

HP Server Automation

for the HP-UX, IBM AIX, Red Hat Enterprise Linux, Solaris, SUSE Linux Enterprise Server, VMware, and Windows® operating systems

Software Version: 9.0

User Guide: Application Automation

Document Release Date: June 2010

Software Release Date: June 2010



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.

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

© Copyright 2000-2010 Hewlett-Packard Development Company, L.P.

Trademark Notices

Intel® and Itanium® are trademarks of Intel Corporation in the U.S. and other countries.

Java™ is a U.S. trademark of Sun Microsystems, Inc.

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

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Adobe® is a trademark of Adobe Systems Incorporated.

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.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

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.

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

Contents

1 Service Automation Visualizer	19
Overview of HP Service Automation Visualizer	19
The SAV and SA Clients	19
SAV Platform Support	20
Overview of SAV Features	20
SAV Usage Examples	21
Launch SAV	21
Discover and Map Business Applications on Servers	21
View Related Networking and Storage Information	22
Define Business Application Definition	22
Troubleshoot Problems and Take Action	23
How SAV Works	23
Data Collection and Display	24
SAV Business Application	24
Launching SAV	30
Launching SAV from Servers, Devices, or Device Groups	31
Launching Business Applications from the SA Client Library	32
Launching SAV from Search Results — SA Client or SAR Client	33
Launching SAV from Generated Reports — SA Client or SAR Client	33
Launching Storage Essentials from SAV	34
SAV User Interface	34
SAV Toolbars	35
Menus and Menu Options	39
Adding and Removing Devices in SAV	39
SAV Maps	41
Tiers Map	42
Server Map	43
Network Map	46
Storage Map	48
Viewing Storage and SA Permissions	50
SAN Map	51
SAV Infrastructure Pane	52
Symbols Used in Maps	53
SAV Properties	56
Exporting Properties Information to .csv	57
Tiers Tree: Tiers, Process Family, Signature Properties	57
Devices Tree: Server and Network Device Properties	61
Storage and SAN Properties	68
SAN Link Properties	72

SAV Options	73
Virtualization Settings	73
Scan Time-Out Preference	74
Discovery Settings	74
Reset All Settings	74
Accessing Servers and Devices From SAV	74
Opening the Device Explorer	75
Opening a Remote Terminal	75
Opening a Global Shell	75
Running Scripts on Devices	76
Creating Business Application Definitions	77
Business Application Templates	77
Creating Business Application Contacts	78
Sending Email to Business Application Contacts	79
Business Application Tiers	79
Cutting and Copying a Tier	80
Pasting a Tier	80
Application and Storage Signatures	80
Signature Evaluation Order	81
Creating an Application or Storage Signature	83
Editing Signatures	84
Deleting Signatures	84
Cutting and Copying Signatures	84
Pasting a Signature	85
SAV Business Application Management	85
Opening a Business Application	85
Saving a Business Application	85
Saving a Business Application as an Application Template	86
ACLs and Server Pool Configurations	86
Viewing ACLs	86
Comparing ACLs — Two Devices In Same Snapshot	87
Comparing ACLs — Same Device Between Two Snapshots	87
Comparing Server Pool Configuration	88
Viewing Server Pool Configuration	88
Comparing Snapshots	89
Creating a Snapshot	89
Opening a Snapshot	90
Scheduling a Snapshot	90
“Source” and “Comparison” Snapshot	90
Comparison Types	91
Comparing Snapshots	93
Significant Scan Result Difference Heuristics	94
Filtering SAV Data	95
Creating a Data Filter in SAV	96
Filter Criteria	97
SAV Scan Error Messages	98
SAV Platform Support	101

Supported Platforms in SAV	101
2 Audit and Remediation	105
Overview of Audit and Remediation	105
Audit and Remediation Examples	105
Audits	107
Audit Policies	107
Audits and the Compliance View	107
Snapshots	108
Terms and Concepts	108
Audits	110
Audit Comparison Types	110
The Auditing Process	111
Audit Elements	112
Creating an Audit	113
Saving an Audit as Audit Policy	115
Viewing Server Audit and Snapshot Usage	116
Configuring an Audit	117
Audit Sources: Server, Snapshot, or Snapshot Specification	118
Server Objects Used in Audits and Snapshots	121
Audit and Remediation Rules	123
Configuration Rules: Expected (Target) and Remediation Values	124
Configuring Specific Audit and Snapshot Rules	126
Configuring Application Configuration Rule	126
Configuring COM+ Rule	131
Configuring Custom Scripts Rule	132
Configuring the Discovered Software Rule	134
Configuring the File Rule	135
Configuring Hardware Rule	138
Configuring IIS Metabase Rule	139
Configuring Internet Information Server Rule	140
Configuring IIS 7.0 Rule	141
Configuring Local Security Settings Rule	143
Configuring Registered Software Rule	145
Configuring Storage Rule	146
Configuring Windows .NET Framework Configurations Rule	147
Configuring Windows Registry Rule	148
Configuring Windows Services Rule	149
Configuring Windows/UNIX Users and Groups Rule	150
Configuring Compliance Checks	151
Renaming Compliance Checks	153
Searching for Compliance Checks from the Audit/Snapshot Specification Window	153
Managing Compliance Checks	154
Editing Compliance Check Properties	154
Creating Custom Compliance Check Categories	155
Restoring Compliance Checks to Defaults	156
Showing Deprecated Checks	157

File Inclusion and Exclusion Rules	157
Inclusion and Exclusion Rule Types	158
Example: Including all .txt Files in a Snapshot or Audit	159
Example: Including Only File a in a Snapshot or Audit	160
Example: Including last temp.txt file and exclude all else	160
File Rule Overlap	161
Parameterizing Filenames for SA/Custom Attributes	162
Using Environment Variables in Pathnames	163
Audit Rule Exceptions	164
Rules That Cannot Have Exceptions	164
Considerations When Applying Exceptions to Device Groups	164
Adding a Rule Exception to an Audit	164
Editing or Deleting a Rule Exception	165
Audit Policies	166
Linking vs. Importing Audit Policies	166
Rule Overlap with Multiple Linked Audit Policies	167
Creating an Audit Policy	167
Linking and Importing Audit Policies	168
Locating an Audit Policy in the Folder Library	171
Exporting an Audit Policy to HTML or CSV	171
Running an Audit	172
Running an Audit from the Library	172
Running an Audit on a Server from All Managed Servers	173
Re-running an Audit from Audit Results	173
Scheduling an Audit	174
Scheduling a Recurring Audit	174
Editing an Audit Schedule	175
Viewing a Completed Audit Job	176
Remediating Audit Results	176
Remediating Rules with Inherited Values	177
Accessing Audit Results	178
Audit Results Window	178
Remediation Methods: All, By Server, or By Rules	180
Viewing and Remediating Audit Results Differences	185
Viewing Audit Results with Exceptions	187
Searching for Audits	188
Deleting Audits	188
Deleting Audit Results	189
Archiving Audit Results	189
Snapshots	190
Snapshot Specification and Snapshot	190
Snapshot Used in an Audit	190
Audit Policies and Snapshot Specifications	191
Snapshot Specification Elements	191
The Snapshot Process	192
Creating a Snapshot Specification	193
Creating a Snapshot Specification from a Server	194

Creating a Snapshot Specification from the Library	194
Configuring a Snapshot Specification	194
Configuring a Snapshot Specification	195
Configuring Snapshot Specification Rules	196
Saving a Snapshot Specification as an Audit Policy	196
Running a Snapshot Specification	197
Scheduling Snapshot Jobs	198
Scheduling a Recurring Snapshot Job	198
Viewing and Editing a Snapshot Job Schedule	199
Deleting a Snapshot Job Schedule	200
Locating Snapshots	200
Searching for Snapshots	200
Archiving Snapshots	203
Deleting a Snapshot Specification	203
Deleting a Snapshot	204
Copying Objects from a Snapshot to a Server	204
Copying Objects to a Server from a Snapshot	205
3 Server Compliance	207
Overview of Server Compliance	207
Compliance Dashboard Usage: Proactive and Reactive	208
Compliance Terms and Concepts	210
Server Compliance Dashboard Categories	210
Compliance Dashboard Statuses	211
Compliance Status Thresholds — Policy, Server, and Group	213
Changing Device Group Compliance Settings	214
Viewing Compliance Dashboard in the SA Client	216
Viewing Individual Server Compliance	216
Viewing Compliance for Multiple Servers	219
Viewing Group Compliance in the Device Group Explorer	222
Adding and Removing Compliance View Columns	225
Filtering By Compliance Status	226
Refreshing For Latest Compliance Information	227
Setting Automatic Compliance Check Frequency	228
Scanning for Compliance	229
Exporting Compliance View Information	230
Compliance Dashboard Remediation	231
Group Compliance Remediation	231
Compliance Remediation for Servers	233
Audit Compliance	234
Audit Compliance Status	234
Audit Compliance Remediation	235
Software Compliance	237
Software Compliance Status	237
Software Compliance Remediation	238
Patch Compliance	241

Patch Compliance Status	241
Application Configuration Compliance.....	244
Application Configuration (App Config) Compliance Status	244
4 SA Client Reports.....	247
Overview of SA Client Reports	247
Reports Features	247
HP Server Automation Client Reports	248
User Permissions	249
Launching the Reports Feature	249
Reports Display	249
Running a Report.....	252
Modifying Report Parameters	252
Report Results Restriction	252
Report Results	253
Graphical Report.....	253
List Report.....	254
Exporting a Report	255
Printing a Report	256
5 Software Management	257
Policy-Based Software Management.....	257
SA Software Management Features	258
Overview of Software Policies	259
Managing Software Resources Using a Software Policy.....	260
The Software Management Process	261
Installing/Uninstalling Software without a Software Policy.....	262
The Install or Uninstall Software Windows	262
Install or Uninstall Software	263
Install Software Using a Software Policy	267
Attach a Software Policy to a Server	268
Attach a Server to a Software Policy	269
The Remediation Window	270
Open the Remediate Window.....	271
Remediating Policies.....	271
Uninstall Software Using a Software Policy	276
Detach the Software Policy from the Managed Server or Device Group.....	276
Remediate the Software Policy to Remove Software	277
Overview of Software Templates.....	277
Overview of Running ISM Controls	278
Open the Run ISM Control Window	278
Running ISM Controls	279
Software Policy Compliance	281
The Software Policy Compliance Scan.....	282
Software Policy Reports.....	282

6 Patch Management for Windows	287
Overview of Patch Management for Windows	287
Patch Management for Windows Features	287
Library	289
Patch Management for Windows Prerequisites	290
Windows Server 2008 Patch Management Support	291
Microsoft Patch Database	293
SA Integration	294
Support for Windows Patch Testing and Installation Standardization	294
Supported Windows Patch Types	295
Supporting Technologies for Patch Management	295
Windows Hotfixes	296
Searching for Patches and Policies	297
Roles for Windows Patch Management	297
Patch Management Process	298
Patch Properties	301
Patch Dependencies and Supersedence	302
Viewing Windows Patches	303
Editing Windows Patch Properties	303
Importing Custom Documentation for a Patch	304
Deleting Custom Documentation for a Patch	304
Finding Vendor-Recommended Windows Patches	304
Finding Servers That Have a Windows Patch Installed	305
Finding Servers That Do Not Have a Windows Patch Installed	305
Importing a Patch	305
Automatically Importing Windows Patches	306
Exporting a Windows Patch	308
Exporting Windows Patch Information	308
Deleting a Patch	309
Policy Management	310
Patch Policy	310
Patch Policy Exception	311
Precedence Rules for Applying Policies	312
Remediation Process	312
Remediating Patch Policies	313
Setting Remediate Options	314
Setting Reboot Options for Remediation	315
Specifying Pre and Post Install Scripts for Remediation	315
Scheduling a Patch Installation for Remediation	316
Setting Up Email Notifications for Remediation	317
Previewing a Remediation	317
Verifying Patch Policy Compliance	319
Creating a Patch Policy	319
Deleting a Patch Policy	319
Adding a Patch to a Patch Policy	320
Removing a Patch from a Patch Policy	320
Attaching a Patch Policy to a Server	320

Detaching a Patch Policy from a Server	321
Setting a Patch Policy Exception	321
Finding an Existing Patch Policy Exception	322
Copying a Patch Policy Exception	322
Removing a Patch Policy Exception	323
Patch Compliance.	323
Patch Compliance Scans.	323
Ways to Start a Patch Compliance Scan	323
Starting a Patch Compliance Scan Immediately	324
Refreshing the Compliance Status of Selected Servers	324
Viewing Scan Failure Details.	324
Patch Compliance Icons	325
Patch Compliance Levels	325
Patch Compliance Rules.	325
Patch Compliance Reports	326
Patch Administration for Windows.	327
Setting the Patch Availability	327
Importing the Microsoft Patch Database.	328
Selecting Windows Products to Track for Patching	328
Scheduling a Patch Compliance Scan	329
Setting the Patch Policy Compliance Level.	329
Importing Windows Patch Utilities	330
Exporting Windows Utility Files	330
Editing the Customized Patch Policy Compliance Level	330
Locales for Windows Patching.	331
Supported Locales.	331
Overview of Locale Configuration Tasks.	331
Configuring the SA Core for Non-English Locales	331
Selecting the Locales of Patches to Import	332
End User Requirements for Non-English Locales	333
Patch Installation.	333
Installation Flags	334
Application Patches	335
Service Packs, Update Rollups, and Hotfixes	335
Installing a Windows Patch	335
Setting Windows Install Options	336
Setting Reboot Options for a Windows Patch Installation	337
Specifying Install Scripts for a Windows Patch Installation	338
Scheduling a Windows Patch Installation.	338
Setting Up Email Notifications for a Windows Patch Installation	339
Previewing a Windows Patch Installation.	339
Viewing Job Progress of a Windows Patch Installation	340
Patch Uninstallation	341
Uninstallation Flags.	342
Uninstalling a Windows Patch.	343
Setting Uninstall Options	343
Setting Reboot Options for a Windows Patch Uninstallation	344

Specifying Install Scripts for a Windows Patch Uninstallation	345
Scheduling a Windows Patch Uninstallation	345
Setting Up Email Notifications for a Windows Patch Uninstallation	346
Previewing a Windows Patch Uninstallation	346
Viewing Job Progress of a Patch Uninstallation	347
7 Patch Management for Solaris	349
Overview of Solaris Patching.	349
Policy Based Patch Management	349
Quick Start to Solaris Patching.	350
Solaris Patch Workflow - Patching Specific Servers	351
Solaris Patch Workflow - Installing Specific Patches.	352
Solaris Patching Features	354
Supported Operating Systems for Solaris Patching.	355
Using Solaris Patch Policies	355
Patch Installation Order.	356
Using Patch Policies with OS Provisioning.	356
Solaris Patch Policy Management Tasks	356
Creating a Solaris Patch Policy	357
Ways to Open a Solaris Patch Policy	358
Editing Solaris Patch Policy Properties.	359
Adding Solaris Patches to a Patch Policy	360
Removing Patches from a Solaris Patch Policy	362
Resolving Solaris Patch Dependencies	362
Adding Custom Attributes to a Solaris Patch Policy	366
Deleting Custom Attributes from a Patch Policy	366
Viewing the History of a Solaris Patch Policy	367
Viewing all the Software Policies Associated with a Solaris Patch Policy.	367
Viewing OS Sequences Associated with a Solaris Patch Policy	367
Viewing Servers Attached to a Solaris Patch Policy.	367
Locating Solaris Patch Policies in Folders.	368
Solaris Patch Management Tasks	368
Running the solpatch_import Command.	369
Initializing the Solaris Patch Database.	369
Maintaining the Solaris Patch Database.	370
Finding the Right Solaris Patches	371
Automatically Importing a Solaris Patch or Patch Cluster	374
Manually Importing a Solaris Patch or Patch Cluster	374
Exporting a Solaris Patch or Patch Cluster.	375
Ways to Open a Solaris Patch	376
Viewing Solaris Patch or Patch Cluster Properties	376
Editing Solaris Patch or Patch Cluster Properties	380
Viewing Vendor Documentation for a Solaris Patch or Patch Cluster.	380
Importing Custom Documentation for a Solaris Patch or Patch Cluster	381
Viewing the Contents of a Solaris Patch Cluster	382
Viewing Patch Clusters Associated with a Solaris Patch	382
Viewing all Software Policies Associated with a Solaris Patch or Patch Cluster	382

Viewing all Patch Policies Associated with a Solaris Patch or Patch Cluster	383
Viewing Servers Associated with a Solaris Patch or Patch Cluster	383
Deleting a Solaris Patch or Patch Cluster	383
Installing Solaris Patches	384
Installing Solaris Clusters	384
Installing Manual Patches	385
Installing Solaris Patches Using a Solaris Patch Policy	385
Attaching a Solaris Patch Policy to a Server	385
Attaching a Server to a Solaris Patch Policy	386
Remediating a Server Against a Solaris Patch Policy	386
Troubleshooting Patch Installation	389
Uninstalling Solaris Patches	390
Solaris Patch Compliance	391
Performing a Solaris Patch Compliance Scan	393
Patching Solaris Zones	393
Viewing Solaris Zones	393
Installing Patches Using Offline Volumes	394
8 Patch Management for Unix	395
Overview of Patching Unix Systems	395
Tracking Patches on Managed Servers	397
Support for Unix Patch Testing and Installation Standardization	397
Viewing Patches in the SA Client	397
Searching for Patches	398
Patch Management Roles for Unix	399
Patch Management for Specific Unix Operating Systems	400
Supported Unix Versions and Patch Types	400
Underlying Technologies for Patch Management on Unix	400
AIX Patches	401
Solaris Patches	402
HP-UX Patches	402
Uploading Unix Patches into the SA Library	402
Unix Patch Information	403
The Patch Properties View	404
The Contents View	405
The Patch Depots View - HP-UX Only	405
The Patch Products View - HP-UX Only	405
The Patch Clusters View - Solaris Only	405
The LPPs/APARs View - AIX Only	406
The Software Policies View	406
The Patch Policies View	406
The Servers View	406
Viewing and Editing Unix Patch Properties	406
Finding Servers That Have a Unix Patch Installed	406
Exporting a Patch	407
Deleting a Patch	407
Using Software Policies to Manage Patches	408

Patch Compliance Reports	408
Patch Administration for Unix	409
Setting the Default Patch Availability	409
Patch Installation	410
Installation Flags	410
Application Patches	411
Installing a Unix Patch	411
Setting Unix Install Options	412
Setting Reboot Options for a Unix Patch Installation	413
Specifying Install Scripts for a Unix Patch Installation	413
Scheduling a Unix Patch Installation	414
Setting Up Email Notifications for a Unix Patch Installation	415
Previewing a Unix Patch Installation	415
Viewing Job Progress of a Unix Patch Installation	416
Patch Uninstallation	417
Uninstallation Flags	417
Uninstalling a Unix Patch	418
Setting Uninstall Options	418
Setting Reboot Options for a Unix Patch Uninstallation	419
Specifying Pre and Post Install Scripts for a Unix Patch Uninstallation	420
Scheduling a Unix Patch Uninstallation	420
Setting Up Email Notifications for a Unix Patch Uninstallation	421
Previewing a Unix Patch Uninstallation	421
Viewing Job Progress of a Patch Uninstallation	421
9 Script Execution	423
Overview of Script Execution	423
Script Execution Features	423
Script Execution Process	424
Types of Scripts	424
Managing Scripts	426
Creating a Script	426
Opening a Script in the SA Client	429
Editing Script Properties	430
Viewing All the Software Policies Associated with a Script	431
Viewing Script Version History	431
Locating Scripts in Folders	432
Exporting a Script	432
Renaming a Script	432
Deleting a Script	432
Executing Scripts	433
Ways to Open the Run Script Window	433
Running a Server Script (Saved Script or Ad-Hoc Script)	434
Running an OGFS Script	438
10 Running SA Extensions	443
Methods of Running Extensions	443

Run Extensions on Managed Servers	444
Running Extensible Discovery on Managed Servers	447
Running Extensible Discovery from the OGSF	448
Adding Scripts to Extensible Discovery	448
Scripts Provided with Extensible Discovery	449
Software Policies Provided with Extensible Discovery	449
Writing Your Own Scripts for Extensible Discovery	450
Adding Your Own Scripts to Extensible Discovery	451
Upgrading Your Scripts in Extensible Discovery	453
Removing Your Scripts from Managed Servers	454
Output from Extensible Discovery Scripts	454
Comparing Custom Fields and Custom Attributes	455
Creating and Managing Custom Fields	456
Data Types in Custom Fields	456
Creating a Custom Field with the Custom Field Management Web Extension	457
Deleting a Custom Field with the Custom Field Management Web Extension	458
11 Operating System Provisioning	459
SA OS Provisioning Prerequisites	459
Supported Operating Systems and Media for OS Provisioning	459
Permissions	461
Network Setup for OS Provisioning	461
Hardware Preparation	461
OS Build Plan Requirements	462
The OS Provisioning Process	463
Overview of the OS Provisioning Process	463
SA OS Provisioning-supplied CD Boot Images	464
Booting Servers Remotely	464
Booting from a CD	465
Network Booting a Linux or VMware ESX Server using PXE	465
Booting a Red Hat Enterprise Linux Server in a Non-DHCP Environment	466
Booting a Red Hat Enterprise Linux Itanium 64-bit Server in a Non-DHCP Environment	468
Network Booting a Windows Server Using PXE, WinPE, and WinPE/OGFS	470
Booting a Windows Server in a Non-DHCP Environment	472
Network Booting a Solaris Server	474
The Manage Boot Clients (MBC) Option	475
Specifying OS Provisioning Tasks	480
OS Build Plans	480
Affect of the OGFS Agent on Server Lifecycle	481
What are OS Build Plans?	481
Baseline OS Build Plans	482
OS Build Plan Requirements	483
Copying a Baseline OS Build Plan	483
Viewing/Modifying an OS Build Plans	485
Minimum Baseline OS Build Plan Modification	488
OS Sequences	489

OS Sequence Contents	489
Defining an OS Sequence	490
OS Sequence Prerequisites	491
Installing (Provisioning) an Operating System	494
The Unprovisioned Servers List	494
The OS Installation Profile	495
Using an OS Build Plan for OS Provisioning	495
Using an OS Sequence for OS Provisioning	498
Model Base Packages Functionality	499
Model Base Packages Script Usage	500
Reprovisioning a Managed Server	500
Advanced SA OS Provisioning Architecture	501
OS Provisioning Components	501
Build Customization Scripts	503
Booting and Provisioning Using WinPE-OGFS Boot Images	503
How the OS Build Agent Locates the Build Manager	503
Loading OS Build Agents	504
Verifying That a Server is Ready for Operating System Installation	505
Recovering when an OS Build Agent Fails to Load	505
Index	507

1 Service Automation Visualizer

Overview of HP Service Automation Visualizer

HP Service Automation Visualizer (SAV) allows you to manage the operational architecture and behavior of distributed business applications in your IT environment by displaying detailed application information in physical and logical drawings.

SAV enables you to scan selected servers or devices in your data center so you can visualize all aspects of your business applications and how it interacts with other components on your network. SAV gives you the ability to create signature-based definitions of your Business Applications and storage components, and models them in Tiers. This provides a detailed and comprehensive picture of how all a business application's components interact.

SAV's detailed picture of your business application includes all related physical and virtual servers, network and storage devices, and physical and logical connections between any of them. When you better understand a business application's processes and interrelationships, you can understand how the business applications are distributed and you are likely to be more effective in troubleshooting errors when they occur.

SAV enables you to take snapshots of a business application (on a one time or recurring basis) and compare the results, so you can view and compare differences in your business application at a specific point in time. You can compare two snapshot results to see changes that have occurred and remediate any differences in the results.

You can also view compliance information in SAV, so you can monitor server and device compliance levels and troubleshoot those that are out of compliance.

SAV is tightly integrated with features in Server Automation (SA), as well as the Network Automation (NA) and Storage Essentials (SE). The kinds of data you can visualize and tasks you can perform in SAV, however, depends upon the SA products you are licensed to run or have configured to work with SA and SAV.

The SAV and SA Clients

The Server Automation Visualizer (SAV) Client is a separately licensed product that requires SA in order to run.

In order to visualize networking information with Network Automation (NA) inside of 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.

In order to visualize SAN objects, such as arrays, switches, volumes in SAV, Storage Essentials (SE) version 6.1.1 or later is required and the Server Automation SE Connector component must be installed and configured on your SA core.

You can also visualize servers and devices in search and report results from inside the Service Automation Reporter (SAR), which is also a separately licensed product.

If you have not purchased SAV, NA, SE or SAR, but would like to, contact your sales representative.

SAV Platform Support

For the current list of support OS platforms and hardware architecture supported by SAV, see [SAV Platform Support](#) on page 101.

Overview of SAV Features

SAV enables you to perform the following tasks:

- Discover, map, and visualize the process families, connections, dependencies, and storage of multi-tiered business applications
- Visualize business applications that run on virtual servers, showing virtual servers in relationship to their hypervisors, as well as virtual switches and port groups (VMware ESX only)
- Visualize business application information in multiple physical and logical layouts, such as an application view, a server view, a network view (including virtual network devices), a storage and SAN view that displays logical and physical storage connections, and an infrastructure view that provides detailed inventory and infrastructure information related to objects scanned.
- Visualize Oracle database instances, including their tablespaces and connection to database files (including redo logs).
- Organize recognized application and storage signatures into multi-tier applications to create a logical view that can be analyzed to verify correct operation
- Map business application process families to application and storage signatures and highlight them with custom color schemes
- Create, schedule and compare snapshots of your Business Applications and all the data captured in them
- Filter business application snapshots to find exactly the data you are looking for
- Create and share business application templates that represent an ideal application definition
- Run scripts on devices or the Global File System (OGFS) on a one-time or scheduled basis
- Troubleshoot and resolve problems by launching the Device Explorer, Network Device Explorer, Global Shell, Remote Terminal, and NAS interface to perform in-depth analysis or to perform actions on the systems under investigation
- Export maps to .gif, .jpg, and .svg files
- Export tables (Properties and Infrastructure tabs) to .csv

SAV Prerequisites

In order to scan and visualize devices and relationships in SAV, the following requirements must be met:

- Server Agent version 7.0 or greater to scan and visualize managed servers from a SA core. The exceptions to this requirement is VMware ESXi Server, which does not require an SA agent for scanning and visualizing those servers in SAV
- SA core that is configured to connect to Storage Essentials (SE). For information, see the Storage Visibility and Automation documentation.
- NA 7.0 or greater server in order to scan network devices and connections

Supported Operating Systems

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

For more detailed information on SAV platform support, see [SAV Platform Support](#) on page 101.

SAV Usage Examples

Understanding how SAV functions within the context of a real datacenter is best illustrated with some general usage examples:

- [Launch SAV](#)
- [Discover and Map Business Applications on Servers](#)
- [View Related Networking and Storage Information](#)
- [Define Business Application Definition](#)
- [Troubleshoot Problems and Take Action](#)

Launch SAV

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


To gain a better picture of the application, he opens the SA Client, selects a group of servers (some of which use remote storage) that the application runs on, and launches SAV.

For information on how to launch SAV, see [Launching SAV](#) on page 30.

Discover and Map Business Applications on Servers

SAV scans the selected servers and discovers all applications, signatures, processes and process families, files systems, local and remote storage, database connections, and any other connections related to all the applications running on the selected servers. SAV displays detailed “maps” of the applications (processes and process families) and servers and connections associated with them, as well as any related network relationships and remote storage and SAN connections.

The application administrator examines this information and sees a short list of items that contains two servers, one of which is his server, and two network devices. Looking at the Server Map, he selects the box that represents his server, and a properties pane opens to display more detailed information about the server. He notices that the server has virtual machine-related information, so he concludes that his business application may be running on a virtual machine instance.

He then clicks **Show Virtual/Physical Containment Relationships**  on the SAV toolbar, and now the map shows that his server is a virtual machine running on a hypervisor. He double-clicks the hypervisor server and once it expands, he sees his server within it. He now understands that his business application runs on a VMware virtual machine (VM), and has visibility into the physical host (hypervisor) on which the VM runs. He also notices that the hypervisor server is connected to a SAN disk array, and he can see the connection between the file system on the server and the storage device.

