list, and in the server lists in the SA wizards.

Table 8-4: Server Status Icons in SA (continued)

SERVER ICON	DESCRIPTION
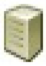	**Unmanaged** Indicates the server is unmanaged, which means it does not have an Server Agent installed on it. For unmanaged virtual servers, this means that someone created a VMware Virtual Machine (VM) or Solaris local zone outside of SA and thus does not have an Server Agent installed on it. For virtual servers, this state could also mean that the VMware VM is has not yet been provisioned, or that your user belongs to a group that does not have permissions to perform operations on this virtual server. For more information on installing an Server Agent on an unmanaged server, see Chapter 4, “Agent Management” on page 105 of this guide.
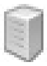	**Deactivated** Indicates a server that was deactivated in HP Server Automation so that it is currently not managed and is no longer reachable. Appears in the Manage Servers list and in the server lists in the SA wizards (however, it is not selectable in the wizards).
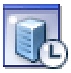	**Scheduled** Indicates a server that is scheduled for an operation (install software, uninstall software, and so forth). In the SA Client, appears in the Job Logs list. In the SAS Web Client, appears in the My Jobs panel in the home page and in the list in the My Jobs page.
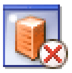	**Error** Indicates a managed server on which an error occurred while HP Server Automation was installing or uninstalling software. Appears in the My Jobs panel in the home page and in the list in the My Jobs page.

Table 8-4: Server Status Icons in SA (continued)

SERVER ICON	DESCRIPTION
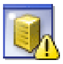	**Warning** Indicates a managed server on which a warning occurred while HP Server Automation was installing or uninstalling software. Appears in the My Jobs panel in the home page and in the list in the My Jobs page.
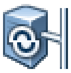	**Application Configuration Out of Sync** Indicates a managed server on which the configuration file on the server is out of sync with the Application Configuration Template (SA model). Appears only in the Application Configuration feature and the server list in the SA Client.
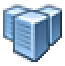	**Static Device Group** Indicates a static server group. The same states that apply to single servers apply to groups. See "Server Groups" on page 282 in this chapter for information about the different types of server groups.
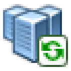	**Dynamic Device Group** Indicates a dynamic server group. The same states that apply to single servers apply to groups.
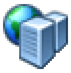	**Public Static Device Group** Indicates a public and static server group. The same states that apply to single servers apply to groups.
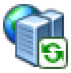	**Public Edenic Device Group** Indicates a public and dynamic server group. The same states that apply to single servers apply to groups.

Figure 8-9 shows the Server Properties columns. Table 8-5, Table 8-6, and Table 8-7 describe the Status, Stage, and Use properties for managed servers.

Figure 8-9: Server Properties Columns in the Manage Servers List

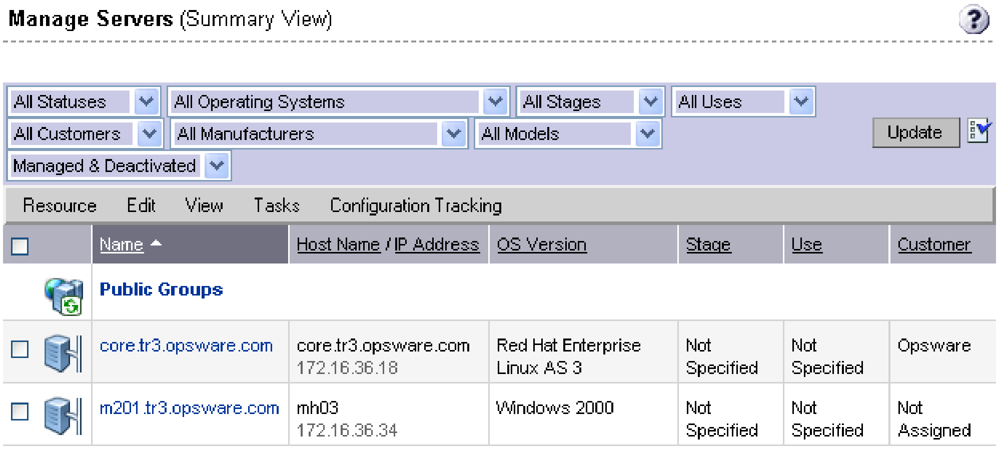

Manage Servers (Summary View)

□	Name	Host Name / IP Address	OS Version	Stage	Use	Customer
	Public Groups					
□	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	Red Hat Enterprise Linux AS 3	Not Specified	Not Specified	Opsware
□	m201.tr3.opsware.com	mh03 172.16.36.34	Windows 2000	Not Specified	Not Specified	Not Assigned

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

Status (short for Agent Status) is set automatically by HP Server Automation.

HP Server Automation toggles each server between OK and Not Reachable by reachability checks.

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

Table 8-5: Values for the Status Property for Managed Servers

STATUS VALUE	DESCRIPTION
OK	Server is manageable by HP Server Automation. Represented as text (OK) in the properties page for a server. Represented as an icon in the Manage Servers and Server Pool lists:

Table 8-5: Values for the Status Property for Managed Servers (continued)

STATUS VALUE	DESCRIPTION
Not Reachable	Server is unreachable by HP Server Automation due to an error (for example, it cannot connect to the SA core); automatically set by HP Server Automation. Represented as text (Not Reachable) in the properties page for a server. Represented as an icon in the Manage Servers list: If you want to discover reasons why the managed server is unreachable, you can run a Communication Test. See “Agent Reachability Communication Tests” on page 129 in Chapter 4 for more information.

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

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

Your SA administrator can change the values for the Stage property. By default, HP Server Automation is installed with the following Stage values.

Table 8-6: Values for the Stage Property for Managed Servers

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

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

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

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

Table 8-7: Values for the Use Property for Managed Servers

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

Server Management Tasks Related to the Server Life Cycle

Managing servers in HP Server Automation involves the following standard tasks:

- Bringing a new server into HP Server Automation so that it appears in the Server Pool

  See *SA User's Guide: Application Automation* for more information about OS Provisioning.

- Installing an operating system on a server

  See *SA User's Guide: Application Automation* for more information about OS Provisioning.

- Installing a patch

  See *SA User's Guide: Application Automation* for more information on installing patches.

- Installing a software

  See *SA User's Guide: Application Automation* for more information on installing software.

- Reprovisioning a server with a new OS

  See *SA User's Guide: Application Automation* for more information about OS Provisioning.

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

You can reprovision servers built by HP Server Automation and SA-managed servers by using this feature.

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

- Deactivating a managed server so that HP Server Automation no longer manages it.

  See "Tasks Associated with Deactivating a Server" on page 324 in this chapter for more information.

You accomplish server management tasks by using the following menus in the Manage Servers list:

- Resource: Allows to create new server groups, view properties of servers.
- Edit: Allows you to add servers to a group, delete server groups, remove servers from a group, deactivate servers.
- View: Allows you to view the summary, hardware, and software information of servers and server groups.
- Tasks: Allows you perform tasks such as running scripts, custom extensions and running communication tests.
- Configuration Tracking: Allows you to edit, and reconcile tracking policies.

Changing the Use and Stage Values for Servers

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

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

   Or

   Search for the server that you want to deactivate.

   See "Using the Search Feature" on page 244 in Chapter 7 for more information. See "Server Searching by IP Address" on page 260 in Chapter 7 for more information.

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

3. Choose **Edit ➤ Usage** from the menu above the Manage Servers list. A window prompts you to select different values, as Figure 8-10 shows.

Figure 8-10: Edit Server Popup Window

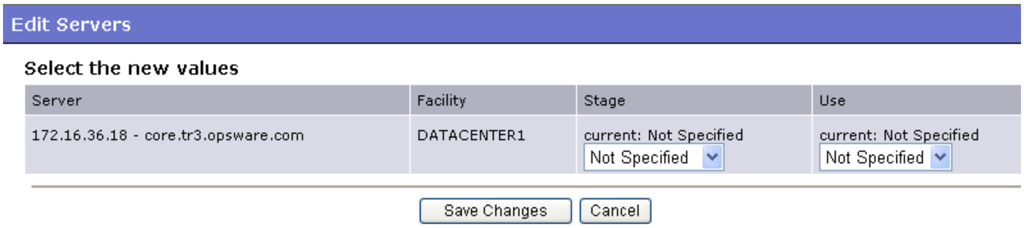

Edit Servers

Select the new values

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

4. Select the Use and Stage values from the lists.

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

Editing the Properties of a Server

You can edit a server only if you have permission in the SAS Web Client to access the customer to whom the server is associated.

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

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

   Or

   Search for the server that you want to edit.

   See “Using the Search Feature” on page 244 in Chapter 7 for more information. See “Server Searching by IP Address” on page 260 in Chapter 7 for more information.

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

3. Change any of the following properties for the server:

   - To change the name that appears in the SAS Web Client, edit the text in the Name field.

   - To change the description of the server, edit the Notes field.

- To change the customer associated with the server, select a different customer from the list. Your SA administrator defines the options for customer selections. Contact your SA administrator if the list does not contain the customer that you want to associate with this server.

You cannot change the customer associated with a server when the server is part of a CDR service, synchronization, or sequence. See the *SA Policy Setter's Guide* for more information.

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

  See "Server Properties" on page 314 in this chapter for more information.

- To change whether configuration tracking is enabled or disabled for the server, select a value from the list.

  See "Configuration Tracking" on page 429 in Chapter 12 for more information.

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

To change the custom attributes of the server, select the Custom Attributes tab. The Manage Servers: Custom Attributes page appears. See "Managing Custom Attributes" on page 339 in this chapter for more information.

Tasks Associated with Deactivating a Server

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

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

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

See "Agent Reachability Communication Tests" on page 129 in Chapter 4 for more information.

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

- Remove custom attributes from the server.
- Delete any configuration tracking policies from the server that are associated with backups.
- Set the server life cycle value to Deactivated.

You cannot deactivate a server when it is part of a CDR service, synchronization, or sequence. See the *SA Policy Setter's Guide* for more information.

Deactivating a Server

Perform the following steps to deactivate a server:

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

   Or

   Search for the server that you want to deactivate.

   See “Using the Search Feature” on page 244 in Chapter 7 for more information. See “Server Searching by IP Address” on page 260 in Chapter 7 for more information.

2. Select the servers that you want to deactivate.
3. Choose **Edit ➤ Deactivate Server** from the menu above the Manage Servers list. A confirmation dialog box prompts you to confirm the deactivation.
4. Click **OK**. The Manage Servers list refreshes and the server appears with a deactivated icon.

Deleting a Server from HP Server Automation

When you want to remove all record of a server from HP Server Automation, you can delete it.

You must deactivate a server before you can delete it from HP Server Automation.

When you delete a server from HP Server Automation, it has these effects:

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

Perform the following steps to delete a server:

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

   Or

   Search for the server that you want to delete.

   See "Using the Search Feature" on page 244 in Chapter 7 for more information. See "Server Searching by IP Address" on page 260 in Chapter 7 for more information.

2. Select the servers that you want to delete.
3. Choose **Edit ➤ Delete Server** from the menu above the Manage Servers list. A confirmation dialog box prompts you to confirm the deletion.
4. Click **OK**. The Manage Servers list refreshes and the server disappears from the list.

Server Management Jobs

This section provides information about server management jobs within HP Server Automation and contains the following topics:

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

My Jobs

The My Jobs information is available only on a per-user basis. You cannot log in as an SA administrator to see the jobs that other SA users have run. The My Jobs information appears in two places in the SAS Web Client:

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

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

My Jobs Display Information

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

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

My Jobs in the SAS Web Client

The My Jobs feature in the SAS Web Client provides information about the following Command Engine scripts:

- OS provisioning
- CDR requests
- Distributed script execution
- Custom Extensions

Jobs in SA Client

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

- Creating snapshots
- Pushing or Auditing an Application Configuration
- Auditing Servers
- Creating Server Snapshots
- Remediating policies
- Running OS Sequence
- Any jobs scheduled to be run at a future date

Jobs in the SA Client appears in the Jobs Logs feature window.

Viewing Job Details in the SAS Web Client

Perform the following steps to view the job details:

1. From the SAS Web Client home page, click the link in the My Jobs panel for the job that you want to view, as Figure 8-11 shows.

Figure 8-11: My Jobs Panel in the SAS Web Client Home Page

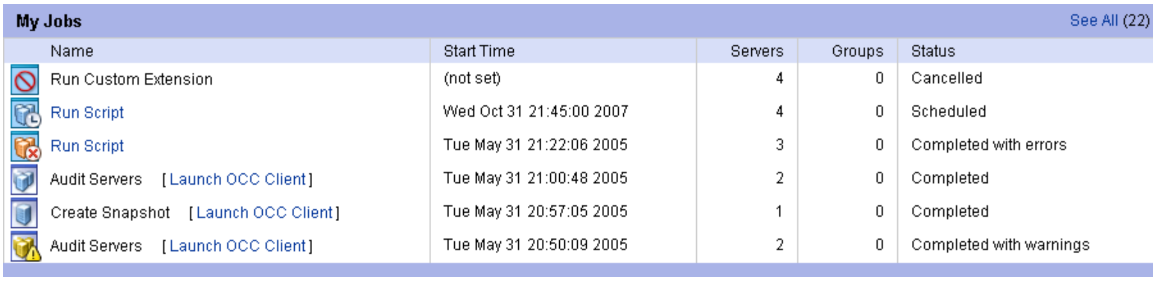

My Jobs See All (22)

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

Or

From the navigation panel, click My Jobs and then click the link for the job to open a window that shows the details of the job, as Figure 8-12 shows.

Figure 8-12: My Jobs Page Accessed from the Navigation Panel

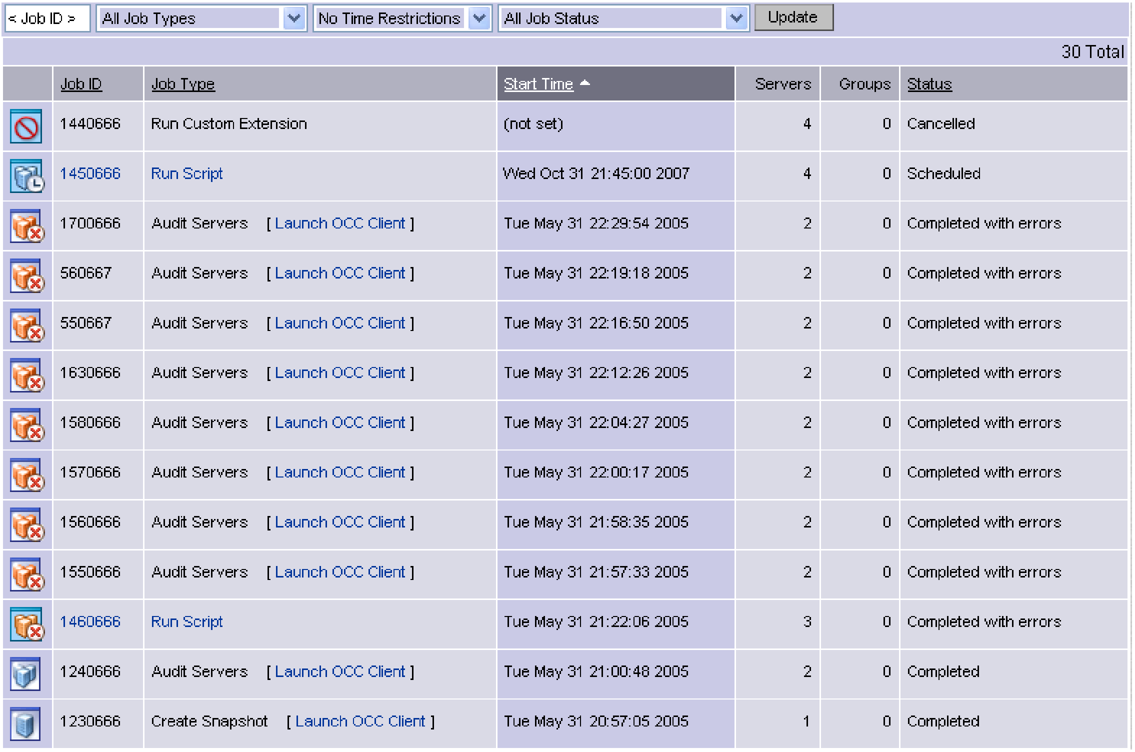

< Job ID > | All Job Types | No Time Restrictions | All Job Status | Update

30 Total

	Job ID	Job Type	Start Time ▲	Servers	Groups	Status
	1440666	Run Custom Extension	(not set)	4	0	Cancelled
	1450666	Run Script	Wed Oct 31 21:45:00 2007	4	0	Scheduled
	1700666	Audit Servers [Launch OCC Client]	Tue May 31 22:29:54 2005	2	0	Completed with errors
	560667	Audit Servers [Launch OCC Client]	Tue May 31 22:19:18 2005	2	0	Completed with errors
	550667	Audit Servers [Launch OCC Client]	Tue May 31 22:16:50 2005	2	0	Completed with errors
	1630666	Audit Servers [Launch OCC Client]	Tue May 31 22:12:26 2005	2	0	Completed with errors
	1580666	Audit Servers [Launch OCC Client]	Tue May 31 22:04:27 2005	2	0	Completed with errors
	1570666	Audit Servers [Launch OCC Client]	Tue May 31 22:00:17 2005	2	0	Completed with errors
	1560666	Audit Servers [Launch OCC Client]	Tue May 31 21:58:35 2005	2	0	Completed with errors
	1550666	Audit Servers [Launch OCC Client]	Tue May 31 21:57:33 2005	2	0	Completed with errors
	1460666	Run Script	Tue May 31 21:22:06 2005	3	0	Completed with errors
	1240666	Audit Servers [Launch OCC Client]	Tue May 31 21:00:48 2005	2	0	Completed
	1230666	Create Snapshot [Launch OCC Client]	Tue May 31 20:57:05 2005	1	0	Completed

The My Jobs page displays the operations that you performed.

2 Click **View Details** to see detailed information about the job. The My Jobs information contains a build log for the job. This build log contains any error messages that HP Server Automation generates. See Figure 8-13.

Figure 8-13: Communication Test Details Page in the My Job Window

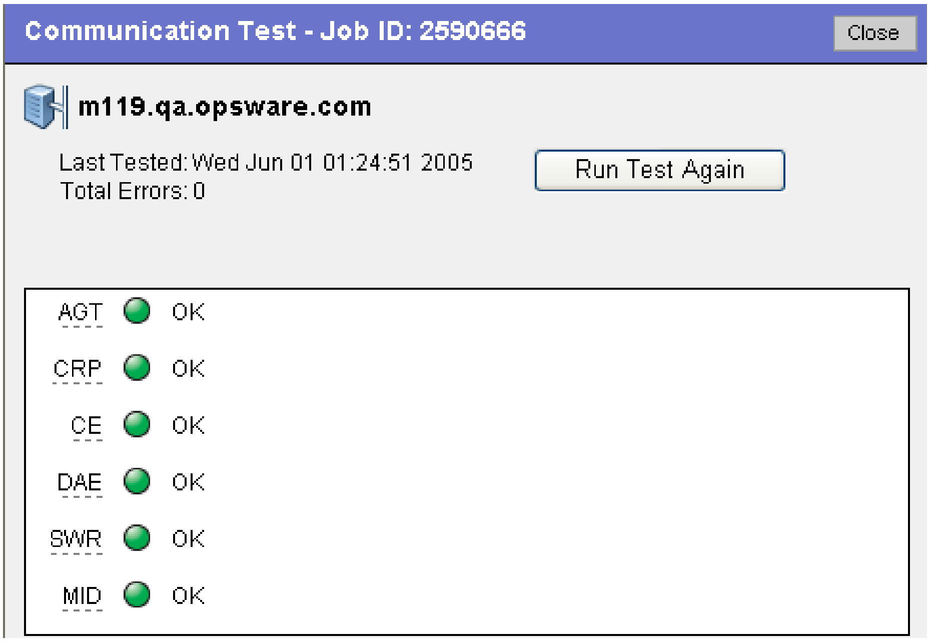

Communication Test - Job ID: 2590666

m119.qa.opsware.com

Last Tested: Wed Jun 01 01:24:51 2005
Total Errors: 0

AGT ● OK
CRP ● OK
CE ● OK
DAE ● OK
SWR ● OK
MID ● OK

Viewing My Job Details from inside the SA Client

You can also view job details from inside the SA Client. If the job was run from inside the SA Client, you will see a link next to the job named Launch SA Client, as shown in Figure 8-14.

Figure 8-14: My Jobs with link to SA Client

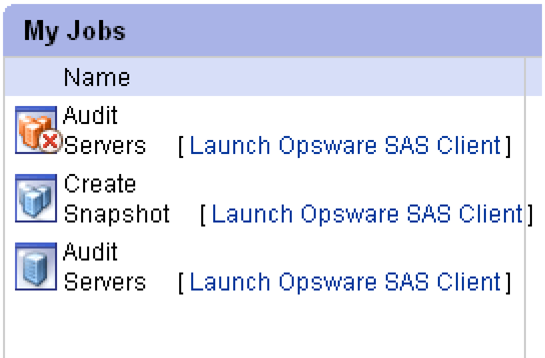

My Jobs

Name

Audit Servers [Launch Opsware SAS Client]
Create Snapshot [Launch Opsware SAS Client]
Audit Servers [Launch Opsware SAS Client]

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

Server Management Scheduling and Notification

This section contains the following topics:

- Scheduling a Job
- Sending Email Notification

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

Scheduling a Job

Perform the following steps to schedule server management tasks:

1. In the Schedule and Notify page of an SA Wizard, choose the Run Now option to execute the operation immediately or choose the Specify Time Option to schedule the operation at a later date and time. See Figure 8-15

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

2. When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

   - **Option 1**: Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.

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

Figure 8-15: Scheduling a Job in an SA Wizard

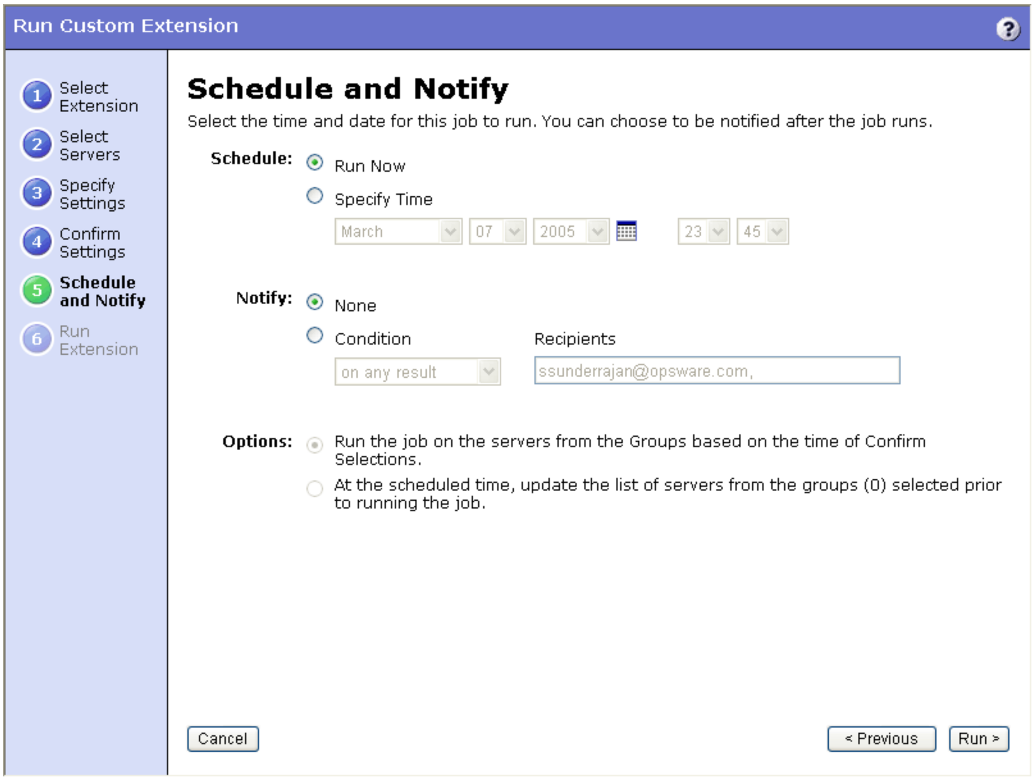

Run Custom Extension

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

Schedule and Notify

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

Schedule: ● Run Now
○ Specify Time
March | 07 | 2005 | 23 | 45

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

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

Cancel | < Previous | Run >

3 Additionally, you can view a scheduled job in the My Jobs page and change the date and time for the job to run, or cancel the job entirely. (Click the name of a scheduled job to open a window to change the date or time that the job will run or cancel it.)

Sending Email Notification

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

You have the option of sending the notification

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

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

To send an email about the job details, choose the Condition option on the Schedule and Notify page and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field. See Figure 8-16.

Figure 8-16: Notifying about a Job in the SAS Web Client

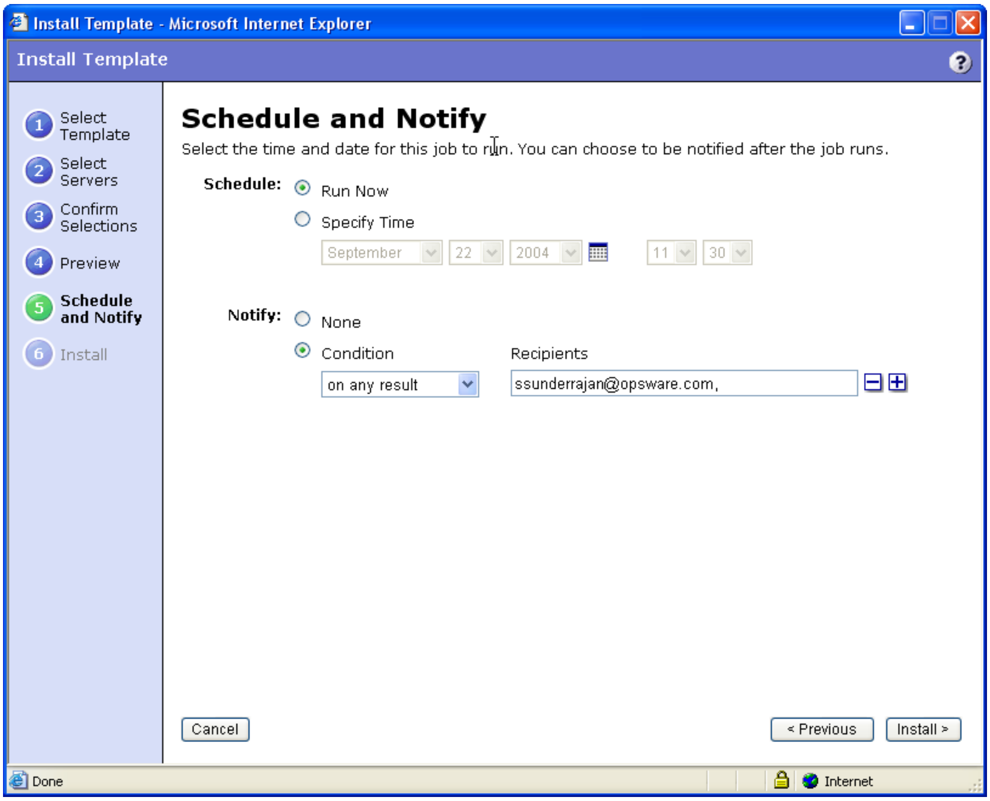

Schedule and Notify

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

Schedule: ● Run Now
○ Specify Time
September 22 2004 11 30

Notify: ○ None
● Condition
on any result

Recipients: ssunderrajan@opsware.com,

Time-Outs for Server Management Jobs

Table 8-8 describes the time-out values that apply to the server management operations in HP Server Automation.

Table 8-8: Time-Outs in HP Server Automation

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

Customizing the Monitor Time-Out Duration

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

During remediate, a periodic heartbeat occurs between the Server Agent and the Command Engine to ensure that the agent is still responsive. This setting controls the maximum amount of time that can pass between these heartbeat messages. Typically, you only need to increase this setting if you install software that reboots the server, and the time that it takes for the server to reboot and for the agent to restart exceeds the default value.

Custom Fields for Servers

This section provides information on custom fields for servers within HP Server Automation and discusses the following topics:

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

Overview of Custom Fields for Servers

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

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

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

Typed Data Supported for Custom Fields

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

- **Number**: The value provided must be a long integer.
- **Short String**: A text string that must be less than or equal to 4000 characters.

- **Long String**: A text string with no length restrictions. Custom fields with this type cannot be used in search or in the rules for dynamic server groups.
- **Date**: The value will be verified for correctness.
- **File**: Indicates a file attachment.
- **URI**: A Uniform Resource Identifier string, which is validated as a URI.

HP Server Automation validates the value you enter in a custom field based on the data type specified for the field.

Uses of Custom Fields

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

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

Changing the Value for a Custom Field

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

   Or

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

2. Click the server name. The Manage Servers: Properties page appears. Scroll to the Custom Fields section of the page. The custom fields that appear in the properties page are specific to your operational environment. Figure 8-17 shows an example of the types of custom fields that can appear for a server.

Figure 8-17: Custom Fields That Appear for a Server

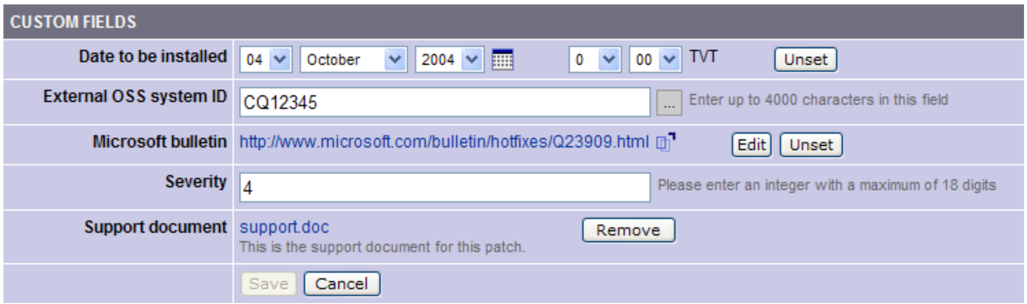

CUSTOM FIELDS

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

3. To change a value in a field requires a number or text, enter the value in the field.

4. To add a file, click **Browse** and select the file from the dialog box.

5. To remove a file, click **Remove**.

6. To specify a URL, enter the URL in the field or click **Edit** to change an existing URL.

7. To add or change a date, click the ▦ icon and select the date or enter the date in the field by selecting the appropriate day, month, and year from the drop-down lists.

8. Click **Save**.

Custom Attributes for Servers

This section provides information on custom attributes for servers within HP Server Automation and contains the following topics:

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

Overview of Custom Attributes for Servers

Users often need to store specific miscellaneous information in the SA Model Repository to facilitate server or application installation and configuration, scripting, or other purposes.

The SAS Web Client provides a data management function by allowing users to set custom attributes for servers. These custom attributes include setting miscellaneous parameters and named data values. Users can write scripts that use these parameters and values when performing a variety of functions, including network and server configuration, notifications, and CRON script configuration.

Custom attributes can be accessed by software packages at installation time to configure settings that might be unique to the installation.

For information about how to set custom attributes required by the software running on a specific server, contact the group responsible for packaging your applications, as Figure 8-18 shows.

Figure 8-18: Custom Attributes Set for a Server

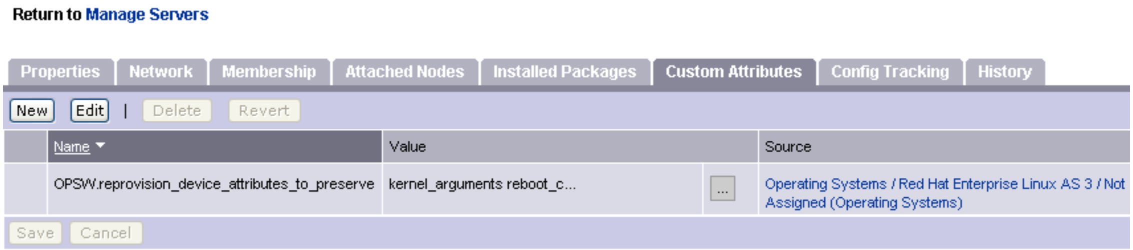

Name	Value	Source
OPSW.reprovision_device_attributes_to_preserve	kernel_arguments reboot_c...	Operating Systems / Red Hat Enterprise Linux AS 3 / Not Assigned (Operating Systems)

Managing Custom Attributes

Do not edit or remove custom attributes without verifying that the change you are making does not impact other users or critical SA operations.

To set custom attributes that affect a specific server, use the Manage Servers list. After locating the server and displaying the server properties, select the Custom Attributes tab. The SAS Web Client displays the currently defined custom attributes for the selected server.

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

See "Adding Server Custom Attributes" on page 340 in this chapter for more information.

You can also set custom attributes to a software policy. See the *SA Policy Setter's Guide* for more information.

Additionally, you can set custom attributes that affect every server associated with a specific customer or for every server in a facility. Navigate to the customer or facility where you want to set attributes. Select Environment ➤ Customers or Facilities in the navigation panel, and click the correct name in the list. Select the Custom Attributes tab, and add attributes for all the servers associated with the customer or located in the facility. When you use this option, you define custom attributes at a customer-specific or facility-specific level. The procedure to add custom attributes to a customer or facility is the same as adding custom attributes to individual servers.

Additionally, you can add custom attributes for a server group by viewing a server group, and then selecting the Custom Attributes tab for that group. The procedure to add custom attributes to a server group is the same as adding custom attributes to individual servers.

Adding Server Custom Attributes

For information about how to set custom attributes required by the software running on a specific server, contact the group responsible for packaging your applications.

Perform the following steps to add a custom attribute for a server:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server for which you want to add custom attributes.

   Or

   Search for the server for which you want to add custom attributes.

2. Click the display name of the server. The Manage Servers: Server Properties page appears.

3. Select the Custom Attributes tab. The Manage Servers: Custom Attributes page appears.

4. Click **New**.

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

6. Click **Save**.

See “Using the Search Feature” on page 244 in Chapter 7 for more information. See “Server Searching by IP Address” on page 260 in Chapter 7 for more information

Editing Server Custom Attributes

If you want to change the name of a custom attribute entry, you need to create a new custom attribute and delete the old custom attribute.

Perform the following steps to edit custom attributes for a server:

1. From the navigation panel, click Servers ➤ Manage Servers or Server Pool. The Manage Servers page appears. Browse the list to find the server for which you want to edit custom attributes.

   Or

   Search for the server for which you want to edit custom attributes.

2. Click the display name of the server. The Manage Servers: Server Properties page appears.
3. Select the Custom Attributes tab to change the custom attributes of the server. The Manage Servers: Custom Attributes page appears.
4. Click the attribute name link for the custom attribute that you want to change.
5. Update the value of the custom attribute.
6. Click **Save** to save your changes. The Manage Servers: Custom Attributes page reappears with the updated value.

See “Using the Search Feature” on page 244 in Chapter 7 for more information. See “Server Searching by IP Address” on page 260 in Chapter 7 for more information.

Deleting Server Custom Attributes

Perform the following steps to delete custom attributes for a server:

1. From the navigation panel, click Servers ➤ Manage Servers or Server Pool. The Manage Servers page appears. Browse the list to find the server from which you want to remove custom attributes.

   Or

   Search for the server from which you want to remove custom attributes.

See “Using the Search Feature” on page 244 in Chapter 7 for more information. See “Server Searching by IP Address” on page 260 in Chapter 7 for more information.

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

See “Using the Search Feature” on page 244 in Chapter 7 for more information. See “Server Searching by IP Address” on page 260 in Chapter 7 for more information.

Service Levels

This section provides information about service levels within HP Server Automation and contains the following topics:

- Overview of Service Levels
- Adding a Service Level to the SAS Web Client
- Editing a Service Level
- Ways to View the Service Level for Servers
- Assigning a Server to a Service Level
- Removing a Server from a Service Level

Overview of Service Levels

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

You can also create service levels to indicate the Service Level Agreement (SLA) for the servers that your IT organization manages. For example, you might create service levels to denote Silver, Gold, and Platinum services.

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

Adding a Service Level to the SAS Web Client

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

1. From the navigation panel, click Environment ➤ Service Levels. The Service Levels page appears.
2. Navigate the hierarchy of service levels until you reach the point in the hierarchy where you want to add a new service level, as Figure 8-19 shows.

Figure 8-19: Service Level Hierarchy

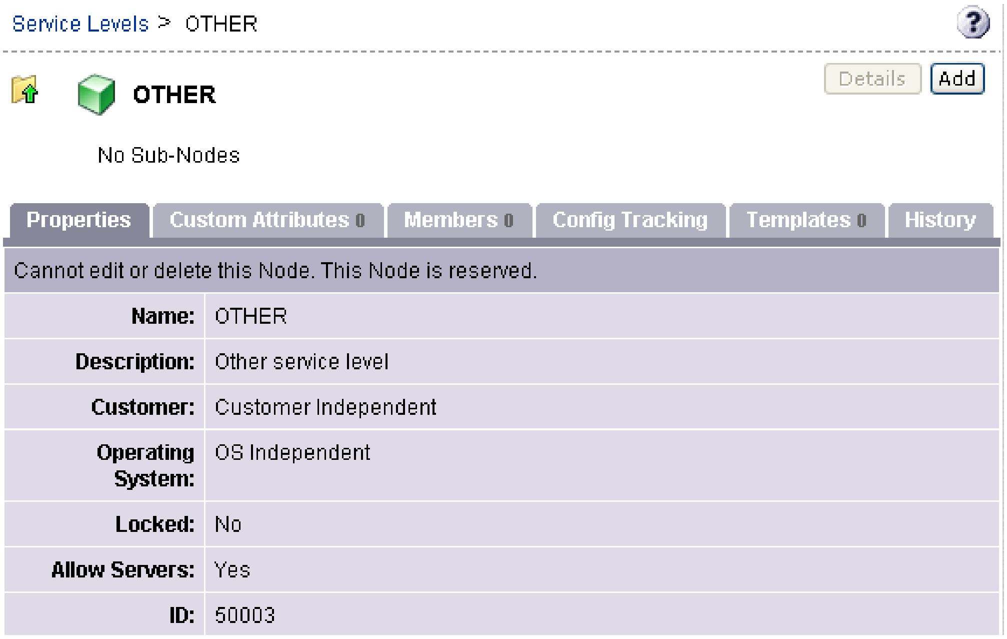

Service Levels > OTHER

OTHER

No Sub-Nodes

Cannot edit or delete this Node. This Node is reserved.

Name:	OTHER
Description:	Other service level
Customer:	Customer Independent
Operating System:	OS Independent
Locked:	No
Allow Servers:	Yes
ID:	50003

3. Click **Add**. The Service Levels page refreshes and the ADD SUB-NODE TO Service Levels form appears in the page.
4. Enter a name for the service level (required), and (optionally) enter notes and a description for the service level.

5. Click **Save**. The service level is added to the hierarchy of service levels. The Edit Service Level page appears, where you can change the properties of the service level, such as the customer association.

Editing a Service Level

Perform the following steps to edit a service level:

1. From the navigation panel, click Environment ➤ Service Levels. The Service Levels page appears.
2. Navigate the hierarchy of service levels until you reach the point in the hierarchy where you want to edit an existing service level.
3. Click **Edit** in the Properties tab. The page refreshes and an editable form appears for the service level properties.
4. Make changes to the service level name, description, notes, whether servers are allowed to be assigned to the service level, the associated customers and operating systems.
5. Click **Save**.

Ways to View the Service Level for Servers

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

- If you are browsing the Manage Server list, you can find the service level for a server by locating the value in the Environment column, as Figure 8-20 shows.

Figure 8-20: Service Level Node Appearing in the Software Tab of the Server List

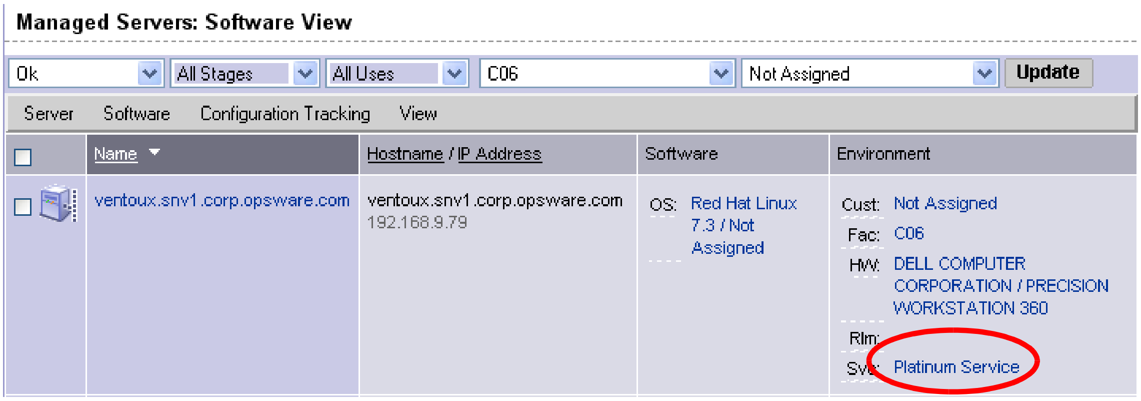

Managed Servers: Software View

□	Name ▼	Hostname / IP Address	Software	Environment
□	ventoux.snv1.corp.opsware.com	ventoux.snv1.corp.opsware.com 192.168.9.79	OS: Red Hat Linux 7.3 / Not Assigned	Cust: Not Assigned Fac: C06 HW: DELL COMPUTER CORPORATION / PRECISION WORKSTATION 360 Rlm: Svc: Platinum Service

- If you searched for the server, click the server name and then select the Attached Nodes tab. You can find the service levels, as Figure 8-21 shows.

Figure 8-21: Nodes Tab That Shows the Service Level to Which a Server is Assigned

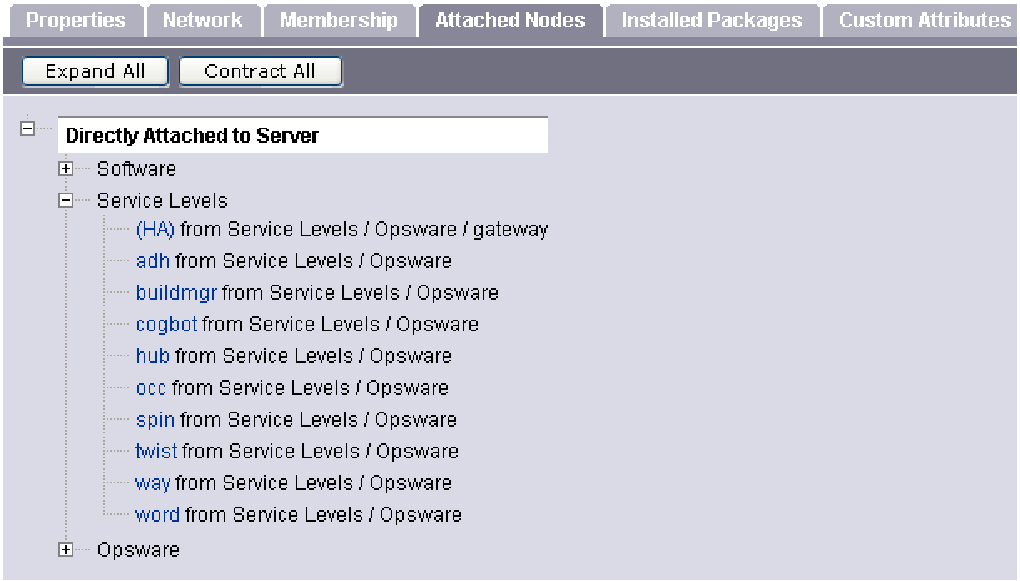

Directly Attached to Server
- Software
- Service Levels
  - (HA) from Service Levels / Opsware / gateway
  - adh from Service Levels / Opsware
  - buildmgr from Service Levels / Opsware
  - cogbot from Service Levels / Opsware
  - hub from Service Levels / Opsware
  - occ from Service Levels / Opsware
  - spin from Service Levels / Opsware
  - twist from Service Levels / Opsware
  - way from Service Levels / Opsware
  - word from Service Levels / Opsware
- Opsware

- To view all the servers assigned to a particular service level, click Environment ➤ Service Levels in the navigation panel. Navigate the hierarchy of service levels until you reach the one for which you want to see which servers are assigned. Select the Members tab, as Figure 8-22 shows.

Figure 8-22: Manage Server Assigned to a Service Level

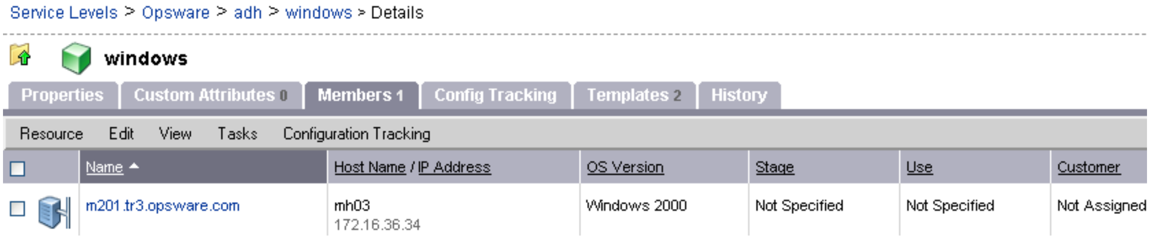

□	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Customer
□	m201.tr3.opsware.com	mh03 172.16.36.34	Windows 2000	Not Specified	Not Specified	Not Assigned

Assigning a Server to a Service Level

Perform the following steps to assign a server to a service level:

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

Or

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

2. Select the servers that you want to assign to a service level.

3. Choose **Tasks ➤ Assign Node** from the menu above the Manage Servers list. A window displays the categories of nodes, as Figure 8-23 shows.