For more information on the SAV maps, see [SAV Maps](#) on page 41.

View Related Networking and Storage Information

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

He takes a closer look at the lines emanating from his server, and notices a prominent, thick black line pointing at some IP address, so he clicks on it. He sees that the line represents 64 connections to another host on port 1433. It looks like the database that he knows his business application uses. He right-clicks the box that his server is pointing at and selects **Add Devices**. A window opens that shows the discovered database selected. He clicks the Add button and his Snapshot refreshes, this time including the new server.

Now he sees that the thick black line is pointing at the new server, and after drilling down, he discovers it pointing specifically at an SQL server process. He continues this until he finds his business application running across and depending on 10 separate servers. He also sees that the SQL server process family runs on two different file systems on the server, which are being stored on a SAN disk array. The connection to the disk array is brown, so he knows this is not a problem with the remote storage device.

Define Business Application Definition

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

His first step is to create the logical tiers of the business application. He selects the Tiers tree and creates four main tiers for the business application: Web, Application, Database, and Storage. He creates sub-tiers for authentication services and integration services. He then defines application signatures to add to each tier, specifying which tier a recognized signature should fall in. For the Storage tier, he creates storage signatures to capture any related storage on any NAS filers or disk arrays.

In order to create reusable application and storage signatures for each tier, he specifies the criteria used to recognize it, including process names, open files, listener ports, command line, environment variables, and so on.


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

Finally, he saves his business application definition so it can be reused by others who want to work with the same business application.

For more information on creating a business application, see [Creating Business Application Definitions](#) on page 77.

Troubleshoot Problems and Take Action

To help keep track of the state of a business application at any given time, the application

administrator continually clicks **Refresh Snapshot**  on the SAV toolbar in order to create new snapshots. Each snapshot can be saved to the SA Client Library or to a local system, which can be used later to compare previous snapshots of the business application with a current state to find any important differences and troubleshoot errors.


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

For more information on snapshots, see [Comparing Snapshots](#) on page 89.

How SAV Works

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


SAV scans a server (or multiple servers) and network and storage devices to gather this information and displays it in the maps and tiers, visualizing all processes and process families, connections, and devices related to the business application. Each SAV session, which can be saved as a Business Application to the SA Client Library (or to a local system), allows you to create, visualize, analyze, define, share, and troubleshoot your business applications.

Clicking **Refresh Snapshot**  allows you to scan the current state of the SAV business application and save it. These scan results (called a "Snapshot") can be compared on a one to one basis using the compare feature (activated by the Compare toolbar button).

Data Collection and Display

SAV scans devices (servers and network and storage devices) and draws maps based on data that is collected in real-time results of a SAV snapshot. Device data is captured directly from servers and then recorded in snapshots. Network device data is scanned and then recorded in scan results by NAS — where it is retrieved by the SAV from the Network Automation data model. Storage data that relates to the selected device is scanned from the SE data model.

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

When you click **Refresh Snapshot**  on the SAV toolbar, SAV creates a new snapshot that captures all the information gathered when you scan a business application (and the servers and devices it runs on) as well as your business application definitions.

SAV uses information gathered from SA, NA, and SE, leveraging the architecture to collect more data on-demand (such as processes that are running, open ports, and the number of users logged in). It also maps business application data to visualize and analyze your operational environment.

SAV collects and displays the following information about managed servers, network, and storage devices:

- Processes and process families (potentially matching application signatures) that are running on managed servers
- TCP and UDP connections between these processes
- Detailed configuration information
- Current runtime information about servers, connections, and processes
- File systems on servers and how they are used by process families, are mapped to fibre channel ports, and are reliant on local and remote storage
- Servers, interfaces, adapters, switches and vSwitches, and switch port connections
- Local and remote storage devices and how they connect to servers, SAN switches, and other storage devices

See [Processes, Process Families, and Extended Process Families](#) on page 27 for an explanation of how SAV interprets this data. See [Filtering SAV Data](#) on page 95 for instructions on how to search the data that was collected by object type, such as by process family, network interface, and so on.

SAV Business Application

A business application is a complex collection of services that typically run across multiple servers, networking (LAN and SAN), and storage devices. A business application in SAV consists of business application definitions (tiers, application and storage signatures, and properties definitions) visible in the Tiers tree, and a collection of maps that visualizes relationships between a business application's signatures, processes (and process families), file systems, storage devices, and external clients and dependencies.

A SAV business application maps to actual instances of business applications that are running on servers that SAV has scanned and displayed. A business application, as seen in the Tiers Map, is a collection of processes running on a managed server that maps to a SAV Application definition, as specified in the Tiers tree. A business application can also include storage devices and how they relate to and connect with process families running on servers.

The SAV business application is further explained in the following sections:

- [Tiers Tree](#)
- [Creating Tiers to Model Business Applications](#)
- [Application Signatures](#)
- [Processes, Process Families, and Extended Process Families](#)
- [Storage Signatures](#)

For information on how to create a SAV application, see [Running Scripts on Devices](#) on page 76.

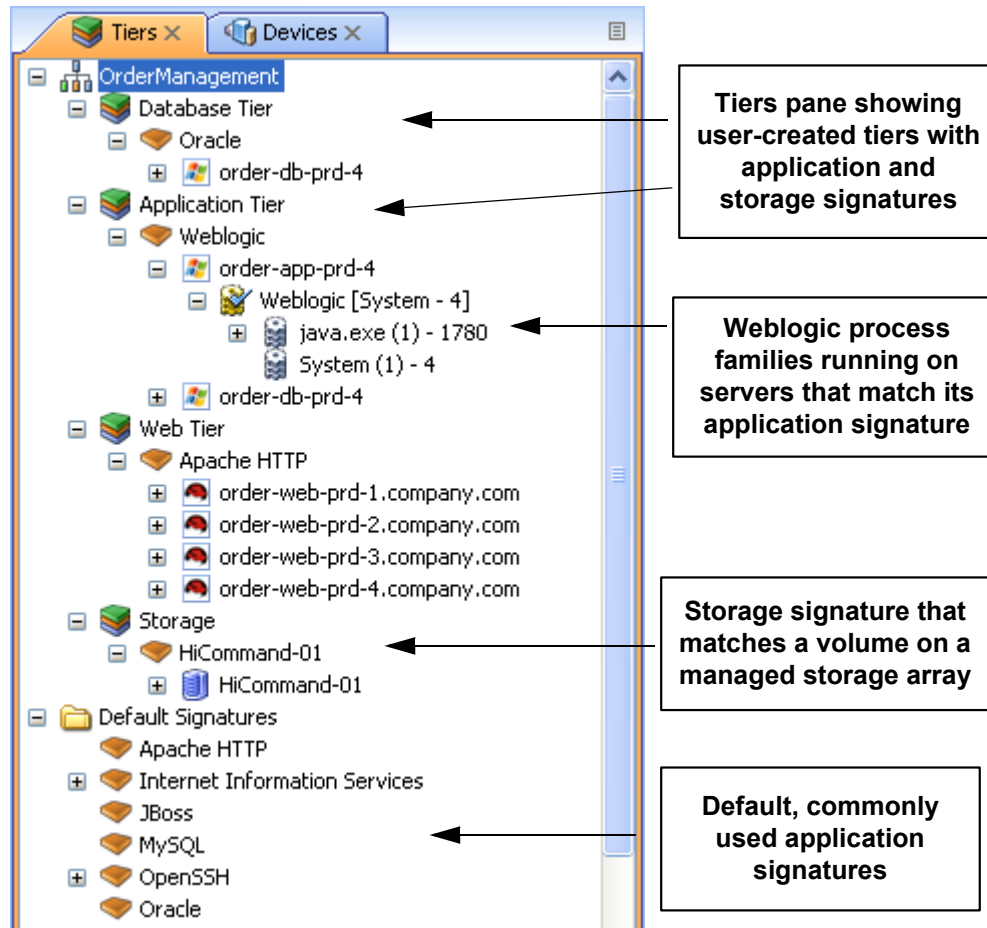
Tiers Tree

The Tiers tree is a logical view of a business application that provides a hierarchical representation of a business application's infrastructure. The Tiers tree provides a different way of looking at what is visually displayed in the Tiers map. The Tiers tree contains tiers and subtiers, which in turn can contain application and storage signatures.

Inside of each application signature is the server or device that the process families run on, and inside of the server or device is the process family itself. A storage signature is similar to an application signature except that its signature matches storage volumes rather than process families (as in an application signature). This allows you to quickly identify and categorize storage dependencies.

See [Figure 1](#) for a picture of the Tiers tree.

Figure 1 Tiers Tree




➤ Default signatures at the bottom of the Tiers tree do not appear in the Tiers map — instead, they are highlighted in the Network Map, Server Map, and Storage and SAN Maps.

If there are no matching process families for a signature, a warning icon ⚠ appears next to it, and the tiers that contain it, in the Tiers tree.

Creating Tiers to Model Business Applications

Creating tiers enables you to model the logical structure of a business application, representing all of its processes and process families as a diagram of elements that run across multiple servers, displaying the connections among them, clients connecting to them, and dependencies to which they connect. Tier definitions can contain a device filter, which restricts the servers whose process families will match the tier's application signatures; or, can contain storage devices whose volumes or filesystems will match the tier's storage signature.


Each application consists of a set of tiers and sub-tiers, such as a Web tier running Apache on Linux, an application tier running WebLogic on Windows, a database tier running Oracle on Solaris, and a storage tier to represent disk arrays and NAS filers.

A tier is represented in the Tiers tree by the  icon, which can contain application signatures, storage signatures, and optional sub-tiers.

For information on how to create business application tiers, see [Business Application Tiers](#) on page 79.

Application Signatures

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

An application signature is represented in the Tiers tree by the  icon. An application signature object consists of a signature and visual display preferences.

A signature is a set of rules that you provide and that SAV uses to identify a process family. This set of rules uses data such as process name, open files, command line, environment variables, connected to port, modules, executable path, and listener port. If SAV discovers the process or process family during a scan according to the signature rule definition, then the process or process family is added to the signature and highlighted in the maps.


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

For more information on creating application signatures, see [Creating an Application or Storage Signature](#) on page 83.

Processes, Process Families, and Extended Process Families


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

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

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

family, for example), it is represented in the Application (or Device) Tree by the  icon.

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


An extended process family is represented by the  icon.


Storage Signatures

If your scanned managed servers are configured with a storage inventory snapshot and the SE Scanner is configured to integrate with your SA core, then you can also model logical storage hierarchies to help visualize and understand how your storage devices and SAN relate to the processes used by your business application.

A storage signature is similar to an application signature except that its signature matches storage volumes rather than process families (as in an application signature). This allows you to quickly identify and categorize storage dependencies.

Storage signatures are created in the Tiers tree. The Tiers Map displays them according to modifiable settings of color and name in the signature's properties.

A storage signature is represented in the Tiers tree by the  icon. A storage signature object consists of the name of the storage volume, LUN name and ID, exported path, and any relevant manufacturer information, such as manufacturer name and model number of the related device.

Storage signatures that do not match any actual volumes in the current scan are flagged with a  warning icon.

For more information on creating storage signatures, see [Creating an Application or Storage Signature](#) on page 83 and [Examples of Regular Expressions](#) on page 98.

Devices Tree

The Devices tree is a logical, tree-based view of top-level information about managed servers, process families, and network and storage devices. This tree hierarchically displays the same top level information that is shown in the Network Map, Server Map, Storage Map and SAN Map.

The Devices tree contains servers, network and storage devices (physical and virtual) as its top nodes. Below the servers are process families and extended process families. Network devices contain VLANs, ports, and port groups (for VMware virtual switches).



VMware ESXi 3.5 hypervisor servers cannot be expanded to view process information in the Devices Tree.

Storage devices in the Devices Tree display the following elements:

- NAS filers and their exported file systems and mapped LUNs used by servers
- SAN arrays and their volumes LUN mapped to servers:
- SAN switches and their fibre channel ports.

The Devices tree also shows virtual devices that were scanned when you launched SAV. These virtual devices can be shown grouped beneath their hypervisor when the Virtualization button is selected. This tree includes the following:

- Virtual servers.
- VMware virtual switches (vSwitches). vSwitches can be expanded to view their port groups.
- Solaris Global zones can be expanded to list all running processes, but this list of processes includes processes on non-global zones not included in the current scan.

Attributes in the Properties pane for Device Tree objects contain the following:

Oracle

- Oracle executable
- Oracle database instance
- Tablespace inside the database

WebLogic

- Applications
- Web Applications
- EJBs

- JDBC Connection Pools

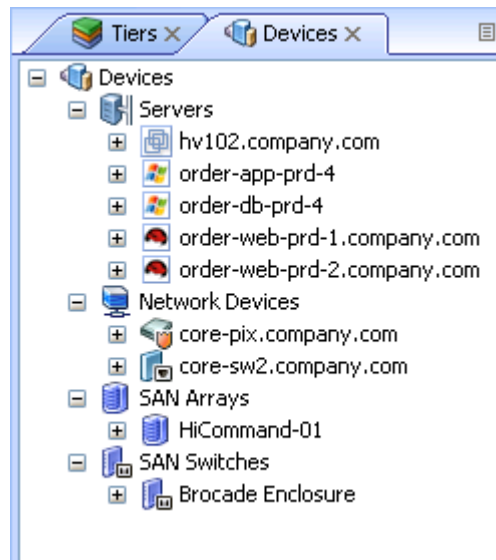
Microsoft IIS

- Web Sites
- FTP Sites
- Bindings

To view online help for these objects, select the object in the Device Tree, then select the Properties tab in the lower left of the SAV window. Then, press F1 on your keyboard.

[Figure 2](#) illustrates the Devices tree, showing servers, network devices, storage and SAN devices.

Figure 2 Devices Tree




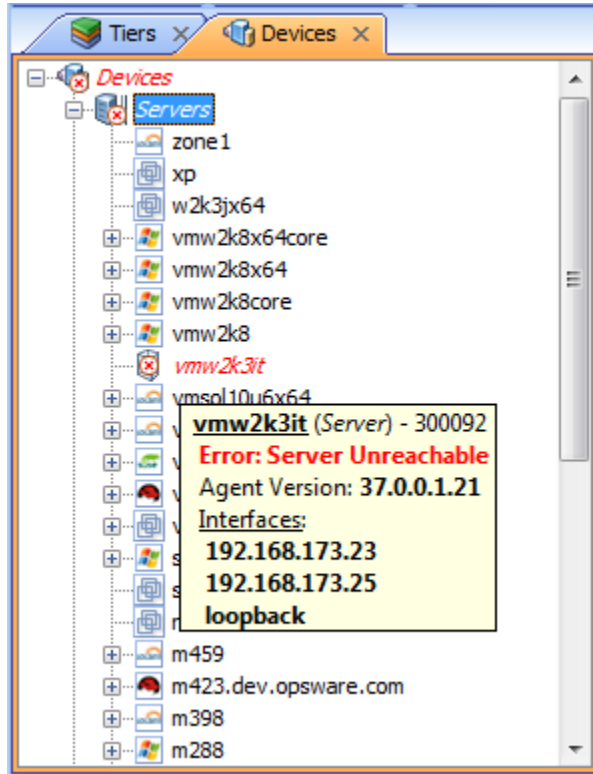
If a server device has an error associated with it, then it appears in the Devices tree with an error icon on it , for example if the server is unreachable by SA. When you move your mouse pointer over the device node in the tree, a tooltip message indicates the nature of the error, as shown in [Figure 3](#).

Figure 3 Devices Tree Server Node with Tooltip Indicating Device Scan Error



For more information on possible device errors, see [SAV Scan Error Messages](#) on page 98.

Launching SAV

You can launch SAV in several different ways:

- From a server or group of servers (using drag and drop or menus)
- From the SA Client Library
- From search or report results
- From a storage or network device (or group of devices)
- From a report or search result inside of the SAR Client

When you launch SAV, it performs an extensive scan of the servers or devices you selected — including all virtual servers and their hypervisors.

For more information about how SAV scans a server or device, see [Data Collection and Display](#) on page 24.

For more information on launching SAV from the SAR Client, see the *SAR User's Guide*.



The Allow Analyze permission is required to use SAV. You also need read access to each managed server that you plan to scan. Write access to each managed server is not required to run the SAV; however, write access is required to perform any actions on the servers, such as opening a remote terminal or running a script.


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

You can launch SAV from inside the SA Client from the following different locations:

- [Launching SAV from Servers, Devices, or Device Groups](#)
- [Launching Business Applications from the SA Client Library](#)
- [Launching SAV from Search Results — SA Client or SAR Client](#)
- [Launching SAV from Generated Reports — SA Client or SAR Client](#)

Launching SAV from Servers, Devices, or Device Groups

To launch SAV servers (virtual servers, or hypervisors), devices (servers, storage devices, or network devices), or groups of devices, perform the following steps:

- 1 Launch the SA Client from one of the following locations:
 - Click the SA Client link in the Power Tools section of the SAS Web Client home page.
 - Double-click the SA Client icon on your desktop (if you installed it on your desktop when you installed the SA Client).
 - Select **Start** menu ► **All Programs** ► **HP Server Automation Client**.
- 2 From the Navigation pane, select the Devices tree.
- 3 From the Device Groups, Servers list, or Storage list, select a device and perform one of the following actions:
 - From the **Actions** menu, select **Open with** ► **HP Server Automation Visualizer**.Or
 - Right-click, and from the menu, **Open with** ► **HP Server Automation Visualizer**.Or
 - From the **Tools** menu, select **HP Server Automation Visualizer** ► **Open Selection**.Or
 - Select the servers and drag them into an open SAV window. After doing this, click **Refresh Snapshot**  on the main toolbar so SAV can scan and display the new device.

After scanning is completed, the SAV application window appears containing the selected device or devices in the Devices tree, Tiers tree, Properties Panes, Server Map, Network Map, Storage Map, Tiers Map, and the Infrastructure pane.

If a SAV scanning process is taking too long, you can cancel. For more information on how to set the scan timeout value, see [Scan Time-Out Preference](#) on page 74.

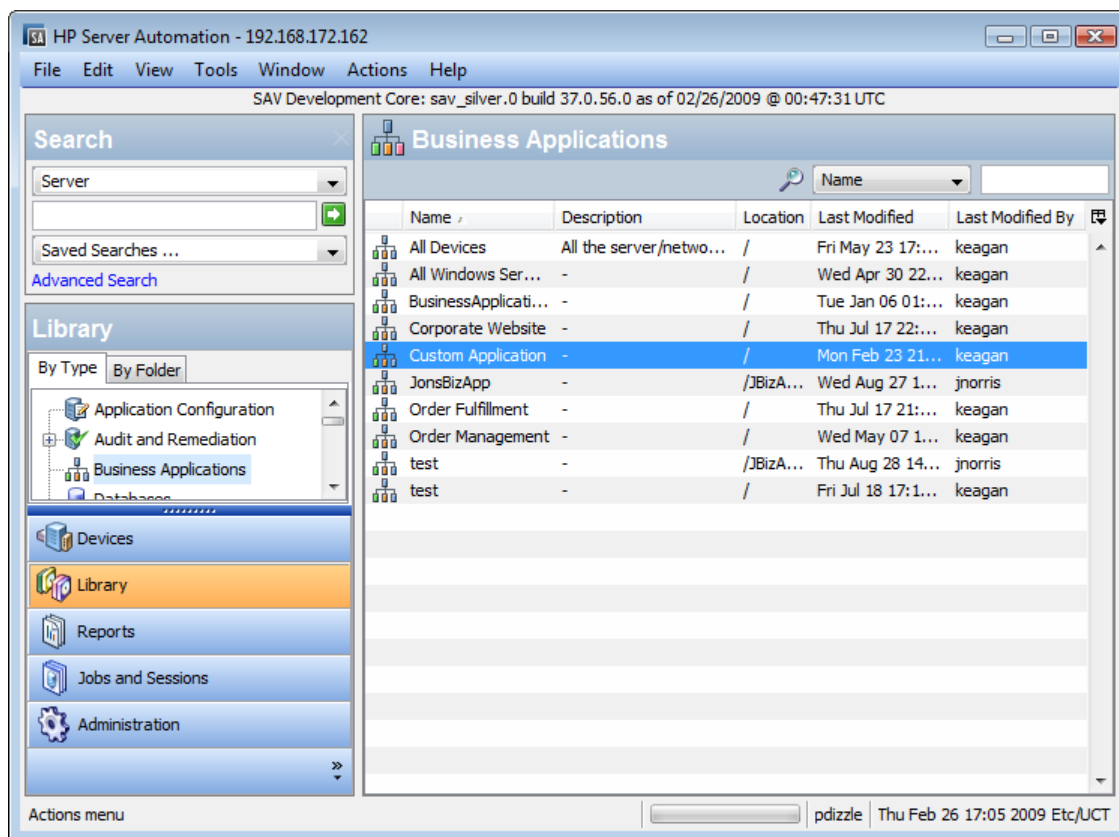


If you have selected virtual servers or a virtual server's hypervisor to open with SAV, you will initially be asked if you want to scan virtualization relationships — in other words, whether or not to scan any virtual and host servers related to the servers that you selected. This could increase the time it takes to complete the scan, depending on how many virtual servers or hypervisors are related to your selected servers. To control virtual server scan settings, see See “Virtualization Settings” on page 73.

Launching Business Applications from the SA Client Library

You can launch Business Applications from the Library. A business application is a complex collection of services that typically run across multiple servers, networking (LAN and SAN), and storage devices. A business application in the SAV consists of business application definitions (tiers, application and storage signatures, and properties definitions) visible in the Tiers tree, and a collection of maps that visualizes relationships between a business application's signatures, processes (and process families), file systems, storage devices, and external clients and dependencies.

Figure 4 Business Applications in the SA Client Library



For more information about creating and saving SAV business applications, see [Creating Business Application Definitions](#) on page 77 and [Saving a Business Application](#) on page 85.


To launch SAV from the SA Client Library, perform the following steps:

- 1 From the Navigation pane, select **Library** ► **By Type**.
- 2 Select the Business Applications object. The Contents pane on the right side shows all SAV business applications you have permissions to see.

- 3 To open a business application and launch SAV, select a business application, right-click, and select **Open**.

Launching SAV from Search Results — SA Client or SAR Client

To launch SAV from search results in either the SA Client or SAR Client, perform the following steps:

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

Or

- From the **Tools** menu, select **Server Automation Visualizer** ► **Open Selection**.

After scanning is completed, the SAV application window appears containing the selected device or devices in the Devices tree, Tiers tree, Properties pane, Server Map, Network Map, Storage Map, Tiers Map, and the Infrastructure pane.

Launching SAV from Generated Reports — SA Client or SAR Client

To launch SAV from report results from either the SA Client or the SAR Client, perform the following steps:

- 1 From the **Start** menu, select ► **All Programs** ► **HP Business Service Automation** ► **HP Server Automation**.
Or, to launch the SAR Client:
From the **Start** menu, select ► **All Programs** ► **HP Business Service Automation** ► **HP Server Automation Reporter**.
- 2 From the Navigation pane, select Reports.
- 3 Expand the Reports, and select a report that will display servers in its results.
- 4 From the report results, drill down and select an individual server or multiple servers, right-click, and select **Service Automation Visualizer**.

After scanning is completed, the SAV application window appears containing the selected device or devices in the Devices tree, Tiers tree, Properties panes, Server Map, Network Map, Storage Map, Tiers Map, and the Infrastructure pane.



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

Launching Storage Essentials from SAV

If your SA core has been configured with Storage Essentials, you can launch Storage Essentials from a SAN array inside of SAV.

To launch Storage Essentials from inside of SAV, perform the following steps:

- 1 From inside the SAV application window, select the SAN map.
- 2 Select a storage array, right-click, and select **Open HP Storage Essentials**.


SAV User Interface

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

- Maps that display the physical and logical and virtual layouts of applications — see [SAV Maps](#) on page 41.
- Trees that display the physical and logical layouts of applications — see [Tiers Tree](#) on page 25 and [Devices Tree](#) on page 28.
- Properties panes that provide granular information about a selected object, signature, process, or connection — see [SAV Properties](#) on page 56.
- Detailed tables for comparison of objects — see [Comparing Snapshots](#) on page 89.
- Dynamic tool bars and detailed tooltips to provide more information about tree and map objects.

[Figure 5](#) shows the types of information that the SAV displays.

Figure 5 Service Automation Visualizer User Interface



SAV Toolbars

The SAV toolbars allow you to open, close, resize, and organize different layout views and trees, as well as execute scripts, launch the Global Shell, create a SAV Snapshot, compare Snapshots, and more.

Depending on the tree and view selected, certain toolbar icons will be unavailable. See [Table 1](#) for a description of SAV's toolbar icons.

Table 1 Toolbar Icons in SAV












Toolbar Icon	Description
Main Toolbar	
	Refreshes the scan results by collecting and displaying new information. Each time you click this button, SAV creates a new SAV snapshot that gets saved as part of the Business Application, and can be used in a snapshot comparison.
	Opens a previously saved .vam or .vat file, or Business Application from the SA Library.
	Saves the current business application (including maps) as a .vam or .vat file in your local file system, in the SA Client Library, or in the Global File System (OGFS). If the business application has not been previously saved, the Save As window displays.
	Prints the selected map. Displays the Print window where you specify page setup (including printing across multiple pages), a title for the printed map, and so on.
	Cuts a selected business application or storage signature or a selected tier in the Tiers tree and saves it to the clipboard.
	Copies a business application component or storage signature in the Tiers tree and saves it to the clipboard.
	Deletes a selected application tier or signature in the Tiers tree or Tiers Map.
	Opens the Device Explorer for the selected device — server, storage device, network device.
	Opens the compliance view of the server (in SA Client) or network device (in NA Client).
	Opens NA for the selected network device (if your core is NA-enabled).
	Opens the Open Remote Terminal window where you select a login ID for a Remote Terminal.

Table 1 Toolbar Icons in SAV (cont'd)



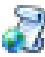




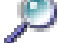





Toolbar Icon	Description
	Launches the Run Script window in the SA Client (for servers) or the NA Client New Task - Run Command Script page (for network devices).
	Opens a Global Shell session.
	Launches the Run OGFS Script window.
	Starts the selected virtual server (VM or Solaris local zone).
	Stops the selected virtual server (VM or Solaris local zone).
	Pauses the selected virtual server (VM only).
	Restarts the selected virtual server (VM or Solaris local zone).
	Allows you to filter the currently loaded scan results to find relevant data. For more information, see Filtering SAV Data on page 95.
	Allows you to compare to SAV snapshots. Clicking once will display the Compare pane at the bottom of SAV window. Click again to hide the Compare pane. For more information on comparing scan results, see Comparing Snapshots on page 89.
	Displays virtualization relationships between virtual server or virtual switches inside of their hosts in the Server map, Storage Map, and Devices Tree.
Properties Pane Toolbar	
	Alphabetize properties of the selected object in the Properties pane.
	Displays additional information about the selected Properties attribute.
	Expands all Properties categories.

Table 1 Toolbar Icons in SAV (cont'd)

















Toolbar Icon	Description
	Collapses all Properties categories.
	Adds a contact to your Business Application. (Available only when you select the top-level node of a Business Application in the Tiers tree.)
	Deletes a contact to your Business Application. (Available only when you select the top-level node of a Business Application in the Tiers tree.)
	Exports Properties pane or table contents (such as Infrastructure or differences) to .csv.
Maps Toolbar	
	Rotates the selected view, toggling it between a vertical and a horizontal orientation.
	Redraws all components in the selected view. Components that have been manually revised will retain their sizing.
	Expands selected tiers in the Tiers tree or closed folders in the selected map. Tiers are expanded recursively down to the business application component that they contain. Managed servers underneath the business application components are not expanded.
	Collapses all tiers in the Tiers tree or closed components in the selected view.
	Opens all tiers in the Tiers tree or signature in the selected map.
	Closes selected tiers in the Tiers tree or folders in the selected map.
	Zooms into the selected view (enlarges display size).
	Zooms out of the selected view (reduces display size).