Figure 8-23: Assign Nodes Popup Window

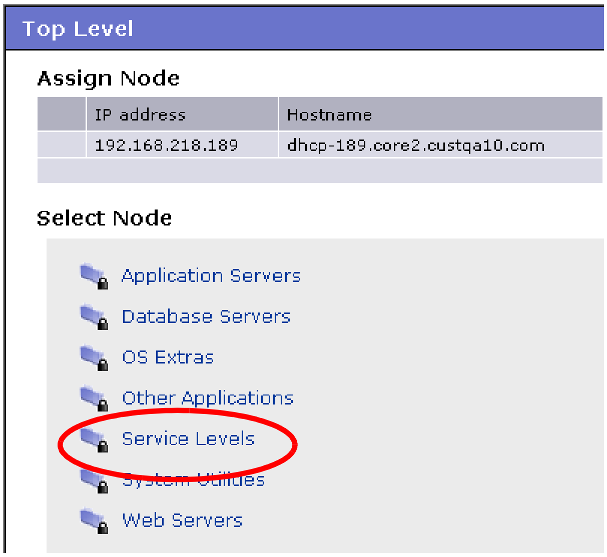

Top Level

Assign Node

	IP address	Hostname
	192.168.218.189	dhcp-189.core2.custqa10.com

Select Node

- Application Servers
- Database Servers
- OS Extras
- Other Applications
- Service Levels
- System Utilities
- Web Servers

4. Click the Service Levels link. The window refreshes to show the service levels created for your operational environment.

5. Navigate to the service level to which you want to assign the server.

6. Click **Assign**. The window closes and you are returned to the Manage Servers list.

Removing a Server from a Service Level

Perform the following steps to remove a server from a service level:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server that you want to remove from a service level.

   Or

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

2. Select the server that you want to remove from a service level.

3. Choose **Tasks ➤ Remove Node** from the menu above the Manage Servers list. A window displays the nodes to which the server is assigned, as Figure 8-24 shows.

Figure 8-24: Remove Nodes Popup Window

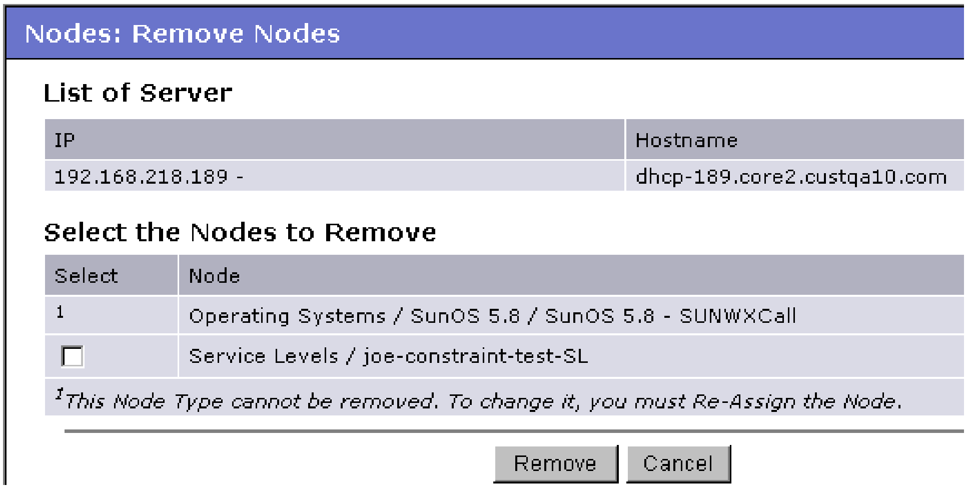

Nodes: Remove Nodes

List of Server

IP	Hostname
192.168.218.189 -	dhcp-189.core2.custqa10.com

Select the Nodes to Remove

Select	Node
1	Operating Systems / SunOS 5.8 / SunOS 5.8 - SUNWXCall
□	Service Levels / joe-constraint-test-SL

[1] *This Node Type cannot be removed. To change it, you must Re-Assign the Node.*

4. Select the service level node from which you want to remove the server and click **Remove**. You are prompted to confirm that you want to remove the server from the service level.

5. Click **Confirm Remove**. The window closes and you are returned to the Manage Servers list.

Network Configuration

This section provides information about network configuration within HP Server Automation and contains the following topics:

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

Overview of the Server Network Configuration

You can use HP Server Automation to automatically configure network settings for a server after you install the OS.

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

HP Server Automation does not support managed servers that have IPv6 addresses.

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

- Host Name
- Domain Name System (DNS) servers
- Management interface (the interface that HP Server Automation should use when managing the server)

  See "Locating the Management IP Address of a Managed Server" on page 278 in Chapter 8 for more information.
- Gateway (the IP address of the default router)
- DNS search domains
- WINS (Windows Internet Naming Service) Servers
- Configuration for each network interface, including whether the interface is configured statically or with a Dynamic Host Configuration Protocol (DHCP) IP address, host name, and subnet mask

You can alter any of these options and then apply the settings to the managed server. HP Server Automation updates the server and reboots it to cause the new settings to take effect.

Configuring Networking for an SA Managed Server

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

Perform the following steps to configure networking for an SA managed server:

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

Or

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

See “Using the Search Feature” on page 244 in Chapter 7 for more information. See “Server Searching by IP Address” on page 260 in Chapter 7 for more information.

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

3. Select the Network tab. The network information for the server displays.

4. Modify any of the following settings to configure the server networking.

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

   - **Host Name**: The host name configured on the managed server. This field only sets the name by which the server knows itself and does not update DNS records for the server.
   - **Management Interface**: Instructs HP Server Automation to use a particular network interface when contacting the server. This is useful, for example, when a server has multiple network interfaces, but not all of them are reachable by the SA core. Designating a particular interface as the management interface allows HP Server Automation to know which interface to use for managing the server.
   - **Gateway**: The IP address of the default router
   - **DNS Servers**: A list of DNS nameserver IP addresses
   - **Search Domains**: A list of DNS domains to search when attempting to resolve host names
   - **WINS Servers**: Set for Windows only; a list of WINS server IP addresses
   - **Interface Configuration** (for each network interface in the system):
     - **DHCP**: If DHCP is enabled for an interface, the system uses DHCP to configure this network interface. In this case, static configuration settings (IP address, host name, and subnet mask) are not relevant for this interface, and the SAS Web Client makes those fields not editable. If DHCP is not enabled, then static settings are required.
     - **IP Address**: The IP address for this interface (unless DHCP is enabled).

- **Host Name**: The local host name for the server. This item is only required for servers running Solaris. Like the Computer Name field, this setting only affects the name by which the managed server knows itself, and does not update DNS records.
- **Subnet Mask**: The IP network mask to use for this interface.

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

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

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

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

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

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

7 Click **OK** to proceed with the configuration changes.

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

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

Details About Changing the Domain for Windows Servers

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

HP Server Automation does not change the domain name of a Windows server because changing the domain name of a server requires password authentication. Changing the domain name of a Windows server is a manual operation. See Figure 8-25.

Figure 8-25: DNS Domain Field Displays in the Network Tab for a Server

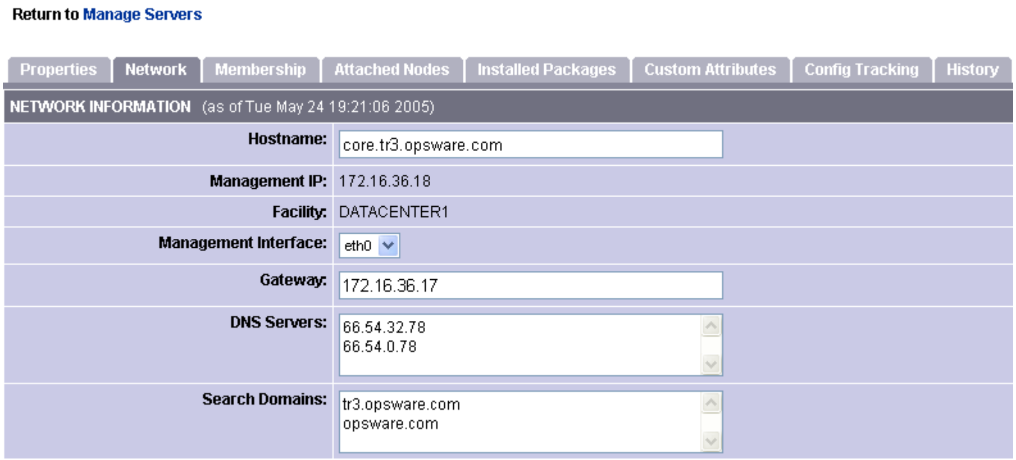

NETWORK INFORMATION (as of Tue May 24 19:21:06 2005)

Hostname:	core.tr3.opsware.com
Management IP:	172.16.36.18
Facility:	DATACENTER1
Management Interface:	eth0
Gateway:	172.16.36.17
DNS Servers:	66.54.32.78 66.54.0.78
Search Domains:	tr3.opsware.com opsware.com

Chapter 9: NA Integration

IN THIS CHAPTER

This section discusses the following topics:

- Overview of NA Integration
- Setup for NA Integration
- Device Information in SA
- Device Groups and NA
- Duplex Mismatch
- Network Device Information in Network Automation
- Network Reports

Overview of NA Integration

Implementing changes in an IT environment often requires a coordinated effort between network administrators, system administrators, and application architects. Together, these members of the IT staff must manage an application environment spanning servers running different operating systems and network devices that include firewalls, load balancers, switches, servers, Web applications, and so on. For example, in some environments, you are required to make changes to network devices in front of an application, such as load balancers, firewalls, switches, and so on.

The NA Integration feature makes this process easier. It enables IT staff members to see how servers are connected to network devices and enables them to closely examine managed servers. With this information, they can determine how all devices are related and coordinate and implement required changes.

Figure 9-1 shows the types of coordinated tasks you can perform by using the NA Integration feature.

Figure 9-1: Overview of Coordinating Tasks Using NA Integration

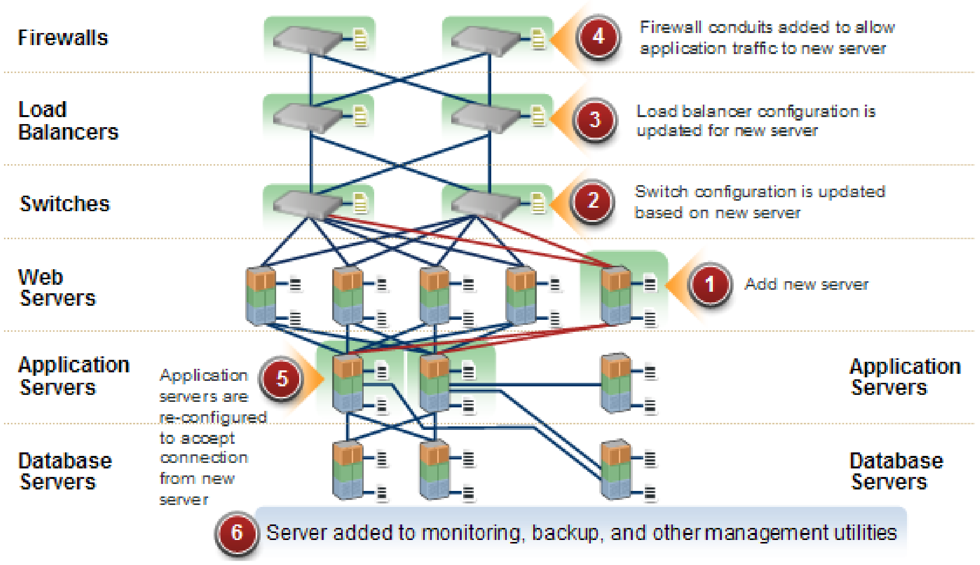
Firewalls
4
Firewall conduits added to allow application traffic to new server
Load Balancers
3
Load balancer configuration is updated for new server
Switches
2
Switch configuration is updated based on new server
Web Servers
1
Add new server
Application Servers
Application servers are re-configured to accept connection from new server
5
Application Servers
Database Servers
Database Servers
6
Server added to monitoring, backup, and other management utilities

This section contains information about how to set up NA Integration, view device details, examine connections between network devices and servers, identify duplex mismatches, and view combined device history information. It also contains information about implementing changes across the environment and generating network reports.

To support an integrated approach to making changes in your environment, such as server reallocation, ensuring compliance across servers and network devices, and detecting and resolving duplex mismatches, the NA Integration feature provides the following SA interface points:

- HP Server Automation (SA)
- Network Automation (NA)
- HP Global Shell (in SA)
- HP HP Service Automation Visualizer (in SA)
- HP Reports (in SA)

NA Integration Features

In the NA Integration feature, you will use several SA and Network Automation features to perform the following tasks:

- Use SA and Network Automation to view summarized and detailed hardware information about managed servers and network devices, and their connections (interfaces and ports).
- Use the Global File System (OGFS) to navigate between managed servers and connected network devices by tracing their associated physical connections, finding network device configurations, and running scripts across servers and network devices.
- Call NA scripts from SA scripts to automate operations across servers and network devices.
- Use features in SA and NA to create diagrams that illustrate managed servers, network devices, and layer 1 connections in your environment.
- Use SA to identify, troubleshoot, and remediate configuration duplex mismatches between managed servers and network devices.
- Use SA to perform actions on device groups that contain both servers and network devices.
- Use SA to review a combined event history log of servers and network devices that records changes made to an application in your environment.
- Use SA to export a combined event history log to a CSV and an HTML file.
- Use NA to directly access additional network device details and event history.
- Use SA to run network reports that identify layer 1 connections and configuration mismatches (duplex compliance).

Reference to *connections* in this documentation refer to *physical connections*, except where noted.

How NA Data is Collected

The NA Integration feature uses the NA Topology Data Gathering and NA Duplex Data Gathering diagnostic tools to collect information about network devices.

NA Topology Data Gathering

The NA Topology Data Gathering diagnostic instructs NA to collect MAC addresses for all switches. MAC addresses are required to discover and add physical connections to the SA data model. For example, when you add another server to a switch, that information will be collected the next time the Network Automation Topology Data Gathering diagnostic runs. You can also run the Network Automation Topology Data Gathering diagnostic or the NA Duplex Data Gathering on demand for specific network devices. See the *NA User's Guide*.

To avoid impairing performance in Network Automation, you cannot run this diagnostic any more frequently than weekly. If you cannot wait a week to refresh the NA data, contact HP Support.

NA Duplex Data Gathering

For network devices, speed and duplex is gathered by the Network Automation Duplex Data Gathering diagnostic, which runs after a device is initially added to NA and then according to a schedule that you define. To ensure that you have the latest speed and duplex information about network devices, SA recommends that you set up a recurring schedule that runs the diagnostic. See "Duplex Mismatch" on page 372 and the *NA User's Guide*.

Setup for NA Integration

When SA and NA are first installed, the SA administrator must perform certain tasks on the core servers to enable integration. For details on these tasks, see "NA Integration" in the *SA Planning and Installation Guide*.

Version Compatibility for NA Integration

The NA Integration feature requires that both Network Automation 6.1 and SA 6.x are installed and configured to enable you to share data across servers and network devices in your environment. See the *SA Planning and Installation Guide*.

The Server Agent 6.x is also required. If an agent older than 6.0 is installed on a server, it cannot return the duplex and speed settings for the network interfaces. Therefore, duplex compliance cannot be calculated for that server.

Resetting the NA Host in the SA Client

Some NA Integration features require that the SA Client (Java) opens the NA Web interface (directly from SA) so that you can access additional details about certain NA events. If your administrator has completed the setup tasks in the *SA Planning and Installation Guide*, but the SA Client is unable to communicate directly with the server running the NA host (server) Web interface, you might need to change the NA option in the SA Client. For example, if a firewall is preventing the SA Client from reaching the NA host, you need to specify the name of a server that is acting as a proxy for the NA host. This will override the default setting. This task must be performed on every desktop running a SA Client that cannot communicate with the NA host.

To reset the NA host in the SA Client, perform the following steps:

1. From the **Tools** menu in the SA Client window, select **Options**.
2. In the Views pane, select HP Network Automation.
3. In the Host field, enter the name of a server that is acting as a proxy for the NA host, such as m208, which is the proxy for the m208.opsware.com NA host.
4. (Optional) Click **Restore Default** to restore the previously saved NA host name.
5. (Optional) Click **Test** to open the Network Automation login window.
6. Click **Save**.

Troubleshooting Tips

To test whether SA is communicating with NA, check the following conditions:

- You can log in to NA with your SA credentials. This verifies that NA can communicate with SA.
- The SA credentials specified in the NA Administrative Settings under External Authentication Type are set to SA. This ensures that NA can look up server MAC addresses.
- The NA Topology Gathering Diagnostic has run successfully. To verify this condition, search for tasks and check their results. This ensures that NA has gathered MAC addresses and tried to look them up in SA.

Device Information in SA

In addition to basic hardware details about managed servers and network devices, the NA Integration feature also reports the following information about network interfaces and network ports:

- On the server side, network interfaces have the following properties: MAC address, subnet mask, interface type, IP address, DHCP setting, connected switch port, speed, and duplex (excluding Windows).
- On the network device side, network ports have the following properties: port name, speed and duplex settings, devices connected, and interface type.

For most devices, auto-negotiation works best when both sides of the connection (server and network device) are set to auto-negotiate mode. For example, a duplex policy could specify that a port should be set to full, half, or auto, and not to full (auto). A full (auto) duplex setting indicates that the port was set to auto-negotiate and it negotiated to full duplex. See the *NA User’s Guide*.

The following tasks describe how you can access detailed hardware information for servers and network devices directly in SA. See “Network Device Information in Network Automation” on page 371 for instructions on how to access hardware information about network devices directly in Network Automation.

Viewing Network Interfaces

To view hardware information about a server, including network interfaces, perform the following steps:

1. From the Navigation pane, select Devices ➤ All Managed Servers.
2. From the View drop-down list, select Hardware.

3 Double-click on a server in the Content pane to display hardware details in the Server Explorer.

Figure 9-2: Hardware View in the Server Explorer

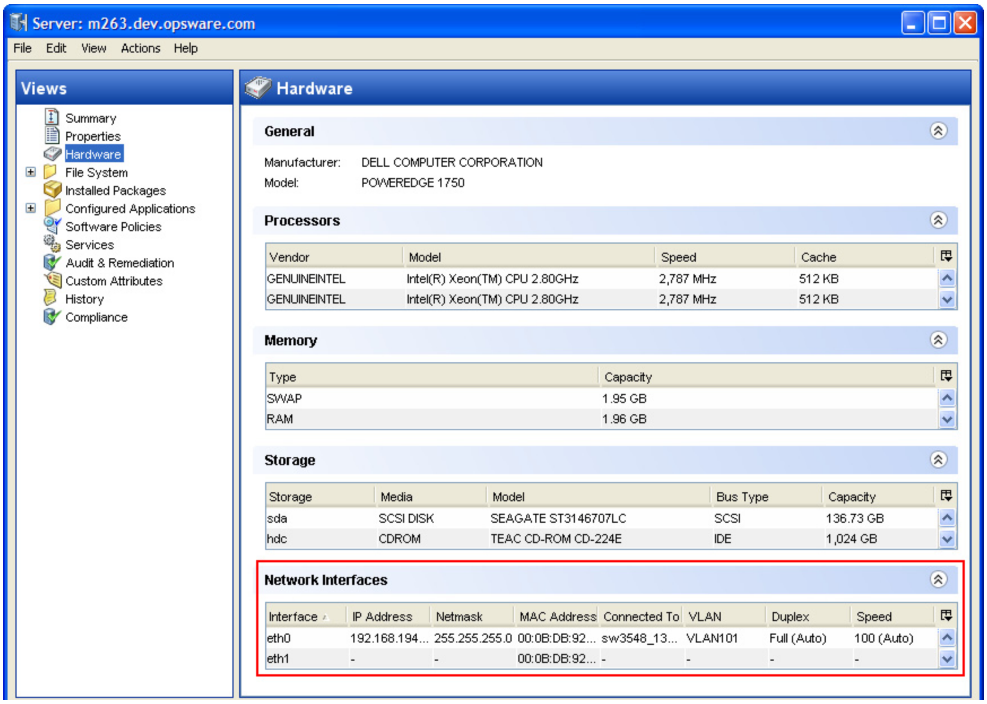

General

Manufacturer: DELL COMPUTER CORPORATION
Model: POWEREDGE 1750

Processors

Vendor	Model	Speed	Cache
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2,787 MHz	512 KB
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2,787 MHz	512 KB

Memory

Type	Capacity
SWAP	1.95 GB
RAM	1.96 GB

Storage

Storage	Media	Model	Bus Type	Capacity
sda	SCSI DISK	SEAGATE ST3146707LC	SCSI	136.73 GB
hdc	CDROM	TEAC CD-ROM CD-224E	IDE	1,024 GB

Network Interfaces

Interface	IP Address	Netmask	MAC Address	Connected To	VLAN	Duplex	Speed
eth0	192.168.194...	255.255.255.0	00:0B:DB:92...	sw3548_13...	VLAN101	Full (Auto)	100 (Auto)
eth1	-	-	00:0B:DB:92...	-	-	-	-

Viewing Network Ports

To view hardware information about a network device, including network ports, perform the following steps:

1 From the Navigation pane, select Devices ➤ Device Groups ➤ Public and then select a group.

2 Double-click on a network device in the Content pane to display the Network Device Explorer.

3 In the Views pane, select Hardware to display information about the selected network device.

Figure 9-3: Hardware View in the Network Device Explorer

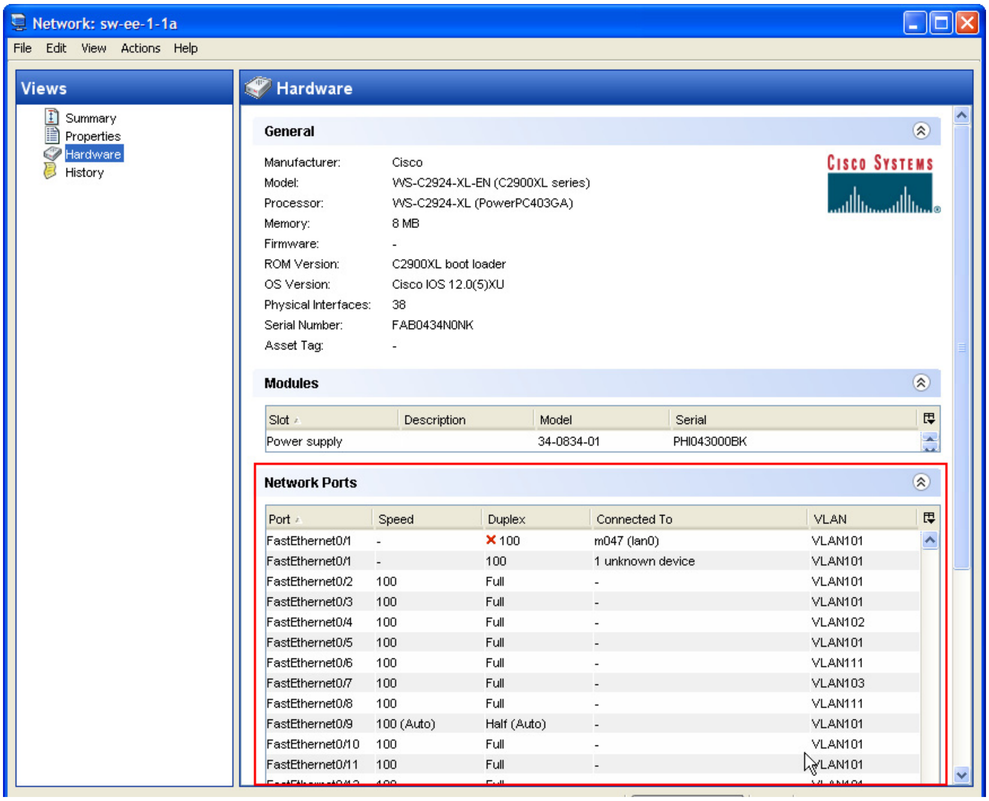

General

Manufacturer:	Cisco
Model:	WS-C2924-XL-EN (C2900XL series)
Processor:	WS-C2924-XL (PowerPC403GA)
Memory:	8 MB
Firmware:	-
ROM Version:	C2900XL boot loader
OS Version:	Cisco IOS 12.0(5)XU
Physical Interfaces:	38
Serial Number:	FAB0434N0NK
Asset Tag:	-

Modules

Slot	Description	Model	Serial
Power supply		34-0834-01	PHI043000BK

Network Ports

Port	Speed	Duplex	Connected To	VLAN
FastEthernet0/1	-	✕ 100	m047 (lan0)	VLAN101
FastEthernet0/1	-	100	1 unknown device	VLAN101
FastEthernet0/2	100	Full	-	VLAN101
FastEthernet0/3	100	Full	-	VLAN101
FastEthernet0/4	100	Full	-	VLAN102
FastEthernet0/5	100	Full	-	VLAN101
FastEthernet0/6	100	Full	-	VLAN111
FastEthernet0/7	100	Full	-	VLAN103
FastEthernet0/8	100	Full	-	VLAN111
FastEthernet0/9	100 (Auto)	Half (Auto)	-	VLAN101
FastEthernet0/10	100	Full	-	VLAN101
FastEthernet0/11	100	Full	-	VLAN101

Connections Between Network Devices and Servers

The NA Integration features are based on layer 2 connections and inferred layer 1 connections. See Figure 9-4 for definitions of the OSI Model layers.

Figure 9-4: OSI Seven Layer Model

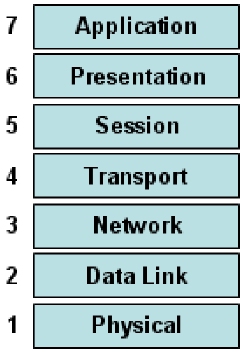

7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Data Link
1	Physical

Data Link Connections

The NA Integration feature includes functionality that detects data link (layer 2) connections and reports on physical (layer 1) and data link connections. These data link connections include switches that are directly connected to a managed server and switches that are indirectly connected through other switches. These connections are discovered by correlating the MAC addresses reported by the device with the known MAC addresses for servers and switches.

Physical Connections

The physical connections are inferred from the data link connections. See "Inferred Physical Connections" on page 363. Physical connections represent direct connections (cables) between server and switches.

In the SA Client, you can see physical connections in the Server Explorer, the Network Device Explorer and in detailed layout diagrams in SAV. In the NA diagramming feature, you can see physical, data link or network (layer 3) connections.

Network Diagrams

You can use the SAV feature in SA and the Diagramming feature in Network Automation to create detailed diagrams that illustrate managed servers, network devices, and layer 2 and layer 1 connections in your environment. You can also export these network diagrams to .gif, .jpg, and .svg files, and then subsequently annotate and use them in other applications.

See the *SA User's Guide: Application Automation* and the *NA User's Guide* for more information about the SAV and the Diagramming tool.

Launching HP Service Automation Visualizer

To access the SAV, perform the following steps:

1. From the Navigation pane, select Devices ➤ All Managed Servers.
2. In the Content pane, select one or more servers.
3. From the **Tools** menu, select **HP Service Automation Visualizer** and then select one of the following options:
   - Select **New** to open the SAV window.
   - Select **Open** to open a previously saved topology.
4. To create and export topology diagrams, see the procedures for using HP Service Automation Visualizer in the *SA User's Guide: Application Automation*.

Launching NA Diagramming

See the *NA User's Guide* for instructions on launching and using the Network Automation Diagramming feature.

Network Directories

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

You can also run three types of NA scripts in the OGFS: command, advanced, and diagnostic. These scripts correspond to the three directories in the OGFS under `/opsw/Scripts/Network`. See "Network Directories" on page 543.

In the OGFS, you can also write scripts in languages, such as Bourne shell and Python, to perform the following tasks:

- Find servers and network devices.
- Find all servers that are connected to a certain switch.
- Find servers with a duplex mismatch.
- Display the network interfaces of a certain server.
- Get the IP addresses of all devices.
- Compare two files to identify changes in a network device's configuration.

- Change device details, such as the snmp-location.

Launching Global Shell

Within the Global Shell feature you can see all of the ports on a switch, view duplex settings in the info file, see all device configurations, compare configuration files, and so on. Every time a network device configuration is changed, a snapshot of it is also created. All actions that you perform on a device are audited.

To access the OGFS in the Global Shell feature, perform the following steps:

1. From the **Tools** menu, select **Global Shell** to launch a terminal window. See "Opening a Global Shell Session" on page 388.
2. To navigate between servers and connected network devices, use the guidelines described in "SA Global Shell" on page 377 and "OGFS Directories" on page 531.

Remote Terminal

The Remote 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 log in to a switch with `rosh` and run the `show vlan` command to view all VLAN details.

See "Remote Terminal" on page 398 and "Logging on to a Managed Server With rosh" on page 392 for more information about using the `rosh` utility.

Inferred Physical Connections

The NA Integration feature also includes functionality that detects and reports on inferred physical (layer 1) connections. These connections are inferred from data (such MAC addresses that are seen by switches), captured, and then added to the SA data model.

These physical connections (inferred layer 1 data) are based on heuristics. In the OSI model, each layer is an abstraction designed to hide the layer below. Therefore, the layer 2 data gathered from devices cannot generate 100% accurate layer 1 data. In particular, layer 1 data may be incorrect if any of the following conditions exist:

- The device does not return the port number where MAC addresses are seen.
- There was no traffic between the devices within a few minutes of when NA gathered the topology data (where MAC addresses are seen).
- There is an unmanaged device between two managed devices.

- There is a hub between two managed devices.

In the SA Client, you can see inferred layer 1 connections by navigating network device directories in Global Shell .

Device Groups and NA

A device group helps you categorize your devices (servers and network devices) in ways that make sense for your organization. For example, you can group devices by customer, facility, usage, application, and so on, and then perform actions on all of the devices in the group.

In HP Server Automation, a device group can contain managed servers *and* network devices, or *only* managed servers. In Network Automation, a device group contains only network devices. You create and edit network device groups only in NA. See the *NA User’s Guide* for more information about using the `rosh` utility.

To monitor an application that is running on multiple servers and relies on multiple network devices in your environment, Hewlett Packard recommends that you model it as a device group that contains all servers and network devices the application runs on. This enables you to troubleshoot the application by using HP Server Automation.

Associating a NA Device Group

When you associate a public device group in SA with a device group in NA, you will be able to monitor information about all servers and network devices that you are interested in. You associate device groups by using identical group names.

Associated device groups have the following requirements:

- The SA device group is public.

- The SA device group is static.

- The names of the associated NA and SA device groups are identical.

To associate device groups in SA and NA, perform the following steps:

1. From the Navigation pane, select Devices ➤ Device Groups ➤ Public.

2. In the Content pane, select a device group.

3. Right-click on the device group and then select Open to display the Device Group Explorer.

4 From the View drop-down list, select Properties.

5 Check the “Associate with a NA device group of the same name” check box to enable this functionality.

6 From the **File** menu, select **Save**.

Combined Device History Log

The combined device history log records events performed on servers and network devices in your environment. These events are recorded in detail as actions performed on a certain date, by a certain user, on a certain server, or on a certain network device.

In many troubleshooting tasks, this type of information is critical because some of these actions (changes) might be the root cause of problems. This log provides detailed information, such as the date the action occurred, the name and type of the device that the action was performed on, and a description of the action, that can help you perform root cause analysis, capacity planning, and compliance remediation tasks. For example, if an application in your environment has suddenly stopped running and you know exactly when it was previously running, you need to examine a combined event history log for the affected servers and network devices, for that time period. This information can help you determine why the application stopped working.

Viewing a Combined Device Event History Log

You can view a detailed list of events that occurred on a server or network device, such as all changes made to an application. You can narrow the time frame of the log display to see changes that occurred daily, weekly, monthly, quarterly, or in a custom range of dates. You can also dynamically filter the display of events by a certain date, device name, device type, event type, or by user name.

You can view a combined device history log for one or more managed servers or for a device group that contains managed servers and network devices.

To view a combined device history log for a device group, perform the following steps:

1 From the Navigation pane, select Devices ➤ Device Group, and then select a device group.

2 In the Content pane, select one or more devices in the group.

3 Right-click and then select View History to list events that occurred on the selected devices.

Figure 9-5: Combined Device Event History

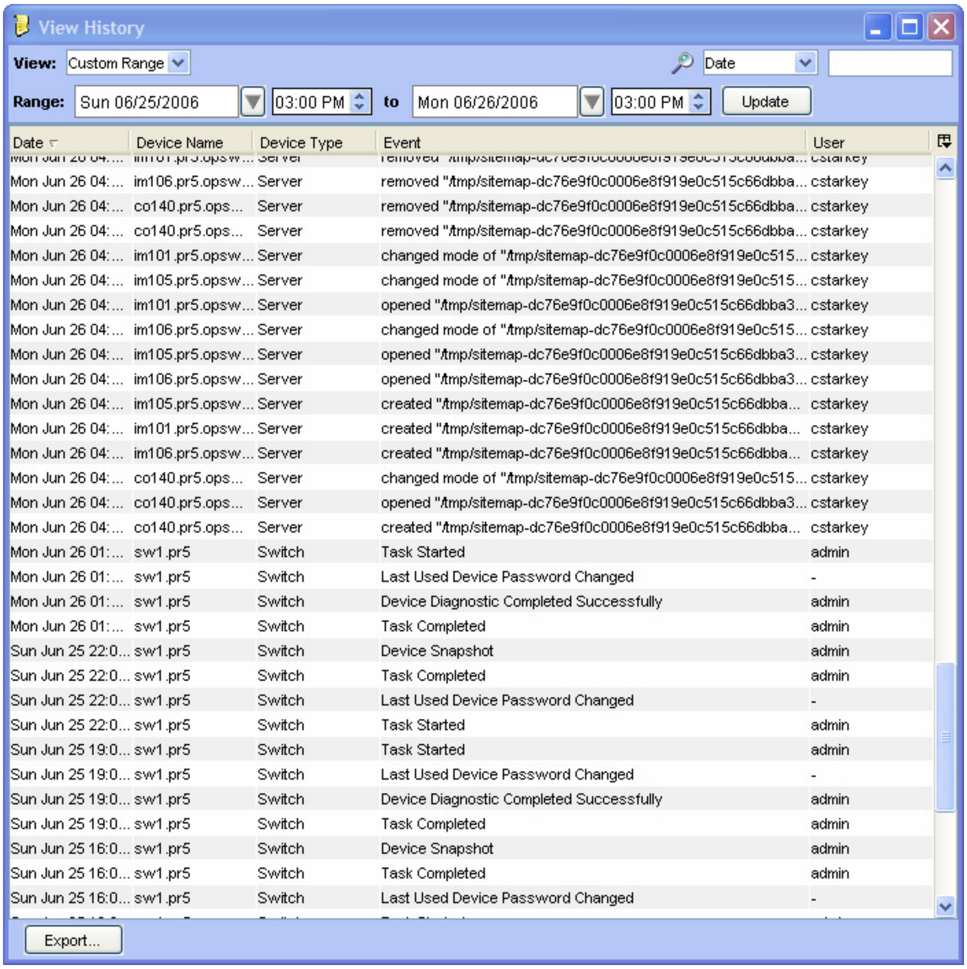

View History

View: Custom Range

Date

Range: Sun 06/25/2006 03:00 PM to Mon 06/26/2006 03:00 PM Update

Date	Device Name	Device Type	Event	User
Mon Jun 26 04:...	im101.pr5.opsw...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	cstarkey
Mon Jun 26 04:...	im106.pr5.opsw...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	cstarkey
Mon Jun 26 04:...	co140.pr5.ops...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	cstarkey
Mon Jun 26 04:...	co140.pr5.ops...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	cstarkey
Mon Jun 26 04:...	im101.pr5.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	cstarkey
Mon Jun 26 04:...	im105.pr5.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	cstarkey
Mon Jun 26 04:...	im101.pr5.opsw...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	cstarkey
Mon Jun 26 04:...	im106.pr5.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	cstarkey
Mon Jun 26 04:...	im105.pr5.opsw...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	cstarkey
Mon Jun 26 04:...	im106.pr5.opsw...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	cstarkey
Mon Jun 26 04:...	im105.pr5.opsw...	Server	created "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	cstarkey
Mon Jun 26 04:...	im101.pr5.opsw...	Server	created "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	cstarkey
Mon Jun 26 04:...	im106.pr5.opsw...	Server	created "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	cstarkey
Mon Jun 26 04:...	co140.pr5.ops...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	cstarkey
Mon Jun 26 04:...	co140.pr5.ops...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	cstarkey
Mon Jun 26 04:...	co140.pr5.ops...	Server	created "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	cstarkey
Mon Jun 26 01:...	sw1.pr5	Switch	Task Started	admin
Mon Jun 26 01:...	sw1.pr5	Switch	Last Used Device Password Changed	-
Mon Jun 26 01:...	sw1.pr5	Switch	Device Diagnostic Completed Successfully	admin
Mon Jun 26 01:...	sw1.pr5	Switch	Task Completed	admin
Sun Jun 25 22:0...	sw1.pr5	Switch	Device Snapshot	admin
Sun Jun 25 22:0...	sw1.pr5	Switch	Task Completed	admin
Sun Jun 25 22:0...	sw1.pr5	Switch	Last Used Device Password Changed	-
Sun Jun 25 22:0...	sw1.pr5	Switch	Task Started	admin
Sun Jun 25 19:0...	sw1.pr5	Switch	Task Started	admin
Sun Jun 25 19:0...	sw1.pr5	Switch	Last Used Device Password Changed	-
Sun Jun 25 19:0...	sw1.pr5	Switch	Device Diagnostic Completed Successfully	admin
Sun Jun 25 19:0...	sw1.pr5	Switch	Task Completed	admin
Sun Jun 25 16:0...	sw1.pr5	Switch	Device Snapshot	admin
Sun Jun 25 16:0...	sw1.pr5	Switch	Task Completed	admin
Sun Jun 25 16:0...	sw1.pr5	Switch	Last Used Device Password Changed	-

Viewing an Event History Log for Servers

To view a combined device history log for one or more managed servers, perform the following steps:

1 From the Navigation pane, select Devices ➤ All Managed Servers.

2 In the Content pane, select one or more servers.

3 Right-click and then select View History to list events that occurred on the selected servers.

Figure 9-6: View History of Servers

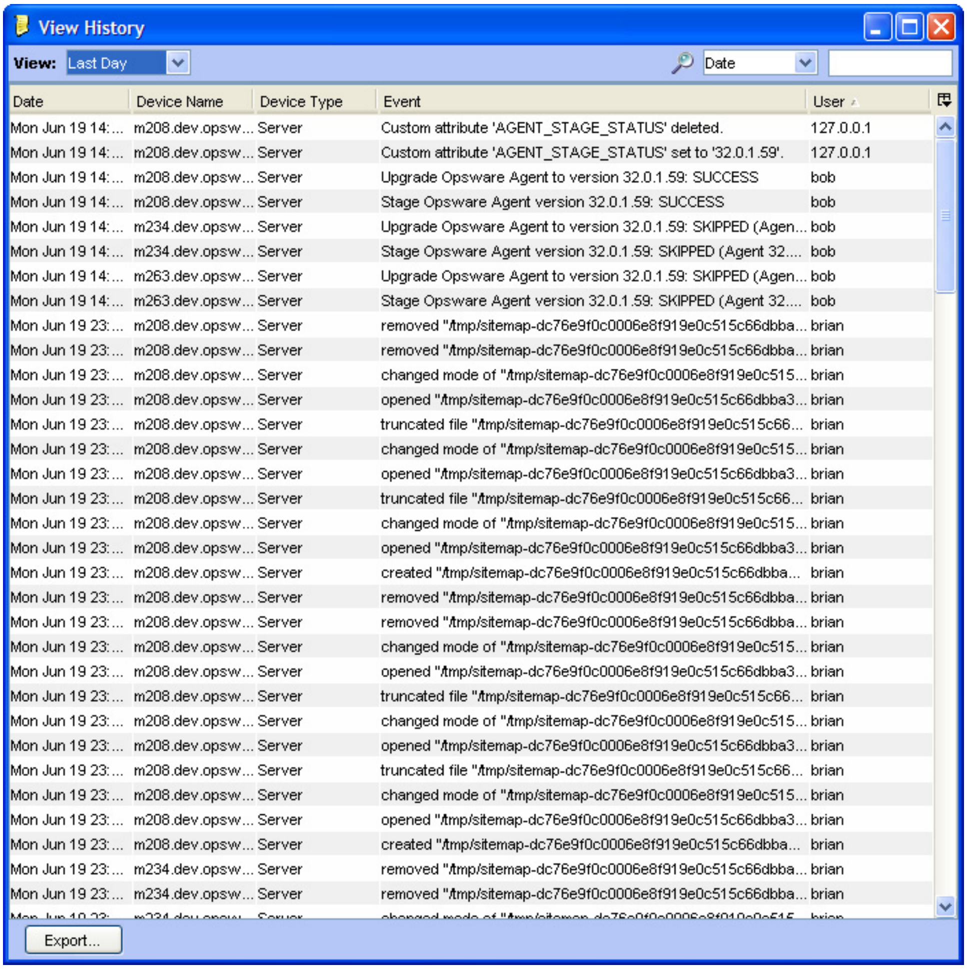

Date	Device Name	Device Type	Event	User
Mon Jun 19 14:...	m208.dev.opsw...	Server	Custom attribute 'AGENT_STAGE_STATUS' deleted.	127.0.0.1
Mon Jun 19 14:...	m208.dev.opsw...	Server	Custom attribute 'AGENT_STAGE_STATUS' set to '32.0.1.59'.	127.0.0.1
Mon Jun 19 14:...	m208.dev.opsw...	Server	Upgrade Opsware Agent to version 32.0.1.59: SUCCESS	bob
Mon Jun 19 14:...	m208.dev.opsw...	Server	Stage Opsware Agent version 32.0.1.59: SUCCESS	bob
Mon Jun 19 14:...	m234.dev.opsw...	Server	Upgrade Opsware Agent to version 32.0.1.59: SKIPPED (Agen...	bob
Mon Jun 19 14:...	m234.dev.opsw...	Server	Stage Opsware Agent version 32.0.1.59: SKIPPED (Agent 32....	bob
Mon Jun 19 14:...	m263.dev.opsw...	Server	Upgrade Opsware Agent to version 32.0.1.59: SKIPPED (Agen...	bob
Mon Jun 19 14:...	m263.dev.opsw...	Server	Stage Opsware Agent version 32.0.1.59: SKIPPED (Agent 32....	bob
Mon Jun 19 23:...	m208.dev.opsw...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	truncated file "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	truncated file "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	created "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	truncated file "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	truncated file "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	opened "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba3...	brian
Mon Jun 19 23:...	m208.dev.opsw...	Server	created "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	brian
Mon Jun 19 23:...	m234.dev.opsw...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	brian
Mon Jun 19 23:...	m234.dev.opsw...	Server	removed "/tmp/sitemap-dc76e9f0c0006e8f919e0c515c66dbba...	brian
Mon Jun 19 23:...	m234.dev.opsw...	Server	changed mode of "/tmp/sitemap-dc76e9f0c0006e8f919e0c515...	brian

4 (Optional) In the View History window, select an option in the View drop-down list to list events by a time period, such as Last Day, Last Week, Last Month, or Custom Range.

5 (Optional) In the View History window, use the search tool 🔎 to dynamically filter the display of events by a certain date, device name, device type, event type, or by user name.

6 (Optional) To identify the device name and type for a certain event, double-click on the event listed in the View History window to display the Event Details window.