Table 1 Toolbar Icons in SAV (cont'd)

Toolbar Icon	Description
	Resizes the selected components in a currently active view to fit within the screen size.
	Resizes all components in the currently active map to fit within the screen size.
	Labels all IPC lines with their associated protocol, such as SSH, HTTP, and so on. Protocols will display in the Server, Network and Tiers Maps. The Show Protocols mode is applied on a per-map basis
	Toggles the Tiers map to display the host server names in the title bar for process families.

Menus and Menu Options

This section discusses the menus and menu options that might not be self-explanatory.

File Menu

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

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

If you want to import a business application template that has already been saved, from the **File** menu, select **Import Template** and import the selected template.

For more information on business application templates, see [Business Application Templates](#) on page 77.

View Menu


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

Adding and Removing Devices in SAV

After you have opened and visualized devices in SAV and created your business application definition, you can add more devices to the initial snapshot in order to see how other devices — servers, network devices, storage devices, and so on — relate to the current state of your business application.

For example, you might have created a business application but you are still not sure about all the storage devices it might be using, or, you might see that there are a few managed servers that are connected to your business application, but that were not initially scanned. You can easily add these to the SAV application window.

Depending on your settings in the Add Devices window that you see when you add a device, SAV will automatically refresh the SAV snapshot and scan the selected devices. If you would rather not have SAV automatically scan any newly added devices, then you can add the devices, uncheck the check boxes in the window, and scan them later by clicking **Refresh**

Snapshot  on the SAV toolbar.

In some cases, you might find that some of the devices you have scanned are not necessary to your snapshot, and so you want to remove them. You can easily do so by selecting the device and selecting to remove it.


There are several potential errors that can occur when a Refresh or a device scan in SAV does not work. For a list of potential errors, see [SAV Scan Error Messages](#) on page 98.

Adding Devices to SAV

To add devices to SAV, perform the following steps:

- 1 From anywhere inside of SAV, either right-click or from the **Application** menu, select **Add Devices**.
- 2 In the Add Devices window, in the left pane you can choose a device category, and the corresponding devices appear in the right pane. You can add servers, device groups, network devices, storage devices, and so on.

The Discovered Dependencies category shows any devices that SAV has discovered to be related to or connected to some of the devices in the existing SAV snapshot.

The Refresh Scan Results option instructs SAV to automatically refresh the SAV snapshot when you click **Add**.
- 3 When you have selected the devices to add, click **Add**. Be sure to save the results or these newly added devices will not be saved in the Business Application.
- 4 If the Refresh Snapshot Results option was not selected, click **Refresh Snapshot**  on the SAV toolbar so SAV will scan the newly added devices. Any devices that have been added to SAV without being refreshed appear as a translucent box in the maps and will not display any properties information.




If you attempt to save or export the Business Application without refreshing the snapshot, a dialog appears asking if you want to save the Business Application. If you want the new information to be included in the Business Application, be sure to refresh the snapshot before saving or exporting.


Removing Devices From SAV

To remove devices from SAV, perform the following steps:

- 1 From inside of SAV, from one of the maps or the Devices pane, select a device, right click or from the Application menu, select Remove Devices.
- 2 You are asked to confirm that you want to remove the selected device. Click **Yes** to remove the device.

- 3 To make sure your SAV snapshot is up to date, click **Refresh Snapshot**  on the SAV toolbar so SAV can update the snapshot and scan the newly added devices, and then save the Business Application.

When a device is removed from a scan, the device and all connections to and from it and all external client IP addresses are removed in the maps, trees and tables, including links to other managed servers in the scan.

When you click **Refresh Snapshot** , these lines to the other managed servers may display as a Client IP or other dependencies.



If you attempt to save or export the Business Application without refreshing the snapshot, a dialog appears asking if you want to save the Business Application. If you want the new information to be included in the Business Application, be sure to refresh the snapshot before saving or exporting.

SAV Maps

SAV provides five visual maps that display physical and logical drawings of managed servers, network and storage devices, and connections in your environment: the Tiers Map, Server Map, Network Map, Storage Map, and SAN Map.

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

- [Tiers Map](#)
- [Server Map](#)
- [Network Map](#)
- [Storage Map](#)
- [SAN Map](#)

In addition to viewing the SAV maps, you can also:


- Show any virtualization relationships in the maps (virtual servers or devices, hypervisors, switches, and more. See [Showing Virtual Server Relationships in the Server Map](#) on page 43.
- Export a map to a .gif, .jpg, or an .svg file. See [Printing a Map](#) on page 55.
- Print a map on single and multiple sheets of paper. See [Printing a Map](#) on page 55.
- View IPC service names — such as HTTP or SMTP — in the Server or Network Maps. See [Showing IPC Service Names in Maps](#) on page 56.

Tiers Map

The Tiers Map displays the logical structure of a business application, including a business application's tiers and the connections between its application and storage signatures, external clients, and other dependencies. By default, this map is initially empty until you create the tiers and define signatures that comprise an application. See [Creating a Tier](#) on page 79.

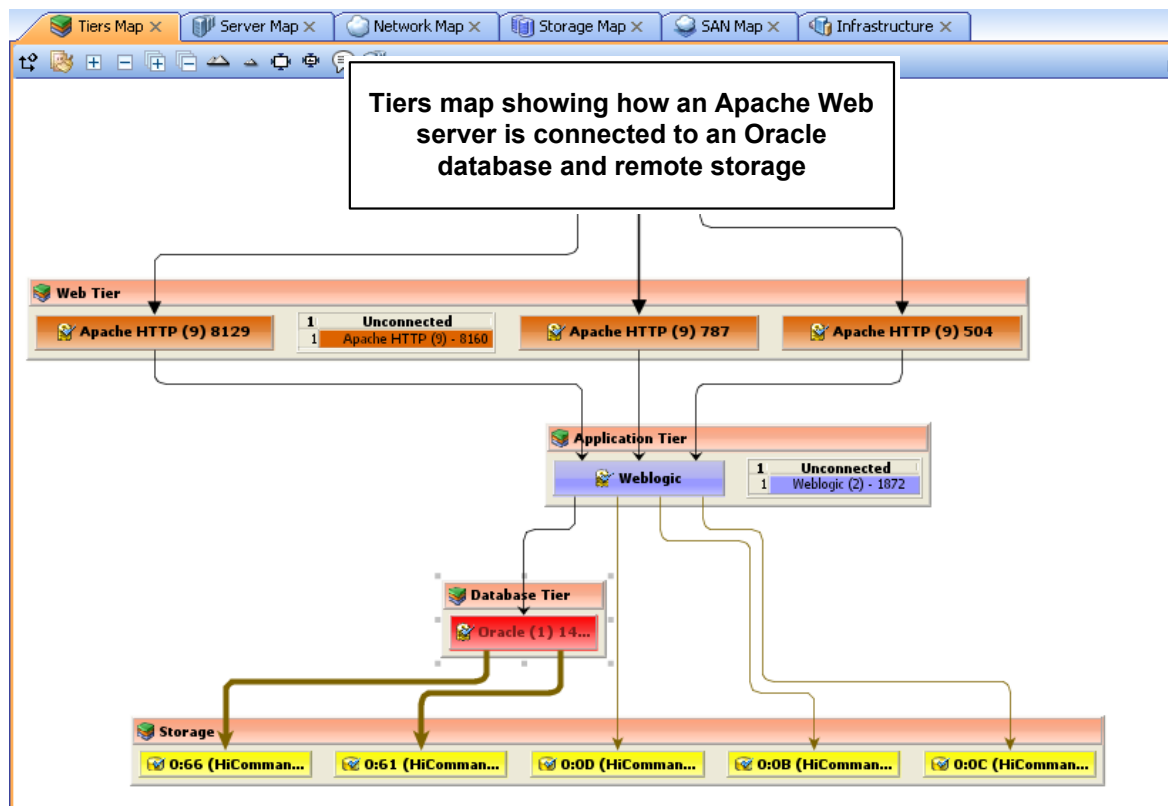
In addition, the Tiers Map shows the external IP addresses (Client IPs) that are connected to the application and the external IP addresses (external dependencies) that the application connects to and depends on.


A storage signature is similar to an application signature except that its signature matches storage volumes rather than process families (as in an application signature). This allows you to quickly identify and categorize storage dependencies. The Tiers Map displays storage signatures according to modifiable settings of color and name in the signature's Properties.

Storage signatures that do not match any actual volumes in the current scan are flagged just like a application signature, with a warning  icon.

Signatures can be categorized within a tier and used to recognize process families as named elements of an application. Tiers that do not have process families are shown in the Tiers Map. You can group them within a tier and make them visually distinct by modifying their color and name. See [Figure 6](#).

Figure 6 Tiers Map



If an application (or storage) signature does not have any process families (or storage devices) associated with it, then the tier object title bar will display a warning icon , for example,



VMware ESXi 3.5 hypervisor servers will not appear in the Tiers map, but any VMs hosted on the servers can be displayed.

Server Map

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

The Server Map shows the external IP addresses (client IPs) that are connected to the application and the external IP addresses (external dependencies) that the application connects to and depends on. If one of these external connections is an SA managed server, then SAV displays the connection as a server.

If you want to include these servers to the scan, right-click the server and select **Add**

Devices. After the device is added, click **Refresh Snapshot**  on the SAV toolbar.

In addition, in both the Server Map and Network Map, a DNS Servers element displays DNS servers in use by all the servers in the scan. If managed server information is known about any of these servers, then you can right-click to Add Devices. No connections will be shown to this DNS servers node.

If you want to see the type of service being used for connections between processes (and the devices they run on) in the Server Map, such as http, ssh, telnet, click the Show IPC Service


Names in Maps  button.



VMware ESXi 3.5 hypervisor servers will not expand to show process information in the Server Map.

Showing Virtual Server Relationships in the Server Map

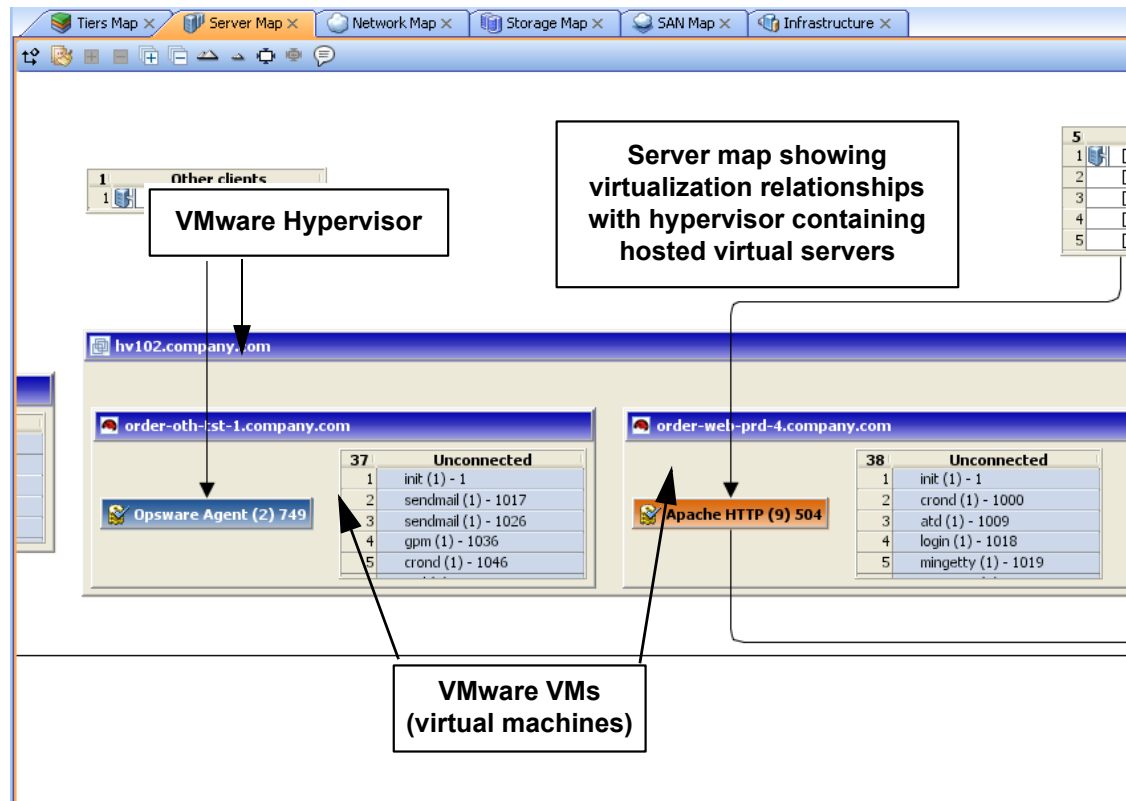
If you scanned virtual servers or devices in SAV, or any devices that contain or are connected to other virtual devices, then clicking **Show Virtual/Physical Containment**

Relationships  will display the relationships between virtual servers and the physical devices that host them in the Server Map (and any other map that shows virtual devices).

Clicking this button shows virtual servers inside of the hypervisor that is hosting them, so you can see which servers are virtual, and which are physical, and how they relate to one another.

Figure 7 shows the Server Map displaying a physical (non-virtual) server (top) and a virtual server (bottom) with the hypervisor name showing in the title bar.

Figure 7 Server Map



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

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

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




VMware ESXi 3.5 hypervisor servers will not display process family and runtime information.

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

For example:

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







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

Starting, Stopping, Suspending, Resetting Virtual Servers

In addition to viewing virtual servers and their relationships, you can also stop, pause, start, and reset virtual servers inside of SAV (but not physical hypervisor servers). Currently, SAV supports VMware ESX 3.0 and Sun Solaris 10 virtual server technologies.

Using either the SAV toolbar, or by right-clicking a virtual server, you can perform the following functions on a virtual server (VMware VM or a Solaris zone):

- **Start Virtual Servers:**  If the virtual server is currently paused, it is resumed. This is enabled if the selected virtual servers are stopped or paused.
- **Stop Virtual Servers:**  Stops selected virtual servers that are running.
- **Suspend Virtual Servers:**  Pauses selected virtual servers that are running. Only available for VMware's VMs, not Solaris local zones.
- **Restart Virtual Servers**  : Restarts if any selected objects are virtual servers that are running.

To start, stop, suspend, or reset a virtual server, perform the following steps:

- 1 From one of the SAV maps that display virtual servers (Servers, Network, Storage, SAN), select a virtual server.
- 2 From either the right-click or from the **Manage** menu select **Start**, **Stop**, **Suspend**, or **Restart Selected Virtual Server**.
- 3 If you want to see virtual servers and their relationships to their hypervisors, click **Show**

Virtual/Physical Containment Relationships



on the SAV toolbar.

Microsoft IIS, Oracle Database, and WebLogic in the Maps

In the maps, you can drill into a Microsoft IIS, Oracle Database, or WebLogic and expand the following objects:

WebLogic process family:

- Web Applications
- EJBs
- JDBC Connection Pools

Oracle Database process family:

- Oracle database instance
- Oracle tablespace
- Oracle data files

Microsoft IIS process family:

- Web Site
- FTP Sites

Network Map

The Network Map displays a physical (and virtual) layout of how the elements of an application connect to each other within the network, including the network interfaces on a server and the devices (switches and vSwitches) to which the server is connected. SAV also displays any firewalls and load balancers in your network environment.

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

The Network Map also shows external IP addresses (client IPs) that are connected to an application and the external IP addresses (external dependencies) that an application connects to and depends on. If one of these external connections is an SA managed server, then SAV displays the connection as a server.

If you want to include these servers to the scan, right-click inside the Network map and select

Add Devices, then click **Refresh Snapshot**  on the SAV toolbar.

Finally, the Network Map displays DNS servers in use by all the servers in the snapshot. If managed server information is known about any of these servers, then you can right-click to Add Devices.

If you want to see the type of service being used for connections between processes (and the devices they run on) in the Network Map, such as http, ssh, telnet, click the Show IPC Service

Names in Maps  button.



Network interfaces for VMware ESXi 3.5 hypervisor servers cannot be expanded to show process information in the Network Map.

Enhanced Layer 1 (L1) Network Graph

The Network Map displays not just the devices that are directly connected to a scanned server, but also the devices that are connected to the devices that are connected to scanned servers. Specifically, the Network Map will display the following physical L1 network information:

- All switches that SAV can detect that are directly connected to scanned servers are shown
- All network devices along the shortest path (based on number of hops) between any 2 servers in the snapshot are shown
- All additional network devices, including:
 - Network devices that were manually added to a business application snapshot
 - Network devices that were shown in a previous snapshot of the business application, but which were not yet manually removed from the Business Application
 - All physical (L1) connections between any two devices that are plugged into each other

Network Speed and Duplex Matching

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


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

ACL and Server Pool Configurations

For network devices (such as firewalls, load balancers, routers, switches), you can view ACL and server pool configuration (load balancers only) information.






For more information on viewing and comparing ACL and server pool configurations, see [ACLs and Server Pool Configurations](#) on page 86.


Virtual Network Devices

If you want to see how virtual network devices are related to the physical devices they run on and are connected to, click **Show Virtual/Physical Containment Relationships**  on the SAV toolbar.

VMware vSwitches are shown alongside virtual machines or other network devices. Connections between them appears as a green matching-duplex ethernet connection.

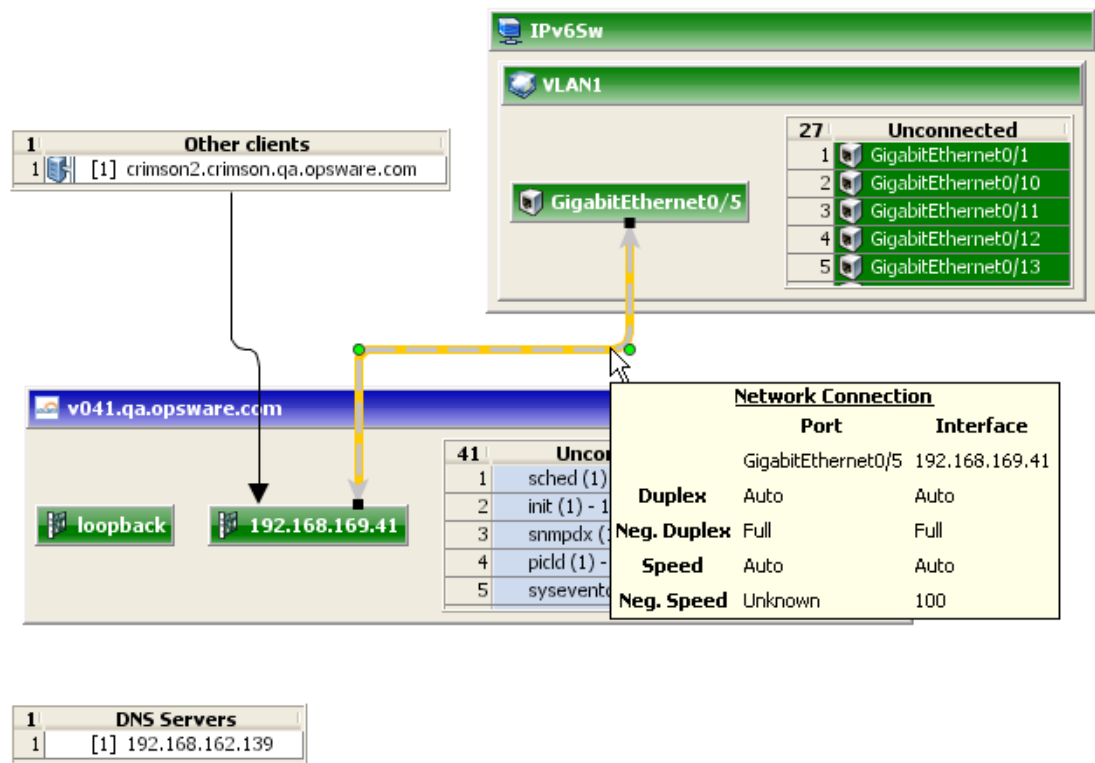
Network interfaces and devices use the following symbols:

-  **Network Device:** (As shown in the Devices tree) A switch or a vSwitch.
-  **Network Interface Card (NIC):** When available, a NIC is shown with its IP address and any processes connected to it.
-  **Listeners:** Process families that are listening on or connected to more than one network interface appear multiple times in the Network Map.
-  **Network Device Port:**
-  **Virtual LAN:**

Network devices without connections are also shown. If it is unknown whether the network device is connected to a server or another network device, a warning icon  appears next to it in the Devices tree. See [SAV Scan Error Messages](#) on page 98.

[Figure 8](#) illustrates network devices (green) as shown in the Network Map, with a switches connected to a network interface on a server, port and MAC address for the switch, and moving the mouse over the connection line displays connection speed and duplex information.

Figure 8 Network Map

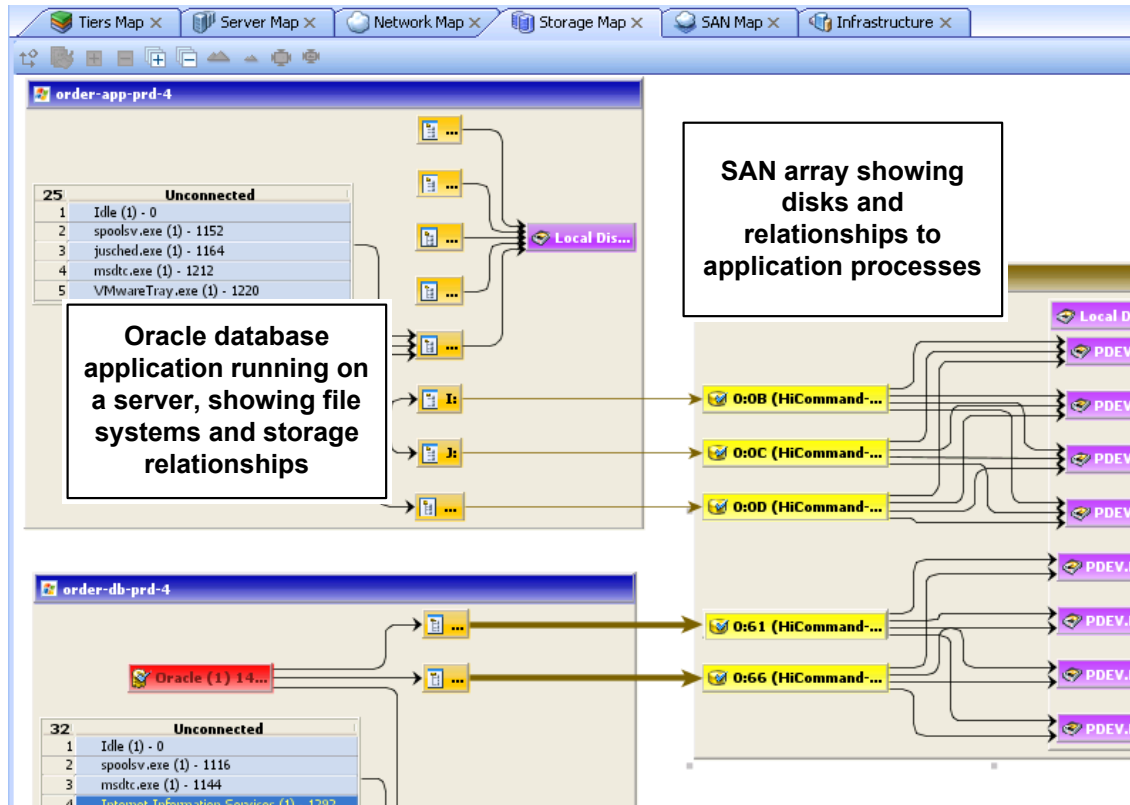


Storage Map

The Storage Map displays storage dependencies for all of the logical storage elements used by servers you have scanned with SAV.

This map provides a graphical view that displays servers, the process families running on them, including the file systems and local or remote storage devices these files are stored on and served from, either through local disks, NAS filers, or fibre channel disk arrays. The map shows connections between process families and their open files and where they files are stored. It also shows on which local (disks) or remote (SAN or NFS) storage those file systems reside. Even multipath SAN connections are displayed, and broken connections shows in red.

Figure 9 Storage Map Showing Oracle Database Mapping to SAN Array



In this map you will see the following data:

- **Servers:** Shows all process families and their relationship to file systems (through all open files) on each server. These file systems are served from NAS filers, local discs, and/or remote volumes from storage arrays. In the Storage Map you can see how specific applications running on servers interact with files systems and where data involved with these process families are stored — locally or remotely.
- **Databases (Oracle):** Shows Oracle database instances, tablespaces, and database files connected to each tablespace. The database files are displayed on the physical file system on which they are stored.
- **File Systems:** Shows file system pathname in the title bar and displays a single list of all open files that live on it. Each file system also shows a connection line between the itself and either local or remote storage.
- **Dependencies:** Shows lines that display how process families use the open files list on each file system on which they have open files. Connection lines also tie NFS-mounted file systems to the exported volumes on their hosting NAS filers, as well as the remote storage array volumes or local disks that they are hosted on.
- **LUN Mapping:** Displays mapping from a server to a disk array. Multipath link (shown in brown) will have thicker lines based on the number of paths to the storage device (thicker lines mean more paths). For links that are multipath, where at least one of the paths is down, the line is colored red.
- **Local Disks:** Groups all local disks into a Disk element, which can be expanded to show other disks.

- **SAN Arrays:** Shows storage disk arrays that are mapped and in use by the servers in the snapshot. Storage arrays that are scanned but are not in use by any of the applications show a special icon to distinguish them from those arrays in use. If you expand the SAN Array, you only see volumes in the array that are LUN mapped to the server. SAN Array backups are also displayed.
- **NAS Filers:** Displays all exported file systems stored on NAS Filers (which appear in a brown box) that are in use by servers included in the snapshot. Those filers that are scanned but are not in use by any of the applications have a special icon to distinguish them from those files that are in use. Other servers and devices that use disk on a filer are displayed by way of an exported file system, and a scroll list of other consumers is shown which points to that disk.

Viewing Storage and SA Permissions

Your user may be able to view some types of storage information in a SAV snapshot even if your user belongs to any groups that do not have permission to see storage devices such as SAN fabrics, arrays, and so on.

Specifically, If your user belongs to one or more groups that have the permission “Manage Business Applications: Read & Write,” then your user will be able to view such devices in a SAV snapshot and objects as fabrics (switches), storage arrays, network devices, and VM info in the SAV snapshot, even if the group does not have individual permissions granted to see those devices and objects.

If your user belongs to one or more groups that do not have “Manage Business Applications: Read & Write,” your user will be able to view SAN fabrics (switches), storage arrays, network devices, and VM info in a SAV snapshot only if the group has those individual permissions granted.