Figure 9-7: Event Details for a Selected Server

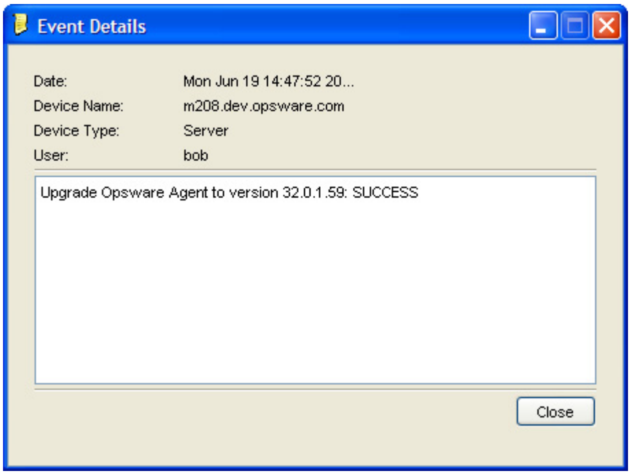

Date: Mon Jun 19 14:47:52 20...
Device Name: m208.dev.opsware.com
Device Type: Server
User: bob

Upgrade Opsware Agent to version 32.0.1.59: SUCCESS

To view the combined device history log for a device group, perform the following steps:

1 From the Navigation pane, select Devices ➤ Public.

2 In the Content pane, select one or more devices in the group.

3 Right-click and then select View History to list events that occurred on the selected network devices.

Figure 9-8: View History of Network Devices

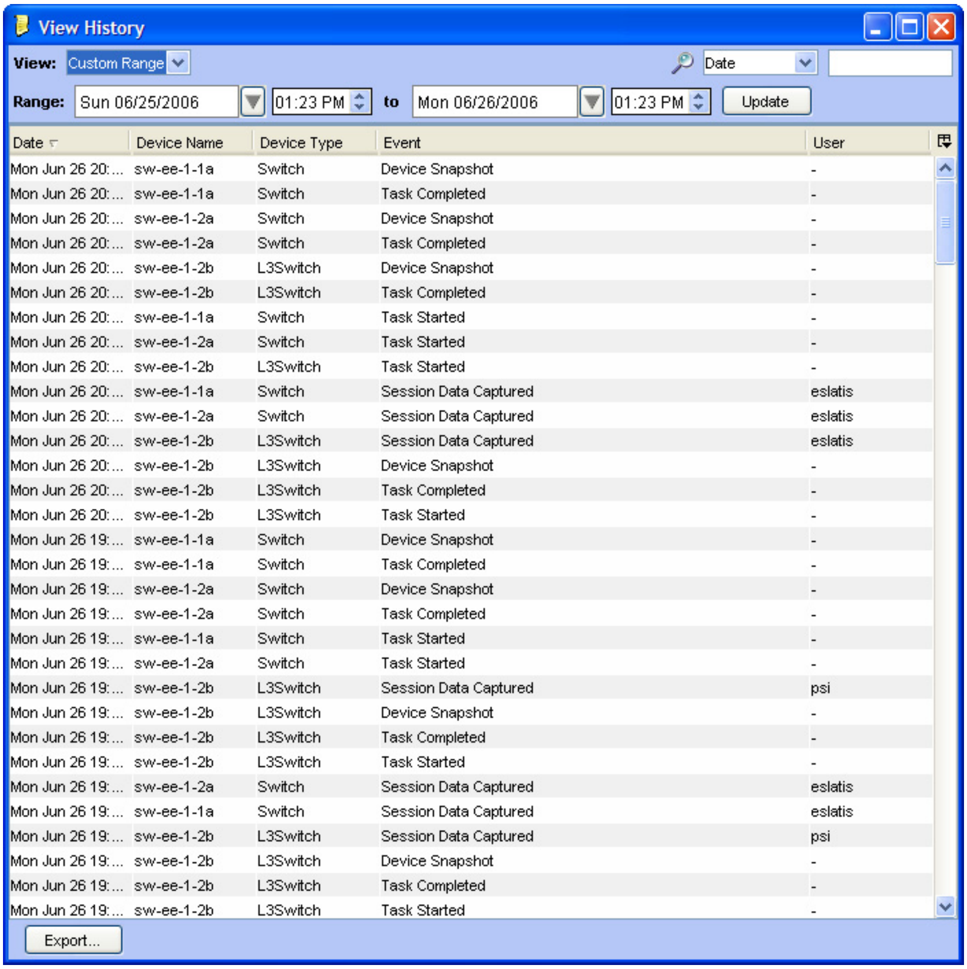

View: Custom Range

Date

Range: Sun 06/25/2006 01:23 PM to Mon 06/26/2006 01:23 PM

Date	Device Name	Device Type	Event	User
Mon Jun 26 20:...	sw-ee-1-1a	Switch	Device Snapshot	-
Mon Jun 26 20:...	sw-ee-1-1a	Switch	Task Completed	-
Mon Jun 26 20:...	sw-ee-1-2a	Switch	Device Snapshot	-
Mon Jun 26 20:...	sw-ee-1-2a	Switch	Task Completed	-
Mon Jun 26 20:...	sw-ee-1-2b	L3Switch	Device Snapshot	-
Mon Jun 26 20:...	sw-ee-1-2b	L3Switch	Task Completed	-
Mon Jun 26 20:...	sw-ee-1-1a	Switch	Task Started	-
Mon Jun 26 20:...	sw-ee-1-2a	Switch	Task Started	-
Mon Jun 26 20:...	sw-ee-1-2b	L3Switch	Task Started	-
Mon Jun 26 20:...	sw-ee-1-1a	Switch	Session Data Captured	eslatis
Mon Jun 26 20:...	sw-ee-1-2a	Switch	Session Data Captured	eslatis
Mon Jun 26 20:...	sw-ee-1-2b	L3Switch	Session Data Captured	eslatis
Mon Jun 26 20:...	sw-ee-1-2b	L3Switch	Device Snapshot	-
Mon Jun 26 20:...	sw-ee-1-2b	L3Switch	Task Completed	-
Mon Jun 26 20:...	sw-ee-1-2b	L3Switch	Task Started	-
Mon Jun 26 19:...	sw-ee-1-1a	Switch	Device Snapshot	-
Mon Jun 26 19:...	sw-ee-1-1a	Switch	Task Completed	-
Mon Jun 26 19:...	sw-ee-1-2a	Switch	Device Snapshot	-
Mon Jun 26 19:...	sw-ee-1-2a	Switch	Task Completed	-
Mon Jun 26 19:...	sw-ee-1-1a	Switch	Task Started	-
Mon Jun 26 19:...	sw-ee-1-2a	Switch	Task Started	-
Mon Jun 26 19:...	sw-ee-1-2b	L3Switch	Session Data Captured	psi
Mon Jun 26 19:...	sw-ee-1-2b	L3Switch	Device Snapshot	-
Mon Jun 26 19:...	sw-ee-1-2b	L3Switch	Task Completed	-
Mon Jun 26 19:...	sw-ee-1-2b	L3Switch	Task Started	-
Mon Jun 26 19:...	sw-ee-1-2a	Switch	Session Data Captured	eslatis
Mon Jun 26 19:...	sw-ee-1-1a	Switch	Session Data Captured	eslatis
Mon Jun 26 19:...	sw-ee-1-2b	L3Switch	Session Data Captured	psi
Mon Jun 26 19:...	sw-ee-1-2b	L3Switch	Device Snapshot	-
Mon Jun 26 19:...	sw-ee-1-2b	L3Switch	Task Completed	-
Mon Jun 26 19:...	sw-ee-1-2b	L3Switch	Task Started	-

4 (Optional) In the View History window, select an option in the View drop-down list to list events by a time period, such as Last Day, Last Week, Last Month, or Custom Range.

5 (Optional) In the View History window, use the search tool 🔍 to dynamically filter the display of events by a certain date, device name, device type, event type, or by user name.

6 (Optional) To display details for a certain event, double-click on the event listed in the View History window to display the Event Details window.

Figure 9-9: Event Details for a Selected Network Device

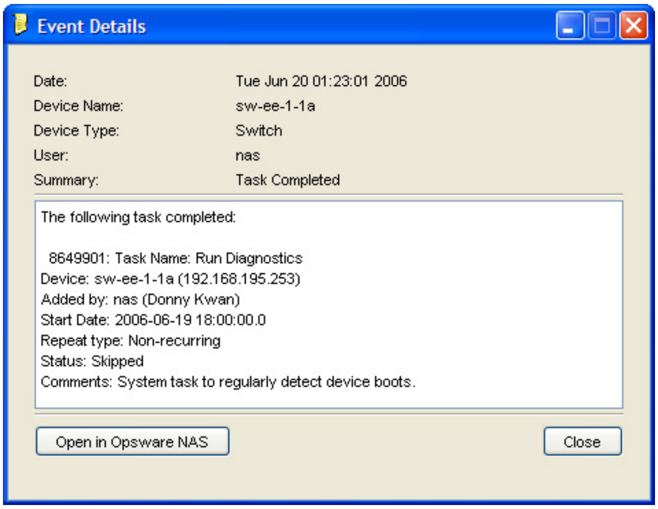

Date: Tue Jun 20 01:23:01 2006
Device Name: sw-ee-1-1a
Device Type: Switch
User: nas
Summary: Task Completed

```
The following task completed:

 8649901: Task Name: Run Diagnostics
Device: sw-ee-1-1a (192.168.195.253)
Added by: nas (Donny Kwan)
Start Date: 2006-06-19 18:00:00.0
Repeat type: Non-recurring
Status: Skipped
Comments: System task to regularly detect device boots.
```

Exporting a Combined Device History Log

If you need to use the log file in different applications, you can export the list of combined device history events to a .csv or an .html file.

To export the combined device history log, perform the following steps:

1 From the View History window, click **Export** to display the Export Dashboard window.

2 In the Look in field, enter the location of where you want to save the file.

3 (Optional) From the Encoding drop-down list, select a character encoding option. The default is Unicode (UTF-8).

4 In the File name field, enter a name for the file.

5 In the Files of type field, select .csv or .html.

6 Click **Export** to save the file in the selected format or click **Cancel** to close this window without saving.

Network Device Information in Network Automation

To help you with troubleshooting tasks that involve network devices in your environment, you can examine additional network device details and network device event history by logging directly in to Network Automation. The NA Integration feature provides a login option to access detailed information about network devices and their event history as recorded in Network Automation.

Viewing a Network Device Directly in Network Automation

If you want to view detailed information about a network device directly in Network Automation, perform the following steps:

1. From the Navigation pane, select Devices ➤ Device Group ➤ Public.
2. In the Content pane, select a network device.
3. Right-click and select **Open with ➤ HP Network Automation** to display the Network Automation login window.
4. Enter your user name and login, and then click **Login** to display device details in NA.

Figure 9-10: Network Device Details in NA

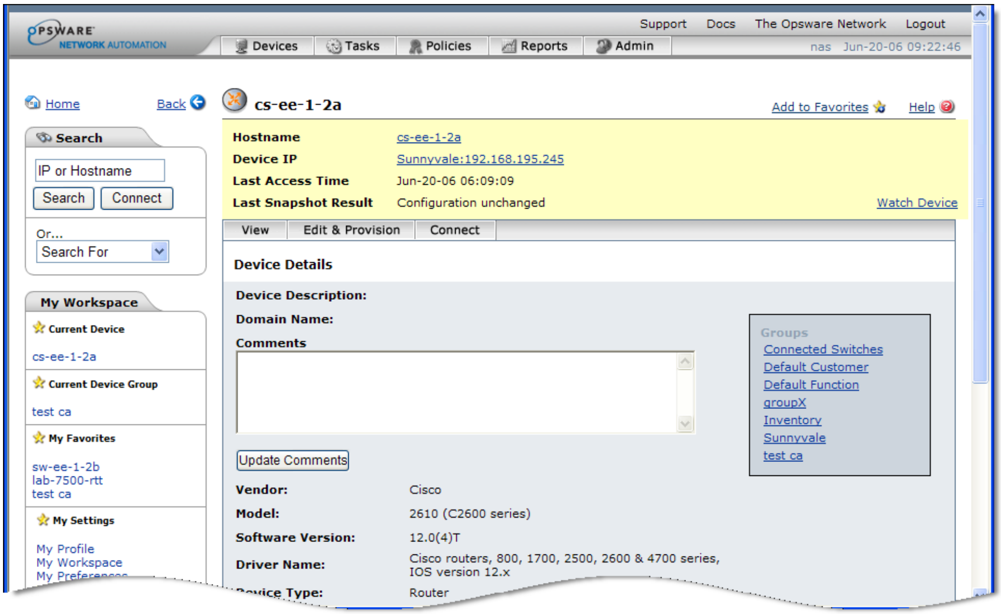

cs-ee-1-2a

Hostname	cs-ee-1-2a
Device IP	Sunnyvale:192.168.195.245
Last Access Time	Jun-20-06 06:09:09
Last Snapshot Result	Configuration unchanged

Device Details

Device Description:

Domain Name:

Comments

Vendor:	Cisco
Model:	2610 (C2600 series)
Software Version:	12.0(4)T
Driver Name:	Cisco routers, 800, 1700, 2500, 2600 & 4700 series, IOS version 12.x
Device Type:	Router

5. Click **Logout** to exit NA.

Viewing Event History Directly in Network Automation

If you want to view event details directly in Network Automation, perform the following steps:

1. In the Event Details window, click **Open with ➤ HP Server Automation** to display the Network Automation login window. See Figure 9-8 on page 369.
2. Enter your user name and login, and then click **Login** to display event details in NA.

Figure 9-11: Event Details for a Network Device in NA

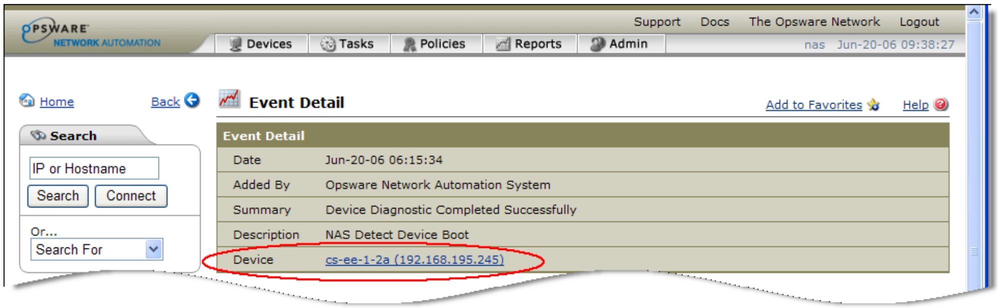

Event Detail

Event Detail	
Date	Jun-20-06 06:15:34
Added By	Opsware Network Automation System
Summary	Device Diagnostic Completed Successfully
Description	NAS Detect Device Boot
Device	cs-ee-1-2a (192.168.195.245)

3. Click on the Device link to view additional information, such as timestamps for when the device was added, the last snapshot, and the last configuration change. See Figure 9-9 on page 370.
4. Click **Logout** to exit NA.

Duplex Mismatch

The NA Integration feature provides automatic detection of duplex mismatches. A duplex mismatch is a configuration mismatch between the speed and duplex of a managed server and a connected network device.

For servers' network interfaces, speed and duplex information is gathered during every hardware registration, which occurs every 24 hours. See "Agent Management" on page 105 for information about hardware registration.

Due to the lack of a device independent method of determining duplex for servers running a Windows operating system, the Server Agent for Windows does not report duplex settings out-of-the-box. A custom script can be added to the Server Agent to collect and report the speed and duplex setting for a certain network interface. For instructions on how to create and integrate the script with the Agent, contact HP Support.

Speed and duplex information for servers is *not* updated when you select **View ➤ Refresh** or press F5 in the SA Client. This data gets updated when the NA Duplex Data Gathering diagnostic runs. See “NA Duplex Data Gathering” on page 356.

For network devices, speed and duplex is gathered by the NA Duplex Data Gathering diagnostic, which runs according to a schedule that you define. To ensure that you have the latest speed and duplex information about network devices, Hewlett Packard recommends that you set up a recurring schedule that runs the diagnostic. See the *NA User’s Guide*.

If the network interface information (speed and duplex) for a server does not match the network port information (speed and duplex) for a connected network device, it is considered to be non-compliant.

In the NA Integration feature, you can see duplex mismatches identified at a top level by using the Dashboard. You can also see duplex mismatches identified by server and network device by using the Server Explorer and Network Device Explorer, respectively.

Viewing Duplex Mismatches in the Dashboard

See the *SA User’s Guide: Application Automation* for information about duplex compliance levels and how they are displayed in the Dashboard.

Viewing Duplex Mismatches by Server

To view duplex mismatches using the Server Explorer, perform the following steps:

1. From the Navigation pane, select Devices ➤ All Managed Servers.
2. In the Content pane, select a server.
3. Double-click on the server to display the Server Explorer.
4. In the Views pane, select Hardware.
5. In the Network Interfaces section, review the Duplex column for detected mismatches. Mismatches are identified by an ❌ icon that precedes the duplex setting (Full, Half, Auto), in the Duplex column.

Viewing Duplex Mismatches by Network Device

To view duplex mismatches using the Network Device Explorer, perform the following steps:

1. From the Navigation pane, select Devices ➤ Device Groups ➤ Public.

2. In the Content pane, select a network device.
3. Double-click on the network device to display the Network Device Explorer.
4. In the Views pane, select Hardware.
5. In the Network Ports section, review the Duplex column for detected mismatches. Mismatches are identified by an ❌ icon that precedes the duplex setting (Full, Half, Auto), in the Duplex column. See Figure 9-2 on page 359.

Network Reports

To help troubleshoot problems that involve physical connections and duplex compliance, you can run and examine several network reports. By using the Reports feature in the SA Client, you can produce the following network reports that identify layer 1 connections and configuration mismatches (duplex compliance) between managed servers and network devices in your environment:

Connections by Network Device

This report lists all physical connections to a selected network device.

Connections by Server

This report lists all physical connections to a selected managed server.

Duplex Compliance (All Servers)

This report groups all managed servers by duplex compliance level to show configuration mismatches between servers and network devices. Click on a section of the chart to display a list of servers in a certain compliance level. Double-click on a server for more details or to perform an action.

Duplex Compliance by Customer

This report lists all managed servers by customer and then by duplex compliance level to show configuration mismatches between servers and network devices. Double-click on a server for more details or to perform an action.

Duplex Compliance by Facility

This report lists all managed servers by facility and then by duplex compliance level to show configuration mismatches between servers and network devices. Click on a section of the chart to display a list of servers in a facility with a certain compliance level. Double-click on a server for more details or to perform an action.

See “Reports” in the *SA User’s Guide: Application Automation* for information about how to run, export, and print these reports.

Chapter 10: SA Global Shell

IN THIS CHAPTER

This section discusses the following topics:

- Overview of the Global Shell
- Global Shell Tutorial
- Global Shell Examples
- Character Encoding for the OGFS
- Global Shell Error Messages
- Remote Terminal

Overview of the Global Shell

The Global Shell is a command-line interface to the Global File System (OGFS). The command-line interface is a Unix shell such as `bash` that runs in a terminal window. The OGFS unifies the SA data model and the contents of managed servers, including files, into a single, virtual file system. You open a Global Shell session from within the SA Client Client or from a direct `ssh` connection in a terminal client on your desktop. With the Global Shell, you can automate repetitive system administration tasks by running scripts across multiple servers in a secure environment.

SA Global File System (OGFS)

The OGFS represents the SA data model as a hierarchical structure of file directories and text files. For example, in the OGFS, the `/opsw/Customer` directory contains details about SA customers and the `/opsw/Server` directory has information about managed servers. The `/opsw/Server` directory also contains subdirectories that reflect the contents (such as file systems and registries) of the managed servers. If you have the required permissions, in the Global Shell, you can view and even modify the file systems of managed servers.

Remote SA Shell (rosh) Utility

The Remote SA Shell (`rosh`) utility enables you to log on to managed servers 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`.

Benefits of the Global Shell

The Global Shell, OGFS, and `rosh` utility, offer the following benefits:

- **Security**: Logging on to managed servers is controlled by the HP Server Automation security framework.
- **Auditing**: Logins and commands on managed servers are recorded in audit log files.
- **Re-use of existing scripts**: Existing native scripts can run on managed servers with the `rosh` utility. For example, you can run .BAT, .vbs, and .sh scripts on managed servers. Scripts written in Unix shells will run within the Global Shell, which supports `bash`, `csh`, and other common shells.
- **Routine maintenance tasks on multiple servers**: By accessing the global view of the OGFS, system administration scripts can run iteratively on groups of servers.
- **Access to the HP Server Automation data model**: Global shell scripts can access information about managed servers, including custom attributes.

Commands Available in the Global Shell

The Global Shell offers the following types of Unix shells:

```
bash (default)
csh
ksh
sh
tcsh
```

Many common Unix commands (too numerous to list here) are available within the Global Shell. To display these commands, in a Global Shell session, list (`ls`) the contents of the following directories:

```
/bin /usr/bin
/opsw/bin
```

The `/opsw/bin` directory contains utilities (such as `rosh`) which are specific to SA. For more information, see "Global Shell Utilities Syntax" on page 519.

Differences Between the Global Shell and Unix Shells

The Global Shell is different from native Unix shells in the following ways:

- **Restricted command set**: Some Unix commands (such as `cron`) are unavailable in the Global Shell. To find out if a command is available, use the `which` command.
- **Limited recursion**: Commands cannot use recursion with the file systems of managed servers. Examples of recursive commands are `find`, `ls -r`, and `rm -r`.
- **SA user**: You log on to the Global Shell as an SA user, not as a Unix user.
- **SA permissions**: The operations that you can perform and the servers that you can access are limited by the SA permissions of your SA user group.
- **Private directories**: The following directories are accessible only by your SA user:

  ```
  /tmp
  /var/tmp
  /usr/tmp
  ```

  For example, the `/tmp` directory seen by the `jdoe` SA user is different than the `/tmp` seen by `tjones`.
- **SA data model in the OGFS**: Stored in the Model Repository, the data model consists of objects such as customers, facilities, and servers. End users manipulate these objects with the SA Client. The OGFS represents the data model in a file system that resembles a Unix file system. Changes to the data model appear as changes in the OGFS, and vice versa.
- **Axis (@) symbol in directory names**: In the OGFS, for example, this symbol appears in the following directories:

  ```
  /opsw/Server/@
  /opsw/Server/@Group
  /opsw/Group/Public/group-name/@
  ```

  The axis (@) symbol represents the end of the filtering criteria for managed servers.

Server Filtering in the OGFS

As you navigate down the OGFS tree, the path grows longer and more specific as fewer servers are visible in the `Server` directory. In the OGFS, the `/opsw` directory contains subdirectories for several types of objects in the SA model space, such as `Server`, `Group`, `Facility`, `OS`, `Application`, `Customer`, and so on.

In the Global Shell interface, you can filter your view of these object types in the `Server` directory by specifying an axis (@) in the path. A path in the SA model space can be a list of filtering criteria that selects objects of a given type. This path begins with the desired object type, such as `/Server`, and each filtering criteria begins with an `@`, such as `@Customer`. An ending `@` denotes the end of the filtering criteria.

Figure 10-1 is graphical representation of related objects (customers and facilities) in a hierarchical `Server` directory. The small boxes represent managed servers. Examples of ways that you can filter this directory immediately follow the diagram.

Figure 10-1: Filtering in the Server Directory

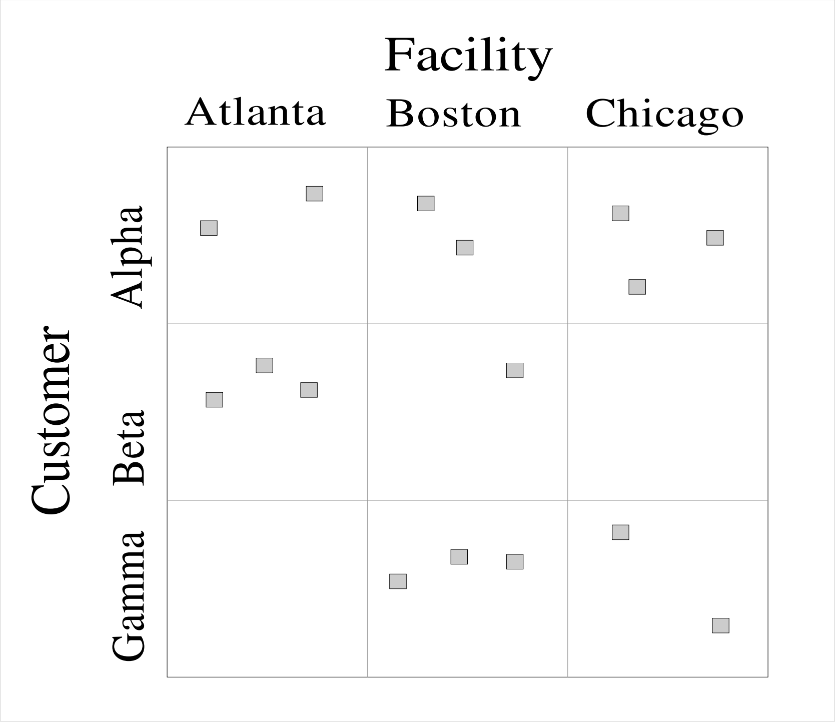

Based on Figure 10-1, the following examples illustrate ways to narrow your search for servers:

- To find all 16 servers, specify the following path:

```
ls /opsw/Server/@
```

- To find servers in the Atlanta facility, specify the following path:

```
ls /opsw/Server/@Facility/Atlanta/@
```

- To find servers that belong to customer Alpha, specify the following path:

```
$ ls /opsw/Server/@Customer/Alpha/@
```

- To find servers in the Atlanta facility that belong to customer Alpha, specify either of the following paths:

```
ls /opsw/Server/@Facility/Atlanta/@Customer/Alpha/@
ls /opsw/Server/@Customer/Alpha/@Facility/Atlanta/@
```

The following paths are filtered away by the OGFS, because they would yield a dead-end. There are no servers belonging to customer Gamma in the Atlanta facility.

```
ls /opsw/Server/@Facility/Atlanta/@Customer/Gamma/@
ls /opsw/Server/@Customer/Gamma/@Facility/Atlanta/@
```

This same filtering logic can be applied to `@Realm`, `@Group`, and `@Application`.

Global Shell Tutorial

This tutorial covers just a few of the highlights of the OGFS and the Global Shell. After completing this tutorial, you will know how to navigate the directories of the OGFS and how to run commands on managed servers from within the Global Shell. Although the tutorial is organized into steps, after performing step 1, you can perform the remaining steps in any order.

Before starting the tutorial, you need the following capabilities:

- You can log on to the SA Client. As you work through this tutorial, you might find it helpful to compare the stdout of the Global Shell with information displayed by the SA Client.
- Your SA user has Read & Write permissions on at least one managed server. Typically assigned by a security administrator, permissions are discussed in the *SA Administration Guide*.
- Your SA user has all Global Shell permissions on the same managed server. For information on these permissions, see “aaa Utility” on page 519.

The example commands in this tutorial operate on a Windows server named `abc.opsware.com`. This server belongs to a device group named All Windows Servers. When trying out these commands, substitute `abc.opsware.com` with the host name of the managed server you have permission to access. Also, replace `jdoe` with your SA user name. If you wish to run the commands on a Unix managed server, replace `ipconfig` with `ifconfig`; and replace `Administrator` with `root`.

Now, let's get started with the tutorial:

1. Open a Global Shell session.

   You can open a Global Shell session from within the SA Client. From the **Actions** menu, select **Global Shell**. You can also open a Global Shell session from a terminal client running on your desk top. For instructions, see "Opening a Global Shell Session" on page 388.

2. Check your session.

   First, enter the `whoami` command, which displays the SA user name for this session:

   ```
   $ whoami
   jdoe
   ```

   You can enter the ps command to view the process status of your Global Shell session. The following `ps` command shows the session is running the default `bash` shell:

   ```
   $ ps
   PID TTY          TIME CMD
    7033 ?        00:00:00 bash
   13712 ?        00:00:00 ps
   ```

   Enter the `uname` command, which displays information about the server running the OGFS component of SA:

   ```
   $ uname -a
   Linux m171.dev.opsware.com 2.4.21-32.ELsmp #1 SMP Fri Apr 15
   21:17:59 EDT 20 05 i686 GNU/Linux
   ```

   If you log on to a Unix managed server with `rosh`, `uname` displays information about that managed server, not the server running the OGFS component. Run the `uname` command when you are not sure if you are interacting with the Global Shell or with the shell of a managed server accessed with `rosh`.

3. Confirm your home directory.

   Every SA user has a home directory in the OGFS. The home directory has a `public/bin` subdirectory where you can store scripts to be executed by other users running Global Shell sessions. Each SA user also has a personal `/tmp` directory for temporary files. You cannot view or modify the `/tmp` directories of other users.

The following commands show some information about the directories of the `jdoe` user:

```
$ cd
$ pwd
/home/jdoe
$ ls -ld /home/jdoe/public/bin
drwxr-xr-x 2 jdoe jdoe 4096 2006-05-17 17:12 /home/jdoe/
public/bin
$ ls -ld /tmp
drwxrwxrwx 3 root root 4096 2006-06-09 23:37 /tmp
```

4 List all managed servers.

The `/opsw/Server` directory of the OGFS contains information about the servers managed by SA. This directory is an example of how the OGFS represents objects (in this case servers) of the SAdata model. Behind the scenes, SA stores this information in a database referred to as the Model Repository.

To view the names of the servers managed by SA, enter the following command:

```
$ ls /opsw/Server/@
abc.ospware.com      m33.opsware.com       gist.opsware.com
pal.opsware.com      hare.opsware.com       qv55.opsware.com
. . .
```

5 Examine server information.

Each managed server has a directory structure containing information about that server. The `attr` subdirectory contains text files that describe the server’s attributes. The attribute name matches the file name and the attribute value is the file contents. The following `cat` command lists the OS version of the managed server named `abc.opsware.com`:

```
$ cd /opsw/Server/@/abc.opsware.com
$ cat attr/osVersion
Microsoft Windows 2000 Advanced Server Service Pack 4 Build
2195 (05-02-2006
```

The `Interface` subdirectory has information about the server’s network interfaces. Here’s an example:

```
$ cat “Interface/Local Area Connection/info”
AdminEnabledFlg:  no
CardIndex:
CardSerialNum:
CircuitId:
Collisions:
ConfiguredDuplex: AUTO
ConfiguredSpeed:  AUTO
```

```
. . .
```

6 List the files of a managed server.

In addition to information about managed servers, `/opsw/Server` contains directories that correspond to the file systems of those servers. If you have the necessary permissions, in a Global Shell session you can access multiple servers from a single virtual file system, the OGFS.

The following command navigates to file system of the `abc.opsware.com` server:

```
$ cd /opsw/Server/@/abc.opsware.com/files
```

The next `ls` command displays OGFS subdirectories that correspond to native users of the managed server. Your security administrator specified these users (login names) when adding OGFS permissions. These are not SA users.

```
$ ls
Administrator LocalSystem
```

Native users might have different views of the managed server's file system. Therefore, under each user, the OGFS presents different file systems for each user. The following `cd` command drills down to the `Program Files` directory as seen by the `Administrator` user on the Windows server.

```
$ cd "Administrator/C/Program Files"
$ pwd
/opsw/Server/@/abc.opsware.com/files/Administrator/C/Program
Files
```

Next, list the files in the `Program Files` directory:

```
$ ls -1
Accessories
Common Files
ComPlus Applications
Internet Explorer
Messenger
. . .
```

Although these files reside in a directory on the managed server's file system, you are in the OGFS, as shown by the preceding `pwd` command. To verify that you are in a Global Shell session (and not in a session running on the managed server), enter the following commands:

```
$ whoami jdoe
$ uname -a
Linux m171.dev.opsware.com 2.4.21-32.ELsmp #1 SMP Fri Apr 15
21:17:59 EDT 2005 i686 GNU/Linux
```

7 Copy a file from the OGFS to a managed server.

By entering the `cd` command, go to your home directory in the OGFS, for example:

```
$ cd
$ pwd
/home/jdoe
```

Next, create a simple text file in your home directory:

```
$ echo “this is text” > myfile.txt
$ cat myfile.txt
this is text
```

Copy the file that you just created to a directory in the file system of a managed server. The following command copies `myfile.txt` to the `C:\temp` directory of the `abc.opsware.com` server:

```
$ cp myfile.txt \
/opsw/Server/@/abc.opsware.com\
/files/Administrator/C/temp/afile.txt
```

Do not copy large files between the OGFS and managed servers. Copy only small files, such as configuration files.

8 Log on to a managed server with `rosh`.

In the preceding steps, you accessed the file system of a managed server from within a Global Shell session. In this step, from the Global Shell you log on to a managed server with `rosh`. After you log in, you interact with the command-line environment (MSDOS or Unix shell) of the managed server.

The following `rosh` command logs in as Administrator to a Windows managed server named `abc.opwsare.com`:

```
$ cd /opsw/Server/@/abc.opsware.com
$ rosh -l Administrator
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-2000 Microsoft Corp.
```

The prompt indicates that you are now in the command-line environment of the managed server. Enter the `ipconfig` and `hostname` commands:

```
C:\WINNT\system32>ipconfig
ipconfig

Windows 2000 IP Configuration

Ethernet adapter Local Area Connection:

        Connection-specific DNS Suffix  . : opsware.com
        IP Address. . . . . . . . . . . . : 192.168.8.217
        Subnet Mask . . . . . . . . . . . : 255.255.254.0
        Default Gateway . . . . . . . . . : 192.168.8.1

C:\WINNT\system32>hostname
hostname
abc
```

Terminate the remote login with the `exit` command:

```
C:\WINNT\system32>exit
```

Enter the `uname` command to verify that you have returned to the Global Shell session:

```
$ uname -a
Linux m171.dev.opsware.com 2.4.21-32.ELsmp #1 SMP Fri Apr 15
21:17:59 EDT 2005 i686 GNU/Linux
```

9 Create a script that runs across servers.

A Global Shell script can iterate within the OGFS and run the `rosh` command to execute native commands on multiple servers. The example script shown in this step iterates through the servers of the public device group named All Windows Servers. On each server, the script runs the `ipconfig` command with the `rosh` command. In this example, substitute your SA user name for `jdoe`.

First, return to your home directory in the OGFS:

```
$ cd
$ cd public/bin
$ pwd
/home/jdoe/public/bin
```

Next, run the `vi` editor:

```
$ vi
```

In `vi`, insert the following lines to create a `bash` script:

```
#!/bin/bash
# This is simple_iterate.sh.
# Change jdoe to your user name.

OUTFILE="/home/jdoe/public/bin/ipconfig_all.txt"
rm -f $OUTFILE

cd "/opsw/Group/Public/All Windows Servers/@/Server"

for SERVER_NAME in *
do
   echo ---- $SERVER_NAME
   echo ---- $SERVER_NAME >> $OUTFILE
   rosh -n $SERVER_NAME -l Administrator \
           "ipconfig" >> $OUTFILE
done
# Last line in simple_iterate.sh.
```

Save the file in `vi`, naming it `simple_iterate.sh`. Quit `vi`.

Change the permissions of `simple_iterate.sh` with `chmod`, and then run it:

```
$ chmod 755 simple_iterate.sh
$ ./simple_iterate.sh
---- abc.ospware.com
---- gist.opsware.com
---- hare.opsware.com
---- m33.opsware.com
. . .
```

As the script runs, it echos the name of each server to `stdout`, and redirects the output of the `ipconfig` command to the `ipconfig_all.txt` file. Enter the `more` command to view the contents of `ipconfig_all.txt`:

```
$ more ipconfig_all.txt
---- abc.ospware.com
Windows 2000 IP Configuration
Ethernet adapter Local Area Connection: . . .
```

10 Learn more.

Here are a few suggested tasks for learning more about the OGFS and the Global Shell:

- Explore the folders and contents under `/opsw/Library`, comparing them with the Library windows of the SA Client.
- If you have NA installed, navigate to the `/opsw/Net*` (network) directories. For descriptions of these directories, see “Network Directories” on page 543.
- On Windows servers, examine the `registry` and `complus` directories under `/opsw/Server/@/`*`server-name`*.
- List the files in the `method` directory, also under `/opsw/Server/@/`*`server-name`*. These files are the executables of the SA Command-Line Interface (OCLI), which enables you to perform SA functions from within the Global Shell. To learn how to run the CLI methods, see the *SA Platform Developer’s Guide*.

Global Shell Examples

The examples in this section use `bash`, the default shell of a Global Shell session. These are relatively simple examples. For more complex examples, including those with method invocations and search filters, see the *SA Platform Developer's Guide*.

Opening a Global Shell Session

You can open a Global Shell session with an `ssh` client or from within the SA Client. When you open a session, the working directory is `/home/`*`user-name`*.

To open a Global Shell session with an `ssh` client, perform the following steps:

1. On a host that is not an SA core server or a managed server, open a terminal window.

2. In the terminal window, enter an `ssh` command with the following syntax:

   ```
   ssh -p 2222 user-name@ogfs-host
   ```

   To use this command, port 2222 must be open on the firewall that protects the OGFS server. The *user-name* is your SA user (login) and the *ogfs-host* is the host name (or IP address) of the core server running the OGFS. The SA user name is not case sensitive. After you enter the `ssh` command, the OGFS prompts for the password of the SA user.

To open a Global Shell session from within the SA Client, from the **Actions** menu, select **Global Shell**.

Finding Servers in the OGFS

List the names of all servers that you can manage with SA:

```
ls /opsw/Server/@
```

List the IDs of the servers:

```
ls -a /opsw/.Server.ID
```

Find all servers in the `.opsware.com` domain by specifying the wildcard character (*):

```
ls -d  /opsw/Server/@/*.opsware.com
```

List all servers in the Atlanta facility:

```
ls /opsw/Server/@Facility/Atlanta/@
```

List the servers in the public device group named Alpha:

```
ls /opsw/Group/Public/Alpha/@/Server
```

List the servers in the All Windows Servers group, first with backlashes to escape the spaces in the group name, then by enclosing the option in single quotes:

```
ls /opsw/Group/Public/All\ Windows\ Servers/@/Server
ls '/opsw/Group/Public/All Windows Servers/@/Server'
```

List the servers of the Widget customer:

```
ls /opsw/Customer/Widget@/Server
```

The following two commands display the same output, the servers in the Atlanta facility that belong to the Green customer:

```
ls /opsw/Server/@Facility/Atlanta/@Customer/Green/@
ls /opsw/Server/@Customer/Green/@Facility/Atlanta/@
```

Getting Server Information from the OGFS

List the Opsware ID and of the server named `m256.opsware.com`:

```
cd /opsw/Server/@/m256.opsware.com
cat self:i ; echo
```

(The preceding `echo` command is optional. It generates a new line character, which makes the output easier to read. The semicolon separates `bash` statements entered on the same line.)

List the name of the server with an Opsware ID of `340039`:

```
cat  /opsw/.Server.ID/340039/self
```

By iterating through the server names with a `for` loop in `bash`, display the platform (operating system) name for each server:

```
cd /opsw/Server/@
for SERVER_NAME in *
do
   cat $SERVER_NAME/attr/platform
done
```

Display the amount of RAM in the server named `abc.opsware.com`:

```
cd /opsw/Server/@/abc.opsware.com
grep Quantity Memory/RAM/info
```

Display the network interfaces of a the server named `blizzard.opsware.com`:

```
cd /opsw/Server/@/glengarriff.snv1.dev.opsware.com/Interface
for INTERFACE_NAME in *
do
   echo ............. $INTERFACE_NAME ............
```

```
      grep  Interface "$INTERFACE_NAME/info"
      echo ""
   done
```

Browsing a Server's File System or Registry

List all of the files the `C:\ Program Files` directory of the Windows server named `abc.opsware.com`:

```
cd /opsw/Server/@/abc.opsware.com/files/Administrator
ls C/Program\ Files
```

List the registry keys of the `abc.opsware.com` server:

```
cd /opsw/Server/@/abc.opsware.com/registry/\ Administrator/
ls *
```

List the contents of the `/var` directory on the Unix server named `m256.opsware.com`:

```
/opsw/Server/@/m256.opsware.com/files/root
ls var
```

Managing Custom Attributes

On the server `abc.opsware.com`, create a custom attribute named `MyGreeting` with the value `hello there`:

```
cd /opsw/Server/@/abc.opsware.com/CustAttr
echo -n "hello there" > MyGreeting
cat MyGreeting
```

Execute `runit.bat` on the `abc.opsware.com` server, passing the value of the `My Test` custom attribute as a command-line parameter for `runit.bat`:

```
cd /opsw/Server/@/abc.opsware.com
TESTPARAM=`cat CustAttr/"My Test"`
rosh -l Administrator "C:\temp\runit.bat $TESTPARAM"
```

- When you create or edit custom attributes within the OGFS, HP Server Automation preserves leading and trailing whitespace characters in custom attribute values.

Copying Files Within the OGFS

Do not use the techniques in this section to copy large files. The OGFS is not designed distribute large amounts of data. However, you can use these techniques for copying small files (such as configuration files) to and from managed servers.

Copy `myfile.txt` from the home directory in the OGFS of the `jdoe` user to the `C:\temp` directory of the Windows server named `abc.opsware.com`:

```
cp  /home/jdoe/myfile.txt \
/opsw/Server/@/abc.opsware.com/files/Administrator/C/temp
```

Copy `myfile.txt` from the home directory in the OGFS of the `jdoe` user to the `/tmp` directory of the Unix server named `m25.opsware.com6`:

```
cp /home/jdoe/myfile.txt \
/opsw/Server/@/m256.opsware.com/files/root/tmp
```

Copy the `C:\temp\myfile.txt` file from the `abc.opsware.com` server to the `m344.opsware.com` server:

```
cp /opsw/Server/@/abc.opsware.com/files/\
Administrator/C/temp/myfile.txt \
/opsw/Server/@/m344.opsware.com/files/\
Administrator/C/temp/myfile.txt
```

Copying Files Between the OGFS and a Development Server

You can securely copy files between the OGFS and a server that is not part of HP Server Automation. To copy the files, perform the following steps:

1. On a host that is not an SA core server or a managed server, open a terminal window.

2. In the terminal window, enter either the `scp`, `sftp`, or `rsync` command and specify port 2222, your SA user name, and the host running the OGFS.

The following three `scp` examples perform the same operation: They copy the file `myscript.sh` from the local machine to the file `/home/jdoe/myscript.sh` in the OGFS. The SA user is `jdoe` and the host running the OGFS is 192.168.166.178.

```
scp -P 2222 myscript.sh jdoe@192.168.166.178:myscript.sh
scp -P 2222 myscript.sh jdoe@192.168.166.178:/home/jdoe
scp -P 2222 myscript.sh jdoe@192.168.166.178:
```

The following example copies `myscript.sh` from the home directory of `jdoe` in the OGFS to the local machine:

```
scp -P 2222 jdoe@192.168.166.178:myscript.sh myscript.sh
```

The following `sftp` example copies `myscript.sh` from the local machine to the OGFS:

```
sftp -oPort=2222 jdoe@192.168.166.178
Connecting to 192.168.166.178...

Opsware Global Shell
```

```
jdoe@opsware's password:
sftp> put myscript.sh
. . . .
```

The following rsync example transfers files from /path on the local machine to /other/path in the OGFS:

```
rsync -av -e "ssh -p 2222" /path \
jdoe@192.168.166.178:/other/path
```

Logging on to a Managed Server With rosh

The next three rosh commands perform the same operation: logging on to the Windows server named abc.opsware.com as the Administrator user. After logging on, the current working directory on the remote shell is the default working directory of the Administrator Windows user. These rosh commands require different options, depending on the current working directory in the OGFS. For example, the first rosh command does not require the -n (server name) and -l (user) options because the option values can be inferred from the current working directory of OGFS. The options of the following three rosh commands differ because of the current working directory:

```
cd /opsw/Server/@/abc.opsware.com/files/Administrator
rosh
. . .
exit
. . .
cd /opsw/Server/@/abc.opsware.com
rosh -l Administrator
...
exit
. . .
cd /home/jdoe
rosh -n abc.opsware.com -l Administrator
...
exit
```

The next rosh command logs into the Unix server named m256.opsware.com as the root user with the current working directory of /tmp:

```
rosh -n m256.opsware.com -l root -d /tmp
```

Running OGFS Scripts on Managed Servers With rosh

The next sequence of commands create a .BAT script in the OGFS and then run the script on a Windows managed server. Created with echo statements, the myfile.bat script resides in the OGFS under /home/jdoe/public/bin. (Note that myfile.bat does not reside in the file system of the managed server.) The myfile.bat script contains three commands: ipconfig, cd, and dir. The rosh command runs myfile.bat on the server named abc.opsware.com as the Administrator Windows user. The following commands create a local .BAT script and run it remotely with rosh:

```
cd /home/jdoe/public/bin

echo ipconfig > myfile.bat
echo "cd c:\temp" >> myfile.bat
echo dir >> myfile.bat

rosh -n abc.opsware.com -l Administrator -s ./myfile.bat
```

Create a script named who.sh in the /home/jdoe/public/bin directory of the OGFS and then run who.sh on the server named m256.opsware.com:

```
cd /home/jdoe/public/bin

echo \#\!\/bin\/sh > who.sh
echo "uname -n" >> who.sh
echo id >> who.sh
echo pwd >> who.sh

rosh -n m256.opsware.com -l root -s ./who.sh
```

Running Native Programs on Managed Servers With rosh

The next two rosh commands run the dir and ipconfig MSDOS commands on the Windows server named abc.opsware.com. Note that the native MSDOS commands are enclosed in quotes. Because the server name and user (login) can be inferred from the current working directory, the first rosh command omits the -n and -l options, as shown in the following code:

```
cd /opsw/Server/@/abc.opsware.com/files/Administrator
rosh "dir & ipconfig"
. . .
cd /home/jdoe
rosh -n abc.opsware.com -l Administrator "dir & ipconfig"
```

Run the ipconfig command on abc.opsware.com and redirect the output to a file in home directory of jdoe in the OGFS:

```
rosh -n abc.opsware.com -l Administrator "ipconfig" \
> /home/jdoe/ipconfig_ouptput.txt
```

On the Unix server named m256.opsware.com, run the uname and ls commands as root:

```
rosh -n m256.opsware.com -l root "uname -a; ls /tmp"
```

Within a for loop in bash, run the MSDOS ipconfig command on each server in the All Windows device group:

```
cd /opsw/Group/Public/All\ Windows\ Servers/@/Server

for SERVER_NAME in *
do
   echo ............. $SERVER_NAME ............
   rosh -n $SERVER_NAME -l Administrator "ipconfig"
   echo ""
done
```

Character Encoding for the OGFS

To support international environments, the OGFS can display information in different character encodings such as Shift-JIS (Japanese) and EUC-KR (Korean). You can control the encoding of Global Sessions in the following ways:

- To specify the encoding of your Global Shell sessions, in the Terminal and Shell Preferences of the SA Client, select an item from the Encoding drop-down list.
- To change the encoding of an active Global Shell session, run the swenc command with the -e option.

If you change the encoding of an active session, you must also change the encoding of the terminal application for that session. This procedure varies according to the terminal application.

Terminal Application Configuration

The terminal application that is hosting a Global Shell or Remote Terminal session must be configured to use the same encoding expected by the session. If the encodings do not match, the data might be displayed incorrectly.

When the SA Client launches the terminal application, it composes the command specified by the Terminal Client field in the Terminal and Shell Preferences. If the Terminal Client field includes the %e substitution variable, the SA Client replaces %e with an encoding name. For Global Shell sessions, this encoding name is specified by the Encoding field of the Terminal and Shell Preferences. For Remote Terminal sessions, this encoding name is the value of the Encoding field in the Properties section of the Device Explorer. If the terminal application does not support the encoding that replaces the `%e` variable, or if the `%e` variable is not specified, you must change the encoding manually in the terminal application after it starts.

Data that Cannot Be Displayed

Data that cannot be displayed might be from a managed server (such as the contents of files) or it might be the name of an object in the SA model. If the session's encoding does not support the data to be displayed, the data often appears as question marks. (However, it might appear as other characters such as exclamation points.) The session attempts to display this data with the current encoding. Usually, this data cannot be accessed. To access this data, select a compatible encoding for the session.

Objects in the SA model, such as Facility and Server, appear as file names in the OGFS. If these file names contain characters that cannot be represented by the encoding of the session, they are displayed as question marks, appended by the Opsware ID of the object. In the following example, the IDs are 10002 and 11002:

```
New York
Paris
Montr?al~10002
??~11002
```

If the model object does not have an ID, such as a custom attribute, then the session attempts to display the name with the current encoding.

LANG and LC_CTYPE Environment Variables

Many Unix commands (such as `ls`) rely on the character encoding, which is determined by the `LANG` or `LC_CTYPE` environment variables. In a Global Session, if the encoding is changed with the `swenc` command, the system attempts to reset these variables.

HP Server Automation determines the new value of the `LANG` variable with the following process:

1. The value of `LANG` is generated by combining the language of the user's profile in the SAS Web Client with the current session encoding. For example, if the language is English and the session encoding is UTF-8, `LANG` is set to en_US.utf8.

2. The value determined by the preceding step is compared with the set of valid locales on the OGFS server (according to the output of `locale -a`). If the value is a valid locale, `LANG` is set to this value.
3. If the value is not a valid locale, the system attempts to find a valid locale that specifies the user's language without the encoding. For example, if the user's language is English and the session encoding is EUC-JP, and this combination does not form a valid locale, `LANG` is set to en_US. If no matching locale is found, `LANG` is left unspecified.

The new value of the `LC_CTYPE` variable is determined in the following order:

1. HP Server Automation attempts to find a valid locale that matches the session encoding, regardless of the language.
2. If a valid locale is found, `LC_CTYPE` is set to this locale. For example, if the session encoding is EUC-JP, the `LC_CTYPE` variable is set to ja_JP.eucjp.
3. If no matching locale is found, `LC_CTYPE` is left unspecified.
4. If HP Server Automation cannot set the `LANG` or `LC_CTYPE` variables with the preceding process, you should set them manually.

Transcoded Data in a Managed Server

Transcoding is the conversion of data from one character encoding to another. HP Server Automation automatically transcodes some of the data between Global Shell sessions and other sources of data. For example, the file names of managed servers are transcoded, but the contents of the files are not. To see which data is transcoded, see Table 10-1. To display the transcoding mode of the current Global Shell session, enter the `swenc` command with no options.

Table 10-1: Data Transcoded for Global Shell Sessions

DATA	TRANSCODING
Objects in the SA model space, such as `Facility`, `Customer`, and `Server`. These objects are stored in the SA Model Repository (database) in UTF-8.	Between UTF-8 and the session encoding.
File and directory names of managed servers.	Between the managed server encoding and the session encoding.

Table 10-1: Data Transcoded for Global Shell Sessions (continued)

DATA	TRANSCODING
Meta-data of managed servers, such as user names and registry key names.	Between the managed server encoding and the session encoding.
File contents of managed servers.	None
Contents of Windows registries, services, COM objects, and IIS metabases.	None
`rosh`: Contents of saved scripts executed on managed servers (a `rosh` push operation).	None
`rosh`: Ad-hoc scripts executed on managed servers.	None
`rosh`: Command-line arguments to programs executed on managed servers.	None
`rosh`: Data streams (such as stdin and stdout) of programs executed on managed servers.	None
`rosh`: Data streams of `rosh` jump operations.	None
OGFS `home`, `tmp`, and `var/tmp` directories.	None

Disabling the Transcoding of Managed Server Data

On Unix servers, file and directory names can contain characters in arbitrary encodings. When accessed through the OGFS, file and directory names are transcoded by HP Server Automation. If the encoding of the names does not match the default encoding of the managed server, the transcoded data might be unusable.

You can disable transcoding with the `swenc` command. To turn transcoding on or off, use the `-T` option:

```
swenc -T {on | off}
```

If transcoding is disabled, file and directory names are passed unmodified from the managed servers. Therefore, you must manually configure the encoding of the terminal application to display the names correctly.

Windows servers store their file system data internally in the UTF-16 encoding. Because the encoding is known, transcoding is performed correctly and does not need to be disabled. Therefore, the `-T` option of the `swenc` command has no effect on Windows servers. For more information, see “swenc Syntax” on page 529.

Global Shell Error Messages

The Global Shell feature provides the file system error messages that are described in Table 10-2

Table 10-2: Global Shell Errors

ERROR	DESCRIPTION	ACTION
Input/output error	Your session has exceeded the time-out limit or the Agent is not running.	Start a new session or check the status of the Agent.
Operation not permitted.	No password was found.	Verify that you have a valid password.
Permission denied.	You are not allowed to view a directory. This does not mean that the directory does not exist on a given server. See the *SA Administration Guide* for more information.	Verify that you have `readFileSystem` permissions.
RFS Specific error	You do not have permissions on the managed server. For example, this error will occur if you are trying to perform an operation on a managed server and you do not belong to the Administrators group that has the required permissions assigned to it.	You must have a set of permissions to perform operations on managed servers. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Remote Terminal

With HP Server Automation, you can log on to a managed server in a terminal window in two ways:

- The Remote Terminal feature, as described in this section.
- The `rosh` command, entered in a Global Shell session.

The Remote Terminal feature opens a terminal window for Unix servers or an RDP client window for Windows servers. Unlike a Global Shell session, a Remote Terminal session does not provide access to the OGFS.

You can specify your terminal and RDP client preferences for Remote Terminal and Global Shell sessions in the Set Preferences window. See “Terminal and Shell” on page 97 in Chapter 3 for more information.

Prerequisite for a Remote Terminal

This feature requires SA login permissions on the managed server. See the *SA Administration Guide* for more information. You cannot log into a remote terminal with a user name or password that contains multi-byte characters.

In order to open a Remote Terminal session on a Windows managed server running Windows Firewall, you must set an exception that is not enabled by default. On the managed server, run Windows Security Center, then select Windows Firewall ➤ Exceptions ➤ Remote Desktop.

This feature must be able to establish a loopback connection on the machine running the SA Client. Some firewall and VPN utilities have settings that prevent loopback connections.

Opening a Remote Terminal

In the Device Explorer of the SA Client, perform the following steps:

1. Select a managed server.
2. Open the managed server.
3. From the **Actions** menu, select **Remote Terminal**.

Chapter 11: Code Deployment and Rollback

IN THIS CHAPTER

This section discusses the following topics:

- Code Deployment Process
- Services, Synchronizations, and Sequences

You must have specific permissions to deploy code and content by using the SAS Web Client. Contact your SA administrator to obtain the necessary access rights. For more information, see the Permissions Reference appendix in the *SA Administration Guide*.

The Code Deployment and Rollback (CDR) feature is not supported on a VMware ESX Hypervisor.

Code Deployment Process

This section provides information on the code deployment process within HP Server Automation and contains the following topics:

- Deploying Code
- Uploading Code and Content to Staging
- CDR Operations and Directories
- CDR Features
- CDR Permissions
- Accessing CDR

Deploying Code

The Code Deployment & Rollback (CDR) feature in the SAS Web Client provides tools for deploying new and updated code and content to your operational environment.

The following figure shows the architecture and process for updating a typical server hosted in an SA managed environment.

Figure 11-1: Typical Code and Content Update in the SA *Managed Environment*

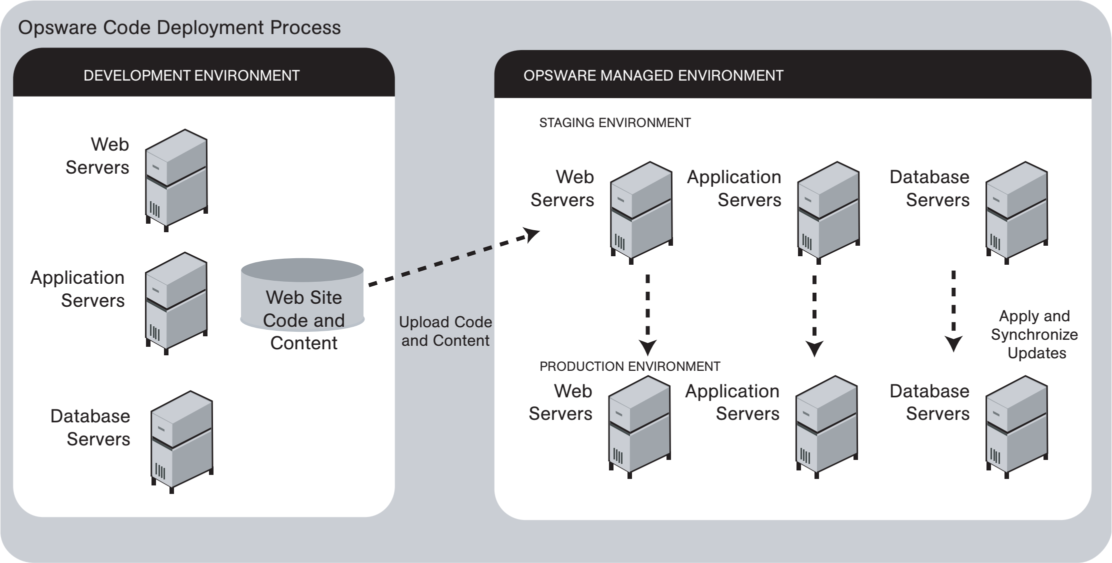

The deployment process involves performing the following high-level tasks:

1. Determining your application code and content deployment requirements and defining the CDR services, synchronizations, and sequences that you need to support them.

   - Services are defined for each different type of web server or application server applications (for example, WebLogic Server) that is installed on the staging and production hosts in your environment.

   - Synchronizations are defined for each service so that you can update files between the source location and one or more destination production hosts that are running the same service.

   - Sequences are optional but can simplify deployment by grouping a collection of service operations and synchronizations that can be performed as a single task.

2. Uploading new or updated code and content to your staging environment.

3. After performing any necessary testing, cutting over to the changed code and content on the staging environment.
4. As necessary, performing CDR service operations, such as backing up code and content from your live site.
5. Performing CDR operations available to synchronize the updated code and content to your production hosts in the SA managed environment.
6. To simplify subsequent deployments of new code and content, defining sequences that specify a series of service operations and synchronizations you want to perform as a single action.

The code and content deployment process that you follow might be different depending on the architecture of your operational environment and your deployment requirements.

Uploading Code and Content to Staging

Before you use CDR to push code and content, you must upload new or updated files to your SA staging environment. You can use content management tools, such as OpenDeploy, scp, or rsync over SSH, to do that.

The following figure shows an example of a typical development environment and how your uploaded code and content move to the staging environment.

Figure 11-2: How Code and Content Move to the Staging Environment

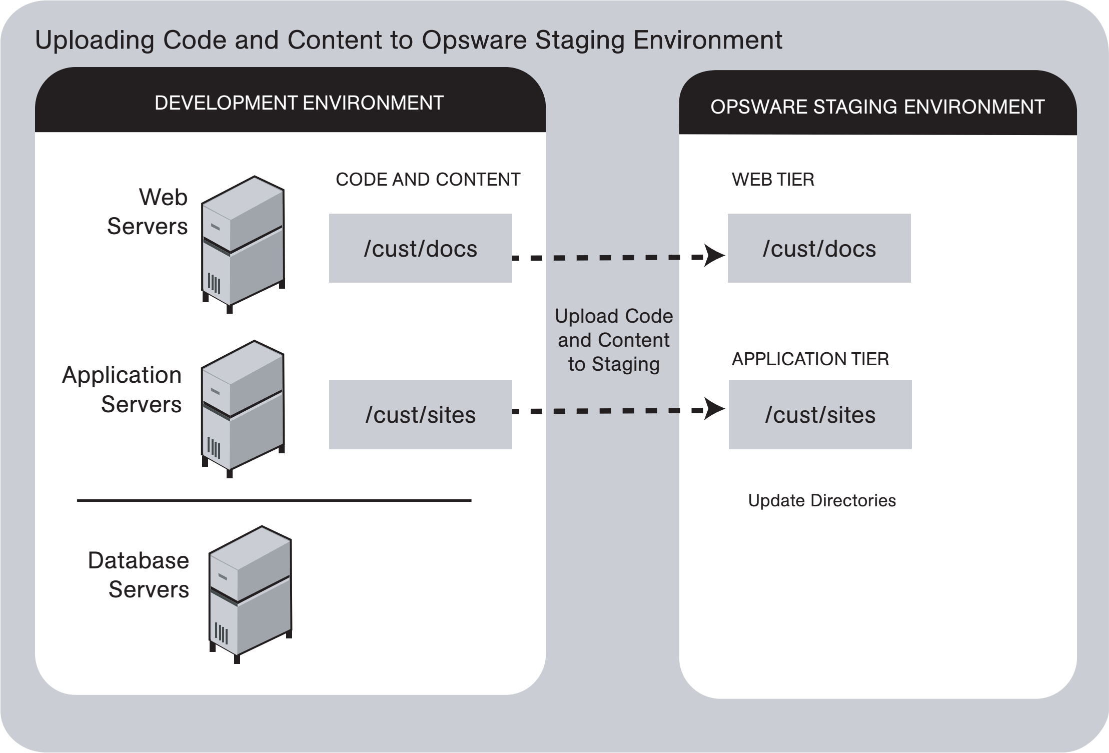

After you upload the files and test your changes, you can synchronize updates to the production hosts running your managed environment. You can run specific synchronizations and perform other service deployment operations by selecting CDR menu options available from the SAS Web Client navigation panel.

CDR Operations and Directories

After you upload updated code and content to your SA-managed staging environment, you can use the CDR operations to cutover to new code and content, perform host synchronizations, and perform other service operations.

CDR uses the following directories to synchronize and cutover code and content for specified hosts:

- **Live directory**: The directory that stores the actual code and content required to run a live site.

- **Update directory**: The directory written to by CDR synchronizations. Stores only the files that changed between the source host Live directory and the Live directories of the destination hosts.
- **Site Previous directory**: This directory holds all the changes necessary to revert the Live directory back to the state it was in before the last cutover. Like the Update directory, the Site Previous directory only stores the files that changed between the current Live directory contents and its previous state.
- **Site Backup directory**: This directory stores a complete backup of the site. The directory is populated when the user issues a Backup service operation.

When you cutover to new code and content, CDR determines the differences between the new code and content in the current Update directory and the Live directory for your site. The files that are different are synchronized to the Live directory. When you synchronize source and destination hosts, CDR moves modified files from the Live directory on a source host to a directory on a destination host.

You cannot use CDR to automate database pushes. However, you can configure CDR so that you can synchronize modified database script files on different hosts.

CDR Features

CDR offers the following features:

- Provides a single tool for deploying code (such as ASP, JSP, and JAR files) and site content (such as HTML, JPEG, GIF, and PDF files). Using a single tool is helpful when the code and content for your site are intermingled.
- Provides direct control over code and content pushes by making it possible to decide what information to update and determine when and how to perform updates.
- Provides flexibility to accommodate frequent updates to staging and production hosts by enabling more frequent pushes in a shorter period of time.
- Allows verification of file changes between staging and production host directories by creating a manifest of updated files. You can verify changes before cutting over to new code and content.
- Provides administrative service operations, including starting and stopping services, and backing up, restoring, and rolling back code and content to return your site to the previous version.

- Lets you push incremental updates to your site so that only files that have changed are pushed to specified locations on staging or production hosts.
- CDR uses the same authentication and navigation that you use in accessing other information and performing other site operations from the SAS Web Client.

CDR Permissions

As with all other features in HP Server Automation, the links that you see on the SAS Web Client Home page and the links that you see in the navigation panel are based on the permissions that you have in combination with the customer you are associated with.

If you do not have permissions for CDR, you cannot see the Code Deployment links on the navigation panel, the link called Deploy Code in the Tasks panel of the SAS Web Client home page appears in italics, and it is not an active link.

If you have CDR permissions to no more than one customer, when you expand the Code Deployment section in the navigation panel, you can see a link called Set Customer. Click that link to view the links to the specific Code Deployment functions that you have permissions for in combination with that single customer.

If you have CDR permissions to more than one customer, you can see a link called Select Customer. Click that link to display a page that shows the customers you are associated with. Select the customer you want to work with. The CDR Home Page appears, with links to the specific Code Deployment functions that you have permissions for. These links are the same functions that you can find in the navigation panel under Code Deployment.

The navigation instructions and screen captures in this chapter show what a user with permissions to all code deployment functions and access to only one customer can see. Consequently, because your permissions and customers might be different, the available menu selections and features that you see might likewise differ.

Accessing CDR

Perform the following steps to access CDR:

1. If necessary, click the Code Deployment link in the navigation panel to expand the list of CDR options.
2. Click the CDR Home link. The CDR Home Page for [*customer name*] appears, as the following figure shows.

Figure 11-3: Code Deployment Home Page

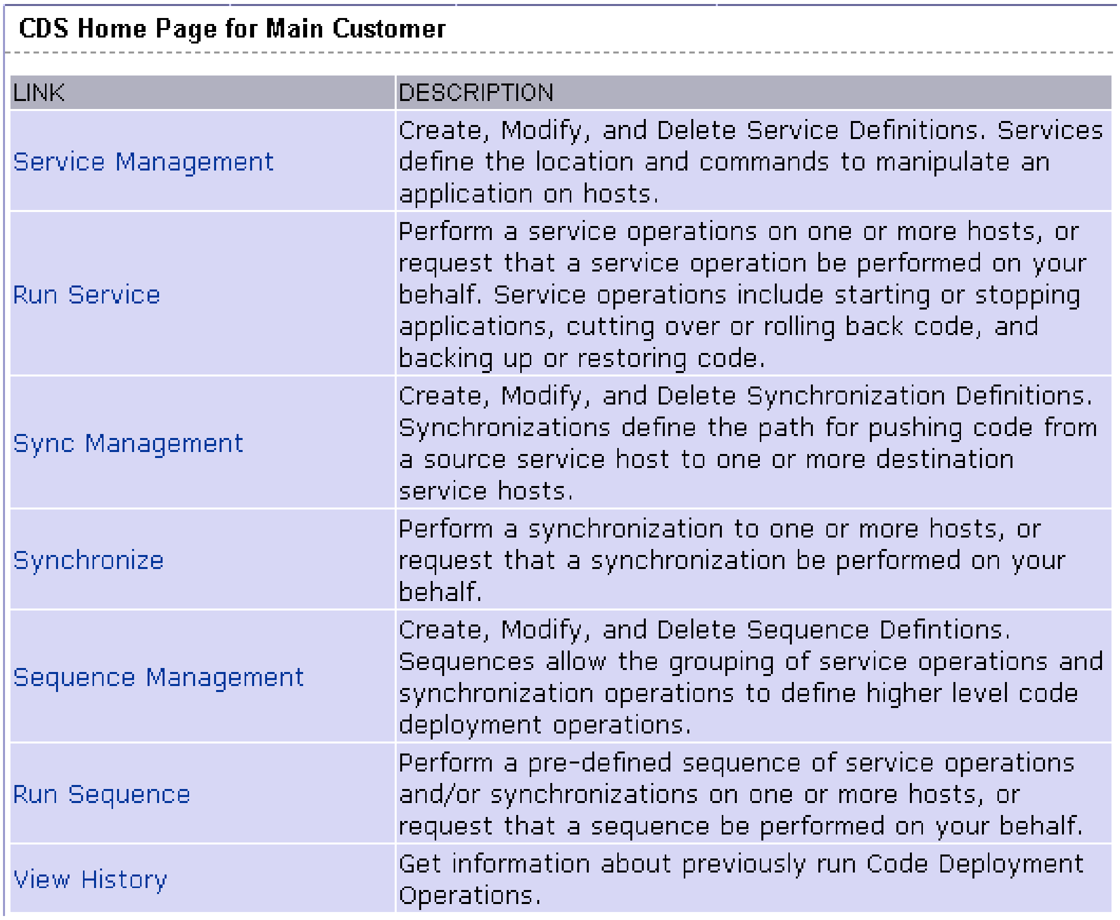

CDS Home Page for Main Customer

LINK	DESCRIPTION
Service Management	Create, Modify, and Delete Service Definitions. Services define the location and commands to manipulate an application on hosts.
Run Service	Perform a service operations on one or more hosts, or request that a service operation be performed on your behalf. Service operations include starting or stopping applications, cutting over or rolling back code, and backing up or restoring code.
Sync Management	Create, Modify, and Delete Synchronization Definitions. Synchronizations define the path for pushing code from a source service host to one or more destination service hosts.
Synchronize	Perform a synchronization to one or more hosts, or request that a synchronization be performed on your behalf.
Sequence Management	Create, Modify, and Delete Sequence Defintions. Sequences allow the grouping of service operations and synchronization operations to define higher level code deployment operations.
Run Sequence	Perform a pre-defined sequence of service operations and/or synchronizations on one or more hosts, or request that a sequence be performed on your behalf.
View History	Get information about previously run Code Deployment Operations.

Depending on your access permissions, the following CDR options appear:

- **Service Management**: Create, modify, or delete service definitions that define the location and commands to manipulate an application on hosts associated with each application instance running in your operational environment.

- **Run Service**: Perform a service operation or request that one be performed.
- **Sync Management**: Create, modify, or delete synchronization definitions associated with code pushes.
- **Synchronize**: Perform a synchronization or request that one be performed.
- **Sequence Management**: Create, modify, or delete sequences of operations.
- **Run Sequences**: Perform a selected sequence or request that one be performed.
- **View History**: View information stored in an operations log to determine the status of particular deployment operations, and whether they completed successfully.

3. Choose the CDR operations that you want to perform, selecting options from the navigation panel or from the CDR home page.

Services, Synchronizations, and Sequences

This section provides information on performing services, synchronizations, and sequences within HP Server Automation and contains the following topics:

- Overview of Performing Services, Synchronizations, and Sequences
- Synchronization of Site Code and Content
- Performing Synchronizations
- CDR Cutover Operation
- CDR Service Operations
- Starting and Stopping Host Services
- Backing Up Code and Content
- Restoring Code and Content from a Previous Version
- Rolling Back Code and Content to the Previous Version
- Accessing Service Operations in CDR
- Performing Service Operations by Service Name
- Performing Service Operations by Host Name
- Performing Sequences
- Processing Code Deployment Requests from Users

- Status View of Previous Operations

Overview of Performing Services, Synchronizations, and Sequences

After you upload updated code and content to your HP Server Automation staging environment, you use CDR to cutover to new code and content, perform host synchronizations, and perform other service operations.

When you cutover to new code and content, CDR determines the differences between the new code and content in the current Update directory and the Live directory. The files that are different are synchronized to the Live directory. When you synchronize source and destination hosts, CDR moves modified files from the Live directory on a source host to a directory on a destination host.

The code and content footprint that CDR maintains on a host is larger than the storage space for the code and content files of the live site. The amount of storage used beyond the actual size of your site increases with the number of files that you modify and back up.

Synchronization of Site Code and Content

CDR uses the following directories to synchronize and cutover code and content for specified hosts:

- **Live directory**: The directory that stores the actual code and content required to run a live site.
- **Update directory**: The directory written to by CDR synchronizations. Stores only the files that changed between the source host Live directory and the Live directories of the destination hosts.
- **Site Previous directory**: This directory holds all the changes necessary to revert the Live directory back to the state it was in before the last cutover. Like the Update directory, the Previous directory only stores the files that changed between the current Live directory contents and its previous state.
- **Site Backup directory**: This directory stores a complete backup of the site. The directory is populated when the user issues a Backup service operation.

CDR can synchronize updates from the Live directory of a source host to either the Update or Live directories on destination hosts. Decide whether you want to synchronize changed files on the source host to Update or Live directories on the destination hosts:

- **Synchronize to Update directories**: CDR determines files that have changed by comparing files in the destination host Live directory and the source host Live directory. Updated files are stored in the Update directories of destination hosts.
- **Synchronize to Live directories**: CDR updates changed files to the Live directory on destination hosts bypassing the Update directory.

If you choose to synchronize directly to Live directories, the Rollback operation does not function properly. Therefore, choose this option only for synchronizations that are not likely to impact site stability.

When you choose to synchronize to Live directories, back up the Live directories first and run the Restore operation to return your site to the previous version. See "Rolling Back Code and Content to the Previous Version" on page 417 in this chapter for more information.

Figure 11-4 shows an example of synchronization between hosts and how you synchronize updates for each service.

Figure 11-4: Service Synchronization of Hosts Using CDR

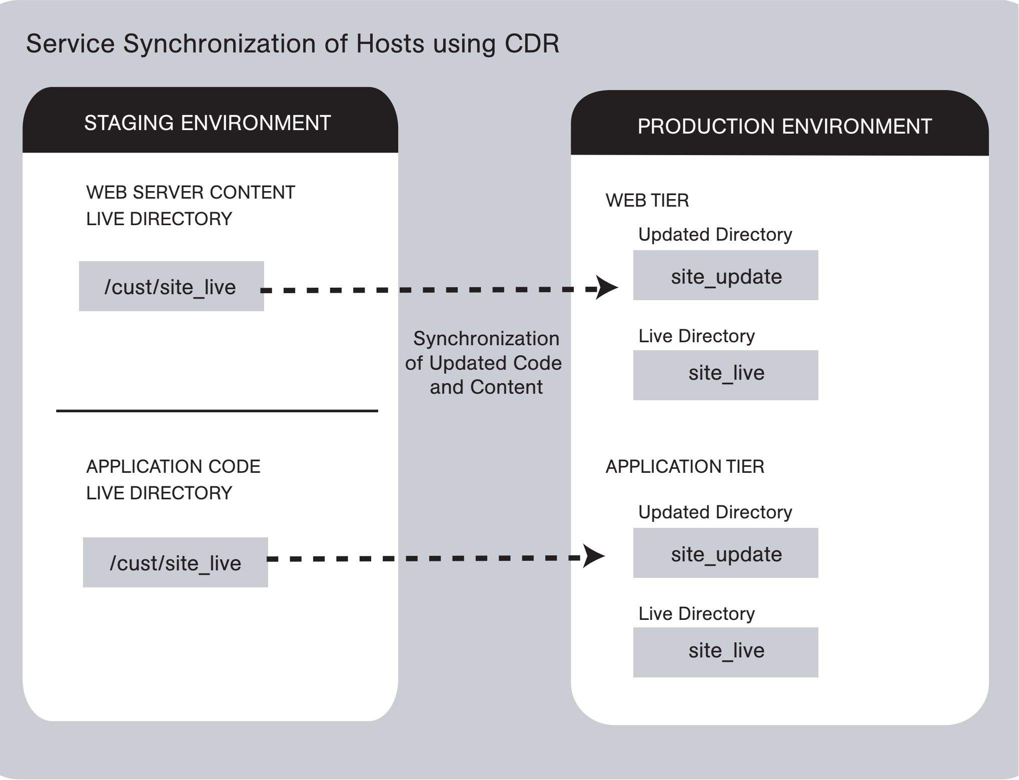

Contact your SA administrator for a description of services and synchronizations set up for your site and the specific operations that you need to run when updating code or content for a particular host service.

Performing Synchronizations

Perform the following steps to synchronize updated code and content from one source host to one or more destination hosts:

1 Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2 Click the Synchronize link.

A page appears that displays the synchronizations that you can perform.

3. Select the synchronization that you want to perform by clicking the link. The CDR Synchronize for [customer name] page appears, as Figure 11-5 shows.

Figure 11-5: The CDR Synchronize Page

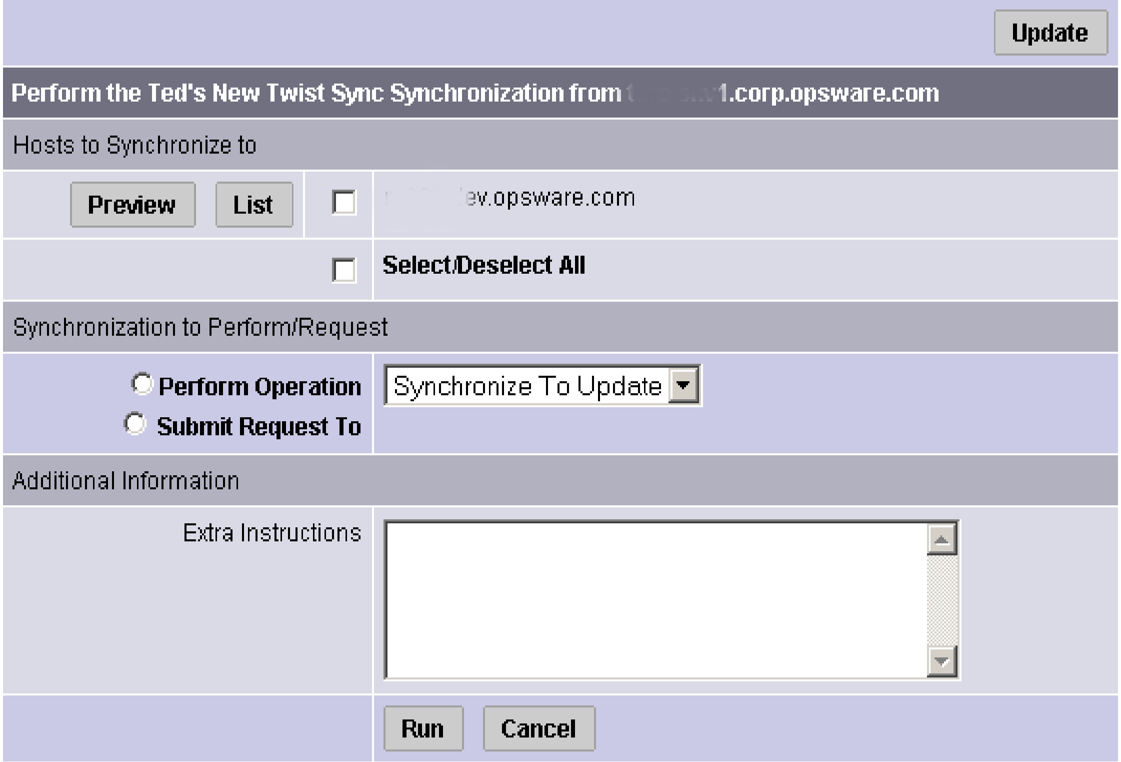

Update

Perform the Ted's New Twist Sync Synchronization from [redacted]1.corp.opsware.com

Hosts to Synchronize to

Preview List □ [redacted]ev.opsware.com

□ Select/Deselect All

Synchronization to Perform/Request

○ Perform Operation ○ Submit Request To | Synchronize To Update ▾

Additional Information

Extra Instructions

Run Cancel

4. Select one or more of the displayed host names on which you want to perform synchronization. The hosts that you select are the destination hosts.

Choose the Select/Deselect All option to select all or clear all host names.

5. (Optional) To view a list of files that will be created, modified, and deleted on the destination hosts, click **Preview**.

   Or

   (Optional) To view a list of all the files on the destination hosts, click **List**.

6. Select the Perform Operation option to directly perform a selected synchronization.

Or

Select the Submit Request To option to send a request to users specified to receive email notification for service and synchronization requests. When you submit a request, specify any additional instructions that you want to include for the requested synchronization. For example, these might be instructions such as the time that you want the synchronization performed, verification, or other related services to perform.

The Perform Operation option is only visible when you are a member of an SAS Web Client user group allowed to directly perform a synchronization.

7 Choose the type of synchronization that you want performed from the drop-down list:

- Synchronize To Update
- Synchronize To Live

See "Synchronization of Site Code and Content" on page 409 in this chapter for information about these options.

8 To initiate the synchronization or send the request, click **Run**.

CDR Cutover Operation

Perform the cutover operation to make the Update directory and the current live site identical.

When you cutover, CDR performs the following actions:

- Updates the Site Previous directory with files from the Live directory. CDR saves modified files and files-to-be-deleted to the Site Previous directory. The Site Previous directory contains the files necessary to restore the live site to the previous version.

- Determines the differences between the Update directory and the current Live directory. The files that are different are synchronized from the Update directory to the Live directory. See Figure 11-6.

Figure 11-6: Directory and File Updates for Cutover Operations Using CDR

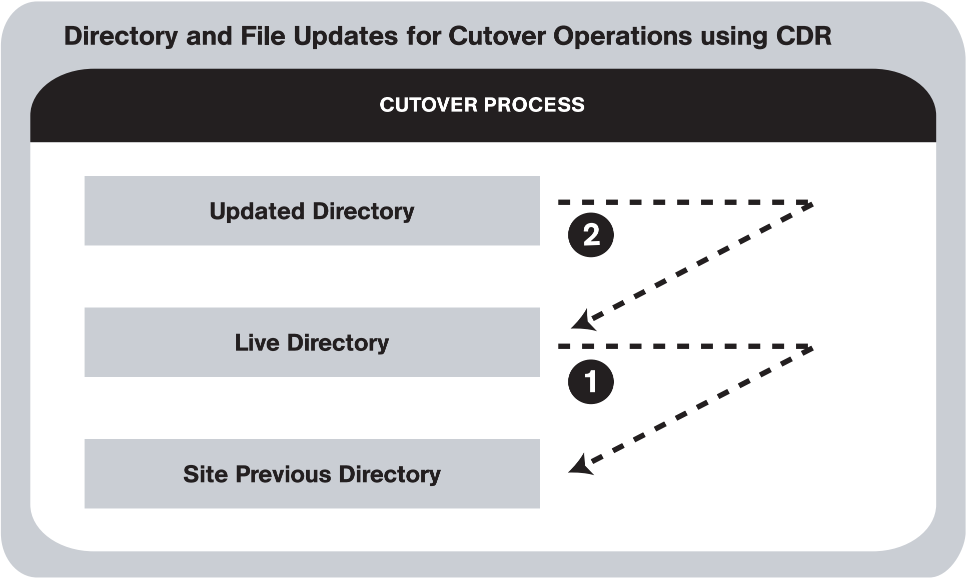

CDR determines file differences between source and destination directories based on file size, modification date and time, ownership, group attributes, and permissions attributes.

By using the cutover process, CDR ensures your ability to rollback to the previous version of your code and content if you experience a problem.

Your SA administrator can configure a CDR service to run scripts before and after cutting over to updated code and content. For example, before and after cutting over, you might distribute content on geographically disperse servers.

See "Synchronization of Site Code and Content" on page 409 in this chapter for information about a description of these directories. See "CDR Service Operations" on page 415 in this chapter for information about how to rollback code and content to the previous version.

CDR Service Operations

In addition to the Cutover operation, CDR provides a number of service operations:

- Starting and stopping host services
- Backing up code and content
- Restoring code and content from a previous version
- Rolling back code and content to the previous version

Performing these operations might be required depending on the type of code and content changes made or the host services that are affected.

The operations that you need to perform are specific to the service (for example, Web server or application server instance) for which you are updating code or content and the particular host.

Starting and Stopping Host Services

Stopping and starting services might be required depending on the type of code and content changes made or the host services that are affected. Discuss your requirements with your SA administrator.

Typically, you only stop and start host services for the hosts in your staging environment. Select members of your staff, or other individuals in your operations center, can stop and start host services for the hosts in your production environment.

Start: Launch a defined service; for example, starting a web or application server instance that is running on a specific host.

Stop: Shut down a defined service; for example, shutting down a web or application server instance before cutting over to new or changed code and content on a specific host.

Backing Up Code and Content

When you use CDR to back up your site, CDR saves the entire contents of the current Live directory for a specific service in the Backup directory. CDR saves the backup copy to the local disk for the host on which you ran the Backup operation. See Figure 11-7.

You can use CDR to keep only one backup copy at a time for a service.

Figure 11-7: Update Directory to Site Backup Directory for Backup Using CDR

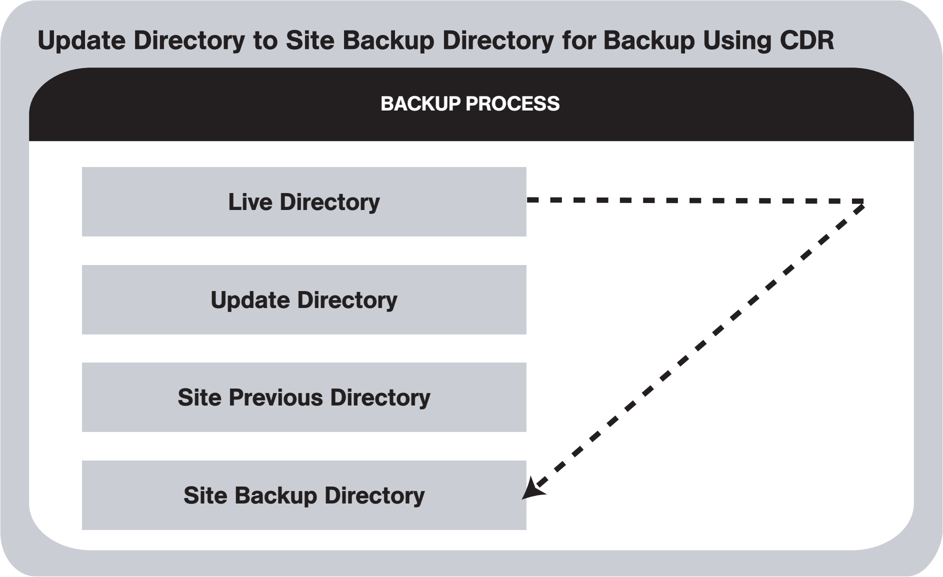

When you run the Restore operation, CDR replaces the Live directory contents with files stored in the Backup directory.

When you reach a high level of site stability, backing up your site is recommended, especially if you plan to make changes to site code and content.

Restoring Code and Content from a Previous Version

The Restore operation restores the previous Live directory by copying the contents of the Backup directory to the Live directory. Restoring code and content from the Backup directory does not change files that are stored in the Update directory. See Figure 11-8.

Figure 11-8: Site Backup Directory to Update Directory for Restoring Using CDR

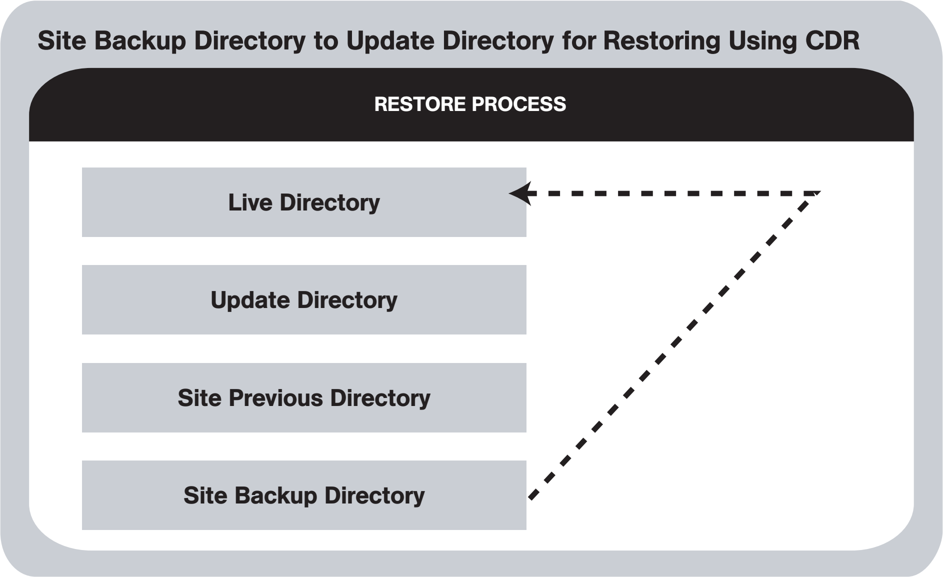

Before you restore code and content, you must have backed up the contents in the Live directory to the Backup directory by performing a Backup operation.

Rolling Back Code and Content to the Previous Version

If you experience a problem after cutting over updated code and content to your production site, you can rollback to the previous version.

Rolling back returns the site to the state it was in prior to the last cutover that you performed. See Figure 11-9.

Figure 11-9: Directory and File Updates for Rollback Operations Using CDR

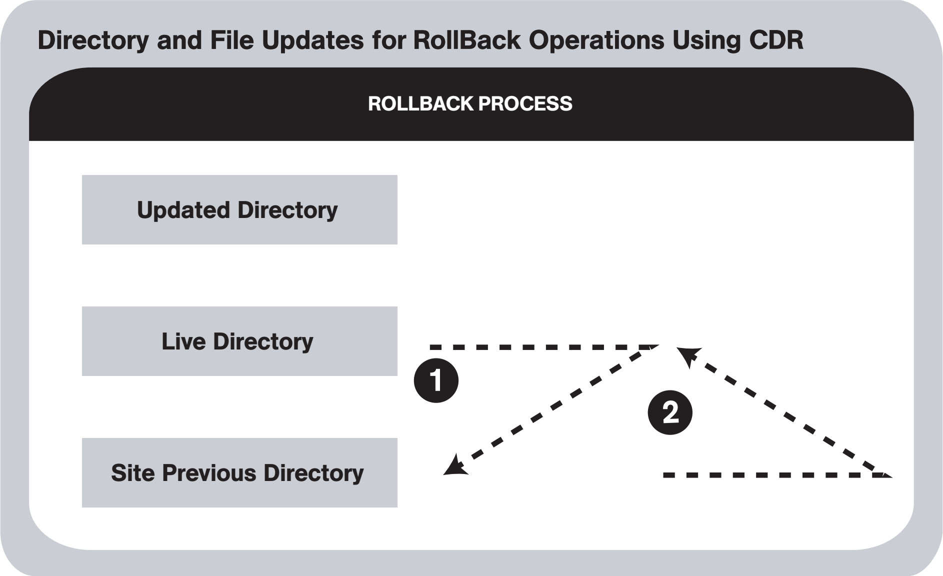

During cutover, CDR updates the Site Previous directory with files from the Live directory. CDR saves modified files and files-to-be-deleted to the Site Previous directory. The Site Previous directory contains the files necessary to restore the live site to the previous version.

During rollback, CDR restores the set of different files (modified files and files that were deleted during cutover) to the Live directory.

If you upload files directly to the Live directory or choose to synchronize directly to Live directories, the rollback operation does not function properly. Under these conditions, back up your Live directory and run the Restore operation to return your site to the previous version.

Accessing Service Operations in CDR

The Service Management option provides a number of service or administrative operations, including starting and stopping host services, backing up, restoring, or rolling back code and content, and cutting over to new site code and content.

To access CDR, your SA administrator must add you as a member of a user group authorized to use CDR.

You have the option of initiating a service either by selecting a service name or by selecting host names. You must select both a service to perform and the hosts on which to perform the service. See Figure 11-10.

Figure 11-10: Run Service Page

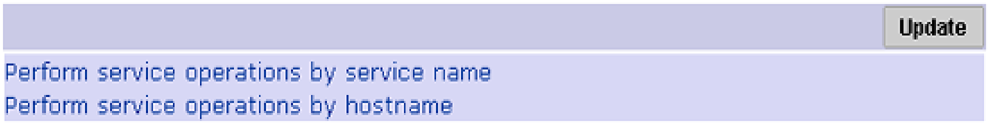

Perform service operations by service name
Perform service operations by hostname

- Initiate a service by selecting the “Perform service operations by service name” option first when you want to perform a service on multiple hosts at the same time. Selecting the service name first shows you all the hosts for which that service is defined.
- Initiate a service by selecting the “Perform service operations by host name” option first when you want to perform a service on a single host or on a specific host where you know the host name. Selecting the host name first shows you all the services that are defined for that host.

Performing Service Operations by Service Name

Perform the following steps to perform service operations by service name:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Select the Service Management option.
3. From the Service Management page, select the “Perform service operations by service name” link.

4 Select a service from a list of services defined for your site. A page that prompts you to select the hosts and the operation that you want to perform appears, as Figure 11-11 shows.

Figure 11-11: Perform Service Operations by Service Name Page

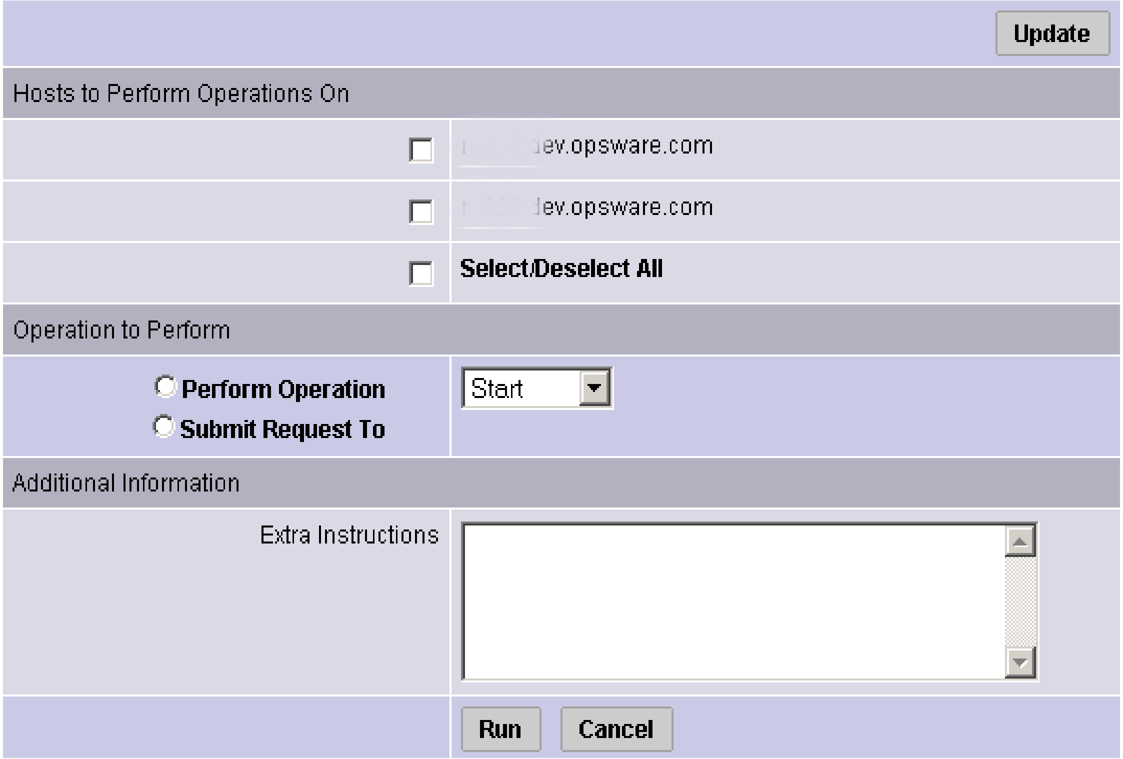

Update

Hosts to Perform Operations On

☐ [redacted]ev.opsware.com

☐ [redacted]ev.opsware.com

☐ **Select/Deselect All**

Operation to Perform

○ **Perform Operation**
○ **Submit Request To**

Start ▾

Additional Information

Extra Instructions

Run Cancel

5 Select one or more of the displayed host names. You can choose the Select/ Deselect All option to select all or clear all host names for the operation that you want to perform.

6 Select the Perform Operation option to directly perform a selected operation.

The Perform Operation option is visible only when you are a member of the CDR user group that allows users to directly perform a service operation.

Or

Select the Submit Request To option to send a request for authorized individuals to perform the operation for you. When you submit a request, specify any additional instructions that might be required to perform the requested operation.

7. Choose the type of operation that you want performed from the drop-down list:

   - Start
   - Stop
   - Cutover
   - Rollback
   - Backup
   - Restore

   See “CDR Service Operations” on page 415 in this chapter for information about these operations.

8. To initiate the operation or send the request, click **Run**.

Performing Service Operations by Host Name

Perform the following steps to perform service operations by host name:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Select the Service Management option.
3. From the Service Management page, select the “Perform service operations by hostname” link.
4. Select a host name from the list of staging and production hosts available for your site. A page that prompts you to select the service and the operation that you want to perform appears, as Figure 11-12 shows.