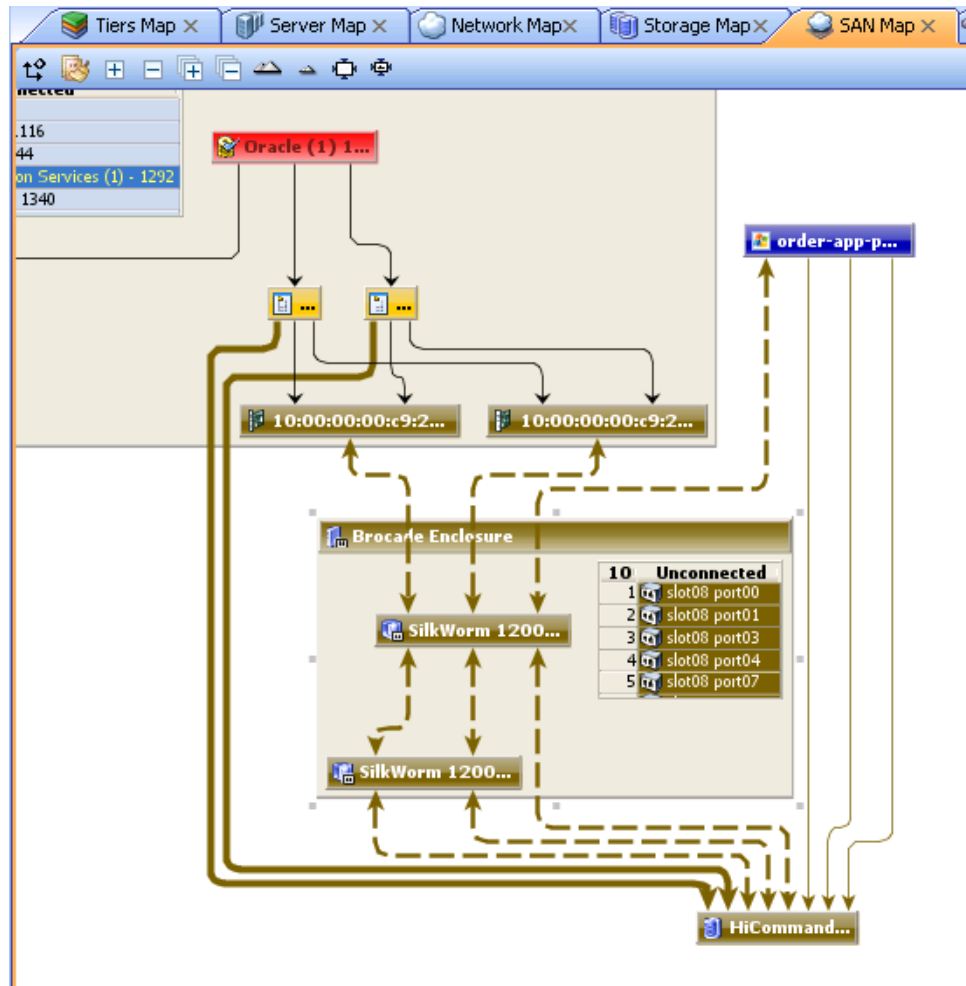
For example, if your user belonged to one or more groups that have the following permission: “Manage Business Applications: Read & Write” but had Manage Fabrics: None, your user would still be able to see fabrics (and SAN switches) in the SAV snapshot.

For more information on setting user group permissions, see the *SA Administration Guide*.

SAN Map

The SAN Map shows a superset of the Storage Map, including a graphical view of the fibre channel storage area network involved in a SAV snapshot, including all servers and their Fibre Channel Adapter (including fibre channel ports), and each adapter's connections to switches in the SAN.

Figure 10 SAN Map with Fibre Channel Switches Routing Storage Traffic from Server to SAN Array



The SAN Map displays the following data:

- Servers and their Fibre Channel Adapters (including fibre channel ports), and each adapter's connections to switches in the SAN.
- All SAN arrays and NAS filers included in the snapshot, with physical connections shown between servers, switches, and storage devices.
- File systems (and all open files) and how they rely on remote or local storage.
- LUN mapping between file systems and storage devices (shown in brown solid lines)
- SAN Switches. Each switch, when expanded, shows all ports being used. Any ports not in use in this snapshot are collected in a scrollable list of Extraneous ports.

- Virtual SAN Switches and Ports. Click **Show Virtual/Physical Containment**

Relationships



on the SAV toolbar to show virtual switches inside of their physical switch parent.

- Logical connections between devices are represented by solid lines. If a connection is down or has a mismatch, the line is red. If the connection is working and has no mismatch, the line is black. Thickness of line indicates the number of connections.
- Actual physical SAN connections (Fibre Channel cables) are shown in brown dashed lines, while LUN mapping connections between a server and a disk array is shown with solid brown lines.

SAV Infrastructure Pane

The Infrastructure pane provides detailed inventory and infrastructure information related to devices you have scanned in SAV. This pane provides a flat list of various objects found in the SAV snapshot and lists detailed properties for each.

Using the drop-down list the Infrastructure pane to filter categories of objects and their properties, to view and sort objects such as, all servers in the snapshot, all SAN arrays in the snapshot, and so on. You can then sort the columns to compare attributes of each item in the list. For example, you can select Servers and look at the load average for all servers discovered in the snapshot, or sort load average between servers. Or, you can view compliance policies are in use by all the servers or network devices in your business application.

Figure 11 Infrastructure Pane in SAV Showing Process Families

Name	Server	Signature	Alias	Group ID	Max. Resident Memory	Max. Virtual Memory
python	vmsles9	SA Agent		2460	19.043 MB	50.766 MB
python.exe	vmnt4	SA Agent		999	22.133 MB	19.34 MB
python	vmsol10u4x64	SA Agent		522	20.68 MB	22.367 MB
python	v019	SA Agent		1505	18.336 MB	54.145 MB
python.exe	vmw2k8	SA Agent		999	31.059 MB	24.754 MB
python	vmsol10u3x64	SA Agent		485	20.5 MB	21.602 MB
python.exe	vmw2k8core	SA Agent		999	28.465 MB	23.609 MB
python	vmsles10x64	SA Agent		2995	28.234 MB	92.84 MB
python.exe	vmw2k8x64core	SA Agent		999	36.598 MB	30.652 MB
python2	vmsles9	Python		16657	2.602 MB	3.703 MB
python2	vmsol10u4x64	Python		16605	3.109 MB	4.492 MB
python2	v019	Python		28283	2.664 MB	5.824 MB
python2	vmsol10u3x64	Python		12518	3.039 MB	4.383 MB
python2	vmsles10x64	Python		31414	3.496 MB	17 MB
sshd	vmsles9	OpenSSH		2661	1.758 MB	4.594 MB
sshd	vmsol10u4x64	OpenSSH		488	1.277 MB	3.438 MB
sshd	v019	OpenSSH		1531	680 KB	3.566 MB
sshd	vmsol10u3x64	OpenSSH		461	1.297 MB	3.273 MB
sshd	vmsles10x64	OpenSSH		3128	1.422 MB	26.426 MB
java	vmsol10u4x64	Java		1658	39.535 MB	107.477 MB
java	vmsol10u4x64	Java		537	73.617 MB	164.684 MB
java	vmsol10u3x64	Java		1453	38.051 MB	103.008 MB
java	vmsol10u3x64	Java		512	86.805 MB	170.633 MB
inetinfo.exe	vmnt4	Internet Inform...		999	3.203 MB	1.234 MB
svchost.exe	vmw2k8	Internet Inform...		999	50.609 MB	47.059 MB
inetinfo.exe	vmw2k8	Internet Inform...		999	13.168 MB	7.77 MB
svchost.exe	vmxp			997	4.188 MB	1.664 MB
alg.exe	vmxp			997	3.129 MB	1.012 MB

Depending upon what you have scanned in SAV, the Infrastructure pane displays the following types of object categories:

- Bindings (IIS)
- Compliance Policy
- Database
- Database File
- Disks
- Fibre Channel Adapters
- Fibre Channel Ports
- File Systems
- NAS Filers
- LUN Volumes
- NFS Exported File Systems
- Network Devices
- Network Interfaces
- Network Port
- Process Family
- SAN Arrays
- SAN Switches
- SAN Zone
- Servers
- Tablespace
- VLAN



VMware ESXi 3.5 hypervisor servers will not display runtime state information in the Infrastructure Pane.

Symbols Used in Maps





SAV uses a variety of symbols in the Network Map, Server Map, and Tiers Map, such as lines, arrows, diamonds, and so on. This section explains SAV map symbols in the following topics:

- [Process and Links Connection Symbols](#)
- [Map Process and Network Connection Symbols](#)
- [Exporting a View to .gif, .jpg, or .svg](#)
- [Printing a Map](#)
- [“Source” and “Comparison” Snapshot](#)

Process and Links Connection Symbols

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


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


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


Map Process and Network Connection Symbols


In the SAV maps, lines represent the following connections between devices:


- **Client link:** An internal connection that is labeled by the client IP address.
- **Process link:** A collection of TCP or UDP connections between process families. This link displays processes that provide a network service, such as listening for network connections, and processes that have a connection to another process or server.
- **Layer 1 connection:** A physical link between a server's network or storage interfaces and switch ports/switches. Layer 1 connections are indicated by colored, dashed lines in the map: These connection symbols are also used for virtual server and device connections as well.

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

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






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

 A brown solid line shows LUN mapping connections from a server to a disk array (in the Storage Map)

 A brown dashed line shows physical fibre channel cable connections (SAN and Storage Maps only).

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


Figure 12 Line Thickness and Process Connection Relationship

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

Exporting a View to .gif, .jpg, or .svg

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

To export a map to an image file, perform the following steps:

- 1 From the **View** menu, select **Export View**. (Or, click the Export View  button at the upper right corner of the map.)
- 2 Select a directory where you want the file to be located.
- 3 Enter a file name that includes either .gif, .jpg, or .svg as the file name extension.
- 4 Click **Export View**.





For information about exporting Properties to a .csv file, see [Exporting Properties Information to .csv](#) on page 57.

Printing a Map


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

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

To adjust the map before you print:


- Click Zoom In  or Zoom Out  to increase or decrease the size of the map before you print.
- Enter a title for the map

To print a map, perform the following steps:

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

Showing IPC Service Names in Maps

If you want to see the type of service being used for connections between processes (and the devices they run on) in the Network or Server map, such as http, ssh, telnet, click the Show

IPC Service Names in Maps  on the SAV toolbar. The service name will appear on each connection line.

To show IPC service names in the Server or Network map, perform the following steps:

- 1 Inside of SAV, select either the Server or Network map.
- 2 From the **View** menu, select **Show IPC Service Names in Maps**.

Or

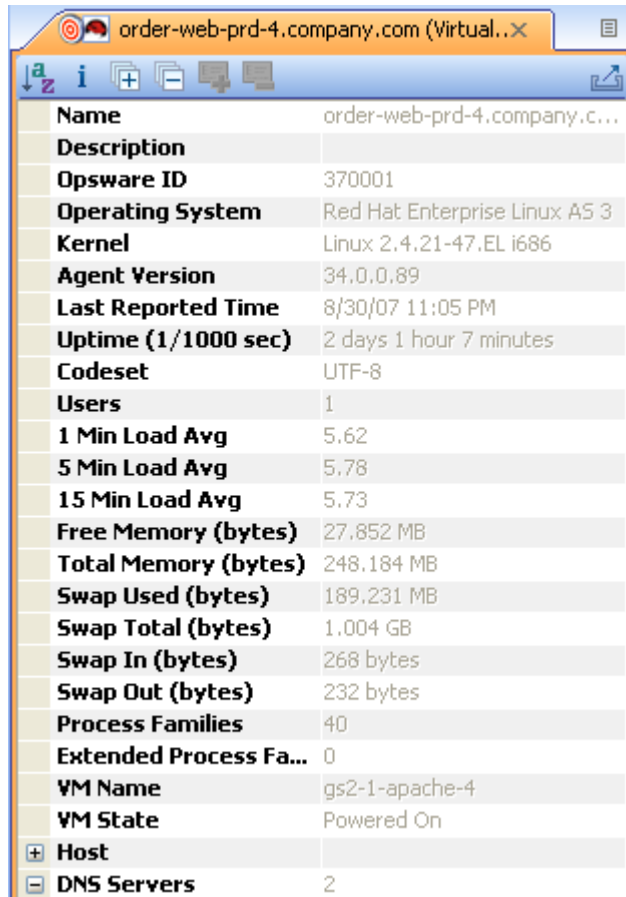
- 3 Click Show IPC Service Names in Maps  on the SAV toolbar

SAV Properties

SAV displays detailed properties for components selected in the Devices tree, Tiers tree, the Infrastructure pane, and any of the maps.

The information displayed in the Properties pane varies depending on the component type, such as a server, network device, storage device, process family, tier, application or storage signature, and links. Depending on the type of object you select, this page includes the number of users logged in, load average, swap usage, memory usage, application components, network devices, network ports, VLANs, tiers, links, and so on. It shows the MAC addresses for each network interface, as shown in [Figure 13](#).

Figure 13 Properties Pane for a Server




Name	order-web-prd-4.company.c...
Description	
Opsware ID	370001
Operating System	Red Hat Enterprise Linux AS 3
Kernel	Linux 2.4.21-47.EL i686
Agent Version	34.0.0.89
Last Reported Time	8/30/07 11:05 PM
Uptime (1/1000 sec)	2 days 1 hour 7 minutes
Codeset	UTF-8
Users	1
1 Min Load Avg	5.62
5 Min Load Avg	5.78
15 Min Load Avg	5.73
Free Memory (bytes)	27,852 MB
Total Memory (bytes)	248,184 MB
Swap Used (bytes)	189,231 MB
Swap Total (bytes)	1,004 GB
Swap In (bytes)	268 bytes
Swap Out (bytes)	232 bytes
Process Families	40
Extended Process Fa...	0
VM Name	gs2-1-apache-4
VM State	Powered On
Host	
DNS Servers	2

Exporting Properties Information to .csv

You can export the information contained in any Properties pane to the .csv format, which enables you to view SAV object properties inside spreadsheet applications.

To export a Properties pane to .csv, perform the following steps:

- 1 From inside of a SAV map or tier, select an object.
- 2 From inside the Properties pane for the selected object, click the Export View  button.
- 3 In the Export to CSV window, choose a location and enter a file name (with the .csv extension), and then click **Save**.

Tiers Tree: Tiers, Process Family, Signature Properties

SAV displays property information about the following SAV elements in the Tiers Tree and maps:

- [Business Application Properties](#) on page 58
- [Process Family Properties](#)
- [Extended Process Family Properties](#)
- [Tier Properties](#)

- [Application Signature Properties](#)
- [Storage Signature Properties](#)



VMware ESXi 3.5 hypervisor servers will not display in the Tiers Tree, although any VMs hosted on the servers will be displayed here.

Business Application Properties

A business application allows you create application tiers and application and storage signatures, so SAV can recognize application processes and storage devices in your data center and display them in the Tiers Map.

The business application Properties pane displays the following information (properties with a * indicate editable fields):

- **Name (*)**: Name of the Business Application
- **Description (*)**: Description of the business application.
- **Tiers**: The number of tiers contained in the business application.
- **Contacts (*)**: User created contacts (Application menu)
- **Custom Fields (*)**: Any Custom field you create is editable.

Process Family Properties

A process family is a collection of processes. The Properties Pane for a process family displays the following information:

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

- **Open files:** The files that the process family currently has open.
- **Software Packages:** The packages associated with the files that the process family has open.
- **Processes:** The number of individual processes in the process family. For each process, the following information is displayed:
 - **PID:** The process ID.
 - **User:** The user ID the process is running as.
 - **Command line:** The command line used to start the process.
 - **Path:** The path to the process binary.
 - **Memory statistics:** The percentage of physical memory consumed by the process, the resident size (in bytes) of the process and the virtual size (in bytes) of the process.
 - **Run time:** The time (in milliseconds) that the process has been running.
 - **CPU Statistics:** The CPU time accumulated by the process and the percentage of CPU consumed by the process since it began.
 - **Environment:** The name and value of each environment variable in the process environment.

Extended Process Family Properties

An extended process family is a collection of process families. If one process family has a listener and another has a connection into that same port, then SAV joins them into an extended family. The Properties Pane for extended process families contains the following information:

- **Name:** Name of the extended process family.
- **Process Families:** Aa list of all process families
- **Incoming connections:** The connections incoming to the process family, grouped by process family (if known, the IP address otherwise) and interface.
- **Outgoing connections:** The connections outgoing from the process family, grouped by process family (if known, IP address otherwise) and interface.

Tier Properties

The Properties Pane for a tier displays the following information (properties with a * indicate editable fields):

- **Name (*)**: The name of the application tier as displayed in the Tiers tree.
- **Subtiers:** The number of subtiers that are currently recognized in the tier.
- **Application signatures:** The number of application signatures currently recognized in the tier.
- **Storage Signatures:** The number of storage signatures currently recognized in the tier.
- **Device Filter (*)**: The devices that are associated with this tier. Only matching devices are filtered for matching application or storage signatures.

Tier Folder Properties

The Properties Pane for a tier displays the following information:

- **Name:** The name of the folder as displayed in the Tiers tree.
- **Application Tiers:** The number of subtiers that are currently recognized in the folder.
- **Application or signatures:** The number of application signatures currently recognized in the folder.
- **Storage Signatures:** The number of storage signatures currently recognized in the folder.
- **Device Filter:** The devices that are associated with the signatures in this folder. Only matching devices are filtered for matching application or storage signatures.

Application Signature Properties

The Properties Pane for an application displays the following information (properties with a * indicate editable fields):

- **Name (*)**: The name of the application component as displayed in the Tiers tree.
- **Alias (*)**: The name of the application component as displayed in the different views. The name will be shown as an alias if not defined.
- **Families**: The number of process families recognized as this application component.
- **Process Name (*)**: The process name filter used to recognize this application component.
- **Command Line**: The command line filter used to recognize this application component.
- **Executable Path (*)**: The executable path filter used to recognize this application component.
- **Open Files (*)**: The open file filter used to recognize this application component.
- **Modules (*)**: Shared libraries that are associated with the process family, DLL files on Windows operating systems and shared object files on Unix operating systems.
- **Environment Variable (*)**: The name and/or value of an environment variable that matches a process family associated with an application signature, where NAME is the name of the environment variable, and VALUE is its value. If you want to find an exact match, use both NAME=VALUE.
- **Ports Connected To (*)**: The port that the server is connected to.
- **Listener Port (*)**: The listen port used to recognize this application component.
- **Background Color (*)**: The background color displayed in the different maps.
- **Foreground Color (*)**: The foreground text color displayed in the different maps.

Storage Signature Properties

The Properties Pane for a storage signature displays the following information:

- **Name (*)**: The name of the storage component as displayed in the Tiers tree.
- **Alias (*)**: The name of the storage component as displayed in the different views. The name is shown as an alias if it is not defined.
- **Remote Volumes**: The number of storage volumes that are defined in this signature
- **LUN ID (*)**: The LUN can be defined with a regular expression to indicate the storage volume the server is connected to.
- **LUN Name (*)**: Name given to the LUN.

- **Exported Path (*)**: Exported path to the LUN.
- **Manufacturer (*)**: Company that manufactured the LUN.
- **Model (*)**: Model name or number of the LUN.
- **Background Color (*)**: The background color displayed in the different maps.
- **Foreground Color (*)**: The foreground text color displayed in the different maps.

Devices Tree: Server and Network Device Properties

SAV displays property information about servers, network devices, and the connections between them:

- [Server Properties](#)
- [DNS Servers Properties](#)
- [Properties for Servers and Devices with Compliance Policies](#)
- [Virtual Server Properties](#)
- [Link Properties for Servers and Network Devices](#)
- [Network Devices Properties](#)
- [Virtual Switch Properties](#)
- [Port Group Properties](#)
- [Network Interface Properties Pane](#)



If your managed server is configured with a storage inventory snapshot and the SE Scanner is configured to integrate with an appropriate SE server, then a server's properties will also include related storage and more detailed file system information..

Server Properties

The Properties Pane for a server displays the following information:

- **Name**: The host name of the server.
- **Server ID**: The SA unique identifier for the server.
- **Operating System**: The operating system of the server.
- **Kernel**: The kernel version of the operating system (when applicable).
- **Agent Version**: The version of the Server Agent that enables the server to be managed and scanned.
- **Last Reported Time**: The most recent time that the Server Agent communicated with the SA core.
- **Uptime**: The length of time the server has been powered on.
- **Codeset**: The character encoding for the server's locale.
- **Users**: The number of users that are currently logged in.
- **Load averages**: 1-minute, 5-minute, and 15-minute load averages. The load average for servers running a Windows operating system displays unknown because it is not supported by Microsoft.

- **Memory usage:** The total free memory.
- **Swap usage:** The total used swap and swap in/out activity.
- **DNS Servers:** All configured DNS servers for the selected server.
- **Virtual Machines/Zones:** If the selected server is a hypervisor (Solaris Global Zone or VMware ESX server), you can expand the list and view all zones (Solaris) virtual machines (VMware ESX). Each virtual machine or zone will display its own server properties. For more information, see [Virtual Server Properties](#) on page 64.
- **Routes:** All configured static routes on the selected server.
- **Interfaces:** The number of network interfaces. For each interface on a server, the following information is displayed:
 - MAC address
 - Broadcast address
 - Subnet mask
 - Device
- **FCAs:** All Fibre Channel Adapters (HBAs) installed on the selected server.
- **File systems:** All files systems in use on the selected server. For each file system, properties include: drive letter, mount point, mount options, type of file system, logical block device used by the file system, amount of free space, percent used, and associated device for each file system.
- **Disks:** All physical disks installed on the selected server.



If you have added a server to SAV but have not yet refreshed the Snapshot, the server will appear grayed out in the maps and in the Devices tab. The properties information will be blank until you initiate a refresh.

DNS Servers Properties



DNS Server properties contain the following information:

- **Name:** Name of the DNS Server.
- **IP Address:** IP address of the DNS Server.
- **Servers:** Provides server properties for each server using this DNS server. For information on individual server properties, see [Server Properties](#) on page 61.

Properties for Servers and Devices with Compliance Policies

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

Each compliance category displays one of the following compliance statuses:

- **Compliant** : The compliance scan ran successfully and the actual server or device configuration matches the criteria defined in the policy.
- **Partial** : The compliance scan ran successfully, but the server or device configuration did not fully pass the compliance criteria defined in the policy.

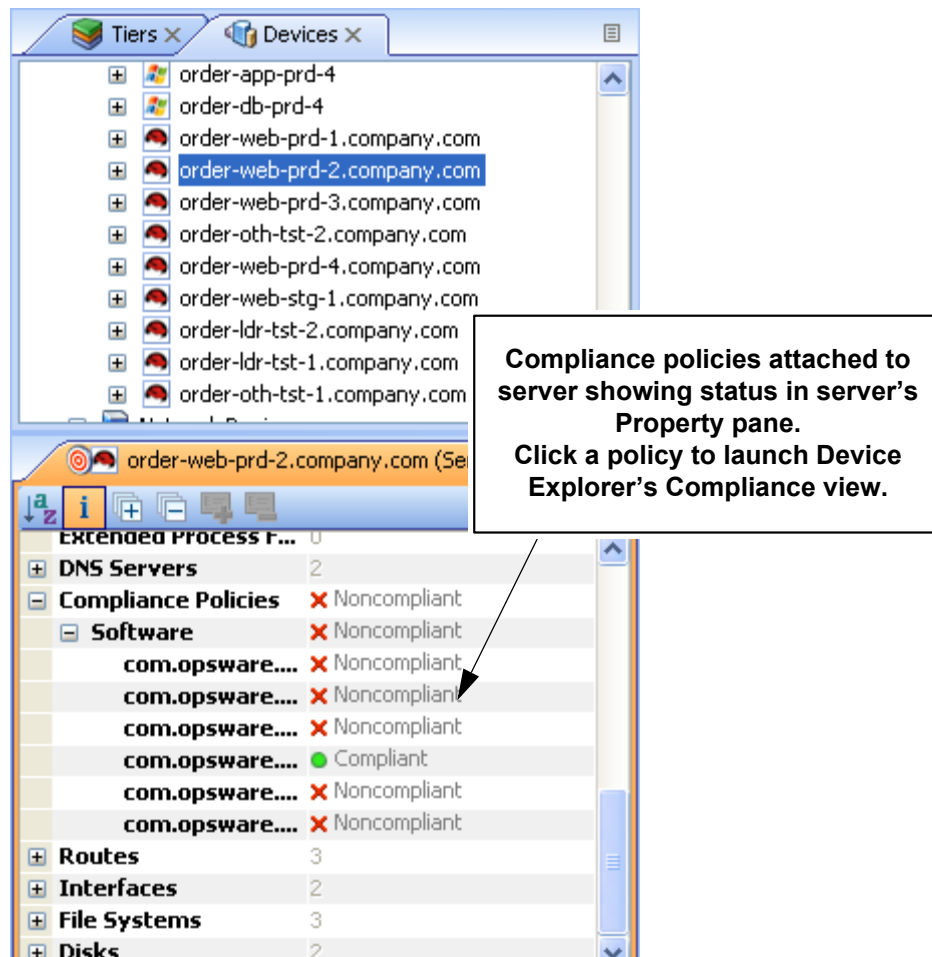
- **Noncompliant** ✖ : The compliance scan ran and the actual server or device configuration did not match the criteria defined in the policy.
- **Scan Failure** ! : The compliance scan was unable to run.
- **Scan Needed** □ : The results are unavailable, perhaps because a compliance scan was never run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

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

For information on launching a Device Explorer, remote terminal, or global shell, see [Adding and Removing Devices in SAV](#) on page 39.


Figure 14 shows a server's properties and lists compliance information about the server. Note that when any compliance policy on the server is non-compliant, then the main compliance policies row shows a non-compliant status, as seen in Figure 14.

Figure 14 Server Properties Compliance Information





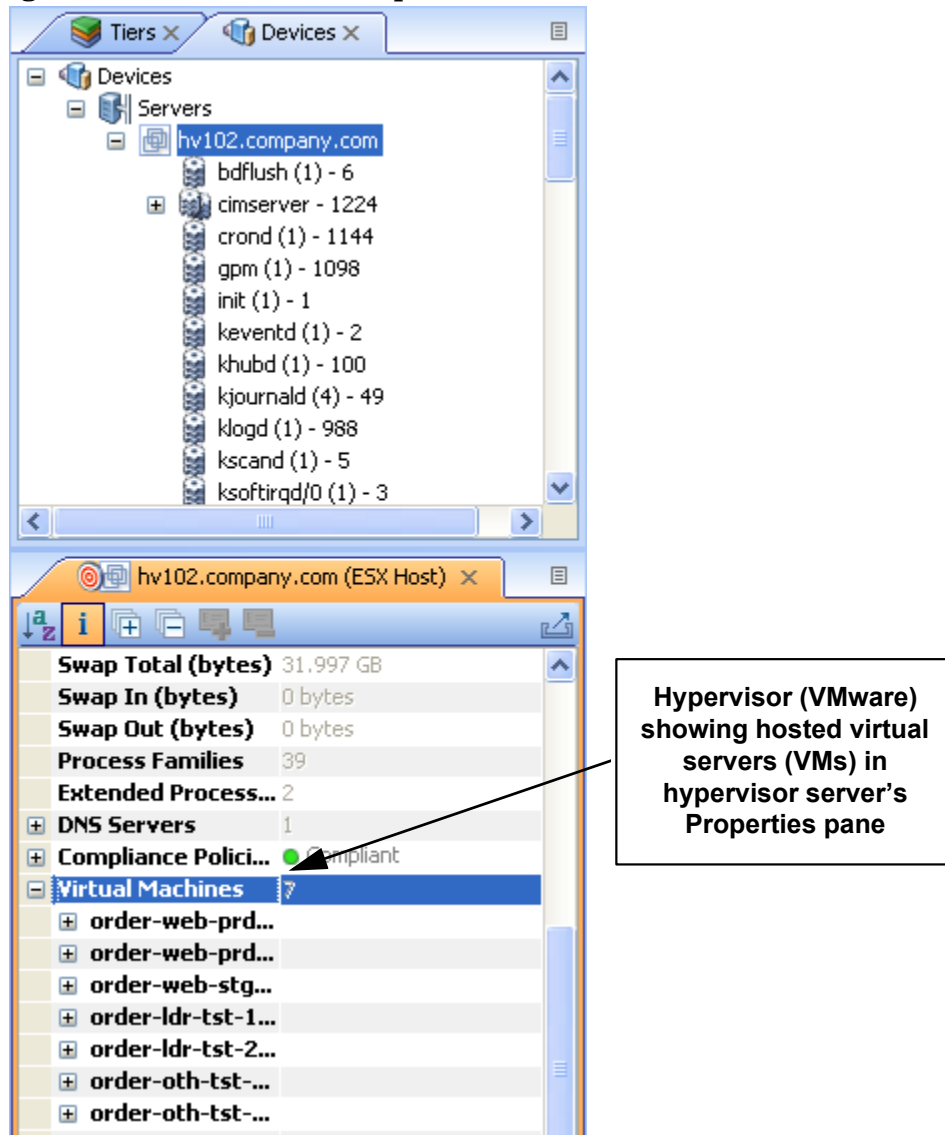
A SAV snapshot is not the same thing as a compliance scan, but they are related. A compliance scan can be run from the SA Client or the NA Client and checks a server or device's compliance status and reports this information to the Compliance Dashboard inside of the SA Client (and by extension, a server's properties in SAV), or to the Policy Compliance page in the NA Client user interface.

The actual compliance state you are viewing in SAV may have changed since you last scanned the server or device. To get the most current information, click **Refresh Snapshot**  on the SAV toolbar. For more information on the Compliance feature, see the *SA User Guide: Application Automation*.

Virtual Server Properties

The properties of a virtual server display all the same information as a physical server except that VMware and Solaris 10 hypervisors show all hosted virtual servers. VMware ESXi 3.5 hypervisor servers will not display in the Tiers Tree, although the VMs hosted on the servers may be displayed here. You can expand each hosted virtual server and view its properties. Conversely, each virtual server will contain in its properties the hypervisor that is hosting it. [Figure 15](#) illustrates this feature.

Figure 15 Virtual Server Properties



Link Properties for Servers and Network Devices

The Properties Pane for links between servers, external connections, and network devices displays the following information:

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

Replication

- **Replication Type:** The type of replication for this selected storage array.
- **Copy Type:** Replication copy type.

- **Status:** Status of the volume replicator.
- **Source:** Source volume of the replication.
- **Target:** Target volume of the replication.

Layer 1 Link Properties

Layer 1 Link properties represents the physical connection between devices and contains the following information:

- **Eff. Duplex:** Effective duplex between the two devices. The duplex represented is the actual duplex. If SAV cannot determine the duplex, then it will use the configured duplex.
- **Eff. Speed:** Effective speed between the two devices. The speed represented is the actual speed. If SAV cannot determine the speed, then it will use the configured port speed.
- **Port:**
 - **Name:** Name configured for the port.
 - **MAC Address:** MAC address assigned to the port.
 - **Duplex:** Configured duplex setting for the port (full, half, or automatic).
 - **Neg. Duplex:** Negotiated duplex, which could be different than actual.
 - **Speed:** Configured speed of the port.
 - **Neg. Speed:** Negotiated speed of the port, which could be different than what has been configured.
 - **Network Interfaces:** Network interfaces connected to this port.
 - **Peer MAC Addresses:** Other MAC addresses connected to the port.
- **Interfaces:** All network interfaces connected to this port.

IPC Links

Inter-process communication (IPC) link properties represent the links between processes in your Business Application and consist of the following attributes:

- **Protocol:** The protocol used for inter-process communication. Usually TCP or UDP.
- **Connections:** The number of connections associated with this link.
- **Destination Address:** IP address of the destination of the link.
- **Source Addresses:** Source IP addresses of the processes connection.

Network Devices Properties

The Properties Pane for network devices (routers, switches, firewalls, load balancers, and so on) displays the following information:

- **Name:** The name of the network device.
- **Last Reported Time:** The date of the last successful snapshot of the network device by NA.
- **Manufacturer:** The vendor that manufactures the network device.
- **Model:** The model number of the network device.
- **Operating System:** The operating system running on the network device.

- **Firmware Version:** The firmware version number for the device.
- **Asset Tag:** The assigned number used for tracking the network device.
- **VLANs:** The total number of VLANs that this network device has.
- **Ports:** The total number of ports that this network device has.
- **ACL Config:** Displays a link to view the ACLs configured for the device. You can view or compare ACL configurations by selecting **View ACL Configuration** from the **Manage** menu. This opens the ACL Configuration dialog for the selected network device(s) that have ACLs.

If you select two network devices, right-click, and select **Compare ACL Configuration**, this opens the Compare ACL Configuration dialog for the selected network device(s) that have ACLs. This allows you to compare ACL configurations between two devices in the same snapshot, or when in comparison mode, between the same device in the two snapshots. For more information, see [ACLs and Server Pool Configurations](#) on page 86

- **Server Pool:** Displays a link to view the sever pool members for the selected device. For load balancers, you can view Server Pool members by clicking the Server Pools link in the Properties pane or selecting **View Server Pool Configuration** from the **Manage** menu.

If you select two load balancers in the Devices tree, right-click, and select **Compare Server Pool Configuration**, you can view the Compare Server Pool Configuration window to see any differences in configuration. For more information, see [ACLs and Server Pool Configurations](#) on page 86

- **Compliance:** For network devices that have compliance policies associated with them, the properties will display its compliance status. For information on Compliance statuses, see [Properties for Servers and Devices with Compliance Policies](#) on page 62.

Network Device Port

Network device port properties display the following information:

- **Name:** Name configured for the port.
- **MAC Address:** The Media Access Control ID assigned to this port.
- **Duplex:** The configured duplex (if it can be collected).
- **Neg. Duplex:** The negotiated duplex (if it can be collected).
- **Speed:** The configured speed in Mbps (if it can be collected).
- **Negotiated Speed:** The negotiated speed in Mbps (if it can be collected).
- **Peer MAC Addresses:** Other MAC addresses connected to the port.

Virtual Switch Properties

Virtual Switch properties displays the following information:

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

Port Group Properties

Port group properties displays the following information:

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

Network Interface Properties Pane

The Properties Pane for a network interface displays the following information:

- **IP Address:** The IP address that is associated with a network interface.
- **MAC Address:** The Media Access Control ID that is associated with a network interface.
- **Subnet Mask:** The subnet that is associated with a network interface.
- **Broadcast Address:** The broadcast address that is associated with a network interface.
- **Device:** The device that is associated with a network interface.
- **Duplex:** The configured duplex (if it can be collected).
- **Negotiated Duplex:** The negotiated duplex (if it can be collected).
- **Speed:** The configured speed in Mbps (if it can be collected).
- **Negotiated Speed:** The negotiated speed in Mbps (if it can be collected).
- **Routes:** Network interfaces and servers will show static route information. VMware ESXi 3.5 hypervisor servers will not show route information.

Storage and SAN Properties

SAV displays property information about the following storage- and SAN-related elements:

- [Server Properties with Storage](#)
- [File Systems](#)
- [LUN IDs](#)
- [SAN Array Properties](#)
- [NAS Filer Properties](#)
- [Fibre Channel Switch](#)
- [SAN Array Disk Volumes](#)
- [Fibre Channel Adapter \(Host Bus Adapter\)](#)
- [Fibre Channel Port](#)
- [Storage Signature](#)

Server Properties with Storage

Servers display the following attribute information in the Properties tab (these elements are visible only in the Storage or SAN Maps):

- **FCA:** Any server that is attached to a SAN storage device will have an HBA (Host Bus Adapter) which SAV labels FCAs (Fibre Channel Adapters). You can expand category in the Properties pane and see a list of all FCAs, each of which can be expanded to display all their ports.
- **File systems:** SAV displays details such as mount options, Type (ext3, NFS), and so on.

- **Disks:** Number of and names of each disk on the server, including type, manufacturer, and disk capacity information.

File Systems

File systems can exist on both servers and NAS Filers (NetApps). SAV displays the following attributes for file systems in the Properties pane:

Mount point/drive letter/export path:

- Capacity
- Free capacity
- Percent used

For server file systems, the following information is displayed:

- Device
- Mount options
- Type (for example, ext3, NFS, and so on)

For NetApp file systems/Qtrees, the following information is displayed:

- Aggregate
- Plexes
- RAID groups

Disk Properties

Disks from SAN arrays, NAS filers, and servers will display the following attributes in the Properties pane:

- Capacity
- Manufacturer
- Model
- Type
- Status
- Serial number
- Firmware version
- Device (only on servers)

SAN Array Properties

SAN Arrays properties include the following information:

- Name
- Description
- Last Reported time
- Manufacturer
- Server ID

- Disks
- Volumes
- Ports

LUN Volume Properties

LUN volume properties include the following information:

- **Name:** Configured name for this LUN.
- **Description:** Description given to the volume.
- **Capacity:** Total space allocated to the volume (bytes).
- **LUN IDs:** All LUN IDs configured for this volume.

NAS Filer Properties

NAS Filers display the following attribute information in the Properties pane:

- **Name:** Name given to the device.
- **Description:** Description of the device.
- **Server ID:** ID that SA uses to identify this device.
- **Hostname:** Host name for the filer.
- **Last Reported Time:** Last time agent on this device reported to SA.
- **Manufacturer:** Manufacturer of the device.
- **Model:** Model number of the device.
- **Operating System Version:** Version of the OS running on the device
- **Serial Number:** Device serial number.
- **Hardware Version:** Device hardware version.
- **Disks:** Number of disks connection to the devices.
- **Ports:** Number of ports in use by the device.
- **Volumes:** Number of LUN volumes connection to the device.
- **Exports:** Number of exports in use by the device.

NFS Export Properties

NFS export properties contain the following information:

- **Export Path:** Path to the remote file system being linked to.
- **Size:** Total capacity of the exported file system.
- **Free Space:** Space available on the exported path.
- **% Used:** Percentage of space being used on the exported file system.
- **Aggregate:**
- **Plexes:** List of all plexes
- **RAID Groups:** List of all raid group (if any) configured on the exported path.

Fibre Channel Switch

Fibre Channel Switches contain the following attribute information in the Properties pane:

- Name
- Description
- Last Reported Time
- Manufacturer
- Model
- Server ID
- Serial number
- Firmware version
- Hardware version
- Ports
- Virtual switches



If a switch is a director-class switch, it could possibly have virtual switch children — these can be expanded as well.

SAN Zones

- Name: The name of the SAN zone.

SAN Array Disk Volumes

Volumes on SAN arrays (and NAS Filers acting as arrays) contain the following attributes:

- Name
- Description
- Capacity
- LUN ID

Fibre Channel Adapter (Host Bus Adapter)

Fibre Channel Adapters (known generically as HBAs) display the following attributes in the Properties pane:

- Node WWN
- Manufacturer
- Model
- Serial number
- Driver version
- Firmware version
- Hardware version
- Ports

Fibre Channel Port

Fibre Channel Ports display the following attributes information in the Properties pane:

- **World Wide Name:** The physical port name that identifies the fibre channel port on a SAN.
- **Description:** Description of the port.
- **Port Number:** The port number that identifies fibre-channel cards and cable connections.
- **Status:** Indicates whether or not the port is open and working.
- **Zonesets:** Zonesets to which the port belongs.
- **Zones:** Zones to which the port belongs.
- **Fabric:** Fabric that contains the fibre channel switch and port.

Storage Signature

Storage signatures match remote storage volumes and display the following information in the Properties pane:

- LUN Name
- LUN ID
- Exported path
- Storage System Manufacturer
- Storage System Model

SAN Link Properties

SAV displays the following links or connections between storage signatures in the SAN Map to illustrate how these signatures are related and connected:

- [Fibre Channel Link](#)
- [LUN Mapping Link](#)
- [NFS Mount Link](#)

Fibre Channel Link

Fibre channel links represent physical fiber cables connecting two ports. The Properties pane displays attributes information for endpoint ports.

LUN Mapping Link

LUN mapping links represent the logical connection between a host LUN volume and the corresponding SAN array volume. LUN mapping links display the following information in the Properties pane:

- Number of paths (indicated by thickness)
- Whether any paths are down (in red)
- Server file system consumer

- Array target volume

NFS Mount Link

NFS mount links represent a dependency of a server file system on a remote NAS Filer export and display the following information in the Properties pane:

- **Mounted To:** Directory of the NFS mount on the server this link originates from.
- **Exported From:** The exported path to the exported file system.

SAV Options

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, including resizing and repositioning all tabbed views.

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

Accessing Servers and Devices From SAV

To help you troubleshoot and take action for server and application errors, SAV gives you easy access to servers and devices through the following methods:

- [Opening the Device Explorer](#)
- [Opening a Remote Terminal](#)
- [Opening a Global Shell](#)

Opening the Device Explorer

To view detailed information about a server or a device (network or storage device) using the Device Explorer, perform the following steps:



SAN switches cannot be opened in the Device Explorer from inside SAV.

- 1 In one of the SAV maps, select one or more servers.
- 2 Right-click and then select **Open in Device Explorer** to open a Device Explorer in the SA Client for each selected server.

See the *SA Users Guide: Server Automation* for information about how to use the Device Explorer.

Opening a Remote Terminal

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



The Open a Remote Terminal feature is not available for VMware ESXi 3.5 hypervisor servers.

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

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

See the *SA Users Guide: Server Automation* for information about using utilities in a Remote Terminal.


Opening a Global Shell

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

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

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

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

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

See the *SA Users Guide: Server Automation* for information about how to use Global Shell.



Running Scripts on Devices

From inside SAV you can run a script, either directly on a selected server or network device (but not on SAN devices), or on the Global File System (OGFS) using the Global Shell — given that your user account has sufficient permissions to run the Global Shell and to perform any operations on servers under SA management.




The Run Scripts feature is not available for VMware ESXi 3.5 hypervisor servers.

There are three possible scenarios in which you can run a script in SAV:


- By selecting a server and clicking **Run Script**  on the SAV toolbar, or selecting **Run Script** from the **Manage** menu. This launches the Run Script Task window.
- By selecting Run Global Shell Script , which launches the Global Shell, and which gives you access to the OGFS.
- Selecting a network device and click **Run Script** on the SAV toolbar, or selecting **Run Script** from the **Manage** menu. This launches the NA interface, where you can log in to NA and run the script on the selected network device.

For more detailed information about the script execution process and how it works, see [Chapter 9, Script Execution](#), on page 423 of this guide. For information on running scripts on network devices, consult the NA online documentation.


To run a script on a server, perform the following steps:

- 1 From inside SAV, select a server from the Devices pane or one of the maps.
- 2 From the **Manage** menu, select **Run Script**, or click Run Script  from the SAV toolbar.
- 3 In the Run Script window, fill out the necessary information and perform the steps to execute the script. For more information on running a script on a server, see [Chapter 9, Script Execution](#), on page 423 of this guide..

To run a global shell script in SAV, perform the following steps:

- 1 From inside SAV, from the **Manage** menu, select **Run Global Shell Script**, or click Run Global Shell Script  from the SAV toolbar.
- 2 In the Run Global Shell Script window, fill out the necessary information and perform the steps to execute or schedule the script execution. For more information on running a global shell script on the OGFS, see [Running an OGFS Script](#) on page 438.

To run a script on a network device, perform the following steps:

- 1 From inside SAV, select a network device from the Devices pane or the Network Map.
- 2 From the **Manage** menu, select **Run Script**, or click Run Script  from the SAV toolbar. This launches the NA web application interface.
- 3 Log in to NA, and on the New Task - Run Command Script page, fill out the necessary information to run or schedule the script execution. For more information on running a script on a network device through NA, consult the NA online documentation by clicking the Help link in the upper right corner of the page.



You must have proper permissions to run global shell scripts and scripts on a device. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Creating Business Application Definitions

A business application definition allows you to transform a data display that contains extraneous and hard-to-understand information into a focused and easy-to-understand view of the relevant data. Based on business application tiers and application and storage signatures that you create, SAV recognizes actual application processes and storage devices and displays them in the Tiers Map according to any visual customizations you make to them.

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

See [Signature Evaluation Order](#) on page 81 for information on the order in which SAV scans application signatures and matches them to processes and process families on servers.

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

- [Business Application Tiers](#) on page 79
- [Creating a Tier](#) on page 79
- [Application and Storage Signatures](#) on page 80
- [Creating an Application or Storage Signature](#) on page 83

Business Application Templates

When you first scan servers with SAV and visualize them, the Tiers pane is empty—it has no business application definitions until you create them. (There are, however, some predefined commonly used default applications built into the product, such as Apache, WebSphere, and so on, contained in the Default Signatures folder.) Once you create and define an application definition with tiers and signatures, you can save the application definition as a template, which can be reused by yourself or others on your team to be automatically be applied to new device scans.

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

Setting a Default Application Template

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

Resetting the Default Application Template

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

Importing an Application Template

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



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

Saving a Business Application as a Template

You can save business application as a template, which makes it available as a generic template that can be used again and shared among other team members.

Note that When you export a Business Application, by design all scan information will be lost, including the servers and devices and their relationships. Business Application definitions and all the components inside of them remain after an export, but turn red to indicate that relationships between live processes and connections have been lost.

To save business application as a template, perform the following steps:

- 1 From the **File** menu, select **Save As**.
- 2 From the Save in drop-down list, select either Opware Global File System or Desktop. (Only business applications can be saved to the SA Client Library. SAV archives and templates can be saved to disk or the Global Filesystem.)
- 3 Enter a name for the business application template, and click **Save**.

Creating Business Application Contacts




In SAV, you can create a list of email contacts — and send emails to contacts on this list — by adding email and contact information to the top level tier of a business application.

You can create groups of contacts, and add to each contact such information as email, instant messenger IDs, phone number, and so on. From the **File** menu, select **Send Email**, and you can email any of the contacts you have added to the Business Application

SAV also will display any email addresses configured on network devices scanned by SAV. (Most network devices have an internal configuration setting such as “sysContact” that allows them to associate an email address for the owner of the device.)

You can add contacts through the properties of the Business Application, and then send emails to any email contacts listed.

To create a new contact for your business application, perform the following steps:

- 1 From the Tiers pane, select the top-level Business Application  icon.
- 2 Select the Properties for the business application, and then click **Add Contact**  (at the top of the Properties pane).
- 3 To enter information for a contact, double-click in the field to the right of each entry. After an entry line is filled, press Return to enter the information. If you want to be able to send emails to a contact, be sure to enter the contact's email address.
- 4 To delete a contact, select the contact in the Properties pane and click **Remove Contact** .

Sending Email to Business Application Contacts

You can send email to any business application contact that has a well-formed email address.

To send an email to a contact, perform the following steps:

- 1 From the **File** menu, select **Send Email**.
- 2 In the Email Contacts window, expand the business application. Each contact that has its name check marked is added to the email. If you do not want to send an email to one of the contacts, select the check mark next to the name.
- 3 Click **Compose** to write and send the email. (SAV launches whatever email client you have installed and configured on your local system.)

Business Application Tiers

Business application tiers provide an architectural framework to organize and display application and storage signatures. You can add, edit, delete, cut, copy, and paste tiers in the Tiers tree. You can paste a tier before or after a selected position in the Tiers tree to rearrange the order. The order of tiers (and the signatures they contain) is significant because it affects the order that the process families are assigned to signatures. (For more information, see [Signature Evaluation Order](#) on page 81.)

If any tiers have application signatures that do not recognize any process families, they and their ancestors are represented with warning icons in the tree and by red title bars in the view. This allows you to quickly identify signatures that should be running but are not.


Creating a Tier

To create an tier in the Tiers tree, perform the following steps:

- 1 In the Tiers tree or map, select either the top-level business application node or a tier, right-click, and select **New Tier**.
- 2 The Properties pane for the tier becomes active and you can edit the tier's properties, such as, give the tier a name.

Deleting a Tier



To delete an tier from the Tiers tree, perform the following steps:

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

Cutting and Copying a Tier


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

To cut and copy an application tier in the Tiers tree, perform the following steps:

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

Pasting a Tier

To paste a tier in the Tiers tree, perform the following steps:

- 1 Select a tier in the Tiers tree and then select the Paste icon  . The tiers that you cut or copied to the clipboard will be appended to the selected tier's children. When you select a signature in the Devices tree, the Paste icon will be disabled.

Application and Storage Signatures

Application and storage signatures are organized and displayed in the Tiers pane inside of the SAV application window. A signature always lives inside of a tier.

Application or storage signature contains the following data:

- A name.
- A signature, which is a set of rules that users provide and that SAV uses to identify a process family or storage mapping. For application signatures, these rules use data such as process name, command line, listen port, environment variable, executable path, and so on. For storage signatures, these rules include either a LUN volume or an NFS File System.
- Object properties such as name, color, and whether this object is used by default each time the user opens up an application.

You can add, edit, delete, cut, copy, and paste signatures in the Tiers tree. You can paste a signature before or after a selected position in the Tiers tree to rearrange the order. The order of signatures (and the tiers that contain them) is significant because it affects the order that the process families and storage mappings are assigned to signatures.

SAV comes with a set of predefined default signatures that recognize a variety of commonly used application process families, such as Apache HTTP, Microsoft IIS, WebLogic, JBoss, Oracle, and so on. So, if your server has any of these applications installed, SAV is able to recognize and display them in the Tiers tree and the maps.

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



Application signatures do not capture process family information on VMware ESXi 3.5 hypervisor servers.

List of Application Signature Discovery Properties

The full list of application signature properties that can be used to find applications and their process families include:

- Process Name
- Command Line
- Connected To Port
- Listener Port
- Executable Path
- Open Files
- Open Modules
- Environment Variable Name and/or Value

The means by which an application signature discovers applications and process families for environment variables allows you to match application signatures by the name of an environment variable, its value, or both.

“Environment variable” is an application signature property that can be added using the following syntax:

`NAME=VALUE`

where `NAME` is the name of the environment variable, and `VALUE` is its value. You can type either name, the value, or both to find matching process families. However, if you want to find an exact match, you must use both `NAME=VALUE`.

Signature Evaluation Order

The order that signatures are recognized in SAV is important because a process family or storage mapping is associated with the first signature that it matches in the Tiers tree. Evaluation order is significant especially when the recognition criteria for a signature matches the same process family or storage mapping found in multiple signatures.

Signatures are evaluated in a depth-first, top to bottom order: signatures in a tier's sub-tiers are evaluated before the tier's signatures (depth -first), and tiers in the Tiers tree are evaluated from top to bottom. Signatures are applied in the order in which they appear in each tier.

After all user-created tiers and signature hierarchies are evaluated, then all of the default SA signatures are evaluated; for example, NA Management System, NA Syslog Server, and so on. After the SA signatures are evaluated, then all of the default predefined signatures are evaluated; for example, Apache HTTP, Internet Information Server (IIS), and so on.


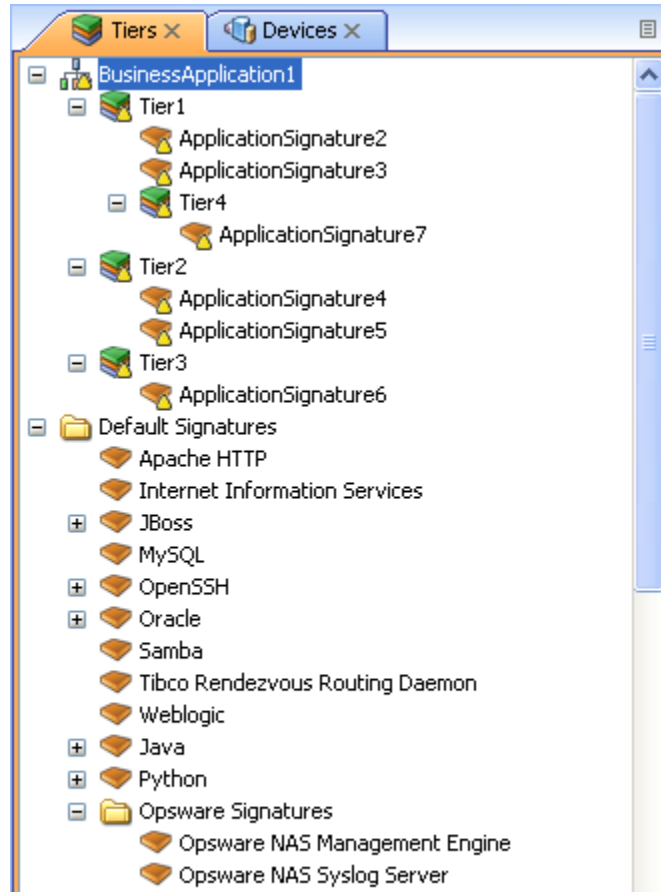
Consider an application definition that has the structure shown in Figure 16. In this image, no processes or process families match the signature definitions, and so the signatures are represented with the  icon.

Figure 16 Tiers Signature Evaluation Order



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


- 1 ApplicationSignature7
- 2 ApplicationSignature2
- 3 ApplicationSignature3
- 4 ApplicationSignature4
- 5 ApplicationSignature5
- 6 ApplicationSignature6
- 7 NA Management System
- 8 NA Syslog Server
- 9 <All remaining SA signatures, top to bottom>
- 10 Apache HTTP

- 11 Internet Information Services
- 12 <All remaining default signatures, top to bottom>

Creating an Application or Storage Signature

To create an application or storage signature in the Tiers tree, perform the following steps:

- 1 Make sure you have already created a Tier into which you want to create an application or storage signature.
- 1 Select the Tiers into which you want to create the signature in either the Tiers tree or map.
- 2 From the **Application** menu, select **New Application Signature** or **New Storage Signature**. The Properties pane becomes active and ready to be filled out for the new signature.
- 3 In the Properties pane for an *application* signature, enter any of the following information:
 - **Process Name:** The name of the process family.
 - **Command Line:** The command line that a signature was started with.
 - **Executable Path:** The path to the executable file of this application component.
 - **Open Files:** The name of an open file.
 - **Modules:** The shared libraries associated with the process family. These include DLLs on Windows and shared object files on Unix.
 - **Environment Variables:** For environment variables, enter a name, the value, or both to find matching process families, where NAME is the name of the environment variable, and VALUE is its value. If you want to find an exact match, you must use both NAME=VALUE.
 - **Connected to Port:** The port the signature is connected to.
 - **Listener Port:** The port on which the signatures are listening.
 - **Alias:** The name of the application component as displayed in the different views.
 - **Background color:** Click to change the background color displayed in the different maps.
 - **Foreground color:** Click to change the foreground text color displayed in the different maps.
- 4 If you created a *storage* signature, enter any of the following information:
 - **Name:** The name of the storage device; for example: HiCommand.
 - **Alias:** Alias for the storage device (if any).
 - **Remote Volumes:** Number of remote volumes on the storage device.
 - **LUN ID:** LUN ID number.
 - **LUN Name:** Name of the LUN.
 - **Exported Path:** Remote exported path for NAS filers.
 - **Manufacturer:** Maker of the storage device.
 - **Model:** Model number.

- **Background color:** Click to change the background color displayed in the different maps.
 - **Foreground color:** Click to change the foreground text color displayed in the different views.
- 5 After each entry, press Return on your keyboard to enter the property.
 - 6 When you have finished defining the signature, click **Refresh Snapshot**  on the SAV toolbar so SAV can update the snapshot and scan your data center to find matching process families and storage devices.

Editing Signatures

To edit a signature in the Tiers tree, perform the following steps:

- 1 In the Tiers tree, select a signature.
- 2 From the Properties pane of the selected signature, double-click in the right-side of the property entry and edit the text.
- 3 Press Return on your keyboard to enter the changes.

Deleting Signatures

To delete signatures from the Tiers tree, perform the following steps:

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



Cutting and Copying Signatures

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




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

To cut and copy a signature tier in the Tiers tree, perform the following steps:

- 1 In the Tiers tree, select a signature.
- 2 From the toolbar, select the  icon or the  icon, or right-click and select **Cut** or **Copy**.
- 3 Then, click **Paste** from the **Edit** menu, or press Control + V.
Or
- 4 You can press the Control button on your keyboard, and then select and drag a signature from one tier to another, creating a copy of the selected signature.

Pasting a Signature

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

- Select a signature in the Devices tree and then select the Paste icon . The signatures that you cut or copied to the clipboard will be appended to the selected tier's signatures.
- Select a signature in the Devices tree and then select **Paste** from the **Edit** menu. The signatures that you cut or copied to the clipboard will be inserted into the selected signature's parent tier *before* the selected signature.
- Select a signature in the Devices tree and then select the **Paste** from the **Edit** menu. The signatures that you cut or copied to the clipboard will be inserted into the selected signature's parent tier *after* (below) the selected signature.



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

SAV Business Application Management


In SAV, a snapshot represents the state of a set of network and storage devices and managed servers, the process families running on those servers, the connections among those process families, local and remote storage devices, file systems, and any external clients and dependencies. Snapshots can be saved as part of a Business Application into the SA Client Library, or as a .vam file to a local or remote file system. A Business Application can contain any number of Snapshots, each of which can be used for Snapshot comparisons.

However, in order to save Business Applications to the Library or OGFS, your user account needs permissions to be able to write to those directories. To obtain the necessary permissions, contact your SA administrator.

Opening a Business Application

After you have launched the SAV, you can open a previously saved Business Application.

To open a Business Application, perform the following steps:

- 1 In the SAV window, select the  toolbar icon or select the **File** menu and then select **Open** to display the Open window.
- 2 In the Look in drop-down list, select the directory on your computer, in the SA Library, or the OGFS where the Business Application was saved.
- 3 Click **Open**.