Figure 11-12: Perform Service Operations by Host Name Page

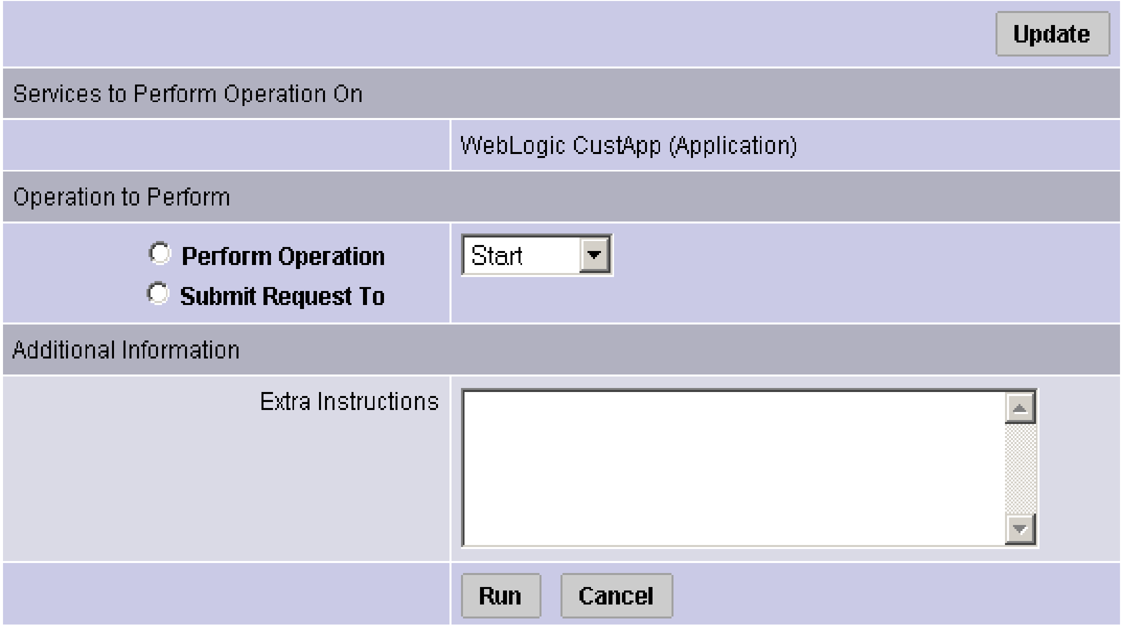

Services to Perform Operation On

WebLogic CustApp (Application)

Operation to Perform

○ Perform Operation
○ Submit Request To

Start

Additional Information

Extra Instructions

5 Select the service for which you want to perform an operation.

6 Select the Perform Operation option to directly perform a selected operation or select the Submit Request To option to send a request to have authorized individuals perform the operation for you. When you submit a request, specify any additional instructions that might be required to perform the requested operation.

The Perform Operation option is visible only when you are a member of the user group that allows users to directly perform a service operation.

7 Choose the type of operation that you want performed from the drop-down list:

- Start
- Stop
- Cutover
- Rollback
- Backup
- Restore

See “CDR Service Operations” on page 415 in this chapter for information about these operations.

8. To initiate the operation or send the request, click **Run**.

Performing Sequences

CDR also allows you to perform service operations and synchronizations that have been set up as a sequence of operations.

Perform the following steps to perform CDR sequences set up for your site:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2. Click the Run Sequence link.

   The CDR Run Sequence for [customer name] page appears that shows the sequences that you can run. See Figure 11-13.

Figure 11-13: CDR Run Sequence Page

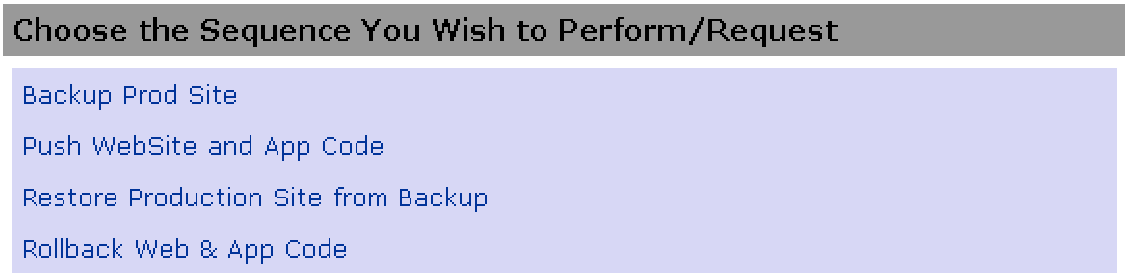

Choose the Sequence You Wish to Perform/Request

Backup Prod Site

Push WebSite and App Code

Restore Production Site from Backup

Rollback Web & App Code

3. Select the sequence that you want to perform by clicking the link. The Run Sequence page appears, as Figure 11-14 shows.

Figure 11-14: Run Sequence Page That Shows Details of a Specific Sequence

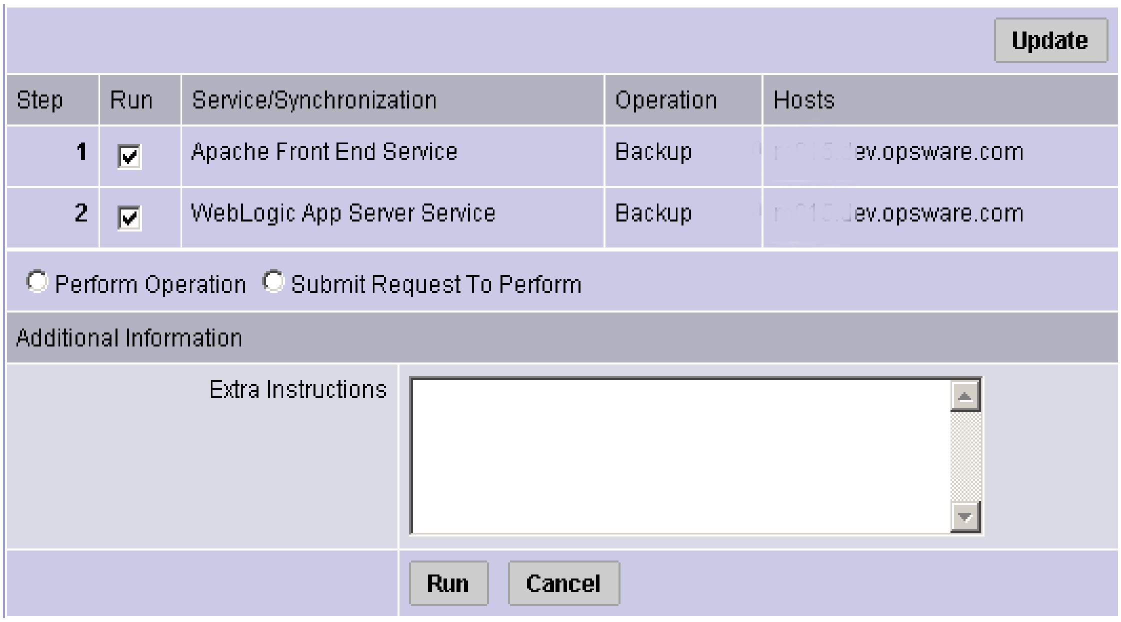

Step	Run	Service/Synchronization	Operation	Hosts
1	☑	Apache Front End Service	Backup	m015.dev.opsware.com
2	☑	WebLogic App Server Service	Backup	m015.dev.opsware.com

○ Perform Operation ○ Submit Request To Perform

Additional Information

Extra Instructions

4. Select the Perform Operation radio button to directly perform the selected sequence.

   Or

   Select the Submit Request To Perform radio button to send a request to the users specified to receive email notification for service, synchronization, and sequence requests. When you submit a request, specify any additional instructions that you want to include for the requested sequence. For example, you might want to include instructions such as the time that you want the sequence to run, verification, or other related services to perform.

> The Perform Operation option is visible only when you are a member of a user group allowed to directly perform a sequence.

5. To initiate the sequence or send the request, click **Run**.

Processing Code Deployment Requests from Users

When CDR users request that a service operation or synchronization be performed on their behalf, an email notification is sent to the individuals assigned to perform the requested service operation or synchronization.

The following message is a typical example of an email notification request.

```
From:    CDR-tool@opsware.com
To:      opscenter@opsware.com
Date:    Tue, 10 Jul 2001 11:25:13 -0700
Subject: Request To Perform Start operation for opsware.com

Please perform the following request

Requestor: jhancock/opsware.com
Request Time: Jul 10, 2001 11:25:13 AM PDT
Requested Service: Demo Apache
Requested action: Start
Perform on the following hosts:
host1.opsware.com
host2.opsware.com

Extra Instructions:
```

In this case, the email message specifies a request from a user, jhancock, to perform a Start operation on two hosts running the Demo Apache service. The subject line provides a summary of the request. In addition, the email indicates the time that the request was sent.

Perform the following steps to process the request:

1. Identify any special instructions that might need to be carried out for the specific request.

2. Log into the SAS Web Client, choose the CDR option, and then select the particular option within CDR to perform the requested operation.

3. When you successfully complete the CDR request, you might want to notify the individual user making the request, and all other involved parties, that the requested operation was completed.

If you encounter problems in completing a request and cannot resolve them, contact your SA administrator for any specific remedies and follow normal escalation procedures defined for your operational environment.

See the *SA Policy Setter's Guide* for more information about troubleshooting tips.

Status View of Previous Operations

CDR maintains a log of operations (service operations, synchronizations and sequences) that were executed. You can view this information to determine the status of particular deployment operations, and whether they have completed successfully. You can also use the My Jobs task area of the SAS Web Client Home page, or the My Jobs link to view this information.

Accessing the Log

Perform the following steps to access the log:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2. Select the View History option.

CDR displays a page providing a list, most recent to oldest, of operations that are either in progress or those that have been completed. This information displays only for the past 60 days. Each page is limited to a list of 10 operations. A Next link displays at the end of the page if there are more than 10 operations to view. Click the Next link to view subsequent operations. A Previous link is available to return to previous pages.

You need to refresh the page to view the status of any operations initiated after you selected the View History link.

A similar page displays after you select the View History link. See Figure 11-15.

Figure 11-15: CDR View History Page

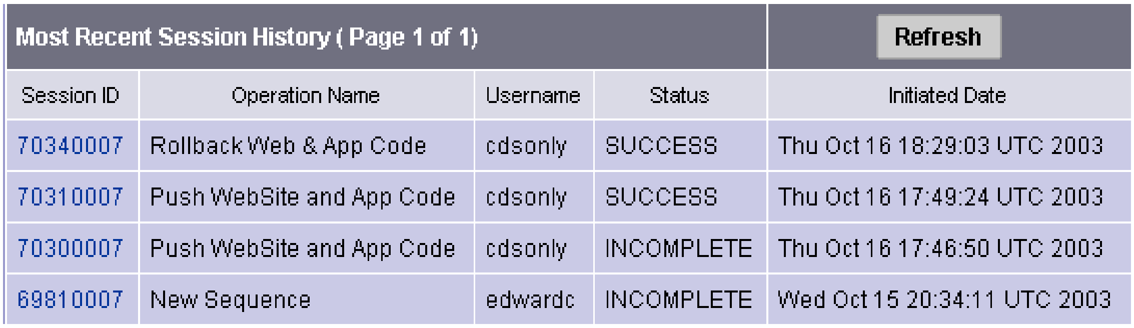

Most Recent Session History (Page 1 of 1)

Session ID	Operation Name	Username	Status	Initiated Date
70340007	Rollback Web & App Code	cdsonly	SUCCESS	Thu Oct 16 18:29:03 UTC 2003
70310007	Push WebSite and App Code	cdsonly	SUCCESS	Thu Oct 16 17:49:24 UTC 2003
70300007	Push WebSite and App Code	cdsonly	INCOMPLETE	Thu Oct 16 17:46:50 UTC 2003
69810007	New Sequence	edwardc	INCOMPLETE	Wed Oct 15 20:34:11 UTC 2003

The individual headings of column information included in the table are:

- **Session ID**: A session is created each time a CDR operation is performed. Click the Session ID to view detailed results of the operation.
- **Operation Name**: The name of the service, synchronization or sequence as specified when they were first defined.
- **User Name**: The user ID of the user that initiated the operation.
- **Status**: This describes the state of the operation at the end of the sequence. The status message varies depending on the type of operation. Single step operations always result in Success/Failure messages while multiple step operations (Sequence operations) could result in Complete with Error, Incomplete, or Success messages. Table 11-3 describes the possible status messages.

Table 11-3: CDR History Status Messages

STATUS TYPE	DESCRIPTION
Abort	This message displays only if a CDR specific script that is executed in order to complete service, synchronization or sequences fails. Such issues should be escalated to your SA Support Representative.
Active	This message displays if an operation is still in progress.
Complete with error	The sequence completed successfully but there were errors reported while the operation was in progress and the user opted to continue rather than cancel the operation.
Failure	The operation (service, sequence or synchronization) did not complete successfully.
Incomplete	The sequence resulted in an error and the user opted to cancel the sequence rather than continue.
Success	The operation (service, synchronization or sequence) was completed successfully and no errors were reported. In the case of a sequence, this message means that all the steps were completed successfully.
Initiated Date	The date on which the operation was initiated.

Chapter 12: Configuration Tracking

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Configuration Tracking
- Configuration Tracking Operations
- Customized Configuration Tracking Policies
- Manual Backups
- Backup History
- Restoration of Backed Up Files

Overview of Configuration Tracking

The Configuration Tracking feature of HP Server Automation allows you to monitor critical configuration files and configuration databases. When SA detects a change in a tracked configuration file or configuration database, the system can perform a number of actions, including backing up the configuration file or sending an email to a designated individual or group. You use configuration tracking policies to identify the files to be tracked and actions to be taken when change is detected. A configuration tracking policy consists of one or more configuration tracking policy entries that specifies the configuration file, the directory of configuration files, or the configuration database that you want to track.

The Configuration Tracking feature is designed for flexibility. For example, you can set configuration tracking defaults for a particular software application, and all attached servers automatically get those defaults. You can also quickly deploy the common configuration tracking defaults to a large number of servers in your SA managed environment or create a specific policy for a single server.

Configuration Tracking allows you to recover from many problems caused by changes to configuration files. Using Automated Configuration Tracking, you can identify which tracked configuration files have changed, thus helping you identify the potential source of a problem. If you back up your configuration files with the Configuration Tracking feature, you can quickly restore the changed configuration files to a previous version.

You can also view a detailed history of all backup activity. This history includes a list of all tracked files that have been backed up and what types of backups occurred. If the backed up configuration files are text-based, you can download the files from the backup history and compare them to determine what specific changes have been made.

The Configuration Tracking feature is not a general-purpose backup solution. Configuration Tracking is designed to monitor text-based configuration files and specific types of configuration databases. The number and size of files that can be monitored on any managed server is limited.

File Types Supported

You can use the Configuration Tracking feature with the following types of files:

- Text-based configuration files
- The COM + Registration Database (Windows 2000)
- The IIS Metabase
- Windows Registry keys

Configuration Tracking Operations

This section discusses the following topics:

- Change Detection
- Types of Actions Performed
- Types of Backups Performed
- Email Automated Configuration Tracking and Logging Actions
- Creating the Email Notification List

Change Detection

All servers that HP Server Automation manages have an Server Agent installed on them. On servers that use Configuration Tracking, every four hours the Server Agent inspects the configuration files and databases that you select to track.

The Server Agent computes an MD5 checksum to determine if the contents of a tracked file have changed. (Any change to the contents of the file results in a change to the MD5 checksum.) If the contents of the file have changed, the action that you specify in the tracking policy is performed. For example, if you create an entry in your tracking policy for the `/etc/passwd` file and select backup as the action to be taken, the file is backed up when the Server Agent discovers a change in the `/etc/passwd` file.

The creation or deletion of a tracked file (or files inside a tracked directory) also counts as change and triggers a policy's action. (There are some exceptions; see the *SA Policy Setter's Guide* for more information.)

Changes to the properties of a tracked file or directory (such as changes to permissions or time stamps) do not count as a change. When a file or directory is backed up, however, its properties are backed up as well.

The first time that a tracking policy is deployed, all targets are considered changed. The Server Agent is encountering the files for the first time, and all of the policy's actions are triggered.

Types of Actions Performed

HP Server Automation can perform the following actions when change is detected in a tracked configuration file or configuration database:

- Back up
- Send email to addresses specified in the policy entry
- Send email to a designated notification group specified by a custom attribute
- Create an entry in the server's standard system log, (the syslog on Unix servers and the event log on Windows servers)

Types of Backups Performed

If you selected backup as the action for a tracked configuration file, the two general types of backups that can occur are incremental backups and full backups.

Incremental Backups

During an incremental backup, only targets that have changed since the last backup (and that have been selected to be backed up) are backed up.

An incremental backup occurs automatically when the Server Agent detects change in a tracked file that is selected to be backed up. (The Server Agent checks for change every four hours.)

Incremental backups also occur before and after you restore a previous version of a backed up configuration file to a server. These backups allow you to rollback the restored files.

Full Backups

During a full backup, all tracked configuration files that were selected to be backed up are backed up, not just the files that have changed.

Once a week, the Server Agent on a server checks to see if any files have changed since the last full backup. If any files have changed, HP Server Automation performs a new full backup. If no files have changed, the full backup does not take place.

You can also force HP Server Automation to perform a full backup on a server by selecting the Perform Manual Backup option. (See “Manual Backups” on page 443 in this chapter for more information.)

See “Backup History” on page 443 in this chapter for information about Backup types.

Backups are stored in the Software Repository until you delete them. You should delete old backups periodically, especially if you are backing up a large number of files that change frequently. See “Deleting Backups” on page 450 in this chapter for information about the procedure for deleting backups.

Email Automated Configuration Tracking and Logging Actions

You can choose to have email sent when a monitored target changes. The following example shows the text of an email generated when a tracked file changed.

```
From: <configurationtracking@yourcompany.com>
Date: Thu Jan 16, 2003  5:40:11 PM US/Pacific
To: <joe@yourcompany.com>
Subject: athena.cust.com: Configuration Tracking CHANGE
notification
```

```
Configuration Tracking has detected a CHANGE event
Host: athena.cust.com
Object: /db/file1l1
```

The email specifies the name of the server and the name of the object that changed. The object can be a file, a directory, or a configuration database.

If you are monitoring a directory target, you receive email about the directory itself and about changes to the files in the directory (except when a file is deleted.) For example, if three new files are created in a directory, you would receive four emails, one for the directory and three for the new files.

If you selected the logging action, an entry is made to the server's standard system log when a change is detected. You select the type of log entry that you want to have written. HP Server Automation uses the following three standard entry types:

- Info
- Warning
- Error

How the entry types are identified is system-dependent. For example, on most systems, Warning entries are identified by the word warning. In some systems, however, a number is used to identify the log entry type.

The following example shows a warning log entry written on a Solaris Server:

```
Jan  8 00:05:25 athena.cust.com Configuration Tracking:
[ID702911
local0.warning] Configuration Tracking: /other/otherfile1 :
Event CHANGE occurred
Jan  8 00:05:25 athena.cust.com Configuration Tracking: [ID
702911
local0.warning] Configuration Tracking: /other/otherfile1 :
Event CHANGE occurred
```

Creating the Email Notification List

Sending email to a server’s backup notification list is one of the actions that you can select in a tracking policy entry. The email notification list is a list of email addresses that you define for the following custom attribute:

```
backup_notification_email
```

This attribute can be set on the server itself or on the customer to which the server is attached. Setting the attribute at the customer level allows you to use the same email notification list for all servers that belong to the same customer (assuming that these servers have all been attached to the same customer and do not have the `backup_notification_list` set on the servers themselves). See the *SA Policy Setter's Guide* for more information about setting custom attributes for customers.

Search Order for Email Notification List Attribute

On a server that has a policy that includes the Email Notification List for Server action selected, the server searches for the backup_notification_email attribute in the following order:

- Server
- Customer

After the custom attribute is found, its value is used (for example, the email address of the notification list) and the search is concluded. If, for example, the backup_notification_email attribute is set on a server, the server's email notification list is used, even if the server is assigned to a customer that has a different backup_notification_email attribute.

Format of Email Notification List

The notification list can contain multiple email addresses. The email address must be formatted as a comma-separated list.

Customized Configuration Tracking Policies

This section provides information on customizing configuration tracking policies within HP Server Automation and contains the following topics:

- Server-Based Entries
- Policies for Customizing Multiple Servers
- Adding or Editing Customized Tracking Policy Entries
- Disabling Customized Tracking Policy Entries
- Enabling Customized Tracking Policy Entries
- Viewing a Server's Tracking Policy
- Reconciling Customized Tracking Policies

- Enabling and Disabling Configuration Tracking on Servers

You can customize the tracking policy for a server or selected group of servers. Ordinarily, a server gets its tracking policy from the policy to which it is attached. In some cases, however, you might need to customize the tracking policy for a particular server or set of servers. If, for example, an application on one server is using an optional configuration file, you can customize the tracking policy for that server so that the optional configuration file is monitored.

Server-Based Entries

A tracking policy entry created for a specific server or selected group of servers is called a server-based tracking policy entry.

Table 12-1 shows the actions that you can perform on server-based tracking policy entries.

Table 12-1: Entry Types

ENTRY TYPE	EDIT	DELETE	ENABLE/DISABLE
Server-based	Yes	Yes	No

To determine if a tracking policy entry server-based, check the Source column in the Track Configurations: Customize Tracking Policies page. Figure 12-1 shows the policy entries for a single selected server.

Figure 12-1: Viewing Configuration Tracking Policy

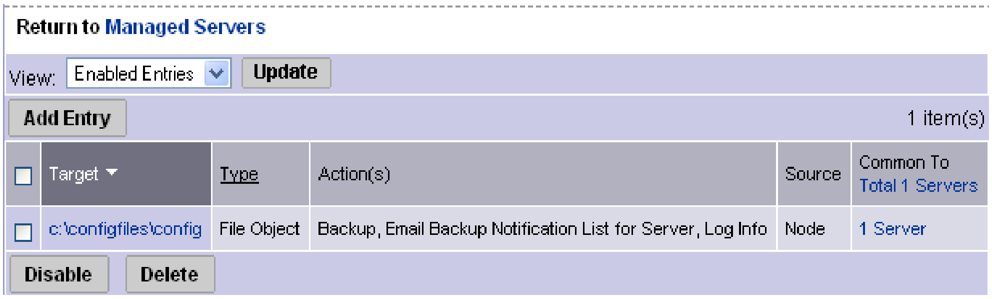

□	Target ▾	Type	Action(s)	Source	Common To Total 1 Servers
□	c:\configfiles\config	File Object	Backup, Email Backup Notification List for Server, Log Info	Node	1 Server

Policies for Customizing Multiple Servers

When you customize a tracking policy, you can select one or multiple servers. Here are some considerations when you customize the policies of multiple servers.

- When you select a group of servers that contain both Unix-based (for example, Solaris, HP-UX, Linux, and so forth) and Windows-based servers, you cannot add tracking policy entries. (You can always add entries when you select a single server.)
- When you select a group of servers made up entirely of either Unix-based servers or entirely of Windows-based servers, you have the option of making new tracking policy entries. The policy entries are added to all servers that you select.
- When you select a group of servers made up entirely of either Unix-based servers or of Windows-based servers, you have the option of editing a tracking policy entry. Any policy entry that can be edited can also be added to all selected servers (if it is not already common to all of them).

It can take 20 seconds or longer to customize polices when you select six or more servers.

Adding or Editing Customized Tracking Policy Entries

Use the following procedure to add or edit tracking policy entries for a server or set of servers running the same general type of operating system. If you select a set of servers running both Windows and Unix operating systems, you cannot add tracking policy entries.

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the server whose policy you want to edit. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. Select the server or set of servers that you want to add server-based tracking policy entries to, as Figure 12-2 shows.

Figure 12-2: Selecting Servers

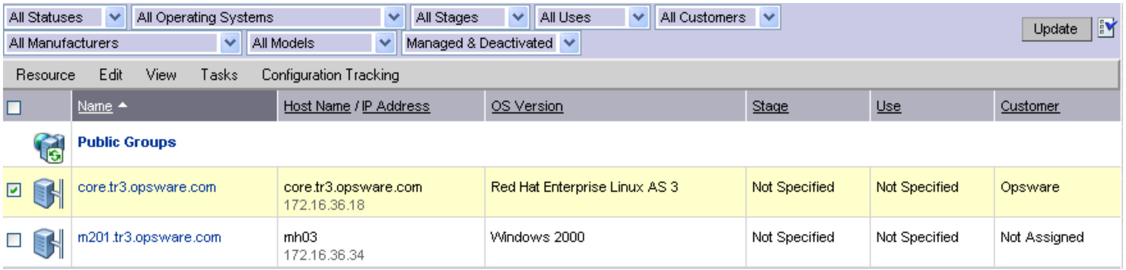

□	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Customer
	Public Groups					
☑	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	Red Hat Enterprise Linux AS 3	Not Specified	Not Specified	Opsware
□	m201.tr3.opsware.com	mh03 172.16.36.34	Windows 2000	Not Specified	Not Specified	Not Assigned

4 From the Configuration Tracking drop-down menu, select Edit Tracking Policies. The Configuration Tracking: Edit Tracking Policies page appears, as Figure 12-3 shows.

Figure 12-3: Configuration Tracking: Edit Tracking Policies Page

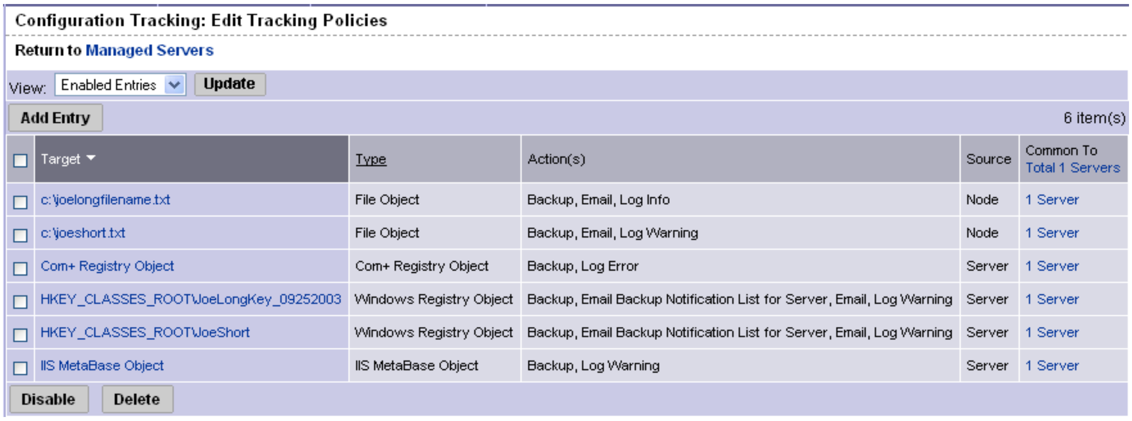

Configuration Tracking: Edit Tracking Policies
Return to Managed Servers
View: Enabled Entries
Update
Add Entry
6 item(s)
Disable
Delete

□	Target ▾	Type	Action(s)	Source	Common To Total 1 Servers
□	c:\joelongfilename.txt	File Object	Backup, Email, Log Info	Node	1 Server
□	c:\joeshort.txt	File Object	Backup, Email, Log Warning	Node	1 Server
□	Com+ Registry Object	Com+ Registry Object	Backup, Log Error	Server	1 Server
□	HKEY_CLASSES_ROOT\JoeLongKey_09252003	Windows Registry Object	Backup, Email Backup Notification List for Server, Email, Log Warning	Server	1 Server
□	HKEY_CLASSES_ROOT\JoeShort	Windows Registry Object	Backup, Email Backup Notification List for Server, Email, Log Warning	Server	1 Server
□	IIS MetaBase Object	IIS MetaBase Object	Backup, Log Warning	Server	1 Server

5 Click **Add Entry** to create a new policy entry. To edit an existing entry, click the link for the entry in the Target field.

6 Define the target, select the type, and select the actions to be performed on the target.

Table 12-2 describes each of the selections that you must make.

Table 12-2: Editing Configuration Tracking Policies

FIELD	DESCRIPTION
Type	**File**: Monitor the file specified in the target field. **Directory**: Monitor all the files in the directory specified in the target field. The following types are available only for Windows servers: **Windows Registry**: Specify key in target field. **IIS Metabase**: Entire Metabase is monitored; do not specify target. **COM + Registration Database**: Entire Registry is monitored; do not specify target.

Table 12-2: Editing Configuration Tracking Policies (continued)

FIELD	DESCRIPTION
Target	If you selected the file type, specify the full path (including the drive letter on Windows servers) of the file that you want to monitor. If you selected the directory type, specify the full path of the directory (including the drive letter on Windows machines) that you want to monitor. You also have the option of monitoring subdirectories. (Select the include subdirectories check box.) If you select the file or directory type, you can use wildcards in the target. (See the *SA Policy Setter's Guide* for more information about Configuration Tracking Policy Targets and Wildcards.) If you selected the Windows Registry type, specify the Windows registry key. This key and all its subkeys are backed up. Use standard syntax for Windows Registry keys (For example, `HKEY_LOCAL_MACHINE\SOFTWARE)` If you selected the IIS Metabase or COM + Registration Database type, you do not specify the target.

Table 12-2: Editing Configuration Tracking Policies (continued)

FIELD	DESCRIPTION
Actions (you can select multiple options)	**Backup**: Back up the specified file, directory, COM + Registration Database, Windows Registry keys, or IIS Metabase. **Email Backup Notification List for Server**: Send an email to the backup notification list for the selected server. (Not available for Windows Registry.) See "Creating the Email Notification List" on page 433 in this chapter for more information. **Email**: Send an email to the address or addresses specified in this field when a change is detected. Use a comma-separated list for multiple email addresses. (Not available for Windows Registry.) **Log**: Add an entry to the server's system log when a change is detected. (Not available for Windows Registry.) You can choose to write the following types of log entries to the server's system log: Info Warning Error

7. Click **Save** to add the entry to the tracking policy.
8. If you want continue to add entries to the tracking policy, click **Add Entry** and repeat this procedure.

Changes do not take effect until you perform a Configuration Tracking policy reconcile.

Disabling Customized Tracking Policy Entries

Perform the following steps to disable customized tracking policy entries:

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Edit Tracking Policies.

The Configuration Tracking: Edit Tracking Policies page appears.

4. Select the tracking policy entry or entries that you want to disable.
5. Click **Disable** to disable the tracking policy entries that you selected.

Changes do not take effect until you perform a Configuration Tracking policy reconcile.

Enabling Customized Tracking Policy Entries

Perform the following steps to re-enable tracking policies:

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Edit Tracking Policies.

   The Configuration Tracking: Edit Tracking Policies page appears.
4. From the **View** menu, select **Disabled Entries** and then click **Update**.
5. Select the tracking policy entry or entries that you want to disable.
6. Click **Enable** to disable the tracking policy entries that you select.

Changes do not take effect until you perform a Configuration Tracking policy reconcile.

Viewing a Server's Tracking Policy

Perform the following steps to view a server's tracking policy:

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Edit Tracking Policies. The tracking policy entries display.
4. Click the Tracking Policy link to display the server's tracking policy.

Reconciling Customized Tracking Policies

When you create or edit any customized tracking policy entries, the entries are not deployed to your servers until you perform a configuration tracking reconcile.

A configuration tracking reconcile is not the same as a standard HP Server Automation remediate. Performing a standard remediate does not deploy your tracking policies to your servers.

A configuration tracking reconcile deploys all configuration polices, not just the customized policies.

Perform the following steps to reconcile the tracking policy to the servers whose policies you customized:

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Reconcile Tracking Policies.

   The Track Configurations: Preview Reconcile page appears and displays the progress of the test reconcile.
4. After the test reconcile completes successfully, click **Reconcile** to perform the actual reconcile operation.

   The Track Configurations: Reconcile page appears.
5. If you want to see more information about the changes made during the reconcile operation, click **View Details**. Otherwise, click **Done**.

Enabling and Disabling Configuration Tracking on Servers

You can enable or disable Configuration Tracking on any server or set of servers in your SA-managed environment.

Disabling Configuration Tracking is not the same as disabling an individual tracking policy entry (See “Disabling Customized Tracking Policy Entries” on page 439 in this chapter for more information”). Disabling Configuration Tracking stops all configuration tracking activity on the selected server.

Disabling configuration tracking, however, does not change a server's tracking policy in any way. If you later re-enabled configuration tracking on the server, the server still has the same tracking policy that it did before you disabled it.

By default, Configuration Tracking is disabled on all managed servers. It is not necessary, however, to manually enable configuration tracking in order to turn on the Configuration Tracking feature. Configuration Tracking is automatically enabled on a server when you reconcile its configuration tracking policies, and you must perform a reconcile in order to deploy tracking policies.

If you want configuration tracking to remain disabled on a server, be careful not to perform a configuration tracking reconcile on the server. (You can, however, still perform a regular reconcile.)

Perform the following steps to enable or disable configuration tracking:

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. Under the Configuration Tracking drop-down menu, select Enable/Disable. A list that displays the state (Enabled/Disabled) of the servers that you selected, as Figure 12-4 shows.

Figure 12-4: Enable/Disable Configuration Tracking Page

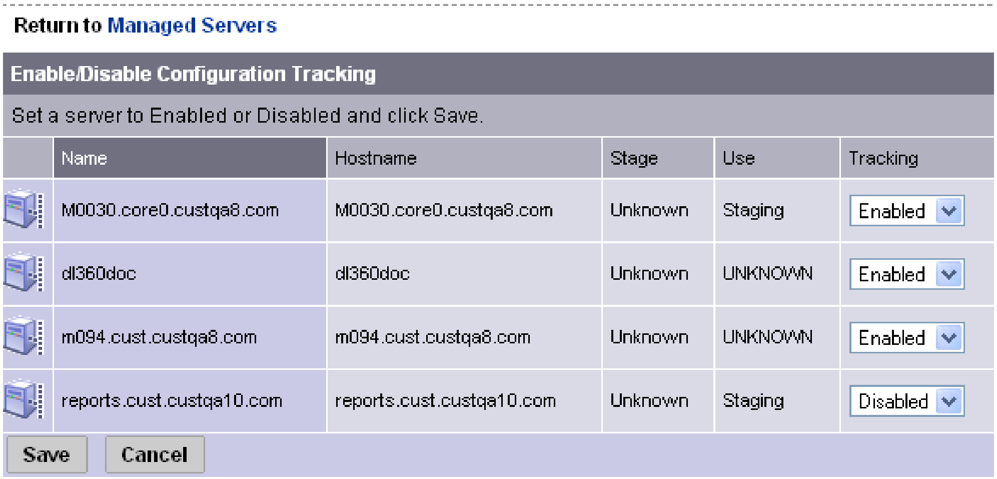

Return to Managed Servers

Enable/Disable Configuration Tracking

Set a server to Enabled or Disabled and click Save.

	Name	Hostname	Stage	Use	Tracking
	M0030.core0.custqa8.com	M0030.core0.custqa8.com	Unknown	Staging	Enabled
	dl360doc	dl360doc	Unknown	UNKNOWN	Enabled
	m094.cust.custqa8.com	m094.cust.custqa8.com	Unknown	UNKNOWN	Enabled
	reports.cust.custqa10.com	reports.cust.custqa10.com	Unknown	Staging	Disabled

4. Under the Tracking field, select Enabled or Disabled to enable or disable configuration tracking on the selected server.
5. Click **Save** to commit the changes.

Manual Backups

On a server or set of servers, you can perform a manual backup of all tracked configuration files and databases for which you have selected the backup action. Manual backups can be useful as a precaution before making changes to configuration files. If a problem arises, you can then immediately restore a backed-up configuration file or database to its previous state.

Manual backups are full backups. All tracked configuration files and databases for which the backup action has been selected are backed up, not just the files that have changed.

Performing Manual Backups

Perform the following steps to perform a manual backup:

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Perform Backup Now.
4. If you do not want to use the default backup name (Manual Backup), type a new name in the backup name field. If you provide a backup name, you can later perform a search for backup names to find this backup point. Backup names do not have to be unique.
5. Click **Start Backup**. The backup progress displays.
6. When the backup is completed, you can click **View Details** to review the list of configuration files or configuration databases that have been backed up.

Backup History

HP Server Automation provides a detailed history of backup activity as well as search capabilities to find backup points by backup name and to find backed up files by file name. This section contains the following topics:

- Viewing the Backup History
- Viewing the List of Backup Events
- Types of Backup Events
- Backup History Search Options
- Backup Info and Backup Manifest
- File Information and File Versions
- Deleting Backups

Viewing the Backup History

Perform the following steps to view the backup history:

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select View Backup History.

   The Track Configurations: View Group Backup History page displays. This page displays the backup activity for the servers that you select. Figure 12-5 shows a sample backup history.

Figure 12-5: Backup History

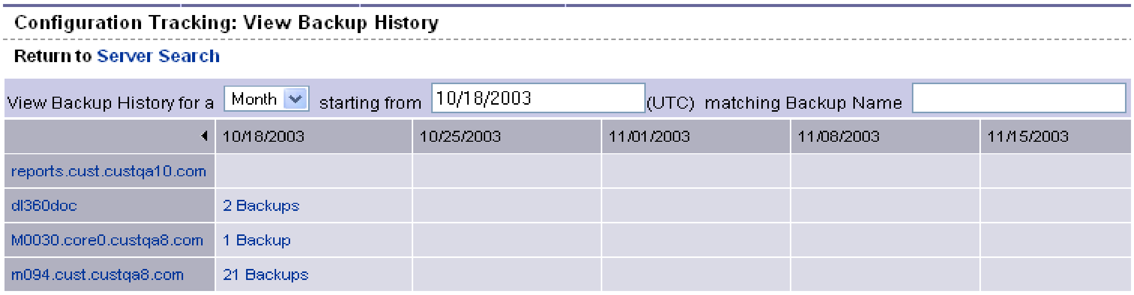

Configuration Tracking: View Backup History

Return to Server Search

View Backup History for a Month ▾ starting from 10/18/2003 (UTC) matching Backup Name

◂	10/18/2003	10/25/2003	11/01/2003	11/08/2003	11/15/2003
reports.cust.custqa10.com					
dl360doc	2 Backups				
M0030.core0.custqa8.com	1 Backup				
m094.cust.custqa8.com	21 Backups				

The number of backups that occurred on a particular date displays as a link in the server's date column.

The number of backups refers to the number of backup events (or backup points) when a particular type of backup took place. It does not refer to the number of files backed up. For example, if a server displays 3 backups in a date column, it could refer to the following three backup events:

- A scheduled incremental backup that backed up four files
- A second scheduled incremental backup that backed up 10 files
- A manual backup that backed up 30 files

Viewing the List of Backup Events

To view the list of individual backup events, click the link that displays the number of backups in the desired date column. Figure 12-6 shows a page with a list of backup events.

Figure 12-6: Backup Events

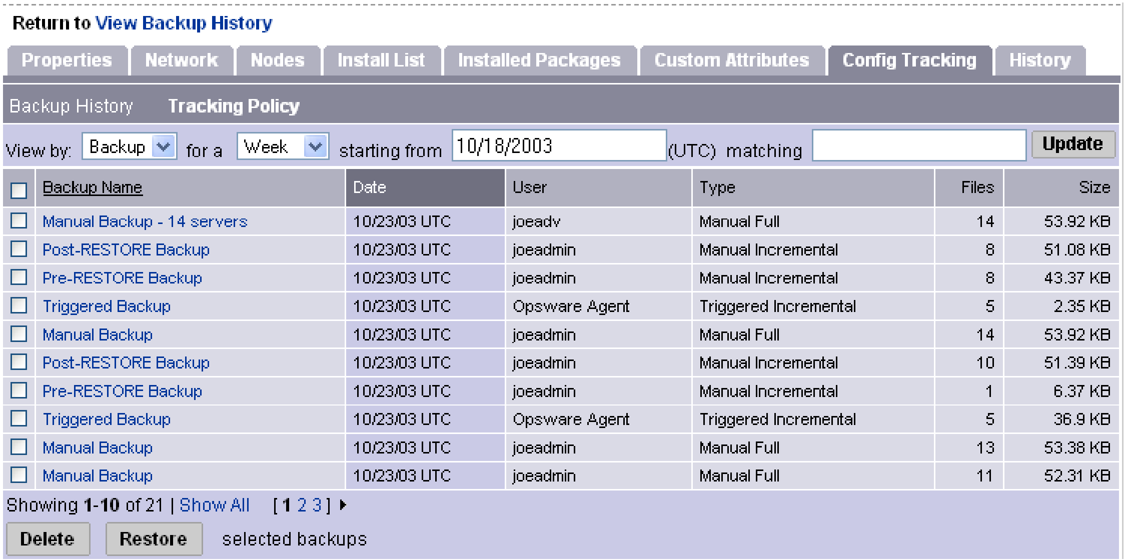

View by: Backup for a Week starting from 10/18/2003 (UTC) matching

□ Backup Name	Date	User	Type	Files	Size
□ Manual Backup - 14 servers	10/23/03 UTC	joeadv	Manual Full	14	53.92 KB
□ Post-RESTORE Backup	10/23/03 UTC	joeadmin	Manual Incremental	8	51.08 KB
□ Pre-RESTORE Backup	10/23/03 UTC	joeadmin	Manual Incremental	8	43.37 KB
□ Triggered Backup	10/23/03 UTC	Opsware Agent	Triggered Incremental	5	2.35 KB
□ Manual Backup	10/23/03 UTC	joeadmin	Manual Full	14	53.92 KB
□ Post-RESTORE Backup	10/23/03 UTC	joeadmin	Manual Incremental	10	51.39 KB
□ Pre-RESTORE Backup	10/23/03 UTC	joeadmin	Manual Incremental	1	6.37 KB
□ Triggered Backup	10/23/03 UTC	Opsware Agent	Triggered Incremental	5	36.9 KB
□ Manual Backup	10/23/03 UTC	joeadmin	Manual Full	13	53.38 KB
□ Manual Backup	10/23/03 UTC	joeadmin	Manual Full	11	52.31 KB

Showing **1-10** of 21 | Show All [**1** 2 3] ▸

Delete Restore selected backups

Types of Backup Events

In addition to information such as the date and time the backup occurred, the list of backup events indicates the type of backup. Table 12-3 describes the type of backup events that can occur.

Table 12-3: Types of Backup Events

TYPE	DESCRIPTION
Triggered Full	The automatic weekly backup of all tracked files for which the backup option was selected. This takes places only if any relevant files have changed since the last full backup.

Table 12-3: Types of Backup Events (continued)

TYPE	DESCRIPTION
Manual Incremental	A backup of all changed files (for which the backup option was selected) that occurs before and after a restoration or a rollback.
Triggered Incremental	An automatic backup that occurs when the Server Agent detects a change to a tracked file for which the backup option was selected. Only changed files are backed up.
Manual Full	A full backup initiated by the user of all tracked files for which the backup option was selected.

Backup History Search Options

By default, you see the backup history for the past week for the servers that you selected.

If you want to display a different date range, you have the following options:

- Display the backup history for different date range by selecting a different option in the "View Backup History for a" box.
- Use the "starting from" field to search for backup history from a past date until the current date.

You can use the matching field to search for a backup name within the selected date range. Wildcards are allowed (the * and ? characters). If you do not type the full name of the backup, you must use a wildcard for any missing parts of the name.

The results of this search show the date and number of backup names on that date that match the * pattern. When you click the link for the number of backups, a list of the matching backup names displays.

Backup Info and Backup Manifest

When you click a backup name (such as Scheduled Backup) anywhere in the backup history, the Backup Info and Backup Manifest tabs display.

Backup Info Tab

The Backup Info tab displays general information about a backup event. This information includes the name of the backup, the date and time of the backup, and the policies that triggered the backup event.

All customized policies are identified as Server Policy, as Figure 12-7 shows.

Figure 12-7: Track Configurations: View Backup Triggered Backup Page

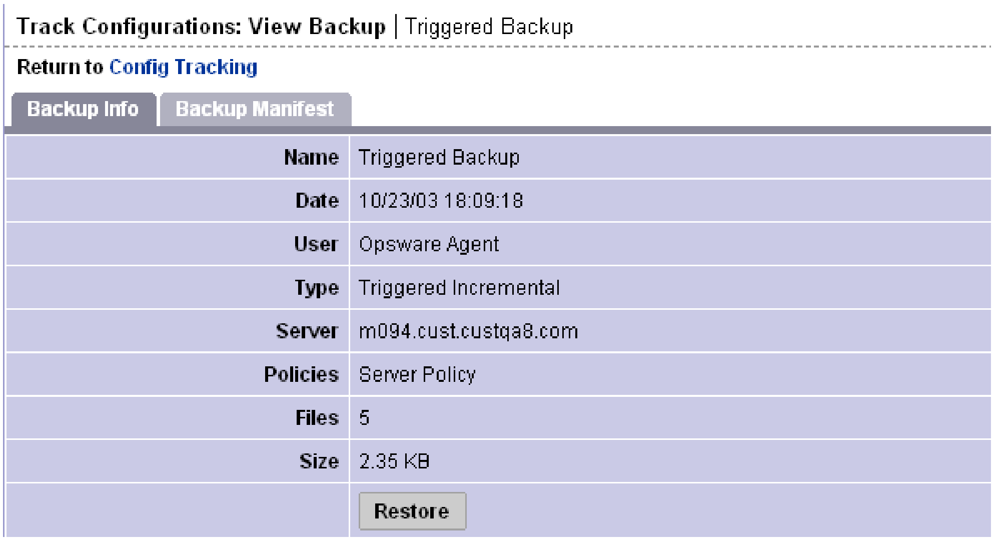

Track Configurations: View Backup | Triggered Backup

Return to Config Tracking

Name	Triggered Backup
Date	10/23/03 18:09:18
User	Opsware Agent
Type	Triggered Incremental
Server	m094.cust.custqa8.com
Policies	Server Policy
Files	5
Size	2.35 KB
	Restore

From the Backup Info tab, you can also place the files that were backed up into the Restore Queue. See “Restoring Backups” on page 454 in this chapter for information about how to use this feature.

Backup Manifest Tab

The Backup Manifest tab displays the list of files that were backed up during the selected backup event, as Figure 12-8 shows.