Saving a Business Application

To save a Business Application, perform the following steps:

- 1 From the **File** menu, select **Save** or **Save As** to open the Save window. (Note that by default, all scan results are selected to be saved.)

- 2 If you chose Save As, in the Save in drop-down list, select the your local computer, the SA Library, or the OGFS and choose where you want to save the Business Application.
- 3 Click **Save**.



If you exit SAV before saving your changes (either application definition changes or Snapshot changes), you are prompted to choose whether you want to save your changes and then exit or exit without saving your changes.

Saving a Business Application as an Application Template

If you would like to save the current application definition as a template, so it can be reused or set to open SAV using that definition, see [Business Application Templates](#) on page 77.

ACLs and Server Pool Configurations

For network devices — such as firewalls, load balancers, routers, switches — you can view Access Control List (ACL) configuration information. For load balancers, you can view server pool configuration (in addition to ACLs).

You also can compare ACL and server pool configurations between two devices in the same snapshot, or when in comparison mode, between the *same* device in the two *different* snapshots.

This section shows you how to perform the following tasks:

- [Viewing ACLs](#)
- [Comparing ACLs — Two Devices In Same Snapshot](#)
- [Viewing Server Pool Configuration](#)
- [Comparing Server Pool Configuration](#)

Viewing ACLs

In SAV, you can view ACLs for network devices such as firewalls, load balancers, routers, switches.

To view ACL configuration information, perform the following steps:

- 1 In the Network Map or Devices Tree, select a network device that has ACLs configured for it, right-click, and select View ACL Configuration. The Access Control Lists window opens.

You can also access the Access Control Lists window by:

- From the Properties pane of the selected device, click the View link in the ACLs property.
 - Select View ACL Configuration from the **Manage** menu, or right-click on the device.
- 2 In the Access Control List window, you can select and copy the text information. Also, from the drop-down list at the top of the window, you can select other devices in the snapshot that contain ACLs and view them. You can also search for specific text strings in the current ACL configuration, highlight strings, and perform other search functions.

- 3 When you are finished viewing ACLs, click **Close**.

Comparing ACLs — Two Devices In Same Snapshot

You can compare ACL configurations between any two devices in the same snapshot, or you can compare ACLs between the same device in two different snapshots. For information on how to compare the same device between two snapshots, see [Comparing ACLs — Same Device Between Two Snapshots](#) on page 87.

To compare ACL configurations for two devices in the same snapshot, perform the following steps:

- 1 From the Devices Tree, expand the Network Devices node.
- 2 Press and hold the Control button on your keyboard, and then select two network devices that have ACLs configured, such as a load balancer, a firewall, router, or LAN switch. (To see if a network device has ACLs configured, select the device and look for ACLs in the Properties pane.)
- 3 From the **Manage** menu (or, right-click), select **Compare ACL Configuration**. In the Compare window, you see two panes side by side, each representing one of the devices you are comparing. At the bottom of each pane is the name of the device.

To indicate the differences for each ACL configuration, the Comparison window uses the following colors:


- **Green:** This indicates that information only exists in device on the right side of the window.
- **Blue:** This indicates that information has been modified.
- **Red:** This indicates that information only exists in device on the left side of the window.
- **Black:** This indicates no changes.


To move through differences between the two configurations, click the arrow buttons at the right top of the window.

- 4 When you are finished viewing the differences, click **Close**.

Comparing ACLs — Same Device Between Two Snapshots

To compare ACLs between the same device in two different snapshots, perform the following steps:

- 1 From the **View** menu, select **Compare**. (Or, click **Compare**  on the SAV toolbar). The Compare pane appears.
- 2 Click **Select**. The Select Comparison Snapshot window opens.
- 3 Select a snapshot you want to compare against the currently loaded snapshot.
- 4 Next, select if you want to show either ANY or ALL of the selected objects (plus their attributes) that are different to show when you compare snapshots.
- 5 From the drop-down list, select an object category for the device type for which you want to compare ACLs.

- 6 As you create comparison rules, SAV automatically displays any or all differences in the Differences pane.
- 7 To view differences in other maps, you can select a map, and all objects that are found to exist will appear normally. All objects that are listed as missing appears grayed out. You can also select other criteria in the Difference pane drop-down list to filter the results of the comparison in more granular detail.
- 1 To close the Compare pane, from the **View** menu, select **Compare** again. (Or, click **Compare**  on the SAV toolbar.)

Comparing Server Pool Configuration

You can compare server pool configurations between any two load balancers in the same snapshot.

(To compare server pool configurations between the same load balancer in two different snapshots, see [Comparing Snapshots](#) on page 89 for more information.)

To compare server pool configurations for two load balancers in the same snapshot, perform the following steps:

- 1 From the Devices Tree, expand the Network Devices node.
- 2 Press and hold the Control button on your keyboard, and then select two load balancers that have a server pool configuration. (To see if a load balancer has a server pool configured, select the device and look for Server Pools in the Properties pane.)
- 3 From the **Manage** menu (or, right-click), select **Compare Server Pool Configuration**. In the Compare window, you see two panes side by side, each representing one of the devices you are comparing. At the bottom of each pane is the name of the device.

To indicate the differences for each server pool configuration, the Comparison window uses the following colors:

- **Green:** This indicates that information only exists in device on the right side of the window.
- **Blue:** This indicates that information has been modified.
- **Red:** This indicates that information only exists in device on the left side of the window.
- **Black:** This indicates no changes.

To move through differences between the two configurations, click the arrow buttons at the right top of the window.

- 4 When you are finished viewing the differences, click **Close**.

Viewing Server Pool Configuration


To view server pool configuration for a load balancer, perform the following steps:

- 1 In the Network Map or Devices Tree, select a load balancer that has server pool configurations, right-click, and select **View Server Pool Configuration**. The Server Pools window opens.

You can also access the Server Pools window by:

- From the Properties pane of the selected device, click the View Server Pools in the load balancer properties.
 - Select View Server Pool Configuration from the **Manage** menu.
- 2 In the Server Pools window, you can select and copy the text information. Also, from the drop-down list at the top of the window, you can select other devices in the snapshot that contain server pool configurations and view them. You can also search for specific text strings in the current server pool configurations, highlight strings, and perform other search functions.
 - 3 When you are finished viewing server pool configuration, click **Close**.

Comparing Snapshots

When you click **Refresh Snapshot**  on the SAV toolbar and then click **Save**, SAV captures and saves all information related to your Business Application. A snapshot including all servers and processes associated with the business application, the current state of all running processes, all local and remote storage devices, and all values and signature definitions you have created in the Tiers tree.

Each time you refresh a snapshot and then save the Business Application, snapshot results are saved within the currently loaded Business Application. You can also schedule snapshots to occur at later time on a one time or recurring schedule.


Each saved snapshot can be used in a one to one comparison between the currently loaded snapshot and a saved one. You can take snapshots from the currently loaded Business Application or from another Business Application to help you determine if any changes have occurred between the current state of the business application and its state as captured in a previously saved snapshot.

When you compare scan results, SAV evaluates certain key objects and their attributes on a one to one basis, and displays any differences in value between those objects. The results of the comparison are displayed in the Differences pane in the SAV window.

Creating a Snapshot

Create a snapshot in SAV any time you want to capture the current state of your business application. Because a data center and all the devices and elements within it are constantly changing, it is a good idea to capture the current state so you can compare the current state of a business application with one you captured in the past.

To create a snapshot, perform the following steps:

- 1 Click **Refresh Snapshot**  on the SAV toolbar.
- 2 From the **File** menu, select **Save**. (Or, click **Save** on the SAV toolbar) A new snapshot has been created.
- 3 To see the snapshot and give it a name, from the **Application** menu, select **Show Snapshots**.
- 4 In the Snapshots window displays all saved snapshots. To rename a snapshot, click the name cell of the list and type a name.

Opening a Snapshot

To view a previous state of a business application, you can load and view a saved snapshot.

To open a saved snapshot, perform the following steps:

- 1 From the **Application** menu, select **Show Snapshots**.
- 2 In the Snapshots window, select a saved snapshot and click **Open**.
- 3 The business application snapshot opens inside of the SAV application window. To delete a snapshot, select it and click **Delete**.

Scheduling a Snapshot

You can automate snapshot creation by scheduling a snapshot at a future point in time, or you can schedule a recurring snapshot to regularly capture the state of your business application.



You can only schedule snapshots for those Business Applications that have been saved to the SA Client Library

To schedule a snapshot, perform the following steps:

- 1 From the **Application** menu, select **Scheduled Snapshots**.
- 2 In the Scheduled Snapshots window, click **New Schedule**.
- 3 Type a name for the snapshot schedule in the Name field.
- 4 In the Scheduled Frequency section, select one of the following snapshot frequency options:
 - **Daily**: Choose to run the snapshot on a daily basis.
 - **Weekly**: Choose a day of the week to run the snapshot.
 - **Monthly**: Choose the months to run the snapshot specification job.
 - **Custom**: In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:


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

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.

- 5 In the Time and Duration section, select a start and end time and day of the month.
- 6 When you are finished filling out the schedule, click **Close**.

"Source" and "Comparison" Snapshot

The currently loaded snapshot in SAV is referred to as the *source* snapshot, while the set of scan results you are comparing against the currently loaded snapshot is called the *comparison* snapshot. When you compare snapshots, you are always comparing the currently

loaded scan result (*source*) with another saved snapshot result (*comparison*). (Remember that to create a new snapshot, you need to click **Refresh Snapshot**  on the SAV toolbar and then click **Save**.)

Comparison Types

SAV displays comparison results based on the following criteria:

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

Object Existence Comparison

Comparing two snapshots helps determine whether or not an object exists between the them. If an object exists in one snapshot, but does not exist in the other, the comparison results display the object with an attribute named “existence” and describes it as either “found” or “missing” on either the source or comparison snapshot.

A process family could be running (and thus, “exist”) when you refresh a snapshot and save the Business Application. But if the process is no longer running, and you refresh the snapshot again and save the Business Application, the results change. When you compare the saved snapshot (the “comparison” snapshot) with the currently loaded snapshot (source), the results of the comparison display in the Differences pane.

The results of the selected row show that on the target snapshot, all the listed LUN volumes are missing, meaning they did not exist in the comparison snapshot, but they exist now in the current snapshot.

Object Attribute Difference

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

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

Table 2 lists all object attributes evaluated during a scan results comparison.

Table 2 Object Attributes Checked for Difference in Snapshot Comparison

Object Category	Object Attributes Compared
Database	Existence, version
Database file	Existence, size (in bytes)
Disk (applies to servers, SAN arrays, and NAS filers)	Existence, firmware version, status
Compliance Policy	Existence, compliance status (compliant, partial, noncompliant, scan failure, scan needed, scanning)
Fibre Channel Adapter	Driver version
Fibre Channel Port (applies to servers, SAN arrays, and NAS filers)	Connected port, existence, fabric, name, zone
File System	Existence, mount options, mount point, significant change in % free space, size (bytes), type
LUN Volume	Capacity (bytes), existence, LUN ID, name
NAS Filer	Existence, hostname, name, operating system version
NFS Export	Existence, export path, significant change in % of free space, size (bytes)
Network Device	ACLs, existence, firmware version, name, operating system, server pools
Network Interface	Broadcast address, connected switch, connected switch port, connected VLAN, duplex, existence, IP address, MAC address, neg. duplex, neg. speed, subnet mask
Network Port	Duplex, existence, MAC address, neg. duplex, neg. speed, speed, VLAN
Process Family	Max CPU utilization, significant change in # of connections, significant change in # of open files
SAN Array	Existence, firmware version, name
SAN Switch	Existence, firmware version, name
Server	Boot time, codeset, DNS server, existence, kernel, name, operating system
Tablespace	Existence, size (in bytes)
VLAN	Existence

“Significant” Object Attribute Differences

SAV also compares a set of attributes by using special heuristics specific to certain attributes, so that differences SAV considers “significant” is shown.

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

Table 3 shows special object attribute differences.


Table 3 “Significant” Object Attribute Differences in SAV

Object	Object Attributes Compared
Server	Load average, percentage of free memory on a server, percentage of free swap memory
Server’s file system	Percentage of free space
Process families	Number of open files, total number of all related connections, total count of process family member connections
NFS Export on NAS Filer	Percentage of free space




For more information on the heuristics used to calculate what is considered a significant difference, see [Significant Scan Result Difference Heuristics](#) on page 94.


Comparing Snapshots

In order to compare two snapshots, you must have at least one saved snapshot. If you select **Save** from the **File** menu, this saves the currently loaded snapshot. If you click **Refresh**

Snapshot  on the SAV toolbar, and then save again, this creates and save a new snapshot.

To compare snapshots in SAV, perform the following steps:

- 1 From the **View** menu, select **Compare**. (Or, click **Compare**  on the SAV toolbar). The Compare pane appears.
- 2 Click **Select**. The Select Comparison Snapshot window opens.
- 3 Select a snapshot you want to compare against the currently loaded snapshot.
- 4 Next, select if you want to show either ANY or ALL of the selected objects (plus their attributes) that are different to show when you compare snapshots.
- 5 From the drop-down list, select an object category to compare and create a comparison rule. You can select any of the object categories, such as file systems, or select All Categories.
- 6 To select another comparison rule, click **Add**  to add another criteria selector. To remove a comparison rule, click **Remove** .
- 7 As you create comparison rules, SAV automatically displays any or all differences in the Differences pane.
- 8 To view differences in other maps, you can select a map, and all objects that are found to exist will appear normally. All objects that are listed as missing appears grayed out. You can also select other criteria in the Difference pane drop-down list to filter the results of the comparison in more granular detail.

- 9 To close the Compare pane, from the **View** menu, select **Compare** again. (Or, click **Compare**  on the SAV toolbar.)

Significant Scan Result Difference Heuristics

When you compare scan results in SAV, specific objects and their attributes are evaluated between each scan result and any differences (and non-existence of objects) is displayed in the comparison results. (For information on the standard set of objects and attributes compared in a scan results comparison, see [Comparing Snapshots](#) on page 89.)

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

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

The following special object attribute differences are compared:

- **Server:** Load average, percentage of free memory on a server, percentage of free swap memory.
- **Server's File System:** Percentage of free space.
- **NAS Filer NFS Exported File System:** Percentage of free space.
- **Process Families:** Number of open files, total number of all related connections, total count of process family member connections.

The heuristics applied to certain attributes in scan results during a comparison are listed in [Table 4](#).

The following variables are used in the expressions:

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

Table 4 Scan Results Comparison Heuristics

Object Attribute	Equation
server — 15 minutes load average	$X > 0.8 * \text{cpu count AND } P > 20\%$ OR $X > \max(1, 0.25 * \text{cpu count}) \text{ AND } P > 100\%$
server — percentage memory free (%)	$N < 0.1 * \text{total mem AND } P > 25\%$
filesystem — percentage free (%)	$N < 0.2 * \text{size AND } P > 10\%$

Table 4 Scan Results Comparison Heuristics

Object Attribute	Equation
process family — open file count (on any member process)	$X > 50 \text{ AND } P > 50\%$
process family — connection count (aggregate across all member processes)	$X > 50 * \text{process count AND } P > 30\%$
process family - connection count (on any member process)	$X > 50 \text{ AND } P > 50\%$

Filtering SAV Data


When you select Filter from the **View** menu, or click **Filter**  on the SAV toolbar, a search control appears above the maps and tabs that allows you to filter according to the following objects (if they were captured as part of the scan) and their attributes, listed in [Table 5](#).

Table 5 Objects and Their Attributes You Can Filter in SAV

Object	Attributes Filtered
Compliance Policy	Name, compliance status (compliant, partial, noncompliant, scan failure, scan needed, scanning)
Databases	Name, status, type, version
Database Files	Name, location, type, path, description, size, free space, % used, status.
Disk (applies to servers, SAN arrays, and NAS filers)	Device, manufacturer, model, serial number, size (bytes), type
Fibre Channel Adapter	Driver version, firmware version, hardware version, model, node world wide name, serial number
Fibre Channel Port (applies to servers, SAN arrays, and NAS filers)	Fabric, name, port number, status, world wide name
File System	% used, device, free space (bytes), mount options, mount point, size (bytes), type
LUN Volume	Capacity (bytes), name
NAS Filer	Hardware version, hostname, manufacturer, model, name, operating system version, Server ID, serial number
NFS Export	% used, export path, free space (bytes), size (bytes)
Network Device	Asset tag, firmware version, manufacturer, model, name, operating system, Server ID, processor

Table 5 Objects and Their Attributes You Can Filter in SAV

Object	Attributes Filtered
Network Interface	Broadcast address, device, duplex, IP address, MAC address, neg. duplex, neg. speed, speed
Network Port	Duplex, MAC address, neg. duplex, neg. speed, speed
Process Family	CPU utilization, command line, connected port, environment variables, listener port, modules, name, open files
SAN Array	Firmware version, manufacturer, model, name, Server ID, serial number
SAN Switch	Firmware version, hardware version, manufacturer, model, name, Server ID, serial number
SAN Zone	Name
Server	1 minute load average, 15 minute load average, 5 minute load average, codeset, free memory (bytes), kernel, name, operating system
Tablespaces	Name, location, description, size, free space, % used, status.
VLAN	Descriptions, ports, VLAN ID


You can filter the current snapshot in according to one or several of the objects in the list, as well as applying operators and attributes, according to certain attributes, and then apply operators to the attributes, depending if the object is a string or a number.



Filter results appear in the Tabs and Maps. In the maps, while all other objects that do not meet the filtering criteria appear grayed out in the maps.

For more information on filtering criteria and regular expressions, see [Filter Criteria](#) on page 97.

Creating a Data Filter in SAV

To filter data that was collected in the currently loaded SAV snapshot, perform the following steps:

- 1 From the **View** menu, select **Filter**. (Or, from the SAV toolbar, click **Filter** ) The Filter pane appears above the maps.
- 2 In the Filter pane rule criteria, choose if you want to show either ANY or ALL of the selected objects (plus their attributes) that meet your filtering criteria.
- 3 From the drop-down list, select an object category to filter in the current snapshot and add criteria to narrow the filter, such as Compliance Policy, File System, Process Family, SAN Array, and so on.
- 4 Using the criteria drop down list, create a meaningful expression. For example, if you chose Disk as a category, you could set Size (bytes) is greater than (>) 5000 (bytes). For more information on expressions, see [Filter Criteria](#) on page 97 and [Examples of Regular Expressions](#) on page 98.

- 5 To select another filter criteria rule, click **Add** . To remove a comparison rule, click **Remove** .
- 6 As you create filter rules, SAV automatically displays any or all results in any of the maps or the Infrastructure pane. All results that meet the criteria appear normally in the maps and the Infrastructure pane. Any results in the snapshot that do not meet the criteria are shown in the maps grayed out.
- 7 From any of the results related to server, you can select the server on which the results were found, right-click, and select **Open Remote Terminal** or **Open Device Explorer** to browse the server.

Filter Criteria

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

Strings Operators

- Contains (default)
- Does Not Contain
- Is
- Is Not
- Starts With
- Ends With
- Matches Regular Expression

Numbers

- == is equal to
- != not
- < less than
- > Greater than
- <= less than or equal to
- >= greater than or equal to

All filtering is performed in case-sensitive mode.

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

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




Examples of Regular Expressions

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

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

SAV Scan Error Messages

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

- **Server Error Icon** : There was an error in gathering information from the server when SAV scanned it (see [Table 6](#) for possible causes for the error).
- **Server Unreachable Error Icon** : The SA core was not able to communicate with the SA Agent installed on the server.
- **Server Unknown** : SAV is unable to scan the server at all, possibly because the server is no longer in the core and under SA management.

It shows these icons before the server name in the Devices tree, Network Map, Virtualization Map, and Server Map. You can move your cursor over the server name to display the detailed error message.



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

Server Scan Errors

Table 6 describes server scan errors and recommended actions.

Table 6 Server Error Messages in SAV

Error	Description	Action
Not Enough Disk Space	A selected managed server does not have enough disk space to perform a scan.	Free up disk space.
Remediation Failed	The Runtime State Server Module failed to remediate on the selected server.	Select the server from the Devices Tree, and then in the property pane. Click the Remediation job number link and the job window from the SA Client opens. Or, select the server, right-click, and select Open Device Explorer to troubleshoot the error.
Scan Timed Out	The scan process has exceeded the time-out limit.	See Scan Time-Out Preference on page 74.
Server Access Denied	By using the OGFS, you are unable to access the server's file system as root (on a Unix server) or as LocalSystem (on a Windows server).	Contact your SA administrator for the required permissions.
Server Capture Failed	The remote capture of data or the transfer of data back to the SA core failed.	Review the log file that is in /tmp/.sitemap/<number> for details in your global shell session.
Server ID Invalid	The server's directory was not found in the OGFS, which means that SA does not know the server exists.	
Server Scan Agent Failed	The driver used to collect data could not be correctly copied to the managed server. This could be caused by a checksum mismatch.	Contact HP Support and provide the log file.
Server Unreachable	The managed server is unreachable by SA. This could be caused if the SA core cannot communicate with the server's agent.	Try again later. If this condition persists, contact your HP administrator.

Table 6 Server Error Messages in SAV (cont'd)

Error	Description	Action
Unknown Scan Error	An unknown error occurred during the scanning process.	Try again later. If this condition persists, contact your HP administrator.
Unsupported Agent for Scan	The SAV does not support the Server Agent version running on a selected managed server.	SA Agent 7.0 or higher is required.
Unsupported OS for Scan	The SAV does not support the operating system running on a selected managed server.	See Supported Operating Systems on page 21.

Network Device Scan Errors

[Table 7](#) describes network device scan errors and recommended actions.

Table 7 Network Device Scan Error Messages in SAV

Error	Description	Action
NA Scan Timed Out	The time needed to gather NA data exceeded the timeout	Scan fewer devices or wait until the NA server can handle this request.
NA Scan Failed	Gathering NA data failed.	Save this snapshot to a Business Application and contact your SA administrator.

Storage Scan Errors

[Table 8](#) describes network device scan errors and recommended actions.

Table 8 Storage Scan Error Messages in SAV

Error	Description	Action
Scan Timed Out	The time needed to gather storage data exceeded the timeout	Scan fewer devices or wait until the SA core server can handle this request.
NAS Scan Failed	Gathering storage data failed.	Export this snapshot to a Business Application and contact your SA administrator.

SAV Platform Support

This section provides information about the operating system platforms and architecture that SAV supports scanning and displaying application (process families), server, and device information.



This list of operating system support for SAV is a subset of the supported platforms for the SA Agent, since in order for SAV to be able to fully scan a server it must be under SA management with an SA Agent. For more information on supported platforms for the SA Agent, see the chapter on server asset tracking in the *SA Users Guide: Server Automation*

Supported Platforms in SAV

For non-Linux and non-VMware operating systems, SAV supports each operating systems kernel out of the box, and assumes no customizations have been made. For a list of non-Linux and non-VMware operating systems and kernels listed supported by SAV, see [Table 9](#).

For Linux and VMware ESX 3 operating systems, there are certain out of the box kernel versions that SAV supports. For information on Linux and VMware operating systems and kernels supported by SAV, see [Table 10](#). Deprecated operating systems are listed with a (*).

Table 9 SAV Supported Operating Systems – Non-Linux/VMware

SAV Supported Operating Systems	OS Versions	Architecture
AIX		
	AIX 4.3* AIX 5.1* AIX 5.2 AIX 5.3 AIX 6.1	POWER 32 bit and 64 bit
HP-UX		
	HP-UX 10.20* HP-UX 11.00* HP-UX 11.11	PA-RISC (direct, nPartition)
	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 (direct, nPartition) PA-RISC (direct, nPartition) PA-RISC (direct, nPartition) PA-RISC and Itanium (direct, nPartition) PA-RISC and Itanium (direct, nPartition)
Sun Solaris		

Table 9 SAV Supported Operating Systems – Non-Linux/VMware (cont'd)

SAV Supported Operating Systems	OS Versions	Architecture
	Sun Solaris 6* Sun Solaris 7* Sun Solaris 8* Sun Solaris 9*	Sun SPARC
	Solaris 10, Updates 1, 2, 3, 4, 5, and 6)	Sun SPARC, 32 bit x86, 64 bit x86 and Niagara Sun Dynamic System Domains are supported on Sun Solaris 6, 7, 8, 9, and 10. Guest logical domains are supported on Solaris 10.
Fujitsu Solaris		
	Fujitsu Solaris 8* Fujitsu Solaris 9* Fujitsu Solaris 10	Fujitsu SPARC
Windows		
	Windows NT 4.0*	32 bit x86
	Windows 2000 Server Family Windows Server 2003	
	Windows Server 2003 x64	64 bit x86 (not Itanium)
	Windows Server 2008	32 bit x86 64 bit x86 (not Itanium)
	Windows XP Professional	32 bit x86
	Windows XP Professional x64	64 bit x86 (not Itanium)

Table 10 SAV Supported Operating Systems – Linux and VMware

SAV Supported Operating Systems	Versions	Kernel	Architecture
Red Hat Linux			
	Red Hat Enterprise Linux 2.1* AS Red Hat Enterprise Linux 2.1*ES Red Hat Enterprise Linux 2.1* WS	2.4.9	32 bit x86
	Red Hat Enterprise Linux 3* AS Red Hat Enterprise Linux 3 ES Red Hat Enterprise Linux 3 WS	2.4.21-x.EL	32 bit x86 64 bit x86
	Red Hat Enterprise Linux 4 AS Red Hat Enterprise Linux 4 ES Red Hat Enterprise Linux 4 WS	2.6.9-x.EL	32 bit x86 64 bit x86 Itanium
	Red Hat Enterprise Linux Desktop 5 Red Hat Enterprise Linux Server 5	2.6.18-8.el5xen	32 bit x86 64 bit x86 Itanium
SUSE Linux			
	SUSE Linux Enterprise Server 8* SUSE Linux Standard Server 8*	2.4.18	32 bit x86
	SUSE Linux Enterprise Server 9	2.4.21, 2.6.5	32 bit x86 64 bit x86
	SUSE Linux Enterprise Server 10	2.6.16.13, 2.6.16.21	32 bit x86 64 bit x86
VMware			
	ESX Server 3 ESX Server 3.0.1 ESX Server 3.0.2 ESX Server 3.0.3	2.4.21-37.0.2.E Lvmnix	32 bit x86 64 bit x86
	VMware ESX Server 3.5	2.4.21-47.0.1.E Lvmnix	32 bit x86 64 bit x86
	VMware ESXi Server 3.5		32 bit x86 64 bit x86

2 Audit and Remediation

Overview of Audit and Remediation

The Audit and Remediation feature allows you to define server configuration policies to help you ensure that servers in your facilities meet policy standards. When servers are found to be out of compliance — not configured the way you want them to be — you can remediate them to conform to your organization's standards.

With Audit and Remediation, you can audit server configuration values based on a live server (or server snapshot), based upon your own custom values, or based on pre-configured audit policies. You can also take server configuration snapshots to capture the current state of a system, so you can compare other servers against a known baseline.

Audit policies allow you to define company or industry-wide compliance standards, which can then be used inside of audits, snapshot specifications, and other audit policies. Referencing audit policies in your audits or snapshot specifications helps ensure that you are up to date with the latest compliance definitions in your organization.

If you have a content subscription to BSA Essentials Subscription Services, you can be kept up to date on the latest industry compliance standards based on the needs of your data center. For example, Subscription Services give you access to regularly updated security best practices, such as the Center for Internet Security (CIS) and Payment Card Industry (PCI), and so on. It also enables access to additional free non-subscription content such as Microsoft Patch Supplement for Server Automation.

BSA Essentials Subscription Services enable you to access the most current regulatory compliance policies (FISMA, Sarbanes-Oxley, etc.) and daily vulnerability alerts. You can also join the content developer communities on the HP Live Network portal to share and access custom-created audit policies and rules.



For information about subscribing to BSA Essentials Subscription Services, contact your sales representative.



SA Audit and Remediation does not support auditing or taking snapshots of VMware ESXi servers.