Figure 12-8: Backup Manifest Tab

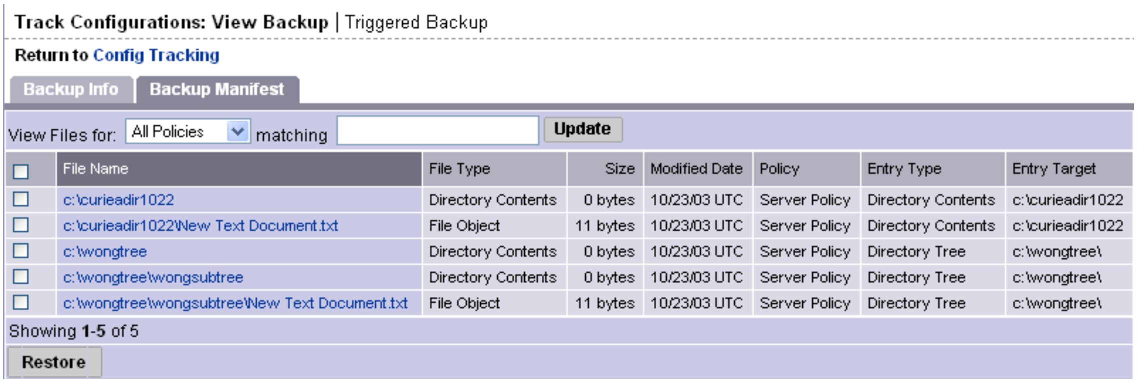

Track Configurations: View Backup | Triggered Backup

Return to Config Tracking

View Files for: All Policies matching

□	File Name	File Type	Size	Modified Date	Policy	Entry Type	Entry Target
□	c:\curieadir1022	Directory Contents	0 bytes	10/23/03 UTC	Server Policy	Directory Contents	c:\curieadir1022
□	c:\curieadir1022\New Text Document.txt	File Object	11 bytes	10/23/03 UTC	Server Policy	Directory Contents	c:\curieadir1022
□	c:\wongtree	Directory Contents	0 bytes	10/23/03 UTC	Server Policy	Directory Tree	c:\wongtree\
□	c:\wongtree\wongsubtree	Directory Contents	0 bytes	10/23/03 UTC	Server Policy	Directory Tree	c:\wongtree\
□	c:\wongtree\wongsubtree\New Text Document.txt	File Object	11 bytes	10/23/03 UTC	Server Policy	Directory Tree	c:\wongtree\

Showing **1-5** of 5

Two types of objects are backed up, as seen in the File Type field.

- File Object in the File Type field indicates that a file has been backed up. You can use the entry to restore the file.
- Directory Contents in the File Type field indicates that a directory object has been backed up. The directory object is the directory itself, and not the contents of the directory. You can use this entry to restore the directory, but you must select the files inside the directory to restore the contents of the directory.

The Entry Type field can have three possible values. The Entry Type identifies the target type in the policy entry that caused the file or directory object to be backed up.

- File Object in the Entry Type field indicates that the file was backed up as the result of a file target type.
- Directory Contents in the Entry Type field indicates that the file or directory object was backed up as result of a directory target type.
- Directory Tree in the Entry Type field indicates that the file or directory object was backed up as a result of a directory target type with the include subdirectories option selected.

File Information and File Versions

When you click a file or directory name anywhere in the backup history (such as in the Backup Info tab), the File Information and File Versions tabs display, as Figure 12-9 shows.

Figure 12-9: File Information and File Versions

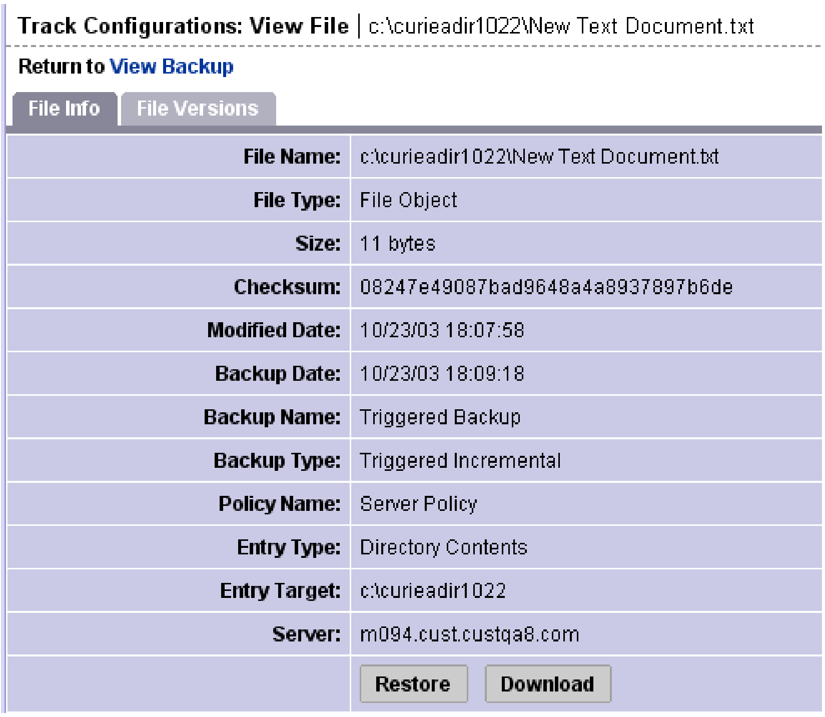

Track Configurations: View File | c:\curieadir1022\New Text Document.txt

Return to View Backup

File Name:	c:\curieadir1022\New Text Document.txt
File Type:	File Object
Size:	11 bytes
Checksum:	08247e49087bad9648a4a8937897b6de
Modified Date:	10/23/03 18:07:58
Backup Date:	10/23/03 18:09:18
Backup Name:	Triggered Backup
Backup Type:	Triggered Incremental
Policy Name:	Server Policy
Entry Type:	Directory Contents
Entry Target:	c:\curieadir1022
Server:	m094.cust.custqa8.com

The File Info tab displays information about the specific file that you selected. The Policy Name refers to the policy that caused the file to be backed up. The Policy Name is the Server Policy if the server's customized tracking policy caused the file to be backed up.

You can place the file that you selected in the Restore Queue by clicking Restore. See "Restoration of Backed Up Files" on page 451 in this chapter for more information.

The File Versions tab displays a list of the backed up versions of the file, as Figure 12-10 shows.

Figure 12-10: File Versions

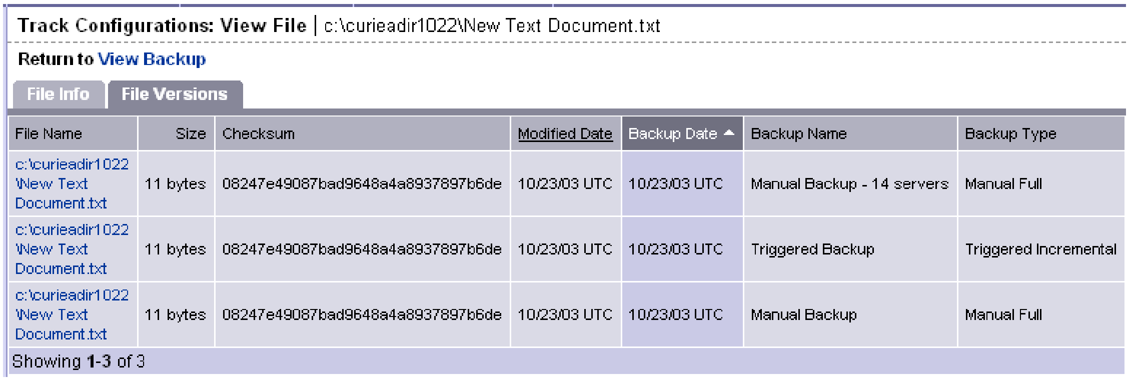

Track Configurations: View File | c:\curieadir1022\New Text Document.txt

Return to View Backup

File Name	Size	Checksum	Modified Date	Backup Date ▲	Backup Name	Backup Type
c:\curieadir1022\New Text Document.txt	11 bytes	08247e49087bad9648a4a8937897b6de	10/23/03 UTC	10/23/03 UTC	Manual Backup - 14 servers	Manual Full
c:\curieadir1022\New Text Document.txt	11 bytes	08247e49087bad9648a4a8937897b6de	10/23/03 UTC	10/23/03 UTC	Triggered Backup	Triggered Incremental
c:\curieadir1022\New Text Document.txt	11 bytes	08247e49087bad9648a4a8937897b6de	10/23/03 UTC	10/23/03 UTC	Manual Backup	Manual Full

Showing **1-3** of 3

It displays the backup name and backup type of the file. If you click any of the files, the File Info for the file is displayed. You can therefore use the File Versions tab to select a different version of the file, and then restore it or download it from the File History.

Deleting Backups

Backups remain in the backup history until you delete them or the server is deactivated. Backups are stored in the Software Repository. You should delete old backup events periodically to reclaim disk storage from the Software Repository.

You can delete entire backup events (identified by a backup name); you cannot delete individual files from the backup history.

See the *SA Administration Guide* for more information about mass deletion of backup files.

Perform the following steps to delete backup points from your backup history:

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. Select the server or set of servers that have backup points that you want to delete.
4. Search for a date that contains the backup points that you want to delete. Click the link with the number of backup events that occurred on that date.

5. Select the check boxes for the backup events that you want to delete from the backup history.

6. Click **Delete**.

   The Delete Backup Confirmation page appears.

7. Click **Delete Backups**. The backups are now deleted.

Restoration of Backed Up Files

You can use the Configuration Tracking feature to restore configuration files or databases that have been backed up. You can restore all files that were backed up during a backup point, or you can select and restore individual files from a backup point. HP Server Automation allows you to rollback a restoration (for example, return your server's tracked files to the state immediately before the restoration).

Overview of Procedures Used to Restore Backed Up Files

This section presents an overview of the procedures you use to restore backed-up configuration files and databases. See “Restoring Backups” on page 454 in this chapter for information about the step-by-step procedure.

To restore backed-up configuration files and databases, view the backup history for a server or set of servers. The backup history displays entries for each date when a backup occurred. As Figure 12-11 shows, the entry for the date displays how many backup points occurred on that date.

Figure 12-11: Backup Points

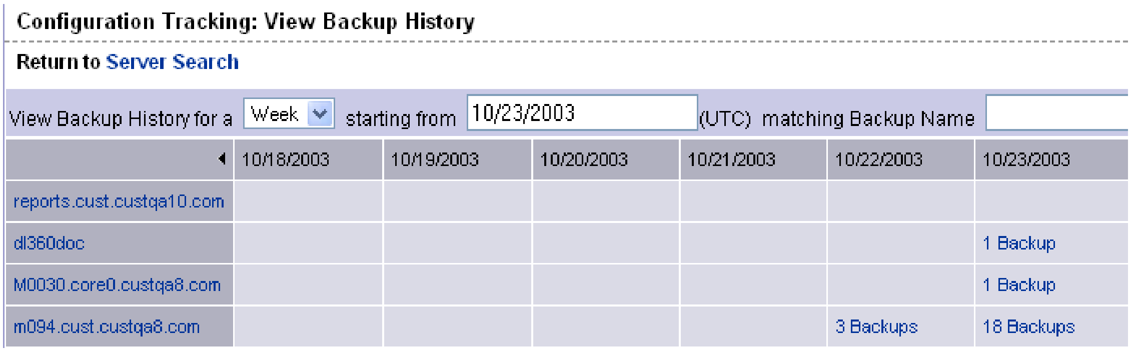

Configuration Tracking: View Backup History

Return to Server Search

View Backup History for a Week starting from 10/23/2003 (UTC) matching Backup Name

◂	10/18/2003	10/19/2003	10/20/2003	10/21/2003	10/22/2003	10/23/2003
reports.cust.custqa10.com						
dl360doc						1 Backup
M0030.core0.custqa8.com						1 Backup
m094.cust.custqa8.com					3 Backups	18 Backups

You select the desired backup date, and you can then select one or more backup points that occurred on that date. Before you restore the backups, you can review all files that were backed up at your selected backup points. You can either choose to restore all backed up files from your selected backup points, or you can select the files individually.

To select the backup, click on the link in the date field that displays the number of backups that were performed on that date.

Restore Queue

This section provides information on the restore queue within HP Server Automation and contains the following topics:

- Storing files in the Restore Queue
- Incremental Backups for Restoration and Rollbacks
- Entries for Directories in the Backup History
- File Not Found Entries
- Restoring Backups
- Rolling Back Restored Files

Storing files in the Restore Queue

When you click **Restore**, the files that you selected are placed in the Restore Queue and the View Restore Queue and Perform Restore page appears. You can then review and select files from all the backup points that you selected.

If you do not want to restore the files at this time, you can perform other actions and the files will remain in the Restore Queue as long as your session is active (even if you leave this page). You can also return to the backup history and select other files to put in the Restore Queue.

When you are ready to restore the backups, return to the Track Configurations: Select a Task page and click the **View Restore Queue** and the Perform Restore link.

Incremental Backups for Restoration and Rollbacks

To make rollbacks possible, HP Server Automation performs two automatic incremental backups. The first backup occurs immediately before the restoration, and the second occurs immediately after the restoration. Similarly, in order to undo a rollback, HP Server Automation performs two automatic incremental backups, one before and one after a rollback.

These backups do not occur if all the files you have selected to restore are identical to the files already on the server. In such a case, the restoration does not in fact change any files on the server, and the rollback option is not available (there are no changes to rollback).

Entries for Directories in the Backup History

As discussed in the "Special Considerations for Directory and Wildcard Targets" in the *SA Policy Setter's Guide*, if you are tracking a directory and the contents of the directory change, the action that you selected is triggered for the directory object itself and for the files that changed. The only exception occurs when a file is deleted from the directory. In that case, the action is triggered for the directory object alone.

If you selected the backup action for the directory target, when the directory contents change, the directory object is backed up as well as the changed files. When a file is deleted, however, only the directory object is backed up.

If you restore files contained in a directory without selecting the entry for the directory object, the directory is re-created on the server if it does not already exist. If you do select the entry for the directory object, however, you can ensure that the directory object is restored with the same properties (such as permissions and time stamp) it had at the selected backup event.

File Not Found Entries

If you are monitoring specific files (as opposed to files monitored in tracked directories and files monitored as the result of wildcard targets), the backup history can contain entries noting "file not found" if the file does not exist on the server, or if the file is later deleted on the server. Similarly, if you are monitoring specific directories (opposed to monitoring directories through wildcard targets), the backup directory can contain "file not found" entries if the directory does not exist on the server, or if the directory is later deleted.

If you select and restore "file not found" entry and the file exists on the server, the file is deleted (it is reverted to the state of the backup event that you selected.)

Exercise caution in restoring entries for deleted directories. If the directory exists on your server and the directory contains files, both the directory and its contents will be deleted.

Restoring Backups

Perform the following steps to restore backed-up configuration files:

1. From the navigation panel in the SAS Web Client, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. Under the Configuration Tracking drop-down menu, select View Backup History.
4. Find the date and server that has the backup points that contain the files that you want to restore. See "Backup History" on page 443 in this chapter for information about finding backup points, viewing backup point details, and viewing lists of backed up files.
5. Select the check box for the backup points that have files that you want to restore.
6. Click **Restore**.

   The files or collection of files from the backup points you selected are placed in the Restore Queue. If you do not want to perform a restore now, you can click the Return to Select a Task link to leave this page. While your session is active, the files remain in the Restore Queue. You can also continue to add more files to the Restore Queue. To return later to the Restore Queue, click the View Restore Queue and Perform Restore from the Track Configurations: Select a Task page.
7. If you are ready to restore files, first review the files in the Restore Queue. If you want to restore all the files in the queue, you can select the check box at the top of the list. If you do not want to restore all of the files in the queue, select the checkboxes for the files that you want to restore.
8. Click **Restore** (or **Restore All** if you want to restore all of the files in the queue).
9. After the restoration is complete, you can click **View Details** for more information about the restored files.

Rolling Back Restored Files

You can rollback any files that you selected to restore. By rolling back the files, you revert the file to the version that was on the server before you performed the restoration. (In other words, you undo the restoration.)

The rollback option is not available if you restored any files that are no longer being monitored by the Configuration Tracking feature. You can choose to restore all changed files or select the files individually.

Perform the following steps to rollback restored files:

1. After you restore the backups, click **Rollback**.

   The Track Configurations: Restore Queue page appears, which contains the list of files that you need to revert in order to rollback the tracked configuration files to their previous state.

   This list might not contain all the files that you selected for your original restore. If any of the files that you selected during your original restore are identical to the files already on the server, they do not need to be rolled back. Only the files that changed display.

2. If you want to rollback all the changed files, click **Restore All**. Otherwise, you can select the individual files that you want to rollback and click **Restore**.

   The Track Configurations: Restore Progress page appears.

Exercise caution in restoring entries for deleted directories. If the directory exists on your server and the directory contains files, both the directory and its contents will be deleted.

Appendix A: Communication Test Troubleshooting

IN THIS APPENDIX

This section discusses the following topics:

- Overview of Agent Communication Tests
- Command Engine to Agent (AGT)
- Crypto Match (CRP)
- Agent to Command Engine (CE)
- Agent to Data Access Engine (DAE)
- Agent to Software Repository (SWR)
- Machine ID Match (MID)
- Common Troubleshooting Tasks

Overview of Agent Communication Tests

The Communication Test performs the following diagnostic tests to determine if an Agent is reachable:

- **Command Engine to Agent (AGT)**: Determines if the Command Engine can communicate with the Agent. The Command Engine is the HP Server Automation component that enables distributed programs to run across many servers. The Command Engine handles the entry of scripts into the SA Model Repository (the script storage location in HP Server Automation) and the versioning of stored scripts.
- **Crypto Match (CRP)**: Checks that the SSL cryptographic files that the Agent uses are valid.
- **Agent to Command Engine (CE)**: Verifies that the Agent can connect to the Command Engine and retrieve a command for execution.

- **Agent to Data Access Engine (DAE)**: Checks whether or not the Agent can connect to the Data Access Engine and retrieve its device record. The Data Access Engine provides an XML-RPC interface to the Model Repository that simplifies interaction with various clients such as the SAS Web Client, system data collection, and monitoring agents on servers.
- **Agent to Software Repository (SWR)**: Determines if the Agent can establish an SSL connection to the Software Repository. The Software Repository is the central repository for all software that the SA system technology manages. It contains software packages for operating systems, applications, databases, customer code, and software configuration information.
- **Machine ID Match (MID)**: Checks that the Machine ID (MID) on the server matches the MID registered in the Model Repository for the Agent.

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

See "Agent Reachability Communication Tests" on page 129 in Chapter 4 for information about how to run Communication Test.

Command Engine to Agent (AGT)

The Command Engine to Agent (AGT) communications test system checks that the Command Engine can initiate an SSL connection to the Agent and execute an XML/RPC request.

The thirteen possible results are:

- AGT – OK
- AGT – Untested
- AGT – Unexpected error
- AGT – Connection refused
- AGT – Connection time-out
- AGT – Request time-out
- AGT – Server never registered

- AGT – Realm is unreachable
- AGT – Tunnel setup error
- AGT – Gateway denied access
- AGT – Internal Gateway error
- AGT – Gateway could not connect to server
- AGT – Gateway time-out

AGT – OK

No troubleshooting necessary.

AGT – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the Command Engine cannot contact the Agent, then no other tests are possible.

What Can I Do If a Test Is Not Run During an AGT Test?

First resolve all tests that failed, and then run the Communication Test again.

AGT – Unexpected error

This result indicates that the test encountered an unexpected error.

What Can I Do If I Get an Unexpected Error?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Hewlett Packard Customer Support.

AGT – Connection refused

This result indicates that the Command Engine is receiving a TCP reset packet when it attempts to connect to the Agent on port 1002. The likely cause is that the Agent is not running. A firewall might also be blocking the connection.

What Can I Do If the Connection is Refused During an AGT Test?

1. Log into the server and confirm that the Agent is running. See “Verifying that an Agent is Running” on page 482 in this appendix for more information.

2. If the Agent is not running, restart the Agent. See “Restarting an Server Agent” on page 483 in this appendix for more information.

3. From the managed server, use netstat to confirm that a socket is in listen mode on port 1002. If not, stop and restart the Agent.

4. From the server itself, use Telnet to connect to the IP address of the server where the Agent is installed and port (1002) that the Agent is listening on. If this does not succeed, stop and restart the Agent.

5. Verify that the Management IP address that HP Server Automation is using to reach the server is the correct address. See “Checking Management IP of a Managed Server” on page 483 in this appendix for more information. If the IP addresses do not match, stop and restart the Agent, then rerun the test.

6. If the previous steps are performed and the test still fails, the problem is likely caused by either a software-based firewall on the server itself or an external firewall blocking the connection.

AGT – Connection time-out

This result indicates that the Command Engine is not receiving any reply packets when it attempts to initiate a TCP connection to the Agent on port 1002. The likely cause is that the server is not running, or that the IP address that HP Server Automation is using to reach the Agent is incorrect. (A firewall might also be blocking the connection.) To check the IP address that HP Server Automation is using to reach the Agent, see “Checking Management IP of a Managed Server” on page 483.

What Can I Do If the Connection Times Out During an AGT Test?

Follow the same steps used to resolve this issue specified in “What Can I Do If the Connection is Refused During an AGT Test?” on page 459.

AGT – Request time-out

This result indicates that the Command Engine is able to successfully complete a TCP connection to the Agent on port 1002, but no response is received from the Agent in response to the XML-RPC request. The likely cause is that the Agent is hung.

What Can I Do If the Request Times Out During an AGT Test?

1. Log into the server and restart the Agent. See “Restarting an Server Agent” on page 483 in this appendix for more information.

2 Check to see whether or not some other process is consistently utilizing an excessive amount of the CPU on the server where the Agent is installed. Also check to see if the system is performing slowly due to a lack of available memory and/or excessive file IO. In any of these cases, the system might be performing too slowly to permit the Agent to respond to the test in a timely manner.

AGT – Server never registered

This test indicates that the server being tested has neither been registered with the Command Engine, nor can it communicate with the Command Engine. The cause of this could be any number of reasons similar to those in the Agent to Command Engine (CE) test. It is also possible (but unlikely) that the Agent was installed but never started.

What Can I Do If the Server Has Not Been Registered with the Command Engine During an AGT Test?

To troubleshoot this error, use the following procedures:

1 Ensure that the Agent is running. For these instructions, See “Verifying that an Agent is Running” on page 482 in this appendix for more information.

2 Ensure that the Agent can contact the Command Engine.

3 If the Agent is in a Satellite facility, ensure that its Gateways are properly configured and that it is properly configured to use those Gateways. See “Checking Network Gateway Configuration” on page 484 in this appendix for more information.

4 If the Agent is not in a Satellite:

- Ensure the host name “way” (no quotes) resolves to its valid IP address. See “Resolving Host Name” on page 484 in this appendix for more information.
- Verify that a connection can be established to port 1018 of way. Use the command “telnet way 1018" (or equivalent).

One (or more) of the above checks will fail. To solve that failure, refer to the corresponding error code for the Agent to Command Engine (CE) test on page 464, or to the realm connectivity and configuration test.

AGT – Realm is unreachable

The Satellite realm where the managed server is located is unreachable. This means that a path of tunnels between the Gateways in the SA core and the realm of the managed server cannot be established.

What Can I Do If the Realm Is Unreachable During an AGT Test?

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact Hewlett Packard Customer Support for assistance in troubleshooting the Gateway network.

AGT – Tunnel setup error

The Command Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do If I Get a Tunnel Setup Error During an AGT Test?

Contact your SA administrator.

AGT – Gateway denied access

The Gateway is working but refused to proxy the connection on behalf of the Agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the Command Engine access to the Agent.

What Can I Do If the Gateway is Denied Access During an AGT Test?

Contact your SA administrator.

AGT – Internal Gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do If There is an Internal Gateway Error During an AGT Test?

Contact your SA administrator.

AGT – Gateway could not connect to server

The Gateway could not establish a connection to the Agent. This might be because the Agent is not running, or because a firewall might be blocking the connection.

What Can I Do If the Gateway Couldn't Connect to the Server During an AGT Test?

If you suspect the Agent is not running, see "Verifying that an Agent is Running" on page 482. To make sure that the Gateway can establish a connection to the IP address of the server where the Agent is installed, try to ping the IP address of the server where the Agent is installed.

AGT – Gateway time-out

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do If the Gateway Times Out During an AGT Test

Ensure that network connectivity is available between the Gateways in the path between the realm of the managed server and the SA core.

Crypto Match (CRP)

This test checks that the SSL cryptographic files that the Agent uses are valid.

The five possible results are:

- CRP – OK
- CRP – Untested
- CRP – Unexpected error
- CRP – Agent certificate mismatch
- CRP – SSL negotiation failure

CRP – OK

No troubleshooting necessary.

CRP – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the Agent cannot be reached, then no other tests are possible.

What Can I Do If a Test Is Not Run During a CRP Test?

First resolve all tests that failed, and then run the Communication Test again.

CRP – Unexpected error

This result indicates that the test encountered an unexpected error.

What Can I Do If I Get an Unexpected Error During a CRP Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Hewlett Packard Customer Support.

CRP – Agent certificate mismatch

This result indicates that the SSL certificate that the Agent is using (cogbot.srv) does not match the SSL certificate that is registered with HP Server Automation for that Agent.

What Can I Do If I Get a Certificate CN Mismatch During a CRP Test?

Use the Recert Agent Custom Extension to issue a new certificate to the Agent. .

CRP – SSL negotiation failure

This result indicates that the Agent is not accepting SSL connections for the SA core. (The SA core is the entire collection of servers and services that provide HP Server Automation services.) The likely cause of this error is that one or more files in the Agent crypto directory are missing or are invalid.

What Can I Do If I Get an SSL Negotiation Failure During an CRP Test?

Run the Server Recert custom extension in the “set allow recert flag only” mode on the server, and then Run the Server Agent Installer with the “-c” switch.

Reinstalling the Agent with the “-c” option (“c” stands for “clean”) removes all certs on the server and also removes the MID file, which forces the Agent to retrieve a new MID from the Data Access Engine.

- See “Agent Reachability Communication Tests” on page 129 in Chapter 4 for information abouthow to install an Server Agent using the “-c” switch.

After you reinstall the Agent, run the test again to check if the Agent is now reachable.

Agent to Command Engine (CE)

This test checks that the Agent can connect to the Command Engine and retrieve a command for execution.

The sixteen possible results are:

- CE – OK

- CE – Untested
- CE – Unexpected error
- CE – Connection refused
- CE – Connection time-out
- CE – DNS does not resolve
- CE – Old Agent version
- CE – Realm is unreachable
- CE – No Gateway defined
- CE – Tunnel setup error
- CE – Gateway denied access
- CE – Gateway name resolution error
- CE – Internal Gateway error
- CE – Gateway could not connect to server
- CE – Gateway time-out
- CE – No callback from Agent

CE – OK

No troubleshooting necessary.

CE – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the Agent cannot reach the Command Engine, then no other tests are possible.

What Can I Do If a Test Is Not Run During a CE Test?

First resolve all tests that failed, and then run the Communication Test again.

CE – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During a CE Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Hewlett Packard Customer Support.

CE – Connection refused

This result indicates that the Agent is receiving a TCP reset packet when attempting to connect to the Command Engine on port 1018. The likely cause is that the Agent is connecting to the wrong IP address. In other words, the Agent does not know the correct IP address of the Command Engine. It is also possible that a firewall might be blocking the connection.

What Can I Do If the Connection is Refused During a CE Test?

1. Check that the name "way" resolves to its correct IP address. For instructions on how to do this, see "Resolving Host Name" on page 484.
2. Check to make sure there isn't a firewall refusing the connection to this IP address.

CE – Connection time-out

This result indicates that the Agent is not receiving any reply packets when it attempts to initiate a TCP connection to the Command Engine on port 1018. The likely cause is that the Agent is connecting to the "wrong" IP address. In other words, the Agent doesn't know the correct IP address of the Command Engine. A firewall might also be blocking the connection.

What Can I Do If the Connection Times Out During a CE Test?

Follow the same steps specified in What Can I Do If the Connection is Refused During a CE Test?.

CE – DNS does not resolve

This result indicates that the Agent cannot resolve the host name "way" to a valid IP address. In other words, the Agent does not know the correct IP address of the Command Engine.

What Can I Do If the Command Engine Name Does Not Resolve During a CE Test?

Log into the server and use a command such as Telnet to confirm that the host name "way" can resolve (for example: "telnet way 1018"). If not, check the DNS configuration of the server to make sure that the host name "way" is configured to its correct IP address. See "Resolving Host Name" on page 484 in this appendix for more information.

CE – Old Agent version

This result indicates that the Agent was unable to contact the Command Engine, but the test was unable to determine the exact cause because the Agent is out of date.

What Can I Do If the Agent is Out of Date During a CE Test?

If this error occurs, it will likely be for one of two reasons: either the host name of the Command Engine ("way") did not resolve, or the connection was refused.

- If you believe that the host name of the Command Engine ("way") did not resolve, then See "CE – DNS does not resolve" on page 466 in this appendix for more information.
- If you determine that the connection was refused, See "CE – Connection refused" on page 466 in this appendix for more information.

Alternatively, you can upgrade the Agent to the latest version (contact Hewlett Packard Customer Support) and re-run the test. See "Agent Reachability Communication Tests" on page 129 in Chapter 4 for information about how to install an Agent.

CE – Realm is unreachable

The Satellite realm where the managed server is located is unreachable. This error means that a path of tunnels between the Gateways in the SA core and the realm of the managed server cannot be established.

What Can I Do if the Realm is Unreachable During a CE Test?

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact your SA administrator for assistance in troubleshooting the Gateway network.

CE – No Gateway defined

The managed server is in a Satellite realm, but its Agent is not properly configured to use a Gateway. Agents located in satellites must use a Gateway to contact the core.

What Can I Do If No Gateway is Defined During a CE Test?

To troubleshoot this error, try the following:

1. Create or open the opswgw.args file on the managed server. The opswgw.args file is located on the managed server at:
   - **Unix/Linux**: /etc/opt/opsware/agent
   - **Windows**: %SystemDrive%\Program Files\Common Files\Opsware\etc\agent
2. Make sure that this file contains a single line as shown:

```
opswgw.gw_list: <gw_ip_address>:<gw_port>,<gw_up_
address>:<gw_port>
```

CE – Tunnel setup error

The Command Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do If A Tunnel Setup Occurs Error During a CE Test?

Contact your SA administrator.

CE – Gateway denied access

The Gateway is working, but refused to proxy the connection on behalf of the Agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the Agent to access the Command Engine.

What Can I Do if the Gateway is Denied Access During a CE Test?

Contact your SA administrator.

CE – Gateway name resolution error

The server running the Gateway in the SA core was unable to resolve the host name "way". It must be able to do this in order to proxy connections on behalf of managed servers in Satellite realms.

What Can I Do if a Name Resolution Error Occurs on the Gateway During a CE Test?

Log into the server where the core Gateway is located and use a command such as ping or host to confirm that the host name "way" can be resolved (for example: "host way").

If you cannot connect, contact your SA administrator so that you can check the DNS configuration of the core Gateway server.

CE – Internal Gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do if an Internal Gateway Error Occurs During a CE Test?

Contact your SA administrator.

CE – Gateway could not connect to server

The Gateway could not establish a connection to the Command Engine. The situation might be because the Command Engine is not running, or because the Gateway is resolving the Command Engine host name (“way”) to the wrong IP address. It is also possible that a firewall might be blocking the connection.

What Can I Do if the Gateway Can’t Connect to Server During a CE Test?

Check that the name “way” resolves to the correct IP address and that the Gateway can establish a connection to port 1018 at that IP. See “Resolving Host Name” on page 484 and “Verifying that a Port is Open on a Managed Server” on page 482 in this appendix.

CE – Gateway time-out

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do if the Gateway Times Out During a CE Test?

Ensure that network connectivity is available between the Gateways in the path between the realm of the managed server and the SA core.

CE – No callback from Agent

The Command Engine was able to contact the Agent, but the Agent did not call back to retrieve its command. However, the Agent reports that it can connect to a Command Engine. This most likely means that the managed server's name resolution mechanism (such as DNS) is not configured to point the server to a different SA core facility than is currently stored for the server by HP Server Automation.

What Can I Do if There is No Callback from Agent?

It is possible that the MID file is missing. If the MID file is missing, it can be recreated easily by creating a file called 'mid' in the correct location which contains the value of the “Server ID” from the Server’s Properties page in the SAS Web Client.

If the MID file is not missing, then ensure that the server's name resolution mechanism (DNS, NIS, and so on) is properly configured so that the host names “spin” and “way” resolve to the appropriate IP addresses or SA core services in the same core, and that those hosts can be reached from the server on ports 1004 and 1018, respectively. If this is the case, it is likely that the server has been redirected to a different core recently, and the Agent has not yet registered, which will cause HP Server Automation to update its records. It this issue remains unresolved for more than 12 hours, contact Hewlett Packard Customer Support.

Agent to Data Access Engine (DAE)

This test checks that the Agent can retrieve its device record from Data Access Engine. The fifteen possible results are:

- DAE – OK
- DAE – Untested
- DAE – Unexpected error
- DAE – Connection refused
- DAE – Connection time-out
- DAE – DNS does not resolve
- DAE – Old Agent version
- DAE – Realm is unreachable
- DAE – No Gateway defined
- DAE – Tunnel setup error
- DAE – Gateway denied access
- DAE – Gateway name resolution error
- DAE – Internal Gateway error
- DAE – Gateway could not connect to server
- DAE – Gateway time-out

DAE – OK

No troubleshooting necessary.

DAE – Untested

This result is returned when a functional area cannot be tested, because of a previous failure that prevents further testing. For example, if the Agent cannot reach the Data Access Engine then no other tests are possible.

What Can I Do If a Test Is Not Run During a DAE Test?

First resolve all tests that failed, and then run the Communication Test again.

DAE – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During a DAE Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Hewlett Packard Customer Support.

DAE – Connection refused

This result indicates that the Agent is receiving a TCP reset packet when attempting to connect to the Data Access Engine on port 1004. The likely cause is that the Agent is connecting to the wrong IP address. A firewall might also be blocking the connection.

What Can I Do If the Connection is Refused During a DAE Test?

1. Check that the name “spin” resolves to its correct IP address. See “Resolving Host Name” on page 484 in this appendix for more information.
2. Check to make sure that a firewall is not refusing the connection to this IP address.

DAE – Connection time-out

This result indicates that the Agent is not receiving any reply packets when it attempts to initiate a TCP connection to the Data Access Engine on port 1004. The likely cause is that the Agent is connecting to the wrong IP address. In other words, the Agent does not know the correct IP address of the Command Engine. A firewall might also be blocking the connection.

What Can I Do If the Connection Times Out During a DAE Test?

Follow the same steps specified in "What Can I Do If the Connection is Refused During a DAE Test?" on page 471.

DAE – DNS does not resolve

This result indicates that the Agent cannot resolve the host name "spin" to a valid IP address. In other words, the Agent does not know the correct IP address of the Data Access Engine.

What Can I Do If the Data Access Engine Name Does Not Resolve During a DAE Test?

Log into the server and use a command such as Telnet to confirm that the host name "spin" can be resolved (for example: telnet spin 1004"). If not, check the DNS configuration of the server to make sure that the host name "spin" is configured to its correct IP address. See "Resolving Host Name" on page 484 in this appendix for more information.

DAE – Old Agent version

This result indicates that the Agent was unable to contact the Data Access Engine, and the test is unable to determine the exact cause, because the Agent is out of date.

What Can I Do If the Agent is Out of Date During an DAE Test?

If this error occurs, it will likely be for one of two reasons: either the host name of the Data Access Engine ("spin") did not resolve, or the connection was refused.

- If you believe that the host name of the Data Access Engine ("way") did not resolve, then see "DAE – DNS does not resolve" on page 472.
- If you determine that the connection was refused, see "DAE – Connection refused" on page 471.

Alternatively, you can upgrade the Agent to the latest version (contact Hewlett Packard Customer Support) and re-run the test. See "Agent Reachability Communication Tests" on page 129 in Chapter 4 for information about how to install an Agent.

DAE – Realm is unreachable

The Satellite realm where the managed server is located is unreachable. This error means that a path of tunnels between the gateways in the SA core and the realm of the managed server cannot be established.

What Can I Do if the Realm is Unreachable During a DAE *Test?*

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact your SA administrator for assistance in troubleshooting the Gateway network

DAE – No Gateway defined

The managed server is in a Satellite realm, but its Agent is not properly configured to use a Gateway. Agents located in satellites must use a Gateway to contact the core.

What Can I Do If No Gateway is Defined During a DAE *Test?*

To troubleshoot this error, try the following:

1. Create or open the opswgw.args file on the managed server. The opswgw.args file is located on the managed server at:

   - **Unix/Linux**: /etc/opt/opsware/agent
   - **Windows**: %SystemDrive%\Program Files\Common Files\Opsware\etc\agent

2. Make sure this file contains a single line as shown:

   ```
   opswgw.gw_list: <gw_ip_address>:<gw_port>,<gw_up_
   address>:<gw_port>
   ```

DAE – Tunnel setup error

The Data Access Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do if a Tunnel Setup Error Occurs During a DAE *Test?*

Contact your SA administrator.

DAE – Gateway denied access

The Gateway is working, but refused to proxy the connection on behalf of the Agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the Agent to access the Data Access Engine.

What Can I Do if the Gateway is Denied Access During a DAE *Test?*

Contact your SA administrator.

DAE – Gateway name resolution error

The server running the Gateway in the SA core was unable to resolve the host name "spin". It must be able to do this in order to proxy connections on behalf of managed servers in Satellite realms.

What Can I Do if There is a Name Resolution Error on the Gateway During a DAE *Test?*

Log into the server where the core Gateway is located and use a command such as ping or host to confirm that the host name "spin" can be resolved (for example: "host spin").

If you cannot connect, contact your SA administrator so that you can check the DNS configuration of the core Gateway server.

DAE – Internal Gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do if an Internal Gateway Error Occurs During a DAE *Test?*

Contact your SA administrator.

DAE – Gateway could not connect to server

The Gateway could not establish a connection to the Data Access Engine. This might be because the Data Access Engine is not running, or because the Gateway is resolving the Data Access Engine host name ("spin") to the wrong IP address. It is also possible that a firewall might be blocking the connection.

What Can I Do if the Gateway Can't Connect to Server During a DAE *Test?*

Check that the name "spin" resolves to the correct IP address and that the Gateway can establish a connection to port 1018 at that IP. See "Resolving Host Name" on page 484 in this appendix for more information and See "Verifying that a Port is Open on a Managed Server" on page 482 in this appendix for more information.

DAE – Gateway time-out

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do if the Gateway Times Out During a DAE Test?

Ensure that network connectivity is available between the Gateways in the path between the managed server's realm and the SA core.

Agent to Software Repository (SWR)

This test checks that the Agent can establish an SSL connection to the Software Repository.

There 16 possible results are:

- SWR – OK
- SWR – Untested
- SWR – Unexpected error
- SWR – Connection refused
- SWR – Connection time-out
- SWR – DNS does not resolve
- SWR – Old Agent version
- SWR - Server identification error
- SWR – Realm is unreachable
- SWR – No Gateway defined
- SWR – Tunnel setup error
- SWR – Gateway denied access
- SWR – Gateway name resolution error
- SWR – Internal Gateway error
- SWR – Gateway Could not connect to server
- SWR – Gateway time-out

SWR – OK

No troubleshooting necessary.

SWR – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the Agent cannot reach the Software Repository, then no other tests are possible.

What Can I Do If a Test Is Not Run During a SWR Test?

First resolve all tests that failed, and then run the Communication Test again.

SWR – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During a SWR Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Hewlett Packard Customer Customer Support.

SWR – Connection refused

This result indicates that the Agent is receiving a TCP reset packet when attempting to connect to the Software Repository on port 1003. The likely cause is that the Agent is trying to connect to the wrong IP address. A firewall might also be blocking the connection.

What Can I Do If the Connection is Refused During an SWR Test?

1. Check that the name "theword" resolves to the correct IP address. For this information, see "Resolving Host Name" on page 484.
2. Check to make sure that a firewall isn't refusing the connection to this IP address.

SWR – Connection time-out

This result indicates that the Agent is receiving a TCP reset packet when attempting to connect to the Software Repository on port 1003. The likely cause is that the Agent is connecting to the wrong IP address. In other words, the Agent does not know the correct IP address of the Software Repository. A firewall might also be blocking the connection.

What Can I Do If the Connection Times Out During an SWR Test?

Follow the same steps specified in "What Can I Do If the Connection is Refused During an SWR Test?" on page 476.

SWR – DNS does not resolve

This result indicates that the Agent cannot resolve the host name “theword” to a valid IP address. In other words, the Agent does not know the correct IP address of the Software Repository.

What Can I Do If the Software Repository Name (“theword”) Does Not Resolve During an SWR Test?

Log into the server and use a command such as Telnet to confirm that the host name “theword” can be resolved (for example: “telnet theword 1003"). If not, contact your SA administrator so that you can check the DNS configuration of the server.

SWR – Old Agent version

This result indicates that the Agent was unable to contact the Software Repository, and the test is unable to determine the exact cause because the Agent is out of date.

What Can I Do If the Agent is Out of Date During an SWR Test?

If this error occurs, it will likely be for one of two reasons: either the host name of the Software Repository (“theword”) did not resolve, or the connection was refused.

- If you think that the host name of the Software Repository (“theword”) did not resolve, then see “SWR – DNS does not resolve” on page 477.
- If you determine that the connection was refused, see “SWR – Connection refused” on page 476.

Alternatively, you can upgrade the Agent to the latest version (contact Hewlett Packard Customer Support) and re-run the test. See “Agent Installation Using ODAD” on page 120 in Chapter 4 for information about how to install an Server Agent.

SWR - Server identification error

Whenever an Agent makes a request of the Software Repository, the identity of the server is validated to confirm that the server should be allowed access to the information requested. This error indicates that the Software Repository was unable to identify the server being tested, or incorrectly identified that server.

What Can I Do If I Get a Server Identification Error?

The Software Repository identifies servers based on the incoming IP address of the request. To troubleshoot this error, try the following:

1. Check the Device Properties tab for the server in the SAS Web Client to see if Network Address Translation (NAT) is in use. If it is, make sure that NAT is statically configured, and that only one server is using the NAT address. If multiple servers are using the same IP address, you will need to reconfigure the NAT device. See "Network Address Translation (NAT) for Managed Servers" on page 277 in Chapter 8 for more information.

2. If the Agent is installed on a cluster, check that each node in the cluster has a unique IP address at which it can be reached. You might have to add static routes to the server to ensure that connections made from that server to the SA core use the unique IP. If NAT is not in use, you can alternately mark the correct interface as the "primary" interface through the Network Configuration tab for the server in the SAS Web Client. See "Network Address Translation (NAT) for Managed Servers" on page 277 in Chapter 8 for more information.

3. The server's IP address might have changed recently. If this is the case, stop and restart the Agent. For instructions on how to stop and start an Agent, see "Restarting an Server Agent" on page 483.

SWR – Realm is unreachable

The Satellite realm where the managed server is located is unreachable. This error means that a path of tunnels between the gateways in the SA core and the realm of the managed server cannot be established.

What Can I Do if the Realm is Unreachable During a SWR Test?

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact your SA administrator for assistance in troubleshooting the Gateway network.

SWR – No Gateway defined

The managed server is in a Satellite realm, but its Agent is not properly configured to use a Gateway. Agents located in satellites must use a Gateway to contact the core.

What Can I Do If No Gateway is Defined During a SWR Test?

To troubleshoot this error, try the following:

1. Create or open the opswgw.args file on the managed server. The opswgw.args file is located on the managed server at:

    - **Unix/Linux**: /etc/opt/opsware/agent

- **Windows**: %SystemDrive%\Program Files\Common Files\Opsware\etc\agent

2 Make sure that this file contains a single line as shown:

```
opswgw.gw_list: <gw_ip_address>:<gw_port>,<gw_up_
address>:<gw_port>
```

SWR – Tunnel setup error

The Data Access Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do If a Tunnel Setup Error Occurs During a SWR Test?

Contact your SA administrator.

SWR – Gateway denied access

The Gateway is working but refused to proxy the connection on behalf of the Agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the Agent to access the Software Repository.

What Can I Do if the Gateway is Denied Access During a SWR Test?

Contact your SA administrator.

SWR – Gateway name resolution error

The server running the Gateway in the SA core was unable to resolve the host name “theword”. It must be able to do this in order to proxy connections on behalf of managed servers in Satellite realms.

What Can I Do if a Name Resolution Error Occurs on the Gateway During a SWR Test?

Log into the server where the core Gateway is located and use a command such as ping or host to confirm that the host name “theword” can be resolved (for example: “host theword”).

If you cannot connect, contact your SA administrator so that you can check the DNS configuration of the core Gateway server.

SWR – Internal Gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do if an Internal Gateway Error Occurs During a SWR *Test?*

Contact your SA administrator.

SWR – Gateway Could not connect to server

The Gateway couldn't establish a connection to the Software Repository. This error might be because the Software Repository is not running, or because the Gateway is resolving the Software Repository host name ("theword") to the wrong IP address. It is also possible that a firewall might be blocking the connection.

What Can I Do if the Gateway Can't Connect to Server During a SWR *Test?*

Check that the name "theword" resolves to the correct IP address and that the Gateway can establish a connection to port 1018 at that IP address. For more information, see "Resolving Host Name" on page 484 and "Verifying that a Port is Open on a Managed Server" on page 482.

SWR – Gateway time-out

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do if the Gateway Times Out During a SWR Test?

Ensure that network connectivity is available between the Gateways in the path between the realm of the managed server and the SA core.

Machine ID Match (MID)

This test checks whether the MID that the Agent reported matches that recorded in the Model Repository (SA data repository).

You can receive four possible errors from the Machine ID (MID) Communication Test:

- MID – OK
- MID – Untested
- MID – Unexpected error
- MID – MID mismatch

MID – OK

No troubleshooting necessary.

MID – Untested

This result is returned when a functional area cannot be tested, because of a previous failure that prevents further testing. For example, if the Agent cannot reach the Model Repository, then no other tests are possible.

What Can I Do If a Test Is Not Run During an MID Test?

First resolve all tests that failed, and then run the Communication Test again.

MID – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During an MID Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Hewlett Packard Customer Support.

MID – MID mismatch

This result indicates that the MID that the Agent reported does not match the recorded MID in the Model Repository for that Agent. The likely cause is that the Command Engine is running the test against the wrong Agent.