Audit and Remediation Examples

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

- [Enforcing Security Standards With Audit Policies](#)
- [Capturing Golden Server Configurations](#)

Enforcing Security Standards With Audit Policies

Your IT organization likely has security policies you want to enforce, so you can be sure that your servers are configured properly and are safe from security attacks. Policy setters in your organization can build audit policies to enforce these security standards. These pre-defined audit policies can be linked to multiple audits or snapshot specifications, so people who manage live servers can reference the right audit policy to ensure their servers are being audited correctly.

For example, your organization might have a farm of Solaris 10 servers that needs to be kept up to date with the most recent commonly known security vulnerabilities as specified by Common Vulnerabilities and Exposures (CVE). For example, your company wants to make sure your servers are not vulnerable to a known threat to Solaris 10, such as CVE-2009-0168 (CVSS 4.9) which checks for an unspecified vulnerability in `ppdmgr` in Sun Solaris 10 and OpenSolaris `snv_61` through `snv_106`.

By subscribing to the BSA Essentials Subscription Services, you can access an online collection of compliances check that you can use to audit your Solaris 10 servers to make sure that they are not open to this vulnerability. The person in your organization who is responsible for defining compliance standards can build an audit policy that contains the CVE-2009-0168 compliance check.

System administrators who are responsible for managing the Solaris servers can create audits for their servers and link their audit's rules to this audit policy. When an audit links to an audit policy, any changes made to the policy are immediately reflected in the audit, so the person who runs the audits on the servers knows that the audit rules are always up to date. So, if a new CVE update came out for Solaris 10 servers, the policy setter could update the policy, and all audits that link to the policy will have the latest compliance definitions.

Knowing that his audit will always contain the latest vulnerabilities checks, he can schedule the audit to run regularly to check all of the Solaris 10 servers he manages. If the audit results show that any of the target servers do not contain the new CVE security check, those servers can be remediated to fix the problem.

Capturing Golden Server Configurations

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

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

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

Audits

An audit defines a set of rules or configuration values you use to determine if the configuration of a server or group of servers match your organization's desired compliance standards. The rules of an audit can be configured in an ad-hoc manner, or more effectively, reference a pre-configured audit policy that clearly defines the desired state of a server's configuration.

The audit can compare a servers' configuration against the rules defined in the audit, check that a configuration value meets the criteria specified in the audit rule, or simply check to ensure that a specific value does or does not exist. Some rules also allow you run scripts to seek to capture more difficult to ascertain configuration values.

For example, you can define the audit to look for such things as whether or not an IIS Metabase value exist (you may not want it to), to make sure a specific Linux services is set to always be running (a critical service that needs to be running for security reasons), to determine if a certain file system directory does not exceed a certain size limit, or to make sure that the maximum length setting for user passwords has not been exceeded, and so on. You can define what the audit should look for, what values you expect to find on the server, and what value to use to fix them when differences are found.

Once configured, an audit can be run once, scheduled for a future run, or be set to run a regular basis. After an audit is run, its results indicate the extent to which those servers meet the definitions set in the audit's rules. In cases where discrepancies are found, you can remediate those servers and bring them into compliance.

For more information on audits, see [Audits](#) on page 110.

Audit Policies

Audit policies consists of a set of reusable rules that allow you to define the desired state of server configuration, based upon industry standards and the compliance goals set by your organization. Audit policies can be linked to audits, snapshot specifications, and other audit policies. When any changes are made to the audit policies, all references to the audit policy are also updated.

Typically, audit policies are created by *policy setter* users who understand the exact compliance standards that a company wants its servers to meet for a specific configuration domain and platform. Users who manage actual servers can utilized the predefined audit policies by linking them to their audits or snapshot specifications. If any changes are made to the audit policy, the audit that links to it also contains the updated rules. This way, those who audit servers can be sure their audits always reflect the latest policy standards in their organization.

For more information on audit policies, see [Audit Policies](#) on page 166.

Audits and the Compliance View

The Compliance View allows users to view the overall compliance levels for servers in their facility and helps them remediate compliance problems. For more information, see [Server Compliance](#) on page 207.

Snapshots

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

For more information on snapshots, see [Snapshots](#) on page 190.

Terms and Concepts

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

- **Archived Audit Result/Snapshot:** Archiving audit results and snapshots allows you to move them from the audit result or snapshot list but keep them available for historical purposes.
- **Audit:** A set of rules (which may contain individual “checks”) that expresses the desired state of a managed server’s configuration objects — for example, a server’s file system directory structure or files, a server’s Windows Registry, application configuration, and so on. An audit also contains sources (servers, snapshots, snapshot specifications), targets (servers or snapshots), rule exceptions, and a schedule.

An audit’s rules can be linked to an audit policy, which means the rules of the audit policy are substituted for those in the audit. An audit can be run to compare server configuration object values against a baseline server, a server snapshot, or user-defined values, to determine how values differ. When an audit reveals a difference between servers or user-entered values, the user can install software and server objects to remediate the variance.

- **Audit Job:** The process that occurs when you run an Audit. An audit job can be run immediately one time, or on a recurring basis by scheduling the job. When an audit job is finished, it produces an Audit Result.
- **Audit rule types:** An audit can contain both types of the following rules:
 - **Comparison:** A rule that compares a server’s or snapshot’s configurations of a server with other servers or snapshots.
 - **Value-based (user-specified):** A rule that compares one or more set of user-defined values. This type of audit includes an audit that links to an audit policy.
 - **Non-Existence:** A rule that checks for the non-existence of an object to determine if it exists on the target server. If the object exists on the target server, then the user or group rule is out of compliance.
- **Audit policy:** A collection of rules that defines a desired configuration for a server. A policy can be used by an audit in the following ways:
 - **Link:** A linked policy maintains a persistent connection between the audit and the policy. This means that the rules in the audit are exactly those of the audit policy, and if any updates are made to the policy, then the latest changes are also reflected in the Audit to which the policy is linked. When an audit policy is linked to an audit or snapshot specification, the rules are shown inside the audit or snapshot specification as read-only. (However, the rules inside the audit policy are still editable.)

- **Import (replace, non-linked):** When a user imports a policy into an audit, the connection between the audit and the audit policy is no longer maintained, and the user can make changes to the audit without affecting the policy. Conversely, any changes or updates made to the policy will not be reflected in the Audit.
- **Import (merge):** When an audit policy is imported and merged into an audit, the audit policy's rules are added to the rules already present in the audit. No persistent link between the audit and the audit policy is maintained. During the merge, if rules are found to conflict, the newly imported rules from the audit policy will replace the rules in the audit policy.
- **Audit Result:** The results of running an Audit. This shows how a target server or a group of servers' configuration object values match or mismatch the values as defined in the audit.
- **Exception:** A server and specific rules that has been excepted, or disabled, so that when the audit is run, the rule exception is not checked on the selected server thus not considered when determining audit compliance.
- **Compliance:** The degree to which a server configuration conforms to a check or test established in a set of rules defined in an audit, snapshot specification, or audit policy. Compliance in Audit and Remediation is defined by the audit's or snapshot's rules, which specify the values expected of the target servers. If the values on the target server are different than specified in the audit's rules, then the server is considered Non-Compliant.
- **Policy Setter:** A person in an organization who is responsible for defining server configuration compliance standards — the way a server should be configured — and who defines audit policies.
- **Rule:** A check on a particular server configuration object along with a desired value, and optional remediation value. Rules come in two types: server-based, which derive directly from a source server, and user-defined, which are created by a user.

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

- **Server Object:** An object from a server to which an audit or snapshot specification rule can be applied. This can be a value (such as minimum password length) or an object, such as a file or directory, registry entry, Windows Services hardware configuration, and so on. For more information on servers objects used in audits and snapshot specifications, see [Server Objects Used in Audits and Snapshots](#) on page 121.
- **Snapshot:** Shows a picture of how an SA managed server is configured at a certain point in time. A Snapshot is the result of a snapshot specification job that has been run.
- **Snapshot Specification Job:** The process that occurs when you run a snapshot specification. A Snapshot job can be run once, or on a recurring basis by scheduling the job. When a snapshot specification job is completed, it produces a Snapshot.
- **Snapshot Specification:** An object window that allows you to define and create a snapshot. In other words, you can define the rules and servers to take a snapshot of.
- **Target:** The server or servers that you run an audit against or take a snapshot of. The target for an audit can be a server, several servers, a group of servers, or a snapshot. The target for a snapshot can also be other servers.

Audits

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

Audit rules allow you to define and check the state of various configurations or objects and files on a server, such as the state of server's file system, registry settings, installed and registered software (patches and packages), events, software, application configurations, operating system settings, and so on. (For more information on configuring audit rules, see [Configuring Specific Audit and Snapshot Rules](#) on page 126.)

If the configuration or object on the target server is different than the state you defined in the audit rules, the rule is considered Non-Compliant. When you view an audit's results, you can remediate the object configuration to make sure the target server's configuration is in compliance with the desired configuration. (For more information on remediating audit results, see [Remediating Audit Results](#) on page 176.)

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

For more information about viewing the results of audit jobs, see [Finding Information in Job Results](#) on page 113.

Audit Comparison Types

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

- **Comparison:** An audit based on configuration values from a source server or source snapshot specified at the time the audit is created. The source server or server snapshot is also known as a “golden” or reference server. For example, you might want to compare file directories or file contents, registry structures, IIS Metabase entries, or user group settings among servers. Using a snapshot as the source of an audit, you can compare the snapshot with other servers in your facility.

Comparison audits can perform the following types of comparisons:

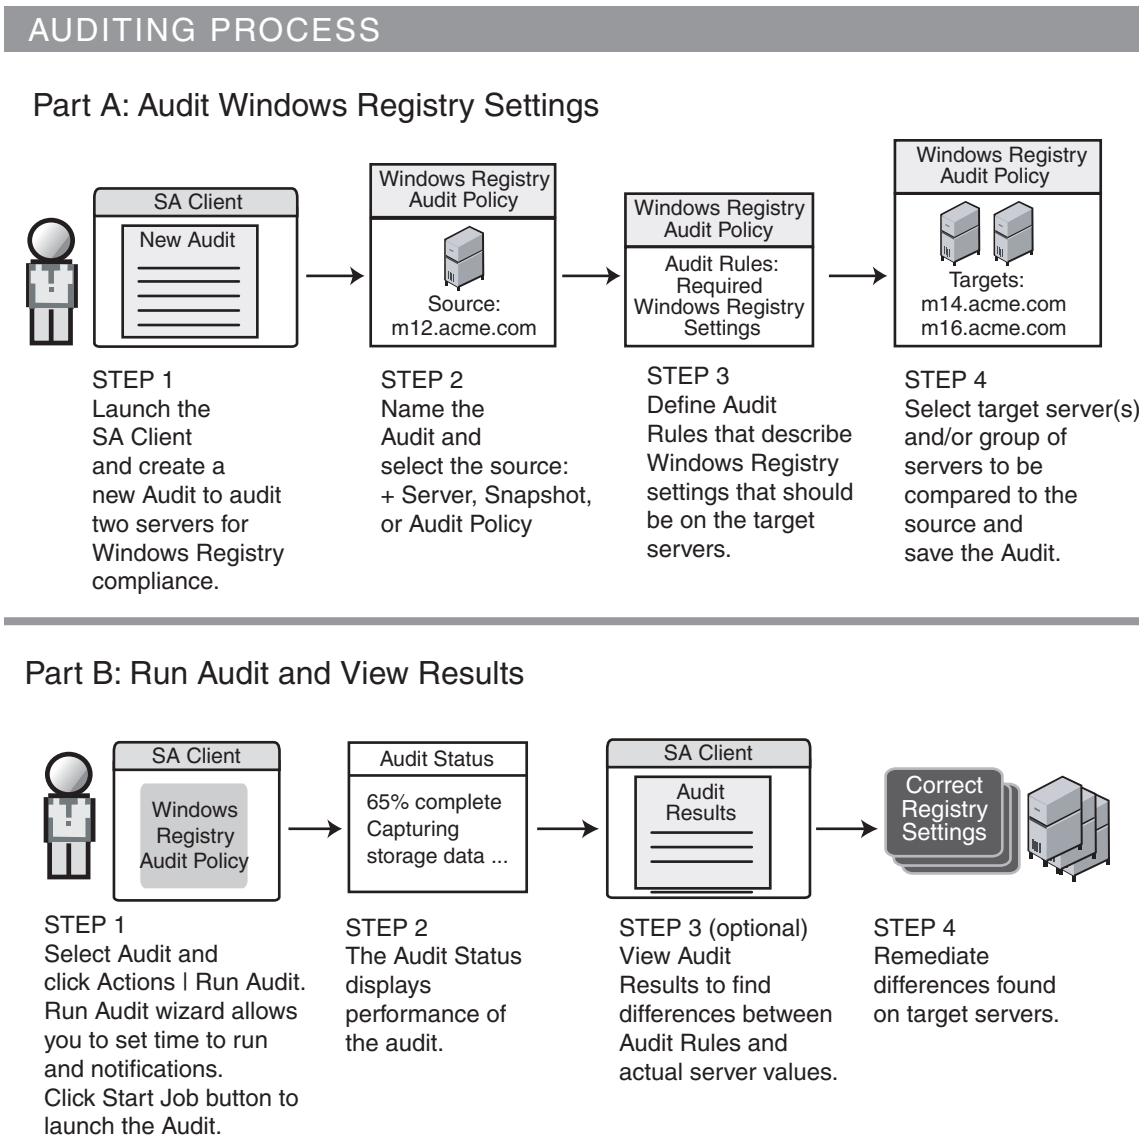
- **Property:** Checks the property of a selected object or object configuration. For example, you could check the release version of a patch on a target server or group of servers, to make sure it matches what you expect to be installed on the targets. You can select this version number based upon a source server or snapshot, or add your own value.
 - **Equivalence:** Checks to determine that a target server configuration is the same between the source server or snapshot of the audit. For example, you could check to see if the target of the audit has the same user group as a group you selected from a source server.
 - **Non-existence:** Checks the target server to determine the non-existence of a server object or configuration. For example, you could check a server to make sure it does not contain a specific COM+ object.
- **User-Defined Value Comparison:** An audit based on custom, user-defined values for each server object (file system, windows services, IIS Metabase, users and groups, and so on). These values can be derived from a source server, or from SA attributes or custom

attributes. This type of audit includes those based on an audit policy. In an audit policy, a user (known as a “policy setter”) pre-defines values for each configuration object based on company or industry compliance standards.

The Auditing Process

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

Figure 17 The Auditing Process



Audit Elements

An audit consists of the following elements:

- **Properties:** The name and description of the audit.
- **Source:** The source of an audit can be a server, a snapshot, or no source at all. (However, some rules require a source.) Choosing a server as the source for an audit allows you to select server objects from that server as the basis of your audit. Choosing a snapshot as the source of an audit allows you to use the configuration values of the snapshot. Choose a snapshot specification as the source allows you to audit a server against itself over time.

For example, if you took a snapshot of a server, then used that snapshot specification as the source of the audit, every time you run the audit, you can compare the original state of the server against the server's actual configuration over time (using a recurring audit schedule). If you choose no source, then you can define only your own custom values for the audit or snapshot.

- **Rules:** A check on a particular server object with a desired value and an optional remediation value. For example, you might check to see if this server contains a specific Windows Service, and if found, determine if the service is turned off. For a description of server objects and rules, see [Server Objects Used in Audits and Snapshots](#) on page 121.
- **Targets:** The servers that the audit will check for compliance. You can choose as many servers and groups of servers as needed for an audit or snapshot.

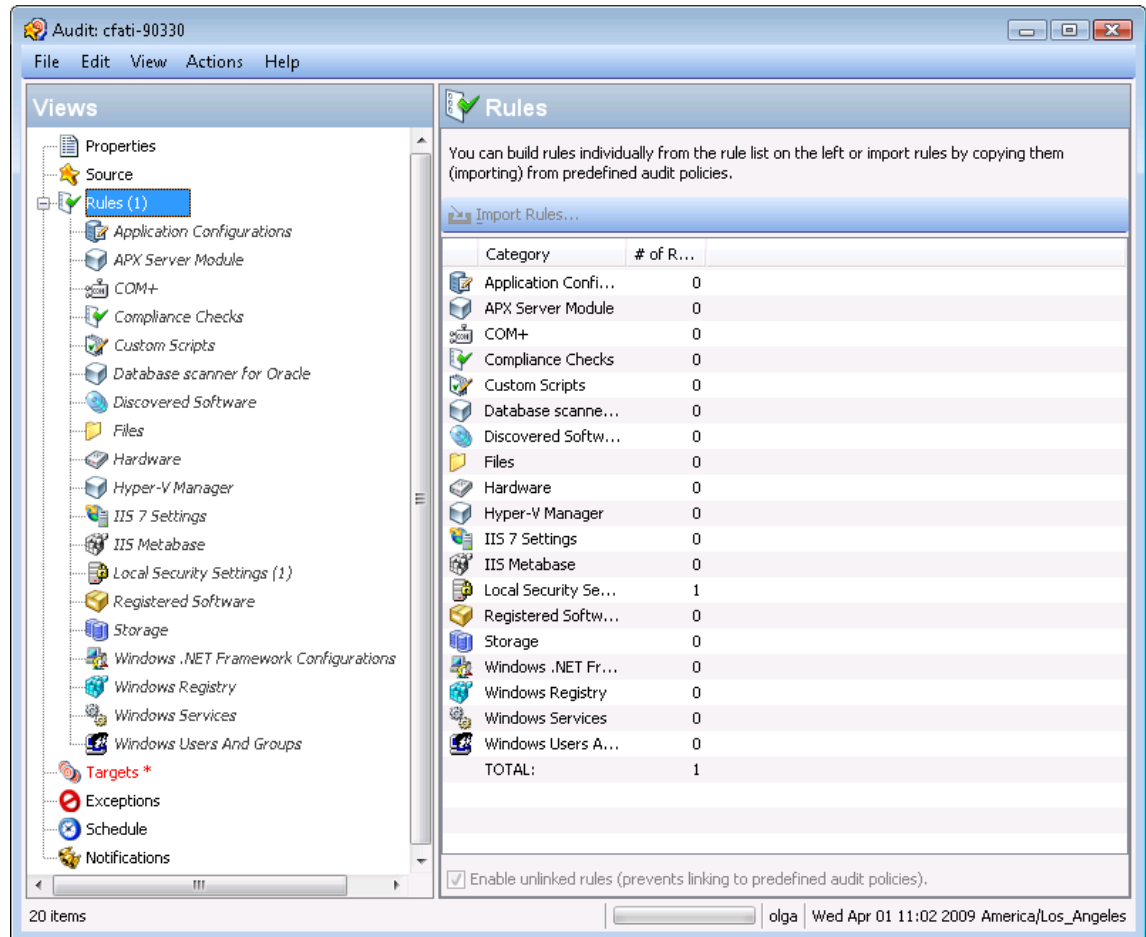


VMware ESXi servers cannot be the target of an audit or snapshot.

- **Exceptions:** Servers and specific rules that will not be checked for compliance when the audit is run.
- **Schedule:** You can run an audit on a onetime basis, or on a recurring schedule. Audits that run on a recurring schedule appear as a single compliance column in the Compliance Dashboard.
- **Notifications:** You can send emails when the audit has finished running, and base the notification on the success, failure, or the completion of an audit job.

To configure an audit, select server configuration objects and then apply rules to those objects in order to define their desired configuration state. For example, [Figure 18](#) shows an audit that has defined four rules. These rules will determine if target server configurations match the rules in the audit. For more information about viewing the results of audit jobs, see [Finding Information in Job Results](#) on page 113.

Figure 18 Audit Window Showing Elements of an Audit



Creating an Audit

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

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

- From a managed server, using the selected server as the source of the audit. You can choose to run the audit on a single server or a group of servers.
- From the Device Groups list, choosing a group of servers as the target at the audit.
- From the Library, by creating a new audit.
- From a snapshot, by creating an audit based on the snapshot.
- From an audit policy, by creating an audit based on the audit policy.

Creating an Audit from a Server

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



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

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

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

For information on how to configure an audit, see [Configuring an Audit](#) on page 117.

Creating an Audit from a Group of Servers

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

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

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

Creating an Audit from the Library

To create a new audit from the SA Client Library, perform the following steps:

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

Creating an Audit from a Snapshot

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

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

Creating an Audit from an Audit Policy

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

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

Saving an Audit as Audit Policy

You can choose to save an audit as an audit policy, which will save only the rules from the audit and create a new audit policy.

All audit policies you create must be saved to the Library in a folder. You must have permissions to write to the folder you want to save the audit policy to. For more information on folder permissions, see *SA Policy Setter's Guide*, or contact your SA administrator.



For more information on creating, using, linking, and importing audit policies, see [Audit Policies](#) on page 107.



You can also save an audit using the “Save As” function to create a new audit with a new name.

To use Save as to create an audit policy from an existing audit (or create a new audit), perform the following steps:

- 1 From inside the Audit or Snapshot Specification window, from the **File** menu, select **Save As**.
- 2 In the Save As window, enter a name. If you are renaming an audit or snapshot specification, you must use a unique name.
- 3 (Optional) Enter a description.
- 4 From the Type drop-down list, select either Audit or Audit Policy.
- 5 If you selected Audit Policy, from the Location section, click Select.
- 6 Select a folder in the SA Client library to save the audit policy to. (You must have write permissions on the folder to save the audit policy.)
- 7 Click **OK**.

Viewing Server Audit and Snapshot Usage

After you create and run an audit, you can view it from the All Managed Servers list or from the Device Explorer, and see all audits that are associated with a specific server.

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

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

- 1 From the Navigation pane, select **Devices ► Servers ► All Managed Servers**.
- 2 In the Content pane, select a server.
- 3 From the View drop-down list, select Audits or Snapshot Specifications. Notice that the lower Details pane shows information about audit and snapshot usage.
- 4 In the Details pane, if you selected Audits., you can choose one of the following options:
 - **Audit - Server is Target:** Shows all audits where the selected server is the target of the audit.
 - **Audit - Server is Source:** Shows all audits where the selected server is used as the source of the audit.
- 5 From any one of these views, you can select an audit or audit results, and perform actions from the Actions menu. For example, you can open an audit, re-run an audit, and so on.
- 6 If you selected Snapshot Specifications, then the Details pane shows all snapshot specifications that target the selected server.

Viewing a Server's Audit Usage from Device Explorer

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

- 1 From the Navigation pane, select **Devices ► All Managed Servers**.
- 2 In the Content pane, select a server, right-click, and select **Open**.
- 3 In the Device Explorer, from the Views pane, select Management Policies ► Audits.
- 4 In the Content pane, from the Show drop-down list, select one of the following options:
 - **Audit - Server is Target:** Shows all audits where the selected server is the target of the audit.
 - **Audit - Server is Source:** Shows all audits where the selected server is used as the source of the audit.
- 5 From any one of these views, you can select an audit and perform actions from the Actions menu. For example, you can open an audit, re-run an audit, and so on.
- 6 Next, from the Views pane you can select Archived Audit Results to see all audit results associated with this server that have been archived. For more information, see [Archiving Audit Results](#) on page 189.

Configuring an Audit

Configuring an audit or audit policy consists of performing the following general steps:

- Name and describe the audit or audit policy
- Select a source for the audit or audit policy: a server, a snapshot, snapshot specification, or none.
- Configure the audit rules — you have the option of linking to an audit *policy*, specifies that you want to use the rules from an audit policy in your audit (and which disables the ability to configure individual rules), or importing all of the rules of an audit policy into the audit.

For more information, see [Linking and Importing Audit Policies](#) on page 168. For more information about configuring specific rules, see [Configuring Specific Audit and Snapshot Rules](#) on page 126.

- Choose a target server, group of servers, or snapshot to audit
- Add audit rule exceptions (optional)
- Schedule the audit
- Set the Email Notification (optional)
- Save the audit



VMware ESXi servers cannot be the source or the target of an audit or snapshot.

To configure an audit, perform the following steps:

- 1 Create the new audit from one of the methods described in [Creating an Audit](#) on page 113. The Audit window opens.
- 2 Enter the following information for the audit:
 - **Properties:** Enter a name and description for the audit.
 - **Source:** Every audit can use a server, snapshot, or snapshot specification as its source. (Or, you can choose no source and define your own rules.) If you use a server as the source, you can browse the server for values to define the audit's rules. If you choose a snapshot, you will be limited to the rules in the snapshot and the snapshot results when you define the audit rules. If you choose a snapshot specification, then the audit will compare the snapshot taken of the targets of the snapshot specification, and compare those against the targets of the audit. When you choose snapshot specification as the source, the rules in the snapshot are not editable. If you choose no source, you must define your own rules, or choose to link to an audit policy in the rules section. Some rules, however, require a source in order to be defined.
 - **Rules:** Choose a rule category from the list to begin configuring your audit's rules. Each audit rule is unique and requires its own instructions. For information on how to configure individual audit rules, see [Audit and Remediation Rules](#) on page 123.

If you want to use an audit policy to define the rules of your audit, click either Link Policy or Import Policy. When you link an audit policy, the audit maintains a direct connection with the audit policy, and disables the ability to create rules. Once you link a policy, the audit will use only the rules configured in the audit policy. So if any changes are made to the policy, the audit will update with the new changes. If you import an audit policy, the audit will use all the rules defined in the policy but will not maintain a link to the audit policy. For information about audit policies, see [Audit Policies](#) on page 166.

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

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

An asterisk (*) in any of these fields represent all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.
 - **Time and Duration:** For each type of schedule, specify the hour, minute, day of the week, and month for the schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose an end date, select End. From the calendar selector, choose an end date. The Time Zone is set according to the time zone set in your user profile.
 - **Notifications:** Enter email addresses to notify people when the audit job finishes running. You can choose to send the email on both the success and the failure of the audit job (not the success of the audit rules). To add an email address, click Add Notification rule. (This is only relevant if the audit is set to run on a recurring schedule.)

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

Audit Sources: Server, Snapshot, or Snapshot Specification

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

Server as Source for an Audit or Snapshot Specification

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

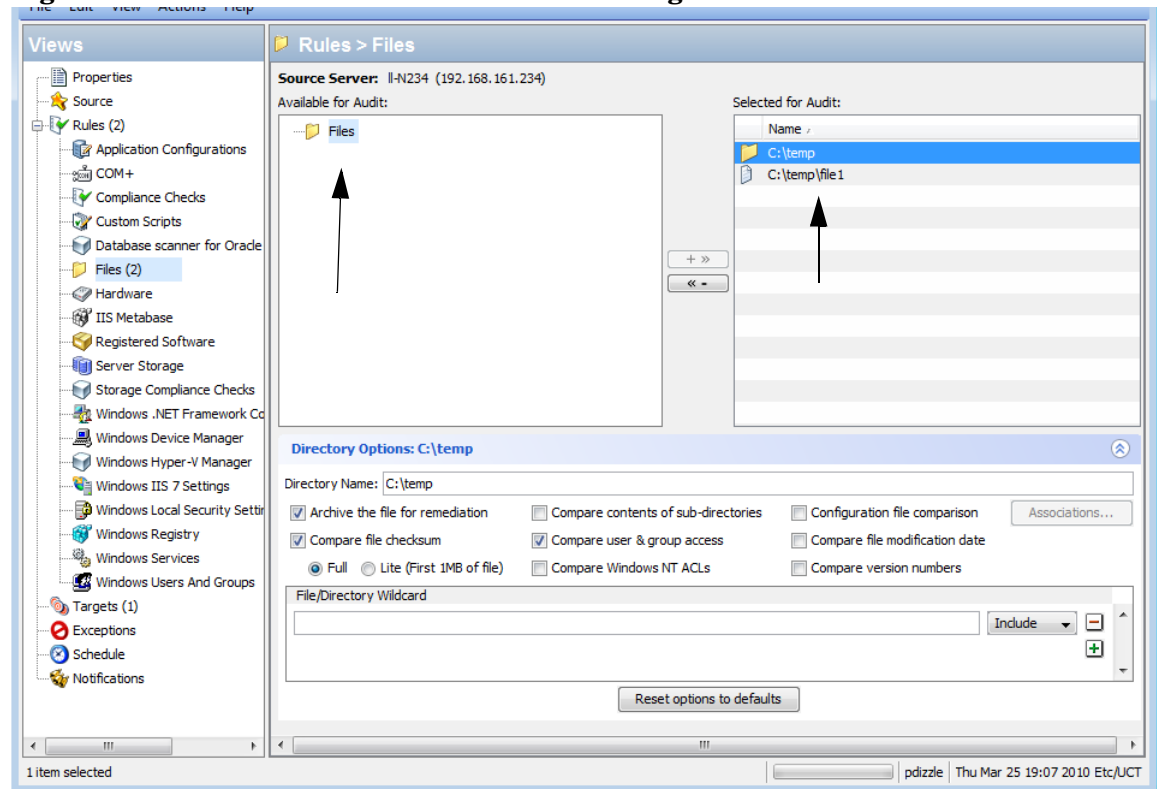
Remember that you can choose several different source servers as you build your audit or snapshot specification rules. In fact, you can choose a different source for each server object rule.