What Can I Do If the MID is Mismatched During an MID Test?

To troubleshoot this error, try the following:

1. Check the Device Properties tab for the server in the SAS Web Client to see if NAT is in use for this server. If it is, make sure that static, 1-to-1 NAT is being used. HP Server Automation requires that all managed servers be reachable on a distinct, consistent IP address, so configurations that assign addresses dynamically or use port-based translation are not supported.

2. If the Agent is installed on a cluster, check that each node in the cluster has a unique IP address at which it can be reached. You might have to add static routes to the server to ensure that connections made from that server to the SA core use the unique IP. If NAT is not in use, you can alternately mark the correct interface as the “primary” interface via the Network Configuration tab for the server in the SAS Web Client.

3. The IP address might have changed recently. If this is the case, stop and restart the Agent. For these instructions, see “Restarting an Server Agent” on page 483.

Common Troubleshooting Tasks

The following list of troubleshooting tasks are common to more than one Communication Test error:

- Verifying that an Agent is Running
- Verifying that a Port is Open on a Managed Server
- Restarting an Server Agent
- Checking Management IP of a Managed Server
- Checking Network Gateway Configuration
- Resolving Host Name

Verifying that an Agent is Running

To verify that an Agent is running on a server, perform the following steps:

1. On Solaris, HP-UX, or AIX, enter this command:

   ```
   /usr/ucb/ps auxwww | grep opsware
   ```

   You should get this result if the Agent is running:

   ```
   /opt/opsware/agent/bin/python /opt/opsware/agent/pylibs/
   shadowbot/daemonbot.pyc --conf /etc/opt/opsware/agent/
   agent.args
   ```

2. On Linux, enter this command:

   ```
   ps auxwww | grep opsware
   ```

   You should get the same result as the preceding step.

3. On Windows, from the Administrative Tools | Services, check to make sure that the `opswareagent` service is running.

Verifying that a Port is Open on a Managed Server

For some errors, you will need to verify that the port is open on the server where the Agent is installed. To do this, perform the following steps:

1. Check if the port is open.

2. On Solaris, HP-UX, AIX, or Linux enter:

   ```
   netstat -an | grep 1002 | grep LISTEN
   ```

   If the port is open on the box, you should get back the following:

   ```
   *.1002    *.*    0      0 24576      0 LISTEN
   ```

3. On Windows, at the command prompt enter:

   ```
   netstat -an | find "1002" | find "LISTEN"
   ```

   If the port is open on the box, you should get back the following result:

   ```
   TCP0.0.0.0:10020.0.0.0:0LISTENING
   ```

4. Confirm that the port is actually open. To do this, from the computer where the Agent is installed, Telnet to port 1002 by using both localhost and the external IP address of the server. Performing the Telnet will help you confirm that a connection refused message is being caused by the lack of an open port on the managed server rather than a problem with networking hardware between the core and the managed server.

Restarting an Server Agent

To restart an Server Agent, log onto the managed server and enter the following commands:

Unix:

```
/etc/init.d/opsware-agent restart
```

Windows:

```
net stop opswareagent
net start opswareagent
```

Checking Management IP of a Managed Server

To check the Management IP of a managed server, perform the following steps:

1. To view the management IP of the managed server, log into the SAS Web Client.

2. From the Navigation panel, click Servers ➤ Manage Servers.

3. From the Manage Servers list, click the display name of the server for which you want to check the Management IP.

4. Select the Network tab of the server’s properties.

5. Check to make sure that the Management IP address matches the IP address of the managed server.

Checking Network Gateway Configuration

To check the network Gateway configuration, perform the following steps:

1. On Solaris, enter this command to check routing table:

   ```
   netstat -rn
   ```

   Your results should look like this:

   ```
   default              192.168.8.1         UG        1   5904
   ```

   where `192.168.8.1` is the IP of the Gateway.

2. On Linux, enter this command to check routing table:

   ```
   route -n
   ```

   Your results should look like this:

   ```
   0.0.0.0       192.168.8.1    0.0.0.0       UG    0     0
   0 eth0
   ```

   where `192.168.8.1` is the IP of the Gateway.

3. On Windows, enter this command to check routing table:

   ```
   route print
   ```

   Your results should look something like this:

   ```
   0.0.0.00.0.0.0192.168.8.1192.168.8.12020
   ```

   where `192.168.8.1` is the IP of the Gateway.

4. In each case, you should also ping 192.168.8.1 (IP) to confirm that you can actually reach the Gateway.

Resolving Host Name

All managed servers (those with agents) must be able to resolve unqualified HP Server Automation service names for the following components:

- spin (Data Access Engine)
- way (Command Engine)

- theword (Software Repository)

If you need to ensure that one of these host names resolves correctly, contact your SA administrator to find out what qualified host name or IP address these service names should resolve to.

1. Try to ping the host. For example, execute the following command if you wanted to resolve the host name, way:

   ```
   ping way
   ```

2. If the host name cannot resolve, you will get the following errors:

   Linux/Solaris/AIX/HP-UX:

   ```
   ping: unknown host way
   ```

   Windows:

   ```
   Ping request could not find host way. Please check the name
   and try again.
   ```

3. If the host name can resolve, you might get back various permutations of these types of messages (OS independent):

   ```
   way is alive
   ```

   or

   ```
   pinging way (ip) with 32 bytes of data
   ```

Appendix B: Agent Utilities

IN THIS APPENDIX

This section discusses the following topics:

- Agent Installation Using the CLI
- Agent Upgrade Tool
- Agent Upgrade Custom Extension

Agent Installation Using the CLI

This section provides information on Agent installation using the Agent Installer CLI and contains the following topics:

- Overview of Agent Installation Using the CLI
- Preparation for Agent Installation
- Preassimilation Checklist
- Installing an Agent by using the Agent Installer CLI
- Agent Installer Options
- Example: Agent Installer Command and Options
- Starting an Agent After Installation
- Verifying Agent Functionality
- Augmenting the Information for a Managed Server
- Uninstalling an Agent on Unix and Windows
- Uninstalling Earlier Versions of Agents on Unix
- Uninstalling Earlier Versions of Agents on Windows

Overview of Agent Installation Using the CLI

When you install Agents on existing operational servers, you should synchronize the local time on the servers with an external time-server that uses a network time protocol (NTP).

Installing Agents on servers makes existing operational servers known to SA so that they can be managed. Assimilating servers into SA is appropriate when many servers are already functioning in the operational environment and need to be managed (for example, when SA technology is initially deployed in a facility).

Installing an Agent on a server with a pre-built OS into SA enables:

- Baseline discovery of the operating system on the server.
- Managing the baseline operating system, including patch management, when the operating system is defined in SA with the Prepare Operating System Wizard.
- Full provisioning and management capabilities for any new applications deployed on the server.

When installed, the Agent registers the server with the Model Repository. SA assigns the server to a generic operating system that corresponds to the operating system that the Agent discovered during the installation. The server is assigned to a placeholder OS node. For each operating system, the SAS Web Client contains a node <*operating_system_version*>/Not Assigned.

The Agent Installer can install Agents when a core is not available to a server. If a newly-installed Agent cannot contact a core, the Agent runs in a dormant mode. While dormant, it periodically attempts to contact the core. When the core becomes available, the Agent performs the initialization tasks, such as hardware and software registration, that usually take place when the Agent is first installed.

The server is tracked in the SAS Web Client. However, the server operating system *cannot* be managed while the server is assigned to the generic operating system node. You must reassign the server to the operating system that was defined with the OS Provisioning feature.

The server is associated with the default facility for the local instance of SA.

If the managed server's IP address does not fall within a specified IP range, the server is associated with the default IP range group (Default). The default group is associated with the customer Not Assigned.

See the *SA Policy Setter's Guide* for information about how servers are associated with customers.

Users install an Agent on each server. Running an Agent on a server allows SA to manage the server. To install an Agent, run the Agent Installer.

The Agent Installer is an application that has the following features:

- Can be invoked from the command line or within a script.
- Installs an Agent.
- Logs its decisions and actions.
- Can be operated unattended because user interaction is not required.

The Agent Installer installs the Agent, retrieves cryptographic material, retrieves configuration information, and writes a configuration file.

Preparation for Agent Installation

It is recommended that you set up a Windows file share to make the Agent Installer for various operating systems available from one place. Setting up a file share allows you to install Agents on servers quickly and easily. If this is not possible, the Agent Installer needs to be moved by using an alternate file transfer mechanism, such as SFTP.

At the completion of the Agent installation process, a managed server is assimilated and the hardware and software data that the Agent discovered is stored in the Model Repository.

To use the Patch Management features on Windows NT 4.0 and Windows 2000 servers, you must install Internet Explorer (IE) 6.0 or later on the server first because a patch utility depends on it. If you do not install IE 6.0 on the server first, the Agent Installer warns you that the Patch Management feature does not work as expected because it checks for IE 6.0 on all Window servers. This prerequisite is not required for Windows 2003, because IE 6.0 is pre-installed for this operating system.

Preassimilation Checklist

Prior to installation, you should perform the following tasks on the server where the Agent is to be installed. Performing these tasks is vital to installing the Agent quickly within maintenance windows.

1. For the Code Deployment & Rollback feature, verify that the following port is accessible between the server you will push code from and the server where you will push code to:

   - `telnet` *<staging_server>* `1002`
   - `telnet` *<production_server>* `1002`

2. Because the Agent runs on port 1002, verify that no other applications are using this port.

   - On a Unix server, enter this command from a terminal window:

     ```
     netstat -an | grep 1002 | grep LISTEN
     ```

   - On a Windows server, enter this command from a terminal window:

     ```
     netstat -an | find "1002" | find "LISTEN"
     ```

3. Check for sufficient disk space for Agent installation on the server.

   The Installer checks for the following amounts of free disk space in these directories:

   - 30MB in `/opt/opsware` (Unix)
   - 100MB in `/var/opt/opsware` (Unix)
   - 30MB in `%SystemDrive%\\Program Files\Opsware` (Windows)
   - 100MB in `%SystemDrive%\Program Files\Common Files\Opsware` (Windows)

   (These default directories can be overridden with parameters at installation time.)

   These space requirements might not be enough. The `vardir` directory is used for dynamic content like logs and downloaded packages. If there is not enough disk space for the packages during remediation, it will fail.

4. On the Solaris operating system, check for legacy sun4m architecture. Currently, the Agent works only for sun4u architecture.

5. For Windows, check the following items:

   - At a minimum, NT 4.0 Service Pack 6a must be installed on the server.
   - Verify that the Windows Registry has the correct settings:

1. Start regedit and locate the following registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control
\FileSystem
```

2. Select the `NtfsDisable8dot3NameCreation` entry.
3. On the **Edit** menu, click **DWORD** and verify that the value is set to 0. The value *must* be set to 0. If necessary (because the value is set to 1), change the value by following your organization's IT policies and reboot the server.

6 To install an Agent on a server running Solaris, you must also install the following Solaris packages:

For Python:

```
SUNWtoo
SUNWtoox
```

For showrev:

```
SUNCadm
SUNWlibC
SUNWlibCx
SUNWadmfw
```

Before you install an Agent on a server, the server must meet certain package and patch requirements that vary by operating system, as Table B-4 shows.

Table B-4: Required Patches and Packages for Agent Installation

SERVER OPERATING SYSTEM	REQUIRED PATCHES
AIX 4.3 AIX 5.1	APAR IY39444 APAR IY39429 NOTE: If AIX 4.3.3.388, 4.4.4.89, or 5.1.0.3 is installed, the Agent Installer displays an error message that indicates the correct APAR to install on the server.
HP-UX (10.20, 11.00, 11.11/11i)	For HP-UX 10.20, PHCO_21018 Additionally, SW-DIST should be upgraded to the HP recommended patch level. You should continue to upgrade this package when HP recommends new versions.

Table B-4: Required Patches and Packages for Agent Installation (continued)

SERVER OPERATING SYSTEM	REQUIRED PATCHES
Linux AS 3.0 Linux WS 3.0 Linux ES 3.0	Red Hat Enterprise Linux 3 Update 3
Solaris 10, 9, 8, 7, and 6	SUNWadmc SUNWcsl SUNWcslr (If available, depending on version) SUNWcsu SUNWesu SUNWlibms SUNlibmsr (If available, depending on version) SUNWswmt It is strongly recommended not to remove packages from the SUNWCreq minimal required install cluster, since many packages are interdependent and operation beyond that of basic SAS functionality may be affected.
Windows 2000	Service Pack 4
Windows NT 4.0	Service Pack 6a

Installing an Agent by using the Agent Installer CLI

The Agent needs administrator-level privileges (root on Unix servers and Local System on Windows servers) to manage a server. Therefore, Agent installation needs to be performed as root on Unix operating systems and as administrator on Windows operating systems.

You can install an Agent on any server listed in “Supported Operating Systems for Managed Servers” in *SA User's Guide: Server Automation*.

Perform the following steps to install an Agent on a server:

1. Log into the server that you want to assimilate by using a remote shell.
2. For Unix operating systems, change the user login to root (`su - root`) and for Windows operating systems, log in as administrator.
3. From the CLI, download the package that contains the Agent Installer from the:

```
/var/opt/opsware/agent_installers
```

direcory in the Core to a directory on the server you want to assimilate.

4 Locate the package `opsware-agent`.

Each operating system and operating system version has different packages for the Agent Installer.

Unix:

```
opsware-agent-<version>-<system_name>-<system_version>
```

Red Hat Linux:

```
opsware-agent-<version>-<system_name>-5CLIENT-<system_
version>
```

Note that Red Hat may use the terms *Client* and *Desktop* interchangeably.

Windows:

```
opsware-agent-<version>-<system_name>-<system_
version>.exe
```

4. Click the package name for the Agent Installer that you want to download. From the Actions menu, select Export Software. The Browse window appears.
5. In the Browse window, provide the location to download the package and Click **Export**. The package is exported to the specified location.

5 From the directory where the Agent Installer was copied, run the Installer by entering the correct executable and options for the installation environment.

See “Example: Agent Installer Command and Options” on page 498 in this appendix for more information.

Agent Installer Options

When you use the Agent Installer CLI, you can include the options that the following table shows to control the way that the Agent is installed on a server.

Table B-5: Agent Installer Options

OPTION	DESCRIPTION
`-f`	Forces Agent installation and removes the target installation directory if it exists. REQUIREMENT: When using the `-f` option, you must run the Agent Installer as root on Unix operating systems and as the administrator on Windows operating systems.
`--logfile`	Specifies the path to the Agent Installer log file. By default, the current directory is set as the path. By default, the log file has the following filename: `opsware-agent-installer-<`*date*`>.log`
`--loglevel <`*level*`>`	Sets the log level for log messages. With this option, specify one of the following levels: `error`, `warn`, `info`, `trace`, or `none`. The level `error` logs the least detail. The level `trace` logs all messages. By default, the log level is set to the log level `info`.
`-o`	Logs all output to `stdout` instead of a log file. This option is invoked automatically if the default log file or the log file passed with the `--logfile` option cannot be created, for example, when running the Agent Installer from non-writeable media, such as a DVD.

Table B-5: Agent Installer Options (continued)

OPTION	DESCRIPTION
`--reconcile <type>`	Reconciles the server against any nodes assigned to the server. The <*type*> can be `full` or `addonly`. `full` – All nodes in a category are selected and reconcile removes software that SA did not install. `addonly` – Software installed outside of SA is not removed. WARNING: When assimilating a server that is already functioning in the operational environment, use caution when specifying the option `--reconcile`. If you specify this option, you might inadvertently uninstall software from the server.
`--rpmbin <path>`	Specifies the path to the RPM binary to use for RPM operations. Use this option, when RPM is already installed on the server, to point the Agent at the RPM binary. Use the `--withrpm` option to install RPM if a usable instance of RPM is *not* already installed. NOTE: It is unnecessary to use this option with the `--withrpm` option.
`-s`	Starts the Agent after installing it. By default, the Agent Installer does not start the Agent.
`--template <ID>`	Assigns the nodes contained in the template to the server. `<ID>` can be an ID or a full name of a template. If this option is specified with the `--reconcile` option, SA assigns the nodes in the template to the server before reconciling the server. WARNING: When assimilating a server that is already functioning in the operational environment, use caution when you specify the option `--template`. If you specify this option, you might inadvertently uninstall software from the server.

Table B-5: Agent Installer Options (continued)

OPTION	DESCRIPTION
`--withmsi`	Installs MSI 2.0 along with the Agent. If MSI 2.0 is already installed, this option has no effect. Works with Windows NT 4.0 Service Pack 6a, Windows 2000, and Windows 2003.
`--withwmi`	Installs WMI 1.5 along with the Agent. If WMI 1.5 is already installed, this option has no effect. Works with Windows NT 4.0 Service Pack 6a.
`--withrpm`	Installs the RPM handler with the Agent. By default, an Agent is not installed with this option. It is recommended that you always include the `--withrpm` option when you install Agents on Solaris servers. NOTE: Use the `--withrpm` option only with the Agent Installers for these operating systems: Solaris 5.6, 5.7, 5.8, and 5.9, and AIX 4.3 and AIX 5.1. On Solaris, RPM 3.0.6 is installed in the directory `/opt/OPSWrpm` and the RPM database is installed in the directory `/var/opt/OPSWrpm/lib/rpm`. On AIX, RPM 3.0.5 is installed in the directory `/opt/freeware` and the RPM database is installed in the directory `/var/opt/freeware/lib/rpm`.
`--workdir` *<path>*	Specifies the path to the Agent Installer temporary working directory. Use this option if the default working directory causes problems with installation.

Table B-5: Agent Installer Options (continued)

OPTION	DESCRIPTION
`--reboot`	During Agent Installation on a Windows server, the Agent Installer copies the ogshcap.dll file to the following location: `%SystemRoot%\system32\ogshcap.dll` If the file is open or is in use, the Agent Installer is unable to copy the ogshcap.dll file. The Agent Installer then informs the user whether to restart the machine and copies the file after restart. You can specify the `--reboot` Installer option in the command line to initiate the reboot at the end of the Agent installation.
`--resetconf(-r)`	Resets Agent configuration file to default the settings.
`--no_anonymous_ssl (-A)`	Disables anonymous SSL. This option applies to dormant Agents only. This option configures the Agent so that browsers cannot connect without a valid certificate.
`--opsw_gw-addr-list`	Specifies the Gateways used during Agent installation.
`--remediate`	Remediates the server against any software policies attached to the server.
`--software_policy <id>`	Attaches server to the software policy *<ID>*.
`--installdir <path>`	Allows you to specify the directory where the Agent will be installed. By default, the Agent is installed in `/opt/opsware/agent` on UNIX, and `%ProgramFiles%\Opsware\agent` on Windows.
`--no_open_fw`	By default, the Agent Installer will modify the Windows Firewall configuration on Windows XP and Windows 2003 (r2) servers to allow the core to contact the managed server on port 1002. If you specify this option, the firewall configuration will not be modified. The server may not be manageable by SA in this case.

Table B-5: Agent Installer Options (continued)

OPTION	DESCRIPTION
`--settime (-t)`	Synchronizes the time on the server on which the Agent is installed with that of the core. NOTE: If the server on which the Agent is being installed is significantly ahead of the clock on the core, then the clock on the managed server is set back in time. Since this can cause problems, do not use the --settime option unless you are sure that this scenario is not a problem in your environment. If a managed server's clock is significantly behind of the clock on the core, the Agent installation might fail. To install an Agent successfully, use the `--settime` option or manually set the time and date on the managed server before retrying the Agent installation.

Example: Agent Installer Command and Options

Enter the following command and options to install the Agent for Solaris 5.7 in the default directories and log the results of the installation in the log file:

```
% opsware-agent-14.2.12.5-solaris-5.7 --logfile opsware-agent-
installer.log --loglevel info
```

Enter the following command and options to install the Agent for Windows NT 4.0 in the default directories and log the results of the installation in the log file:

```
% opsware-agent-14.2.12.5-win32-4.0.exe --logfile opsware-
agent-installer.log --loglevel info
```

Starting an Agent After Installation

If you do *not* include the `-s` option on the command line when you install an Agent, you have to start the Agent on the server manually.

Unix:

```
    /etc/init.d/opsware-agent start
```

Windows:

```
net start opswareagent
```

Verifying Agent Functionality

Perform the following steps to verify Agent functionality:

1. From the navigation panel in the SAS Web Client, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server whose Agent installation you want to verify. If necessary, select the correct customer and facility for the server and click **Update**.

   Or

   Search for the server whose Agent installation you want to verify.

2. Verify that the server appears in the Manage Servers list and has the correct properties.

   See “Using the Search Feature” and “Server Searching by IP Address” in *SA User’s Guide: Server Automation* for more information.

3. If you want to discover reasons why a server is unreachable, you can run a Communication Test. See “Agent Reachability Communication Tests” on page 129 in Chapter 4 for more information.

Augmenting the Information for a Managed Server

Use caution when you augment the discovery process for an Opsware-managed server that is functioning in the operational environment. You might inadvertently install or uninstall software from the server. During the test remediate, verify what software will be uninstalled from the server before you perform the actual remediate operation.

Perform the following steps to augment the information for an Opsware-managed server:

1. Model the OS and other applications running on the server in SA by defining the OS with the Prepare Operating System Wizard and by creating nodes and templates for applications running on the assimilated server.

   See the *SA Policy Setter’s Guide* for more information about Operating System Definitions.

2. Move the server to the appropriate nodes for the OS and installed applications.

The server is tracked in the SAS Web Client; however, the server operating system *cannot* be managed while the server is assigned to the generic operating system node. You must reassign the server to the operating system that was defined with the OS Provisioning feature.

3. Remediate the server.
4. If an IP range group was set up, servers are automatically associated with customers when users install an Agent on the servers. Otherwise, the servers are associated with the Not Assigned customer. To change the customer associated with a server, See "Editing the Properties of a Server" on page 323 in Chapter 8 for more information.
5. To specify the server's use, stage, and state, edit the server's properties. See "Editing the Properties of a Server" on page 323 in Chapter 8 for more information.

Discovery is complete. SA assumes that the server should always be running the specific OS build it has been associated with. Any changes to the OS outside of SA are not captured in the model.

Users can deploy and manage new applications on the server, just as if SA initially provisioned the server. Users can also deploy OS level patches on the server, or rebuild the OS by using the OS build with which the server was associated.

Uninstalling an Agent on Unix and Windows

To uninstall Agents on Windows NT, Windows Scripting Host 5.1 or Internet Explorer 5.5 must be installed on Windows NT.

Perform the following steps to uninstall an Agent on Unix or Windows:

1. Log into Unix as root user. Log into Windows as Administrator.
2. Change directories to any directory other than the Agent's installation directory.
3. *On Unix*, enter the following command:

   ```
   <installation_directory>/bin/agent_uninstall.sh
   ```

   By default, for Solaris and AIX, the Agent Uninstaller will not remove the SA RPM package. For command line options for the agent uninstaller, including how to activate removal of the SA RPM package, See "Agent Uninstaller Options" on page 501 in this appendix for more information.

4. *On Windows*, enter the following command:

```
msiexec /x <installation_directory>\bin\agent_uninstall.msi
```

You can also use the Windows Control Panel's Add or Remove Programs option to remove the Agent.

5. As the uninstall proceeds, the *Unix* platform `stdout` shows the uninstallation progress. The *Windows* uninstall does not show uninstallation progress.

Agent Uninstaller Options

When you use the Agent Uninstaller, you can include the options that Table B-6 and Table B-7 show.

Table B-6: Agent Uninstallation Unix Options

OPTION	DESCRIPTION
`--uninstallerVersion`	Show the uninstaller version.
`--help`	Show this help.
`--no_deactivate`	Do not deactivate the server; by default, the server is deactivated.
`--force`	Do not prompt for confirmation before deactivating the server.
`--delete_opsw_rpm`	Remove the OPSW RPM package (AIX, Solaris only). Use the following commands to remove the RPM package: Solaris: `pkgrm -n OPSWrpm` AIX: `installp -u rpm.rte`

Table B-7: Agent Uninstallation Windows Options

OPTION	DESCRIPTION
`NO_DEACTIVATE="1"`	Do not deactivate the server; by default the server is deactivated.
`FORCE="1"`	Do not prompt for confirmation before deactivating the server.

During Agent Uninstallation on a Windows server, the Agent Installer removes the ogshcap.dll file from the following location:

```
%SystemRoot%\system32\ogshcap.dll
```

If the file is open or is in use, the Agent Installer is unable to remove the ogshcap.dll file. The Agent Installer then prompts the user to restart the machine and removes the file after restart.

Uninstalling Earlier Versions of Agents on Unix

Perform the following steps to uninstall Agents versions 5.1 and earlier:

1. Stop the Agent on the server by running the following command as root:

Linux:

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

Solaris:

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

HP-UX:

```
/sbin/init.d/cogbot stop
```

AIX:

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

2. Deactivate or delete the server by using the **SAS Web Client Server** menu.

3. For Linux servers only, run `chkconfig` to de-register the Agent initialization script:

```
% /sbin/chkconfig -del cogbot
```

4. As root, delete the following files and directories to remove the Agent files from the server:

Linux:

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

Solaris:

```
/etc/init.d/cogbot

/etc/rc2.d/S79cogbot

/etc/rc0.d/K44cogbot
```

HP-UX:

```
/sbin/init.d/cogbot
```

```
/sbin/rc2.d/cogbot
```

AIX:

```
/etc/rc2.d/init.d/cogbot
```

```
/etc/rc.d/S79cogbot
```

All Unix:

```
/opt/OPSW
```

```
/var/lc
```

Uninstalling Earlier Versions of Agents on Windows

Perform the following steps to uninstall earlier versions of Agents on Windows:

1. Stop the Agent by running the following command as administrator:

   ```
   C:\> net stop shadowbot
   ```

2. Deactivate or delete the server by using the **SAS Web Client Server** menu.

3. Deregister the Agent service by running the following command as administrator:

   ```
   C:\> "%SystemDrive%\Program
   Files\Loudcloud\blackshadow\watchdog\watchdog.exe" -x
   ```

4. As administrator, delete the following directories to remove the Agent:

   ```
   "%SystemDrive%\Program Files\Loudcloud"
   ```

   ```
   "%SystemDrive%\Program Files\Common Files\Loudcloud"
   ```

Agent Upgrade Tool

This section contains the following topics:

- Ways to Upgrade Agents
- Prerequisites for Using the Agent Upgrade Tool
- Upgrading the Agent on Managed Servers
- Commands for the Agent Upgrade Tool

- Options for the Agent Upgrade Tool
- Example: Options for the Agent Upgrade Tool
- Example: Commands and Output for Agent Upgrade Tool

Ways to Upgrade Agents

After you upgrade SA running in a facility, you should upgrade the Agents on every managed server to the new version, so that you can utilize the new features in the newly-upgraded core.

SA features continue to work on a managed server even when it is running an older Agent. However, new features in the new versions might not be available for that server.

Refer to the Release Notes for the new version for information about the compatibility of new features with older agents.

You can upgrade the Agents on managed servers in the following ways:

- Use the Agent Installer (a command line interface) to install a new Agent on one server at a time.

  See "Agent Installation Using the CLI" on page 487 in this appendix for information about how to use the Agent Installer.

- Use the Agent Upgrade Tool to upgrade Agents on groups of servers. Running the tool upgrades deployed Agents on managed servers. You can run the script simultaneously on many servers to upgrade large groups of Agents.

The Agent Upgrade Tool has the following characteristics:

- It is a command line interface that provides a flexible mechanism for selecting servers to upgrade, and for monitoring and reviewing upgrade operations.
- You can use it to upgrade many Agents on managed servers simultaneously.
- It runs within your preferred Unix shell, allowing it to leverage the power of standard Unix shells and text processing tools.
- You can use it to upgrade a server in any facility running SA. You can run it from an OPSH shell attached to any SA in any facility.

The OPSH shell is a program that authenticates users in SA, starts the user's normal Unix shell (as specified in the standard password database). Using the OPSH shell allows the user to run the Agent Upgrade Tool in this facility.

Prerequisites for Using the Agent Upgrade Tool

- On a core server of the facility being upgraded, install the OPSH Shell RPM by downloading the `opsh` package from the SAS Web Client. See the *Opsware® SA Content Migration Guide* for instructions on downloading a package.

  Installing the `opsh` RPM places the OGFS shell and the Agent Upgrade Tool in the directory `/opt/OPSWopsh/bin`.

- You need the correct permissions to upgrade Agents. Run the OPSH shell by specifying the SA `admin` user and password to ensure that you have the appropriate permissions. (Contact your SA administrator to obtain the password.)

  When you start an OPSH shell to run the Agent Upgrade Tool, the user name and password are authenticated by SA

- The server where you install the `opsh` RPM must be able to resolve the name `way.<facility-domain>` to the host running the Command Engine in the facility's core. For example, if the Command Engine runs on a host in the `prod.opsware.com` domain, the servers must resolve the name `way.prod.opsware.com`. You specified the `facility-domain` when you installed the core.

Upgrading the Agent on Managed Servers

To upgrade the Agent, perform the following steps:

1. After you install the `opsh` RPM on a core server, enter the following command as root to start the OPSH shell:

   ```
   opsh [username@]facility-domain
   ```

   For example:

   ```
   opsh admin@prod.opsware.com
   ```

   See the preceding section for the name resolution requirement for `facility-domain`. See "Commands for the Agent Upgrade Tool" on page 506 for more information on `opsh`.

2. (Optional) To obtain information about the current Agents running on the managed servers before you upgrade them, enter any of the following commands and options:

   ```
   opsh_agent query server-options
   ```

   (Enter this command if you want to view a report of the Agent versions running on the servers before you upgrade them.)

`opsh_agent verify` *server-options schedule-options* `\`

agent-version

(Enter this command if you want to verify the versions of the Agents running on the managed servers before you upgrade them.)

3. To upgrade Agents on specified servers, enter the following Agent Upgrade Tool commands and options:

   `opsh_agent stage` *server-options schedule-options* `\`

   `[--always]` *agent-version*

   (Enter this command if you want to download the package for the Agent to the managed server before you run the upgrade.)

   `opsh_agent upgrade` *server-options schedule-options* `\`

   `[--always]` *agent-version*

4. (Optional) To review the status of the Agent upgrade, enter the following command and option:

   `opsh_agent review` *session-id*

Commands for the Agent Upgrade Tool

- `opsh [username@]` *facility-domain*

  This command starts an OPSH shell and authenticates the user name against the SA facility running at the specified domain.

  If you do not specify a user name, the currently logged in user name is used. The OPSH shell prompts for a password.

  A new Unix shell (which is attached to the specified SA core-domain) is started. (The password database for the user specifies which Unix shell to use.)

- `opsh_agent query` *server-options*

  This command must be run from an OPSH shell started with the `opsh` command.

  This command queries the reported version of Agents and any staging status for the specified servers by examining data in the Model Repository.

  One line is printed to stdout for each server that shows device ID, IP address, current Agent version, and any staging status.

You can specify the servers by using the `--device, --customer, --facility, and --os` options.

- `opsh_agent stage server-options schedule-options \`

  `[--always] agent-version`

  You must run this command from an OPSH shell started with the `opsh` command.

  This command contacts the Agent on each specified server and instructs it to download the package for the specified version of the Agent from the Software Repository.

  If the download is successful, the staging status is written to the Model Repository for the server.

  To download the package to the server even when this command was entered previously (recorded in the Model Repository), specify the `--always` option.

  One line is printed to stdout for each server that shows the device ID, IP address, and a success or failure indicator.

  You can specify the servers by using the `--server, --customer, --facility and --os` options.

  A session is started and the session ID displays for later review. After the session ID displays, you can type `CTRL-C` and review the session later using the `opsh_agent review` command.

- `opsh_agent upgrade server-options schedule-options \`

  `[--always] agent-version`

  You must run this command from an OPSH shell started with the `opsh` command.

  This command contacts the Agent on each specified server and instructs it to upgrade to the specified version. If the necessary package has not been downloaded on the server already (the `opsh_agent stage` command was entered), the package is downloaded from the Software Repository.

  If the upgrade is successful, the package is removed from the server and the staging status is deleted from the Model Repository.

  To upgrade the Agent even when the specified version of the Agent was already installed on the managed servers, enter the `--always` option. (The Model Repository records when Agents are upgraded on servers.)

One line is printed to stdout for each server that shows the device ID, IP address, and a success or failure indicator.

You can specify the servers by using the `--server, --customer, --facility,` and `--os` options.

A session is started and the session ID displays for later review. After the session ID displays, you can type CTRL-C and review the session later by using the `opsh_agent review` command.

- `opsh_agent verify server-options schedule-options \`

  `agent-version`

  You must run this command from an OPSH shell started with the `opsh` command.

  This command contacts the Agent on each specified server to verify that it is running the specified version.

  One line is printed to stdout for each server that shows the device ID, IP address, the word OLD, NEW, or CURRENT and the actual Agent version running on the server.

  You can specify the servers by using the `--server, --customer`, `--facility,` and `--os` options.

  A session is started and the session ID displays for later review. After the session ID displays, you can enter CTRL-C and review the session later by using the `opsh_agent review` command.

- `opsh_agent review session-id`

  You must run this command from an OPSH shell started with the `opsh` command; although, not necessarily the same OPSH shell from which the original command was started.

  This command attaches to a running `opsh_agent stage, opsh_agent upgrade` or `opsh_agent verify` session running on the Command Engine. It prints the same output to stdout that the original command would have printed if the user had not typed CTRL-C and terminated the command. If the session is complete, it shows the same results that were shown when the session completed.

Options for the Agent Upgrade Tool

Server-options: `--server|-S <svr-spec> --customer|-C <cust-spec>`

`--facility|-F <fac-spec> --os|-O <os-spec>`

If more than one of the `--customer`, `--facility`, or `--os` options is specified, only servers that match all options are selected. Any servers specified by using the `--server` option are added to (or subtracted from) the list specified by combining the

`--customer`, `--facility,` and `--os` options.

Table B-8: Options for the Agent Upgrade Tool

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--server`	`-S`	`<svr-spec>`	Server by device ID, IP address, or system name
`--customer`	`-C`	`<cust-spec>`	All servers associated with the customer specified by the customer ID or name
`--facility`	`-F`	`<fac-spec>`	All servers in the facility specified by the facility ID or name
`--os`	`-O`	`<os-spec>`	All servers running the operating system specified by the OS name

Schedule-options: `--when|-W when-time --until|-U until-time`

Table B-9: Schedule Options

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--when`	`-W`	`<when-time>`	Start time for a stage, upgrade, verify, or test operation in the format: MM/DD/YYY-HH:MM If the `--when` option is used, the operation starts at the specified time, but the command displays a session ID and returns immediately. When you schedule an operation, you use the `review` command to review the results after the operation has run. If the `--when` option is not specified, the operation starts immediately and the command displays the output of the command.

Table B-9: Schedule Options (continued)

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--until`	`-U`	`<until-time>`	End time for a stage, upgrade, verify or test operation in the format: MM/DD/YYY-HH:MM If the `--until` option is specified, the operation stops processing servers at the specified time. Any servers that are not complete are left in a consistent state; this might require that the session run past the specified time.

Miscellaneous options: `--ip|-I --always|-A --parallel|-P --theword|-T`

Table B-10: Miscellaneous Options

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--ip`	`-I`	(N/A)	Display IP addresses instead of host names.
`--always`	`-A`	(N/A)	Always stage or upgrade servers even if the current version is staged or upgraded.
`--parallel`	`-P`	`<Concurrency>`	The maximum of concurrent commands. (Recommended default = 10)
`--theword`	`-T`	`<hostname>`	The host name or IP address to use when contacting the Software Repository from a server.

Example: Options for the Agent Upgrade Tool

The following table provides examples for running the Agent Upgrade Tool.

Table B-11: Examples of Options for the Agent Upgrade Option

EXAMPLE	DESCRIPTION
`--server 1,2`	Selects servers 1 and 2.

Table B-11: Examples of Options for the Agent Upgrade Option (continued)

EXAMPLE	DESCRIPTION
`--facility Y,Z`	Selects all servers (for all customers) in facilities Y and Z.
`--customer -A,-B --facility Z`	Selects all servers in facility Z except those owned by customers A and B.
`--server 1,2,-3,-4 --customer A,B --facility Y,Z`	Select servers 1 and 2 as well as all servers owned by customers A or B which are in facilities Y or Z except servers 3 and 4.
`--server 1,-2 --customer A,B --facility -Y,-Z --os SunOS 5.8`	Select server 1 and all servers owned by customers A or B, except those in facilities Y or Z and which are Solaris 5.8 machines excluding server 2.

Example: Commands and Output for Agent Upgrade Tool

```
# cd /opt/OPSWopsh/bin

# ./opsh admin@prod.opsware.com

admin@prod.opsware.com's password:

#

# ./opsh_agent verify --os "SunOS*" 14a.2.12.18

Session 37802500101L

Device ID Name/IP address Version Result Status Reason

410101L core2-1.prod.opsware.com  14a.2.12.18 CURRENT SUCCESS

^C

Interrupted review of running session 37802500101L

Use review 37802500101L command anytime to review session status

#

# ./opsh_agent review 37802500101L

Session 37802500101L
```

```
Device ID Name/IP address Version Result Status Reason

410101L d033.prod.opsware.com 14a.2.12.18 CURRENT SUCCESS

670101L dhcp-174.prod.opsware.com 14a.2.12.16 OLDER SUCCESS

1460100L emb218-37.manu.opsware.com 14a.2.12.18 CURRENT SUCCESS

20100L f001.manu.opsware.com 14a.2.12.18 CURRENT SUCCESS

10100L f002.manu.opsware.com 14a.2.12.21 NEWER SUCCESS

210100L m022.manu.opsware.com 14a.2.12.18 CURRENT SUCCESS

Session 37802500101L completed.
```

Agent Upgrade Custom Extension

SA includes a custom extension (AgentUpgrade) that you can use to upgrade Agents from the previous three versions.

The Custom Extension feature is accessed through the Run Custom Extension Wizard. This Wizard allows a user to choose a custom extension to run, specify necessary input data for the extension, validate the data required to run the extension, run the extension, and view or download the results from the job. When a user runs a custom extension, the job shows up in My Jobs.

To access the Run Custom Extension Wizard, users must be assigned to a user group that has the permission Wizard: Opsware Extension. When users have this permission, they can run any custom extensions on the servers they have access to in the SAS Web Client.

Any user who has the ability to run custom extensions can run the Agent Upgrade Custom Extension.

The Agent Upgrade custom extension only upgrades agents that are not up-to-date. Therefore, rerunning the custom extension on a group of servers will not affect servers that are already running the most recent version of the Agent.

Upgrading Agents with the Agent Upgrade Custom Extension

To run the Agent Upgrade custom extension, perform the following steps:

1. From the SAS Web Client home page, click the Run Custom Extensions link in the Tasks panel.

Or

From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers list appears. Select the servers on which you want to run a custom extension and choose **Run Extension** from the **Tasks** menu.

The Run Custom Extension Wizard appears.

2 Select the "AgentUpgrade" custom extension and click **Next**.

If you have not already selected servers from the Manage Servers list, the Select Servers page appears.

3 If prompted, select the servers or groups on which you want to run the custom extension and click **Next**. You can find the servers that you want to run a custom extension on by browsing the list or by searching.

The Agent Upgrade Settings page appears as shown in Figure B-1:

Figure B-1: Specify Settings Page for the Agent Upgrade Custom Extension

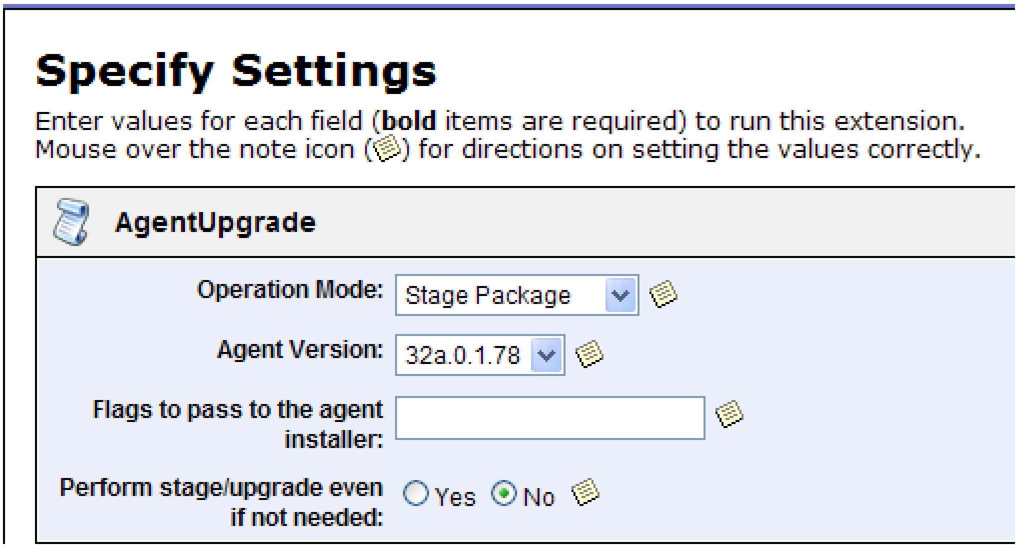

Specify Settings

Enter values for each field (**bold** items are required) to run this extension. Mouse over the note icon (🗒) for directions on setting the values correctly.

AgentUpgrade

Operation Mode:	Stage Package
Agent Version:	32a.0.1.78
Flags to pass to the agent installer:	
Perform stage/upgrade even if not needed:	○ Yes ● No

4 Specify the settings for the custom extension and click **Next**.

- Select the Agent Upgrade action you want to perform:
  - **Stage Package**: Downloads the Agent Installer package (the binaries) but does not upgrade the Agent.

  - **Perform Upgrade**: Upgrades the servers to the specified Agent version by downloading the Agent Installer package (binaries) if not previously staged. If the Agent Installer package was already staged, upgrades the Agent version by using the staged version of package on the server.
  - **Verify Upgrade**: Check whether all the servers are running the selected Agent version.
- Select the Agent version to upgrade to. All agent versions available in the Software Repository are displayed in the list. By default, the latest version available in the Software Repository is selected.
- Enter additional parameters to pass to the Agent Installer. The Agent Upgrade custom extension does not validate the parameters you enter. See “Agent Installer Options” on page 494 in this chapter for information about the parameters you can enter in this field.
- Specify whether the custom extension should upgrade Agents even if the version of the currently running agent is equal to or greater than the version being upgraded to.

The Confirm Settings Screen appears and displays the selected options.

5. Review the values that you entered and the servers that you selected on which to run the Agent Upgrade custom extension. (You can remove servers from the list by deselecting their check boxes. Click the “Show remaining servers” link to display the complete list of selected servers.)

6. Click **Next**.

7. On the Schedule and Notify page, you have the following options:

- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.
- **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

  When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

- **Option 1**: Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

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

8 Click **Run** to start the Custom Extension Wizard.

If you selected to run the job at that time, a progress bar appears for the servers on which the extension is running that shows the progress of the job.

9 When the custom extension finishes running, you can click **View Details** to see the results. The Custom Extension Results window appears. The results window contains a tab for the operation summary and a tab for the operation details.

By default, the Summary tab is displayed.

Figure B-2: Summary Tab in the Details Window of the Agent Upgrade Custom Extension

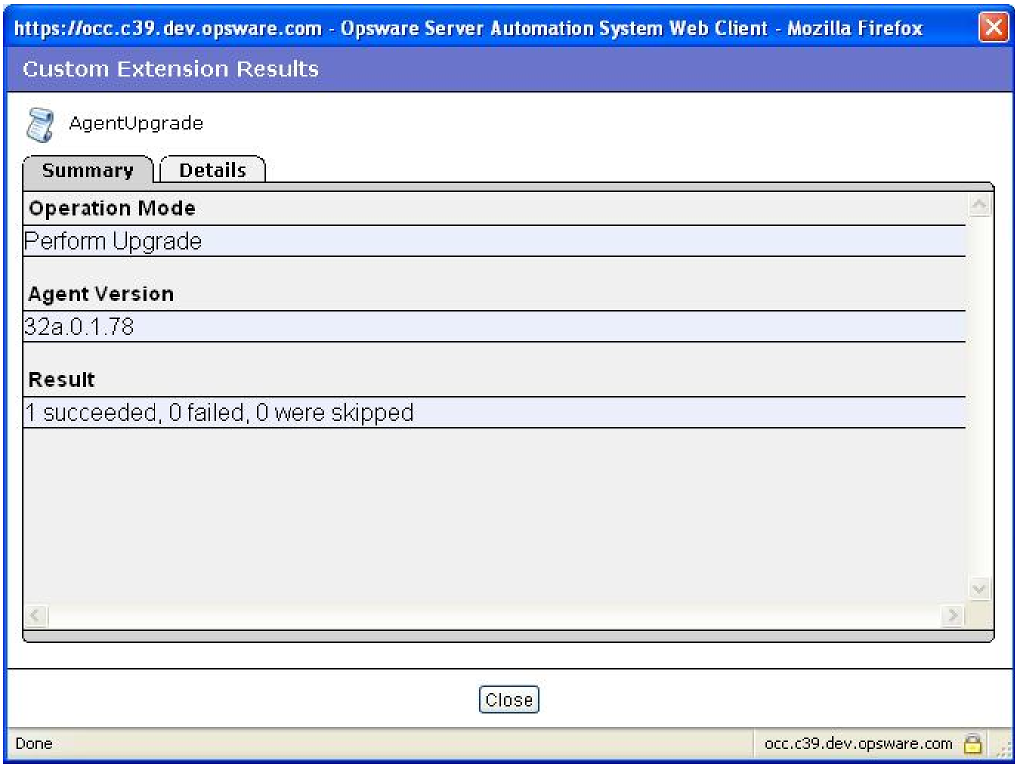

https://occ.c39.dev.opsware.com - Opsware Server Automation System Web Client - Mozilla Firefox

Custom Extension Results

AgentUpgrade

Operation Mode

Perform Upgrade

Agent Version

32a.0.1.78

Result

1 succeeded, 0 failed, 0 were skipped

The Details tab contains a table with one row for each server selected. The table in each tab includes the following columns:

- **Name**: The name of the server
- **Agent Version**: The agent version running on the server after the custom extension completes
- **Reason**: If an operation failed, a text description of the reason the operation failed
- **Result**: (Only shown for Verify operations) Shows the result of the verification:
  - **OLDER**: The agent on the server is older than the selected version
  - **UP TO DATE**: The agent on the server is up to date

- **Status**: The exit status of the command SUCCESS or FAILURE

Figure B-3: Details tab in Results Window for the Agent Upgrade Custom Extension

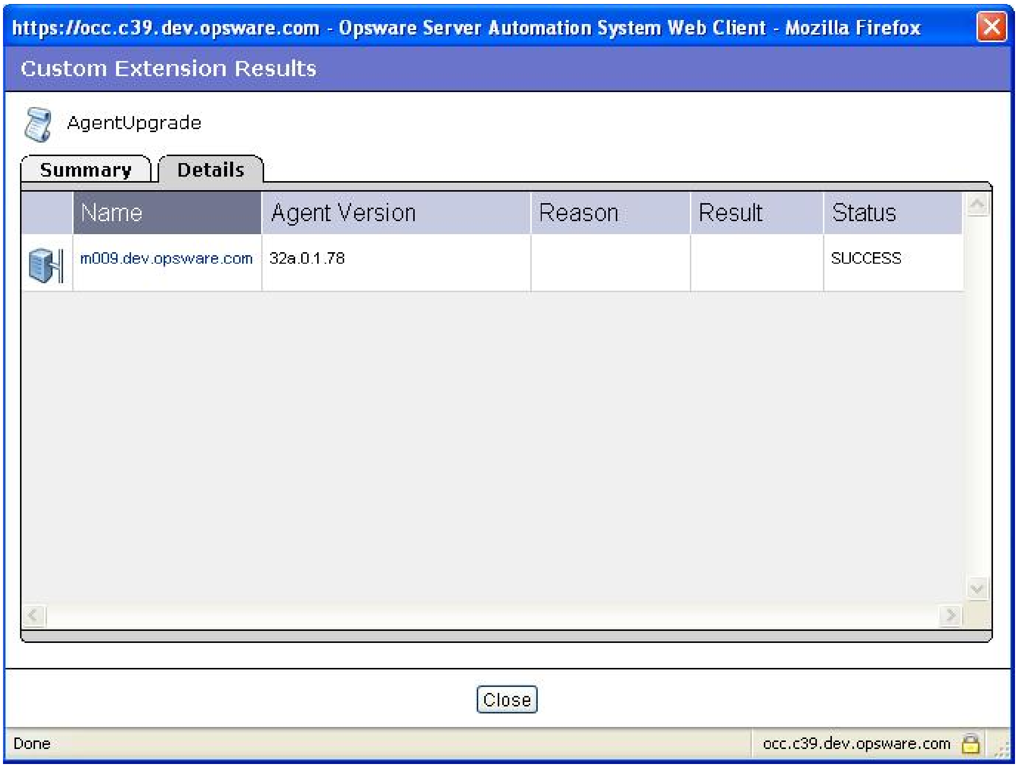
https://occ.c39.dev.opsware.com - Opsware Server Automation System Web Client - Mozilla Firefox
Custom Extension Results
AgentUpgrade
Summary
Details
Close
Done
occ.c39.dev.opsware.com

Name	Agent Version	Reason	Result	Status
m009.dev.opsware.com	32a.0.1.78			SUCCESS

10 When you are done viewing the results, click **Close** to close the window.

After you close the wizard, you can view the progress of the custom extension by viewing My Jobs (accessible from the Home page or the navigation panel). Each custom extension job is identified with the name Run Custom Extension. Click the name link to identify which extension was run.

See “Server Management Scheduling and Notification” on page 332 in Chapter 8 for information about the My Jobs feature.

Appendix C: Global Shell Utilities Syntax

IN THIS APPENDIX

This section discusses the following topics:

- aaa Utility
- rosh Utility
- swenc Utility

aaa Utility

The `aaa` utility grants and revokes permissions for operations that use the OGFS. For example, the `aaa` utility grants permission for the `readServerFilesystem` operation, allowing you to browse a server's file system in the SA Client. To run the `aaa` utility, you must belong to the Administrators user group.

The permissions granted and revoked by the `aaa` utility are stored in the `/opsw/Permissions` directory of the OGFS. For details on the contents of the directory, see "/opsw/Permissions Directory" on page 541.

aaa Syntax

The `aaa` utility has the following syntax:

```
aaa shell-perm (grant | revoke) -o operation [-u user-group]
[-f facility | -c customer | -g device-group [-s | -l login]]
```

Table C-1 describes the command options and Table C-2 lists the operations that can be granted or revoked the `aaa` utility.

Table C-1: aaa Options

OPTION	DESCRIPTION
`-o` *`operation`*	The operation on which to grant or revoke the permission. For a list of allowed values, see the Operation column in Table C-2.

Table C-1: aaa Options (continued)

OPTION	DESCRIPTION
`-u` *`user-group`*	The SA user group that is assigned the permission. This value is inferred from the current working directory if it corresponds to a user group. If it cannot be inferred, specify a user group.
`-f` *`facility`*	The name, ID, or path to a facility, such as: `/opsw/Facility/Chicago` Permission will be granted to all servers in this facility.
`-c` *`customer`*	The name, ID, or path to a customer, such as: `/opsw/Customer/Alpha` Permission will be granted to all servers that belong to this customer.

`-g` *`device-group`*	The name, ID, or path to a public device group, such as:
	`/opsw/Group/Public/Unix Servers`
	Permission will be granted to all servers that belong to this group. To specify the device group by name, omit the following:
	`/opsw/Group/`

`-l` *`login`*	A login account on the servers that are specified by the `-f`, `-c`, or `-g` option. On a Unix server, for example, the *`login`* is the Unix user name. Login accounts with multi-byte characters are not supported.
`-s`	The login account on the servers (specified by `-f`, `-c`, or `-g`) is the same as the SA user name. (Use of `-s` is also referred to as defining a reflexive permission.)

aaa Usage RulesM

The following usage rules and recommendations apply to the `aaa` utility:

- For operations that are performed on a server, one of the -f, -c, or -g options is required.
- As a best practice, when you are granting permissions, use care when you select servers so that you do not capture more servers than you intend. This is particularly important when using the -c or -f option. For example, if you want to grant permission to the loginToServer operation for all servers in the Chicago facility as root, you could use the -f option to select all servers in a particular facility. However, this may also select Windows servers, which is probably not desired since the root user does not typically exist on Windows servers. In this case, you should define a public device group that only includes servers in the Chicago facility which are running a Unix operating system.
- If you specify the -f, -c, or -g option, you must also specify either the -s or -l option. The choice of the -s or -l option depends on the policies of your organization. If users log into managed servers with generic user names (such as root), then you should specify the -l option. If users log into managed servers with individual user names, which are the same as their SA user names, they should specify the -s option.
- The -f and -c options are provided as a convenience; however, in general, it is recommended that you define permissions based on device groups instead.
- The revoke command can only remove a permission that was previously granted. If the permission was not previously granted, the revoke command has no effect.
- The revoke command only removes a permission for a specific user group. If a user has overlapping permissions, revoking permissions from a single user group will not prevent the user from performing that operation. For example, suppose a user belongs to two user groups that both have the launchGlobalShell permission. If this permission is revoked from only one of those user groups, the user still has the launchGlobalShell permission.

aaa Examples

The following example gives all members of the Advanced Users group permission to open a Global Shell session:

```
aaa shell-perm grant -o launchGlobalShell \
-u ‘Advanced Users’
```

The following command allows members of the Advanced Users group to view the file systems as root of all Unix servers:

```
aaa shell-perm grant -o readServerFilesystem \
-u 'Advanced Users' -g 'Public/All Unix Servers' -l root
```

The next example gives all members of the `Unix Admin` user group permission to log in as `root` to all servers in the `Public/Trading Servers` device group:

```
aaa shell-perm grant -o loginToServer -u ‘Unix Admin’\
-g ‘Public/Trading Servers’ -l root
```

The following example allows the `Advanced Users` group to run commands as `root` on servers associated with the `Acme Inc` customer.

```
aaa shell-perm grant -o runCommandOnServer \
-u ‘Advanced Users’ -c ‘Acme Inc’ -l root
```

The next example removes the permission for the `Unix Admin` user group to log into servers that belong to the device group named `Public/Unix Servers`. The command applies to any login, because the `-l` option is not specified.

```
aaa shell-perm revoke -o loginToServer -u ‘Unix Admin’\
-g ‘Public/Unix Servers’
```

The following example allows the `Oracle Users` group to log into servers that belong to the device group `Oracle Servers` as the login `oracle`. For instance, if the SA user `joe` belongs to the `Oracle Users` group, he can log into the servers as the server user `oracle`.

```
aaa shell-perm grant -u ‘Oracle Administrators’ \
-o loginToServer -g ‘/opsw/Group/Public/Oracle Servers’ \
-l oracle
```

Instead of the `-l` option, the next example has the `-s` option, which allows the `Oracle Users` group to log into servers that belong to the device group `Oracle Servers` as the login that matches the SA user name. For instance, if the SA user `joe` belongs to the `Oracle Users` group, he can log into the servers as the server user `joe`.

```
aaa shell-perm grant -u ‘Oracle Administrators’ \
-o loginToServer -g ‘/opsw/Group/Public/Oracle Servers’ -s
```

Global Shell Operations (Permissions)

The actions that an SA user can perform within the Global Shell are determined by the operations specified by the `aaa` utility. Most of these operations, such as `readServerFilesystem`, can be granted on both managed servers and on a login basis. The login is the user name on the managed server, such as the Administrator user on Windows or root on Unix.

A login is not specific to a particular platform (operating system). For example, if the permissions specify that a user can read the file system as root, then root will appear under the files subdirectory, regardless of the platform. The Server Explorer of the SA Client displays the login names that you have been authorized to access that server's file system.

The operations are listed in the `/opsw/Permissions` directory of the OGFS.

Table C-2 identifies and describes the server operations in Global Shell. In the table, the On Server column identifies which operations can be granted for a set of managed servers, and the On Login column identifies the operations that can be granted for specific logins (users).

Table C-2: Global Shell Operations

OPERATION (PERMISSION)	DESCRIPTION	ON SERVER	ON LOGIN
`launchGlobalShell`	Launches the Global Shell.	No	No
`loginToServer`	Opens a shell session on a Unix server. In the SA Client, this is the Remote Terminal feature that opens a terminal window for a Unix server.	Yes	Yes
`readServerComplus`	Reads COM Plus objects as a specific login. In the SA Client, use the Server Explorer to browse these objects on a Windows server.	Yes	Yes
`readServerFilesystem`	Reads a managed server as a specific login. In the SA Client, use the Server Explorer to browse the file system of a managed server.	Yes	Yes
`readServerMetabase`	Reads IIS Metabase objects as a specific login. In the SA Client, use the Server Explorer to browse these objects on a Windows server.	Yes	Yes
`readServerRegistry`	Reads registry files as a specific login. In the SA Client, use the Server Explorer to view the Windows Registry.	Yes	Yes

Table C-2: Global Shell Operations (continued)

OPERATION (PERMISSION)	DESCRIPTION	ON SERVER	ON LOGIN
`relayRdpToServer`	Opens an RDP session on a Windows server. In the SA Client, this is the Remote Terminal feature that opens an RDP client window for a Windows server.	Yes	No
`runCommandOnServer`	Runs a command or script on a managed server using a `rosh` operation, where that command or script already exist. In the SA Client, this is used for Windows Services when you use the Server Explorer.	Yes	Yes
`runTrustedOnServer`	Only for internal use by HP Server Automation. Do not use this operation. HP Server Automation uses this operation for scripts in `/opsw/Script/Shared`, which implement certain HP Server Automation features. These scripts are provided with HP Server Automation and cannot be created or modified by users.	Yes	Yes
`writeServerFilesystem`	Modifies files on a managed server as a specific login. In the SA Client, use the Server Explorer to modify the file system of a managed server.	Yes	Yes

rosh Utility

The Remote SA Shell (`rosh`) command makes a client connection that enables you to remotely run programs on managed servers. You invoke the `rosh` command from within a Global Shell session.

rosh Syntax

For servers, the `rosh` command has the following syntax:

```
rosh (-n server-name | -i server-id)[-d dir] [-l login-name]
```

```
[-s] [-t | -T] [command [arg ...]]
```

For network devices, the `rosh` command has the following syntax:

```
rosh (-n device-name | -i device-id) [-N] [-C comment]
[-L] [-P parameters] [-s] -[-V variables] [command [arg ...]]
```

Table C-3 describes the `rosh` options of the preceding syntax statements.

Table C-3: rosh Options and Commands

OPTION	DESCRIPTION
`-C comment`	A comment for the log of a network script invocation.
`-d` *dir*	Sets the working directory (path) on the remote server. The default is the remote user's home directory.
`-i` *server-id*	Specifies the server by its ID, which must already exist in the `/opsw/.Server.ID` directory.
`-l` *login-name*	Specifies the login name of the remote user who performs operations on a remote server, which must already exist in the `/opsw/Server` directory.
`-L`	The network script should be run line-by-line.
`-m`	Network device mode.
`-n` *server-name*	Specifies the server by its name, which must already exist in the `/opsw/Server` directory.
`-N`	The -i or -n option refers to a network device instead of a server.
`-P` *parameters*	Parameters for a network advanced script.
`-r`	Relays RDP data to a managed server (on Windows).
`-s` *script-name*	Treats a command as the name of a saved script that will be sent to and run on the remote server.
`-t`	Forces the remote session to run in a pseudo terminal (for Unix servers only).
`-T`	Forces the remote session to run without a pseudo terminal (for Unix servers only).
`-V` *variables*	Variables for a network command or advanced script.
`-w` *seconds*	Inactivity timeout.

Table C-3: rosh Options and Commands (continued)

OPTION	DESCRIPTION
`-W` *`seconds`*	Overall timeout.

`command`
`[args . . .]`
	Runs a program or saved script.

rosh Usage Rules

The following usage rules apply to the `rosh` program:

- Specify either the `-n` or `-i` option to log into or run programs on a managed server. These options are mutually exclusive, but if both are specified the `-i` option has precedence.
- If neither the `-n`, `-i,` and `id` options are specified, the managed server can be inferred if your working directory is at or below:

  `/opsw/Server/.../`*`server-name`*`/`

  Or

  `/opsw/.Server.ID/`*`server-id`*`/`
- If `-r` is specified, no other option (excluding `-n` or `-i`) can be specified.
- If `-l` is not specified, the login-name can be inferred if your working directory is at or below:

  `/opsw/Server/.../`*`server-name`*`/files/`*`login-name`*`/`

  Or

  `/opsw/.Server.ID/`*`server-id`*`/files/`*`login-name`*`/`
- If `-s` is specified and *`command`* is a saved shared script with a `setuid` policy, the `login-name` specified by the `-l` option will be overridden. In this case, the `-l` option may be omitted. These scripts are stored in `/opsw/Script/Shared`.
- If your working directory is not below *`server`*`/files/`*`login-name`* and `-d` is not specified, the `cwdpath` defaults to the home directory for *`login-name`*. To default to the home directory, you must specify `-l`.
- For network scripts, if your current working directory is below a network device directory in the OFGS, you do not need to specify the device with `-N`, `-n` or `-i`. The network device is implied by the current working directory.

- For network scripts, if the full path of the script is not specified, `rosh` uses the search path indicated by the `NETWORK_SCRIPT_PATH` environment variable. If this variable is not set, `rosh` searches for the script in these directories:

  ```
  /opsw/Script/Network/Command/
  /opsw/Script/Network/Diagnostic/
  /opsw/Script/Network/Advanced/
  ```

rosh Operations

The `rosh` command establishes a client connection which enables you to remotely run programs on managed servers. The SA Global Shell feature provides the following modes of operation for `rosh`:

- **jump**: This operation starts a shell session in a pseudo-terminal on a managed server. This mode operates when you do not use the `-s` option and when you do not specify a command or a script. You must have the `loginToServer` permission on the managed server to jump.
- **reach**: This is a remote execution of commands that are native to the platform (operating system) of the managed server. This mode operates when you specify a command. You must have the `runCommandOnServer` permission on the managed server to reach.
- **push**: This is a remote execution of a script on a managed server. The script is stored in the OGFS and is sent to the managed server by `rosh`. You must have the `runCommandOnServer` permission on the managed server.

rosh Examples

The following examples illustrate what these operations look like for an SA user named `psi` at this path:

```
/opsw/Server/@/salish.snv1.corp.opsware.com/files/root/etc

[psi@m168 etc](538) $ uname -n; id; pwd
m168.dev.opsware.com
uid=59796(psi) gid=59796(psi) groups=59796(psi)
/opsw/Server/@/salish.snv1.corp.opsware.com/files/root/etc
```

The `rosh` jump command would display the following information about the managed server:

```
[psi@m168 etc](539) $ rosh
[root@salish etc]# uname -n; id; pwd
salish.snv1.corp.opsware.com
```

```
uid=0(root) gid=0(root)
groups=0(root),1(bin),2(daemon),3(sys),4(adm),6(disk),10(wheel)
,12(mail),7(lp),4(adm),9(kmem),6(disk),5(tty),3(sys),2(daemon),
8(mem)
/etc
[root@salish etc]# logout
```

The `rosh` reach command displays the following information about the managed server:

```
[psi@m168 etc](541) $ rosh "uname -n; id; pwd"
salish.snv1.corp.opsware.com
uid=0(root) gid=0(root)
groups=0(root),1(bin),2(daemon),3(sys),4(adm),6(disk),10(wheel)
,12(mail),7(lp),4(adm),9(kmem),6(disk),5(tty),3(sys),2(daemon),
8(mem)
/etc
```

The `rosh` push command displays the following information about the managed server:

```
[psi@m168 etc](544) $ cat /tmp/who.sh
#!/bin/sh
uname -n
id
pwd

[psi@m168 etc](543) $ rosh -s /tmp/who.sh
salish.snv1.corp.opsware.com
uid=0(root) gid=0(root)
groups=0(root),1(bin),2(daemon),3(sys),4(adm),6(disk),10(wheel)
,12(mail),7(lp),4(adm),9(kmem),6(disk),5(tty),3(sys),2(daemon),
8(mem)
/etc
```

The following example runs a script on a network device:

```
$ cd /opsw/Network/@/sw-ee-1-2b
$ rosh -s -C 'Updating device location' \
-V 'Location=Opsware - Sunnyvale' 'Set Location'
run script task 8725081 completed successfully.
Results:
Script 'Set Location for Cisco IOS configuration (for drivers:
Cisco switches, Catalyst 2950, 3550, 3750 & 8500 series, IOS
version 12.x)' completed.
```

swenc Utility

The `swwenc` command enables you to switch the character encoding within a Global Shell session.

swenc Syntax

The `swenc` command has the following syntax:

```
swenc [-e encoding] [-T {on | off}] [-E] [-x] [-c command]
```

Table C-4 describes the `swenc` options.

Table C-4: swenc Options

OPTION	DESCRIPTION
`-c` *`command`*	Executes *`command`* and exits, reverting the session encoding to its previous state.
`-E`	Lists the valid character encodings.
`-e` *`encoding`*	Changes the character encoding of the current session.
`-T {on | off}`	Turns on or off the transcoding of data from the Unix managed server. (Data from Windows servers does not need to be transcoded.) See “Transcoded Data in a Managed Server” on page 396.
`-x`	Prevents the launching of a sub-shell.

swenc Usage Rules

The following usage rules apply to the `swenc` utility:

- If you specify no options, `swenc` displays the character encoding and transcoding mode of the current session.
- Unless you specify the `-x` option, `swenc` starts a new sub-shell, which uses the encoding specified with the `-e` option. To leave the sub-shell and revert to the previous encoding, enter `exit`.

- Changing the encoding with `swenc` affects all processes in the current session, including background processes. If you change the encoding while background processes are running, the background processes might encounter errors.
- The `swenc` command affects only the current Global Shell session. For example, if you run the `rosh` command after the `swenc -e` command, the `rosh` command does not inherit the encoding that you changed with the `swenc` command.
- The `swenc` command does not change the working directory of the session, unless the working directory contains path names that cannot be represented in the new encoding. In this case, the working directory is the user's home directory.
- If you change the character encoding, make sure that the encoding of the terminal application that hosts the Global Shell and Remote Terminal sessions is set properly. To view or change the terminal client, go to the Terminal and Shell Preferences of the SA Client.

Appendix D: OGFS Directories

IN THIS APPENDIX

This section discusses the following topics:

- Directories in the OGFS
- root (/) Directory
- /opsw Directory
- /opsw/Server Directory
- Other Directories Under /opsw
- Network Directories

Directories in the OGFS

The SA Global File System (OGFS) is accessible from within a Global Shell session.

Many directories contain similar files. The `id` file contains the HP Server Automation unique identifier (primary key) for the object represented by the directory. The `attr` directory contains text files that describe the attributes of the managed server. The `info` file is deprecated in HP Server Automation 6. Instead of the `info` file, use the files in the `attr` directory. The `self` file represents this specific server object (instance). The `method` directory contains executables that invoke the methods in the SA API. For details on the executables in the `method` directory, see the *Server Automation Platform Developer's Guide*.

In the directory listings that follow, the italicized text represents variable paths (specific instances of objects in the data model). For example, *server-1* is the name of a specific managed server. The italicized text in parentheses contains comments.

root (/) Directory

At the `root` level of the OGFS is a directory for each SA user. Each user directory and all files and directories under it are visible only to processes in an authenticated session.

Each SA user has the following private directories:

- A home directory which is at /home/*user-name* (SA user name)
- Temporary directories located at /tmp, /var/tmp and /usr/tmp

Each user's home directory contains a public directory (/home/*user-name*/public) which is readable by all other SA users and can be used to share files with other SA users.

The root directory has the following structure:

```
/. (root)
   bin/
   dev/
   etc/
   home/
   lc/
   lib/
   opsw/
   opt/
   proc/
   sys/
   tmp/
   usr/
   var/
```

/opsw Directory

The opsw directory represents the data model of HP Server Automation. For example, the /opsw/Customer directory represents customer objects in the data model. Global Shell scripts and other client applications can navigate within the data model in the OGFS. Except for the api and bin directories, all of the directories under the opsw directory represent objects. The api directory contains executables that invoke methods on the HP Server Automation API. The bin directory contains Global Shell utilities such as rosh and aaa.

For HP Server Automation, the top level of the opsw directory has the following structure:

```
/opsw/
      api/
      Application/
      .Application.ID/
      bin/
      Customer/
```

```
        .Customer.ID/
        Facility/
        .Facility.ID/
        Group/
        .Group.ID/
        Hardware/
        .Hardware.ID/
        Library/
        .Library.ID/
        NetModel/
        NetOS/
        NetType/
        Network/
        .Network.ID/
        OS/
        .OS.ID/
        Realm/
        .Realm.ID/
        Script/
        Server/
        .Server.ID/
        ServiceLevel/
        .ServiceLevel.ID/
```

The ID directories organize objects by their unique SA identifier. These directories are equivalent to those organized by name. For example, if a server with an ID `10043` is named `m44.opsware.com`, then the following directories contain the same information:

```
    /opsw/.Server.ID/10043
    /opsw/Server/@/m44.opsware.com
```

If the NA is installed, the `/opsw` directory also contains several network directories. See "Network Directories" on page 543.

/opsw/Server Directory

The `Server` directory not only contains information about managed servers, but also organizes the servers by their associated objects. For example, the `/opsw/Server/@Group` directory organizes servers by device group.

At the top level, the /opsw/Server directory has the following structure:

```
/opsw/Server/
             @/
                server-1/
                server-2/
```

```
   . . .
@Application/
   Application Servers/
   Database Servers/
   . . .
@Customer/
   customer-1/
   customer-2/
   . . .
@Facility/
   facility-1/
   facility-2/
   . . .
@Group/
   Private/
     private-group-1/
     private-group-2/
   . . .
   Public/
     public-group-1/
     public-group-2/
   . . .
@Hardware/
@OS/
@ServiceLevel/
```

The sections that follow list some of the directory structures under `/opsw/Server`.

/opsw/Server/@ Directory

This directory contains all managed servers. The `files` directory reflects the file system of the managed server. When you modify the file system beneath the `files` directory, you do so as a specific user (such as `login-1`) of the managed server.

The `Interface` directory contains information about the network interface of the server. This directory might contain symbolic links to a network device. See "/opsw/Network Directory" on page 544 for more information on these links.

The paths in the following directory structure are under `/opsw/Server/@`. For example, the full path name of *`server-1`* in the following structure is `/opsw/Server/@/`*`server-1/`*.

```
server-1/
   attr
   ChangeLog
   complus
   CPU/
```

```
       0
       1
       . . .
   CustAttr/
      custom-attribute-1
      custom-attribute-2/
      . . .
   files/
      login-1/
         (file system as seen by login-1)
      login-2/
         (file system as seen by login-2)
      . . .
   info
   Interface/
      network-interface-1/
      . . .
   Memory/
      RAM
      SWAP
   metabase/
      login-1/
         (metabase as seen by login-1)
      login-2/
         (metabase as seen by login-2)
      . . .
   method/
   registry/
      login-1/
         (registry as seen by login-1)
      login-2/
         (registry as seen by login-2)
      . . .
   self
   Storage/
      storage-device-1/
      . . .
server-2/
   . . .
server-3/
   . . .
```

/opsw/Server/@Facility Directory

This directory filters servers according to their facility:

```
facility-1/
```

```
   @ /(all servers in facility-1)
      server-1/
      server-2/
      . . .
      @Group/
         group-1/
            @/ (all servers in both facility-1 and group-1)
               server-1/
               . . .
         group-2/
         . . .
facility-2/
. . .
```

/opsw/Server/@Group Directory

This directory does not contain network devices. This directory filters servers according to their device groups:

```
group-1/
   @/ (all servers in group-1)
      server-1/
      . . .
   child-group-1/
      @ (all servers in child-group-1)
         server-1/
         . . .
   child-group-2/
   . . .
   @Facility/
      facility-1/
         @ (all servers in both group-1 and facility-1)
            server-1/
   . . .
group-2/
   @/ (all servers in group-2)
   . . .
```

/opsw/Library Directory

The Library directory contains information about folders and the objects within folders: application policies, OS sequences, and packages. This directory has the following structure:

```
/opsw/folder-path-1/@/
                      AppPolicy/
                               app-policy-1/
```

```
                                           attr/
                                           method/
                                           self
                               app-policy-2/
                               . . .
                      attr/
                      method/
                      OSSequence/
                               os-sequence-1/
                                           attr/
                                           method/
                                           self
                               os-sequence-2/
                               . . .
                      Package/
                               package-1/
                                           attr/
                                           method/
                                           self
                               package-2/
                               . . .
                      self
/opsw/folder-path-2/
. . .
```

Other Directories Under /opsw

This section lists, in alphabetical order, the object directories under `/opsw` other than the `Server`, `Library`, and `Net*` (network) directories.

/opsw/Application Directory

This directory is deprecated in HP Server Automation 6.0.

This directory represents the SA Software Tree, a hierachical structure for organizing applications:

```
/opsw/Application/
                  Application Servers/
                     application-1/
                        @/
                          CustAttr/
                             custom-attribute-1
                             custom-attribute-2
                             . . .
                          .id
                          info
                         child-application-1/
                            @/
                             . . .
                             grandchild-application-1/
                                @/
                              . . .
                             grandchild-application-2/
                              . . .
                         child-application-2/
                         . . .
                     application-2/
                     . . .
                  Database Servers/
                     application-1/
                     . . .
                  Operating System Extras/
                     application-1/
                     . . .
                  Other Applications/
                     application-1/
                     . . .
                  System Utilities/
                     application-1/
                     . . .
                  Web Servers/
                     application-1/
                     . . .
```

/opsw/Facility Directory

Typically, a facility identifies the geographical location of a data center, such as a city or building. This directory contains information about facilities and the servers they manage:

```
/opsw/Facility/
               facility-1/
                  @/
                    attr
                    CustAttr/
                       custom-attribute-1
                       custom-attribute-2
                       . . .
                    info
                    method
                    self
                    Server/
                       server-1/
                       server-2/
                       . . .
               facility-2/
                   . . .
```

/opsw/Group Directory

This directory represents device groups. This directory also contains groups for network devices under the following conditions:

- NA is installed.
- The SA group has the NA-associated attribute set.
- The SA and NA group names are the same.

In the following structure, the `child-group` and `grandcdhild-group` reflect nested device groups:

```
/opsw/Group/
            Public/
               group-1/
                  @/
                    CustAttr/
                       custom-attribute-1
                       custom-attribute-2
                        . . .
                  info
                  Server
                      server-1/
                      server-2/
                      . . .
                  child-group-1/
                     @/
                       . . .
                       grandchild-group-1/
                          @/
                          . . .
                       grandchild-group-2/
                     . . .
                  child-group-2/
                  . . .
               group-2/
               . . .
            Private/
               group-1/
               . . .
               group-2/
               . . .
```

/opsw/Library

This directory corresponds to the Library of folders displayed by the SA Client.

```
/opsw/Library/ct
             folder-1/
               @/
                 attr
                 method
                 OSSequence
                 self
                 Software
                 SoftwarePolicy
               subfolder-1/
                 @/
                 . . .
               subfolder-2/
                 . . .
             folder-2/
               @/
               . . .
```

/opsw/OS Directory

This directory contains information about the operating systems defined in SA. It does not contain the actual bits for the operating systems, which are stored in the OS Provisioning Media Server. The `/opsw/OS` directory has the following structure:

```
/opsw/OS/
         os-name-1/
            Not Assigned/
             @/
                CustAttr/
                   custom-attribute-1
                   custom-attribute-2
                   . . .
                info
                Serv-1/er
                    server-1/
                    server-2/
                      . . .
            os-name-2/
               . . .
```

/opsw/Permissions Directory

This directory contains information about the Global Shell permissions. The `info` directory contains a file for each operation (permission) that can be granted with the `aaa` utility. For each user group, the `operations` subdirectory contains a text file corresponding to an operation (such as `launchGlobalShell`) that have been granted to the group by `aaa`. The contents of these text files summarize the parameters of the permissions for that operation and user group.

For example, the `readServerFilesystem` text file is in the `operations` directory of the `Advanced Users` group:

```
/opsw/Permissions/UserGroups/Advanced Users/operations/
readServerFilesystem
```

The `readServerFilesystem` text file contains the following information:

```
Facility: C40(40) Login: sysadmin
Group: Unix Servers (2880040) Login: root
```

In this example, all members of the `Advanced Users` group can read the file system as user `sysadmin` on servers that belong to `Facility C40` (with ID `40`), and as user `root` on servers that belong to the device group `Unix Servers` (with ID `2880040`).

If a user group directory does not contain a text file for an operation, then the group does not have permission to perform that operation. If the text file is empty, the user group has permission to perform the operation, but the operation has no parameters. The `launchGlobalShell` operation, for example, has no parameters.

The `/opsw/Permissions` directory has the following structure:

```
/opsw/Permissions/
                 info/
                      launchGlobalShell
                      loginToServer
                      readServerComplus
                      readServerFilesystem
                      readServerMetabase
                      readServerRegistry
                      relayRdpToServer
                      runCommandOnServer
                      runTrustedOnServer
                      writeServerFilesystem
                 UserGroups/
                           user-group-1/
                              description
                              operations/
```

```
                                                    launchGlobalShell
                                                    loginToServer
                                                    . . .
                                       user-group-2/
                                       . . .
```

/opsw/Realm Directory

A realm is a logical name for a group of IP addresses that can be contacted by a particular set of SA Gateways. Typically, each Satellite and Facility has one or more distinct realms. To find out which managed servers belong to a realm, note the server names below the *realm*/@/Server subdirectory.

The Realm directory has the following structure:

```
/opsw/Realm/
            realm-1/
              @/
                info
                Server/
                   server-1/
                   server-2/
                   . . .
            realm-2/
               . . .
```

/opsw/Script/Shared Directory

This directory contains utility scripts that are included with HP Server Automation. These scripts are not the same as the DSE shared scripts that are accessible with the SAS Web Client. The contents of this directory cannot be changed by end users.

```
/opsw/Script/
             Shared/
                    script-1/
                       description
                       policy
                       source
                       version
                    script-2/
                    . . .
```

/opsw/ServiceLevel Directory

This directory is deprecated in HP Server Automation 6.0.

Service levels are user-defined categories such as Silver, Gold, and Platinum. The `/opsw/ServiceLevel` directory has the following structure:

```
/opsw/ServiceLevel/
                   service-level-1/
                      @/
                        CustAttr/
                           custom-attribute-1
                           custom-attribute-2
                           . . .
                        .id
                        inf-1/o
                            . . .
                       child-service-level-1/
                         @/
                         . . .
                          grandchild-service-level-1/
                            @/
                            . . .
                       child-service-level-2/
                      . . .
                      Server/
                             @/
                              . . .
                   service-level-2/
                   . . .
```

Network Directories

For NA, the top level of the `opsw` directory includes the following subdirectories:

```
/opsw/
      . . .
      NetModel/
      NetOS/
      NetType/
      Network/
      .Network.ID/
      . . .
      Script/Network
```

/opsw/Network Directory

This directory organizes the devices by their model, OS, type, and group. All network devices are in the `/opsw/Network/@` directory. The `Changelog` directory contains time-stamped events for the device. The `Config` directory contains time-stamped configuration files.

At the top level, the `/opsw/Network` directory has the following structure:

```
/opsw/Network/
             @/
                device-1/
                   attr/
                   ChangeLog/
                   Config/
                   info
                   method/
                   Module/
                   Port/
                       port-1/
                          .id
                           info
                           Link/
                       port-2/
                       . . .
                   self
                   Vlan/
                device-2/
                . . .
             @Group/
             @NetModel/
             @NetOS/
             @NetType/
```

The `Port` and `VLAN` directories contain Layer 2 information. Within each subdirectory of the `Port` directory, an `info` file contains duplex information and other data for a specific network port. The `Link` directory contains a symbolic link to a server interface or a port on another network device (if the MAC address to the interface or port is available). In the following example, the `eth0` under `Network` is a symbolic link to `eth0` under `Server`:

```
/opsw/Network/@/sw-ee-1-2b/Port/FastEthernet0_1/Link/eth0
     symbolic link to -->
/opsw/Server/@/m180.mycomp.com/Interface/eth0
```

Similarly, the `Server` directory can have a symbolic link to the correpsonding entry under `Network`:

```
/opsw/Server/@/x.mycomp.com/Interface/eth0/Link/FastEthernet0_7
      symbolic link to -->
/opsw/Network/@/sw-ee-2-4a/Port/FastEthernet0_7/eth0
```

/opsw/Network/@Group Directory

This directory does not contain groups for managed servers. This directory filters network devices according to their groups:

```
group-1/
   @/ (all devices in group-1)
      device-1/
      . . .
   @NetModel/
   @NetOS/
   @NetType/
   child-group-1/
      @ (all devices in child-group-1)
         device-1/
         . . .
   child-group-2/
   . . .
group-2/
   @/ (all devices in group-2)
   . . .
```

/opsw/NetModel Directory

This directory contains subdirectories for vendors, for example, `/opsw/NetModel/Cisco`. The `/opsw/NetModel` directory has the following structure:

```
vendor-1/
          model-1/
             @/
             @Group/
             @NetOS/
             @NetType/
          model-2/
. . .
vendor-2/
. . .
```

/opsw/NetOS Directory

This directory organizes devices by their operating system. The `/opsw/NetModel` directory has the following structure:

```
family-1/
     os-1/
        @/
        @Group/
        @NetOS/
        @NetType/
     os-2/
      . . .
. . .
family-2/
. .
```

/opsw/NetType Directory

This directory organizes network devices by the following types:

```
Firewall
L3Switch
L4to7Switch
Proxy
Router
Switch
unknown
VPN
Wireless Access Point
WirelessAP
```

The `/opsw/NetType` directory has the following structure:

```
type-1/

      @/
      @Group/
      @NetOS/
      @NetType/
type-2/
. . .
```

/opsw/Script/Network Directory

This directory contains utility scripts for network devices. For information about these scripts, see the NA documentation.

```
/opsw/Script/Network/
                     Advanced/
                     Command/
                     Diagnostic/
```

Appendix E: Custom Extensions

IN THIS APPENDIX

This section discusses the following topics:

- SA Custom Extensions

SA Custom Extensions

Hewlett Packard Professional Services can extend functionality for customers by creating custom extensions to HP Server Automation (SA). Custom Extensions (which are custom Command Engine scripts) extend SA functionality to cover specific customer needs.

In HP Server Automation, the Command Engine is a system for running distributed programs across many servers (utilizing the Server Agents running on the servers). HP Server Automation features, such as the Code Deployment feature, use Command Engine scripts to implement part of their functionality

The Custom Extension feature is accessed through the Run Custom Extension Wizard. This Wizard allows a user to choose a custom extension to run, specify necessary input data for the extension, validate the data required to run the extension, run the extension, and view or download the results from the job. When a user runs a custom extension, the job shows up in My Jobs.

When a custom extension is added to one facility, it is automatically propagated to the other facilities in the multimaster mesh of HP Server Automation facilities.

To access the Run Custom Extension Wizard, users must be assigned to a user group that has the permission Wizard: Opsware Extension. When users have this permission, they can run any custom extensions on the servers they have access to in the SAS Web Client.

Running a Custom Extension

To run a custom extension, perform the following steps:

1 From the SAS Web Client home page, click the Run Custom Extensions link in the Tasks panel.

Or

From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers list appears. Select the servers on which you want to run a custom extension and choose **Run Extension** from the **Tasks** menu.

The Run Custom Extension Wizard appears.

2 Select the custom extension that you want to run and click **Next**.

If you have already selected servers from the Manage Servers list, you must ensure that the custom extension that you select can run on the operating systems of the selected servers. Otherwise, an error message appears in this page and you cannot proceed.

Some custom extensions do not require that you select servers from the Server list. For example, the extension might prompt you in the Specify Settings step to enter server host names in a text box. If you already selected servers from the Servers list, an error message appears in the page.

If you have not already selected servers from the Manage Servers list, the Select Servers page appears.

3 If prompted, select the servers or groups on which you want to run the custom extension and click **Next**. You can find the servers that you want to run a custom extension on by browsing the list or by searching.

The Specify Settings page appears. See Figure C-1.

Figure C-1: The Specify Settings Page of the Run Custom Extension Wizard

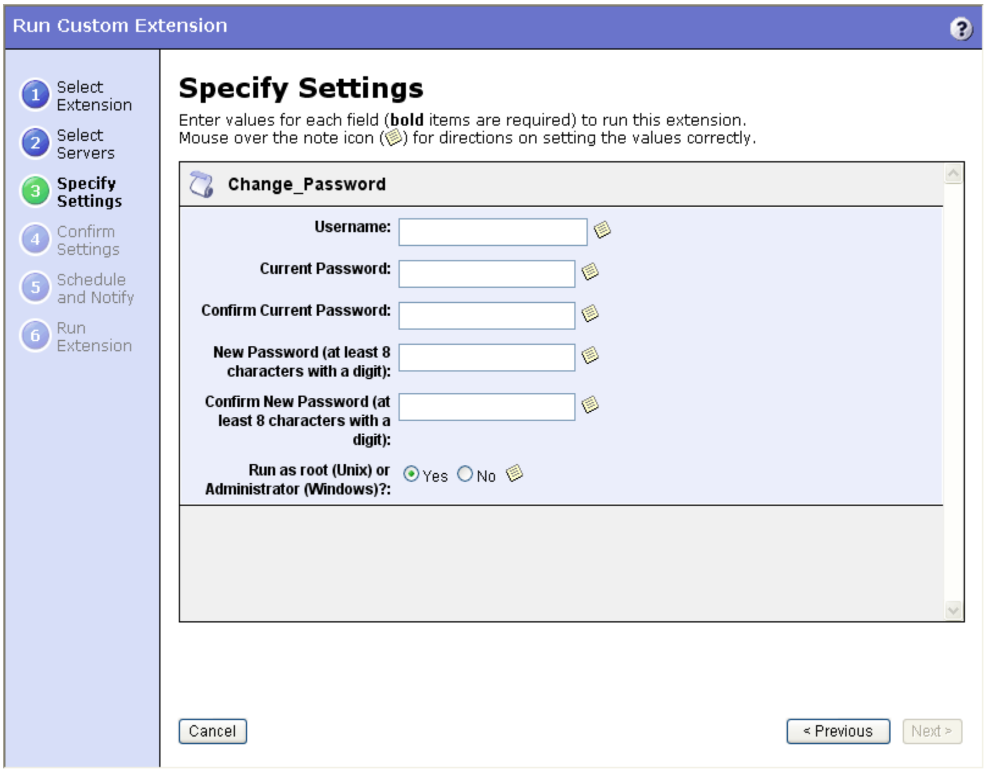

Run Custom Extension

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

Specify Settings

Enter values for each field (**bold** items are required) to run this extension.
Mouse over the note icon (✎) for directions on setting the values correctly.

Change_Password

Username: ______ ✎
Current Password: ______ ✎
Confirm Current Password: ______ ✎
New Password (at least 8 characters with a digit): ______ ✎
Confirm New Password (at least 8 characters with a digit): ______ ✎
Run as root (Unix) or Administrator (Windows)?: ● Yes ○ No ✎

4 Specify the settings for the custom extension and click **Next**.

The settings that appear in the page are unique to the custom extension that is being run. For information about what data to enter in a field, move your mouse pointer over a Note icon.

The Confirm Settings page appears. See Figure C-2.

Figure C-2: The Confirm Settings Page of the Run Custom Extension Wizard

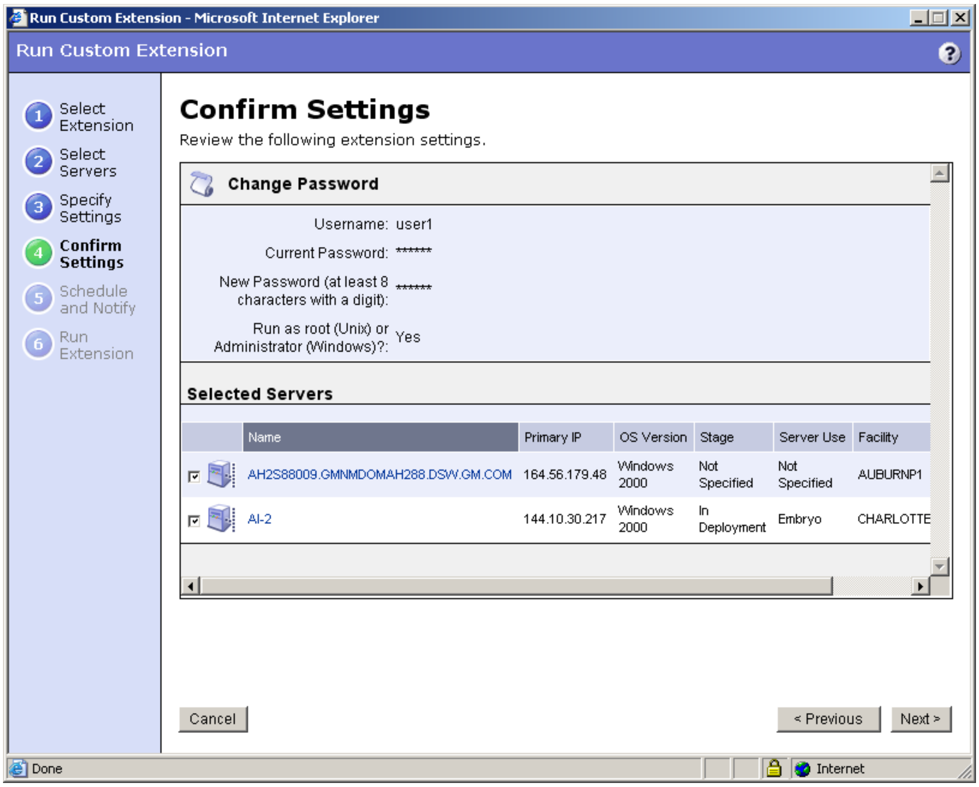

Run Custom Extension - Microsoft Internet Explorer

Run Custom Extension

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

Confirm Settings

Review the following extension settings.

Change Password

Username: user1

Current Password: ******

New Password (at least 8 characters with a digit): ******

Run as root (Unix) or Administrator (Windows)?: Yes

Selected Servers

	Name	Primary IP	OS Version	Stage	Server Use	Facility
☑	AH2S88009.GMNMDOMAH288.DSW.GM.COM	164.56.179.48	Windows 2000	Not Specified	Not Specified	AUBURNP1
☑	AI-2	144.10.30.217	Windows 2000	In Deployment	Embryo	CHARLOTTE

5 Review the values that you entered and the servers that you selected on which to run the custom extension. (You can remove servers from the list by deselecting their check boxes. The list displays the first nine servers on which the custom extension will run. Click the "Show remaining servers" link to display the complete list of selected servers.)

6 Click **Next**.

7 On the Schedule and Notify page, you have the following options:

- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

- **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:

- **Option 1**: Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

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

8 Click **Run** to start the Custom Extension Wizard.

If you selected to run the job at that time, a progress bar appears for the servers on which the extension is running that shows the progress of the job. Depending upon how the custom extension was written, you will see the progress displayed in one of two ways:

- If the custom extension was written to show progress and status for individual servers, you will see individual progress bars for each server. When the custom extension has finished running, you will see the **View Details** button. You can click **View Details** to display detailed error information.
- If the custom extension was written to show progress for all servers, then you will see a single progress bars for all servers. When the custom extension has finished running for all servers, you will see **View Details**. You can click **View Details** to display detailed error information.

If an error occurs while the custom extension is running on servers, the progress bar moves to 100% and an error message appears below the bar.

For any servers where the custom extension has failed to run, you will have to either re-launch "Run Custom Extension" from the home page and pick the servers you want to try again or choose those servers from a Manage Server ➤ My Servers ➤ Server Search list and choose the menu item **Run Custom Extension**.

9 (Optional) When the custom extension finishes running, you can click **View Details** to see the results.

The Custom Extension Results window appears. The tabs in the window can vary depending on how the custom extension was implemented. To download the results to a file, click the download link. When you are done viewing the results, click **Close** to close the window.

10 Click **Close** to end the wizard.

Closing the wizard does not stop the custom extension if it is still running. After you close the wizard, you can view the progress of the running custom extension by viewing My Jobs (accessible from the Home page or the navigation panel). Each custom extension job is identified with the name Run Custom Extension. Click the name link to identify which extension was run.

See "Server Management Scheduling and Notification" on page 332 in Chapter 8 for information about the My Jobs feature.

Index