VMware ESXi servers cannot be the source of an audit or snapshot.

When you choose a server as the source for an audit, [Figure 19](#) shows what you see in the audit or snapshot specification window's Content pane (right side of window):

Figure 19 Server as Source of Audit for Building Audit Rules

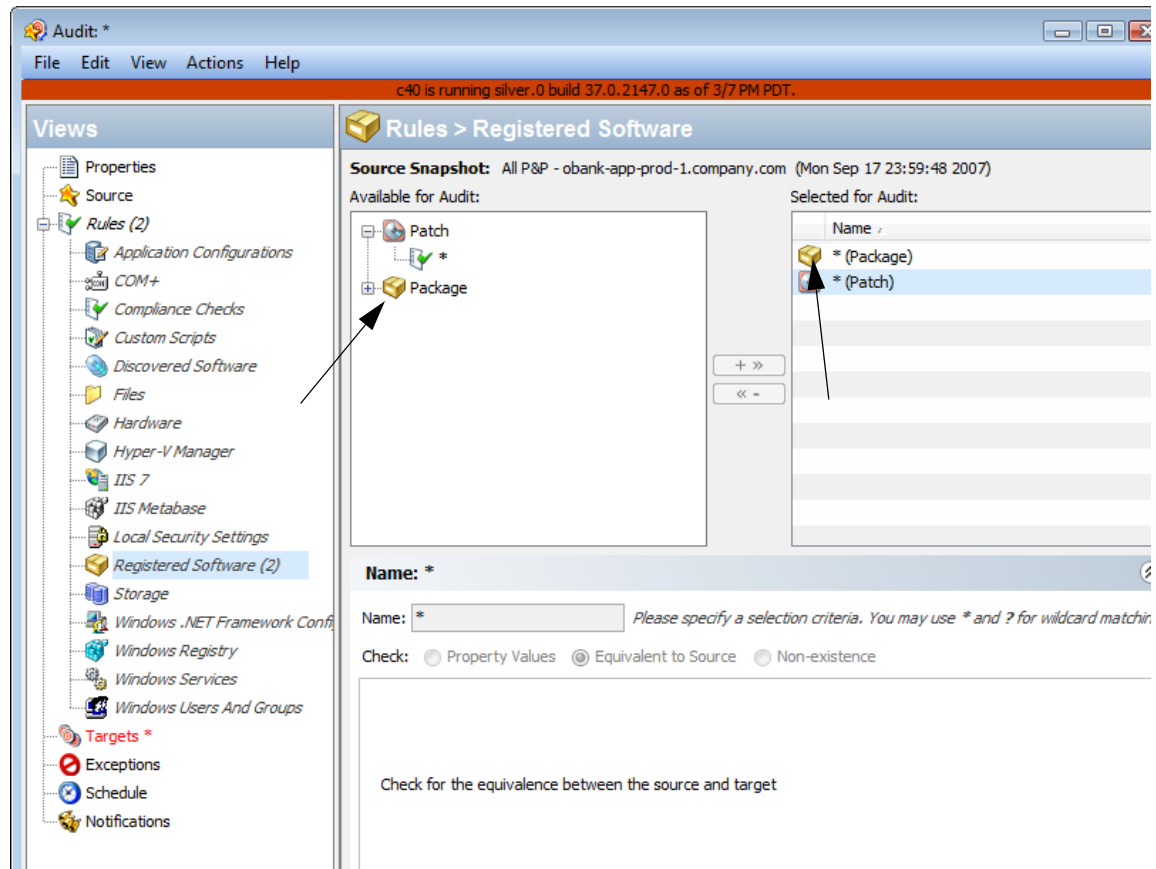


Snapshot as Source for an Audit or Snapshot Specification

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

[Figure 20](#) displays the choices you have for building audit or snapshot specification rules when you use a snapshot as the source. You can choose from the snapshot's results and the snapshot's rules.

Figure 20 Snapshot as Source of Audit: Available Server Objects to Build Audit Rules



Snapshot Specification as Source of Audit — Reflexive Auditing

Choose this option if you want to keep track of a server's configuration over time and monitor any changes that occur. For example, you might want to keep track of an application to make sure that its configuration remains correct over a period of time. If this application runs on several servers, you can create a snapshot specification that defines a desired state of server configuration, and then run the snapshot.

Next, you can create an audit and use the snapshot specification as the source for your audit. Each server that was targeted by the snapshot are now also included as targets of the audit. Next, when you run the audit (either on-demand or on a scheduled basis), each server's current configuration will be compared with the state originally captured from the snapshot. If the snapshot specification that serves as the source of the audit is set to run on a recurring basis, then the audit will compare against the most recently run snapshot. Any changes are displayed in the audit results window.

Rules That Use a Source Value From Source Server

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

- Any of the pre-configured rules that you do not set the value to derive from a source (server or snapshot or snapshot specification)
- Custom Scripts rules that you do not set the compare value to derive from a source (server or snapshot or snapshot specification)

You cannot save an audit that contains rules that require a source and no source has been specified. You must select a source for all comparison checks and for rules that compare against a source value.

Server Objects Used in Audits and Snapshots

Table 11 lists all server objects that you can create rules for inside an Audit or a snapshot specification. Some server object values are captured and audited live and some objects are captured from the Model Repository.

Table 11 Audit and Remediation Server Objects

Server Object	Description	Captured Live and/or from Model Repository
Application Configurations	Contents of application configuration files and their values.	Live
Windows COM+	COM+ objects and component categories.	Live
Custom Scripts	Write your own custom scripts to retrieve information from a server and compare contents. For example, you can run a script to gather output from a custom application and evaluate returned output against values set in the audit. (Python 1.5.2 only for python scripts.)	Live
Discovered Software	Discovered Software provides a signature-based software discovery mechanism for Windows and UNIX managed servers to help you manage applications and software that are not managed by SA.	Live
Files	Contents of files and directories (and subdirectories), user and group access, checksum for files, file modification date, and Windows ACLs (Windows only).	Live
Hardware	CPU, storage devices, and memory.	Model Repository
IIS Metabase	Microsoft IIS Metabase objects and configuration values to snapshot or audit.	Live
IIS 7.0	Microsoft IIS 7.0	Live
Internet Information Server	Real time information about IIS for a Windows server, such as server name, server type, server state, log file path, document file path, and so on.	Live
Local Security Settings	Real time information about security settings, including security settings such as password policy, audit policy, user rights, and security options.	Live

Table 11 Audit and Remediation Server Objects (cont'd)

Server Object	Description	Captured Live and/or from Model Repository
Registered Software	All installed packages or patches actually installed on a source server, whether or not they have been registered by the model repository.	Live
Runtime State	The Runtime State window rule allows you to use time information about run time data for an audit rule, such as DNS servers, Routes, and Processes for every managed server.	Live
Storage	Information related to storage devices and SAN devices and connections in your data center (if your core is storage-enabled). In order to audit and snapshot SAN objects, Storage Essentials (SE) version 6.1.1 or later is required and the Server Automation SE Connector component must be installed and configured on your SA core.	Live
BSA Essentials Subscription Services “compliance checks”	If you are subscribed to BSA Essentials Subscription Services, you have access to many different types of audit rules and their constituent components (also known as “compliance checks”. The exact kind of checks you have access to depend on your subscription, but can include such rules as the latest patch supplements for Microsoft Windows, current regulatory compliance policies (for example, FISMA, Sarbanes-Oxley), user-created rules from the BSA Essentials Subscription Services developer community, daily updated vulnerability content, and so on.	Live
Users and Groups	Compare information about users and groups on servers, such as user name for last login, whether or not CTRL + ALT + DELETE is enabled, and so on.	Live

Table 11 Audit and Remediation Server Objects (cont'd)

Server Object	Description	Captured Live and/or from Model Repository
Windows .NET Framework Configuration	Real time information about Assembly Cache and Configured Assembly List, 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 use information such as assembly name, public key token, codebases, binding policy, file name, file data.	Live
Windows Registry	Select Windows Registry directories or registry key values to capture and compare.	Live
Windows Services	Select Windows services.	Live
Windows Users and Groups	Users and groups information on a Windows Unix servers.	Live



A Windows COM+ category (folder) that does not have any objects will not be included in a Snapshot or Audit, even though SA will display an empty COM+ folder in the Device Explorer.



Audit and Remediation does not support device files or sockets.

Audit and Remediation Rules

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

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

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

Remediation values are implemented manually, after the audit has run, from the Audit Results window. For more information on how to remediate audit results, see [Remediating Audit Results](#) on page 176.

Configuration Rules: Expected (Target) and Remediation Values

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

For example, the Discovered Software rule checks for all registered and unregistered software installed or deployed on a target servers. The Hardware rule allows you to check the CPU, memory, or storage values that exist on target servers. In this case, no extra rule parameters are necessary.

Other rules are more complex and require more advanced configuration, such as specifying an expression that looks for a range of values and specifies remediation that replaces undesired values.

In an audit and audit policy, you can also define what, if any, remediation value you would like the object to have. Remediation values are used only if a server object is found to be different than the desired state — that is, the configuration on the target server is out of compliance with the rules of the audit. Remediation values are implemented manually, after the audit has been run, from the Audit Results window.

An audit rule consists of the following components:

- **Server Object:** This is a specific server configuration that an audit can evaluate, such as a server's file system, application configuration values, hardware information, installed software (patches and packages), Windows Registry entries, and so on. A server object usually consists of several other things that you can check as well. For example, on a Windows server you might want to know if a specific Windows service exists on target servers and whether or not it is enabled.
- **Target Value:** This is value or setting you want to check for on the target server, or the desired value. For example, you might want to determine if a specific directory exists on a server, or if an application is configured properly, or a particular service is enabled, and so on.
- **Remediation Value:** This is the value that you want to change for the server object during remediation, if the target value is not found on the target server. Remediation value is not implemented automatically; you must make the remediation change after the audit has run.

[Figure 21](#) illustrates an audit rule defined for a Windows Service named File Replication.

Figure 21 Custom Audit Rule Configured with Remediation Values

The screenshot shows the Windows Security console with the 'Rules > Compliance Checks' view. A list of rules is shown, with 'Maximum Application Event Log Size' selected. Below the list, the details for this rule are displayed. The 'Check' tab is selected, showing the 'Target Value' section with an operator of '>', a reference of 'Value', and a value of '16777216'. The 'Description' section explains that the rule determines the maximum log file size for the Application Event Log, with CIS and MSFT recommending 16MB (16777216 bytes).

Rules > Compliance Checks

Maximum Application Event Log Size

Check Remediation Properties

Target Value

Operator: Reference: Value:

> Value 16777216

Description

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

In this example (Figure 21), the audit rule has been configured in the following manner:

- **Rules > Compliance Checks:** Lists the selected rule from the BSA Essentials Subscription Services, in this case, Maximum Application Event Log Size.
- **Rules Details**
 - **Description:** Describes value is being checked on the target server. In this case, the audit will check to see if the Application Event Logging file size has not exceeded the CIS and MSFT recommended size limit of 16MB (16777216 bytes).
 - **Target Value:** This is the desired value compared against the value on the server that is the target of the audit. In this example, the rules is configured to determine if the target server's Application Event Logging file size has not exceeded 16777216 bytes. For example, the Target Value parameters in this example have been set to: > Value 16777216.

This instructs the audit to evaluate the target server's Application Event Logging file size and determine if it exceeds 16MB.
- **Remediation Value:** The remediation value determines the action to take if the value on the target server does not match the value you defined in the audit (target value). In this example, the remediation value is set to the CIS and MSFT recommended size limit of 16MB (16777216 bytes). You can remediate this value from the audit results window, after the audit has run, and only if the target server's value fails the rule criteria.

Configuring Specific Audit and Snapshot Rules

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

- [Configuring Application Configuration Rule](#)
- [Configuring COM+ Rule](#)
- [Configuring Custom Scripts Rule](#)
- [Configuring the Discovered Software Rule](#)
- [Configuring the File Rule](#)
- [Configuring Hardware Rule](#)
- [Configuring IIS Metabase Rule](#)
- [Configuring Internet Information Server Rule](#)
- [Configuring IIS 7.0 Rule](#)
- [Configuring Local Security Settings Rule](#)
- [Configuring Registered Software Rule](#)
- [Configuring Storage Rule](#)
- [Configuring Windows .NET Framework Configurations Rule](#)
- [Configuring Windows Registry Rule](#)
- [Configuring Windows Services Rule](#)
- [Configuring Windows/UNIX Users and Groups Rule](#)
- [Configuring Compliance Checks](#)



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



Some SA cores may contain legacy content, specifically, Event Logging, Operating System, and Users and Groups rules with compliance checks. These checks have been integrated into the CIS policies available from the EP.

Configuring Application Configuration Rule

The application configuration audit rule allows you to audit configuration file values on managed servers, to check that those files are configured the way you want them to be.

You can choose from a list of predefined application configuration templates which serve as the basis of comparison for the target configuration file you want to audit. You can also choose from custom application configurations that a user in your organization has created and made available for usage in an audit, snapshot specification, or audit policy.

An application configuration in an audit models the values and structure of an application's configuration file, which allows you to set rules that check the values in actual configuration files on managed servers.

When you choose an application configuration inside an audit, snapshot specification, or audit policy and click **View**, you will see the contents of the configuration file from the source of the audit. All key-value pairs that you are able to add to the audit rule will display.

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

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



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


After you view the contents of the source application configuration file, you can define create your rules by selecting values from the source file and building rules that will be used to check against the target configurations. You can also define remediation values in the event that the audit finds differences between the rules and the target configuration file values.

Creating an Application Configuration Rule

To understand how to configure an application configuration rule, it is useful to look at an example. Your goal is to create an audit rule for a UNIX hosts file (`/etc/hosts`), and then audit a group of servers' `/etc/hosts` files to make sure they contain the correct values.

You know that the UNIX hosts file on a particular “golden” server represents the ideal state of hosts file configuration that you would like other servers to conform to. You can choose that golden server as the source for your audit and borrow the values from that file to construct the rule for the audit. Once you create the rule and save the audit, you can run the audit against a group of servers to see if their `/etc/hosts` files are configured correctly (according to the audit rule).

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

- 1 Create an audit from any one of the methods for creating an audit listed at [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. The source selected for the audit will determine what types of rules, if any, you can create for an application configuration. You must choose a source or you will not be able to configure the application configuration rule.
- 3 In the Audit window, from the View pane, select Rules ► Application Configurations.
- 4 In the content pane of the Audit window (right side of window), click the Add  button to access all available configuration templates.

- 5 In the Select Configuration Templates window, select one or more templates you would like to add to the audit rule, and then click **OK**.
- 6 Select the template you want to configure, and its contents appear below in the template editor.
- 7 Click **View**. (If you cannot view the contents of the configuration file, you might need to enter the correct path in the Filename section.) You see the contents of the configuration file in the File View tab.

For example, if you view a UNIX hosts file, you would see something similar to that shown in [Figure 22](#):

Figure 22 Application Configuration Audit Rule for Hosts File

Rules > Application Configurations

Source Server: m197.qa.opsware.com (192.168.160.197)

Name	Location	Filename
hosts.tpl	/Content/Configurations	/etc/hosts

Rule Details: hosts.tpl

Filename: /etc/hosts View

Contents: File View Rule View

```
# At minimum, this file must contain the name and address for each
# device defined for TCP in your /etc/net file. It may also contain
# entries for well-known (reserved) names such as timeserver
# and printserver as well as any other host name and address.
#
# The format of this file is:
# Internet Address      Hostname      # Comments
# Items are separated by any number of blanks and/or tabs. A '#'
# indicates the beginning of a comment; characters up to the end of
# line are not interpreted by routines which search this file. Blank
# lines are allowed.

# Internet Address      Hostname      # Comments
# 192.9.200.1           net0sample    # ethernet name/address
# 128.100.0.1           token0sample  # token ring name/address
# 10.2.0.2              x25sample    # x.25 name/address
127.0.0.1              loopback localhost # loopback (lo0) name
192.168.160.197 m197.qa.opsware.com
```

Operator: Value Reference: Value Value:

Remediate With: Value

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

- 8 In order to create an audit rule for this configuration file, you need to choose a key-value pair from the hosts file on the source server (the server you choose as the source for the audit).
- 9 To create this rule, first select an IP addresses in the File View tab area, which shows the contents of the file obtained from the source server. In the example in [Figure 22](#), you can select an IP address such as 127.0.0.1. After you select the IP address, the element becomes highlighted in dark blue. This means that the element is ready to have a rule created from it.

(For more information on the color scheme used when configuring an application configuration audit rule, see Table 12 on page 130.)

Once you have selected the IP address in the contents area, notice that the value in the Operator field in the below is set to blank. This means that an operator has not yet been added to the rule. To add the value to the rule, you can either double-click it, or enter the following parameters in the rule expression area below the contents:

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

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

- 10 Next, select a host name in the File View tab area. Notice that the initial IP address you selected in the previous step has turned green. This means that the next rule parameter you set will be paired with the IP address you previously selected.
- 11 In the Rule section, set the following parameters:
 - **Operator:** Choose = (equals).
 - **Reference:** Choose Value. (If you choose a custom attribute here for the rule definition, this custom attribute must also exist on the target servers or the audit for this rule will fail.)
 - **Value:** Choose host.
 - **Remediation:** Choose host. This adds the final part of the rule that will check the target server for the key-value pair of IP address 127.0.0.1 matched with host.

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

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

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

- 13 To configure more application configuration rules, select more application configurations from the Available for Audit section.
- 14 To finish configuring the audit, define other rules and set the target servers, schedule, and notification for the audit.

- 15 Save the audit.
- 16 To run the audit, from the **Actions** menu, select **Run audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Application Configuration Audit Rule Color Scheme

When you first view an application configuration, all elements that can be used to build an audit rule will appear in blue underlined text. After you start selecting and building rules, then the colors will change. [Table 12](#) describes the color scheme used for configuring application configuration audit rules.

Table 12 Application Configuration Audit Rule Color Scheme

Text Color	Description
Blue underlined	This shows all elements in the source configuration file that can be used in a rule.
Highlighted Dark Blue	This shows an element is selected but has no rule has been associated with it.
Highlighted Light blue	This shows all that you add an element to a rule.
Highlighted Medium blue	This shows all that an element is both selected and has a rule associated with it.
Green	<p>This shows all that the element is a primary key and is related to the current selected element. This means that the element will be used in the same rule that the current selected element will be used in.</p> <p>If the currently selected element is given a comparison value (=, contains, matches...) then the other elements with the green text will automatically be given a comparison value of “=”.</p> <p>An example of this would be:</p> <pre>127.0.0.1 localhost</pre> <p>If localhost is selected, then 127.0.0.1 would be green. If localhost is given a comparison value, then 127.0.0.1 will also be given an automatic comparison value, giving you a rule such as:</p> <p>There is an entry where ip is equal to 127.0.0.1 AND hostname is equal to localhost.</p>
Bold	This represents a primary key.
Italicized	This shows a custom attribute or SA attribute.

Configuring COM+ Rule

To configure a Windows COM+ rule, select the source COM+ objects that you want to audit or snapshot on a target server. The COM+ rule also checks Access Control Levels (ACLs) for the selected object as well as any ACLs that are inherited.

COM+ objects are categorized based on attributes of the object, where the COM+ object specifies zero or more categories. The audit or snapshot window displays all COM+ objects in one node in the Rules section of the COM+ object tree. To add a COM+ rule to the audit or snapshot, select it and click the right arrow button.

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

Selecting the “Archive all associated files” option will also include all AccessPermissions and LaunchPermissions associated with the COM+ object in the audit or snapshot rule, including those that are inherited parent COM+ objects.



You cannot audit the COM+ root folder, but can audit as many of the COM+ individual objects or sub categories as you wish.

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

- 1 Create the new audit using one of the methods for creating an audit listed in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► COM+.
- 4 In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a COM+ object or object category.
- 5 Click the right arrow button to move the COM+ object or object category into the Selected for Audit section. All COM+ object or object categories you select will be audited on the target servers or snapshot specification. (You can select individual and COM+ categories for the rule, but you cannot select the root folder to add to the audit rules.)
- 6 You can now choose an option from the bottom of the rule window:
 - Select the Archive all associated files option if you want to be able to remediate COM+ rules in your audit or snapshot results.
 - Select Compare only the file name and not the full pathname if you want the COM+ rule to check only the selected filename and not the full path.
- 7 To finish configuring the audit, define any other COM+ object or object category rules you want and set the target servers, schedule, and notification for the audit.
- 8 If you want to be able to
- 9 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 10 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Creating an Audit Policy](#) on page 167.

Configuring Custom Scripts Rule

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

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

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

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

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

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

- 1 Create the new audit using one of the methods for creating an audit in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User’s and Groups, must have a source.)
- 3 To build a script and define the audit rule, you can choose the following options:

Source

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

Rule Details

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

Success Criteria


- **Output:** Either Exit Code or Standard Output.
- **Operator:** Choose an Operator, such as equals (=), not equals (<>), less than (<()), greater than (>), and so one.
- **Reference:** Choose the source of the script output.

- **Source:** Select this option if you want the rule to execute the script on the source when an audit is run, and gets the value that the script requests. It will then compare that value with the value retrieved from the script that was run on the target server.

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

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

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

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

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

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

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

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

Remediation

- **Type of Script:** Choose from Batch, Python 1.5.2, PowerShell, or Visual Basic (VBS).
 - **Script:** Type or copy and paste the script contents here. Or, click **Import Script** to import a script from your computer.
- (Optional) You can add a remediation script to run if the audit comparison fails. The remediation will not be applied automatically; you can only run the remediation script from the audit results after the audit has run.

For a snapshot, the remediation script you define here can be executed on target servers on an individual server basis.
 - To finish configuring the audit, set the target servers, schedule, and notification for the audit.
 - To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
 - To run the audit, from the **Actions** menu, select **Run audit**. For more information about running an audit, see [Creating an Audit Policy](#) on page 167.

Custom Scripts Example

The following is an example of a custom VB script rule designed to enable a Windows user account and set the user's password.

(Note that this script will only work on Windows OS versions later than Windows NT 4.0. If you wish to enable a user account and set the password on Windows NT 4.0, you will have to do so manually.)

```
strComputer = "."
strAccountName = "red2"
Set objUser = GetObject("WinNT://" & strComputer & "/" & strAccountName )
objUser.AccountDisabled = False
objUser.SetPassword "AiH345^hjq"
objUser.SetInfo
```

Configuring the Discovered Software Rule


The Discovered Software rule provides a signature-based software discovery mechanism for Windows and UNIX managed servers to help you audit and snapshot applications and software that are not managed by SA.

Specifically, the Discovered Software can:

- Discover unregistered software which is not currently managed by SA
- Create an inventory of software that is not installed as part of an OS-registered application, or that was custom built
- Give you the ability to create snapshots of the discovered software on a server and then periodically audits against the snapshots over time
- Enable you to track in-house or custom built software

To configure the discovered software rule, perform the following steps:

- 1 Create a new audit using one of the methods in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source.
- 3 In the Audit window, from the View pane, select Rules ► Discovered Software.
- 4 In the content pane of the Audit window, in the Available for Audit section expand the Software icon. This may take a few moments to load if this is the first time you are loading the rule and you have selected a source for the audit or snapshot.
- 5 Select an element from the list and then click the right arrow button to move the rule object into the Selected for Audit section, which enables you to create a rule for the element.
- 6 For each check you want to configure in the rule, in the lower section of the audit window you can select one of the following rule criteria types:
 - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.

- **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
 - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.
- 7 You can also configure the rule based upon a wildcard search by selecting the Wildcard rule object . When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.
- For example, you could enter simply * which would match everything on the target, P* would match all objects that begin with a capital P, while *P would match all elements ending with uppercase character 'P'.
- After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.
- It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.
- 8 To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy, which enables other users to access the rule set you create in the audit. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 10 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Configuring the File Rule

The file rule allows you to audit and compare files and directories on a target server in the following ways:

- **Directory name:** Displays the absolute path of the selected file or directory
- **Archive the file for remediation:** Archives the entire file. Use this option if you want to be able to remediate and view any files differences found between the rule and the target file.

When selected, this option enables the audit to check for differences of a specified file, according to the differences you specify in the rule. If any differences are found, remediating the differences will copy the source file to the target server and replace the target file with the source.

Keep in mind, however, that selecting this option can potentially create disk space demands on the SA core's database, depending on the size and number of files being compared.

- **Compare file checksum:** Performs a Checksum on the contents of the selected file or files in a directory. You can choose to audit the entire contents of the file, or just the first 1MB of the file.

- **Compare contents of subdirectories:** Includes contents of all subdirectories for a selected file system folder to the audit.
- **Compare user and group access:** Audits the user and group access related to the file and directories.
- **Windows NT ACLs (Access Control List):** Audits the Windows Access Control List for files and directories.
- **Configuration file comparison:** Allows you to use an application configuration to evaluate configuration files on a target server. Selecting this option (and then clicking **Associations**) enables you to utilize a configuration template to compare any differences in values between a source configuration file and one on a target server. For more information on how to use this feature, see [Comparing Files in Audits with Configuration Templates](#) on page 137.
- **Compare file modification date:** Audits the file modification date to use for file or folder comparison.
- **Compare version numbers:** For specific Windows file types — .exe, .dll, .ocx, .olb, .scr, .rll, .sys, .drv, .acm — the author of the file can set a File Version and Product Version. This option, compares these version numbers and if they are different, the rule is considered non-compliant and the actual values on the target file can be viewed in the audit results. (Note that not all files with these extensions always have a product or file version attribute, but some can.)
- **File/Directory wildcard:** Allows you to specify directories and files in the file system you want included in and excluded from the audit. For more information on how this option works, see [File Inclusion and Exclusion Rules](#) on page 157.

For more information on using environment variable and custom attribute to parameterize filenames and paths, see [Parameterizing Filenames for SA/Custom Attributes](#) on page 162.

If you want to use environment variables in file name PATH on Unix, see [Using Environment Variables in Pathnames](#) on page 163 for more information.

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

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



If you are checking ACLs for the File rules, and the user and group ACL does not exist, then after the audit is run and after remediation, if user and group does not exist on target a temporary user and group will be created as unknown name. The next time you run Audit it shows up as unknown, which shows name other than the source user. For more information on remediation, see [Remediating Audit Results](#) on page 176.

To configure file rules, perform the following steps:

- 1 Create the new audit using one of the methods in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)

- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Files.
- 4 In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a folder or file to create a rule for.
- 5 Click the right arrow button to move the folder or file into the Selected for Audit section. All folders or files that you select will be used to audit or snapshot the target server.
- 6 In the Selected for Audit section, select a folder or file to apply a rule to.
- 7 In the Directory Options section, select file system rule options to apply to the selected folder or file. If you would like to reset the original settings of the source file system, select the Reset options to match those of the File System option.
- 8 (Optional) For folders, you can select a File/directory Wildcard option to specify files and directories that you want to include or exclude from the audit.

Click the **plus (+)** button to add a new rule, or click the **minus (-)** button to remove a rule. For more information on how to enter files and directories and how this affects the audit, see [File Inclusion and Exclusion Rules](#) on page 157.
- 9 (Optional) If you want to use an application configuration to compare configuration files, select Configuration file comparison, and click **Associations**.
- 10 In the Edit AppConfig Associations window, from the Installed AppConfig Templates, select a template you would like to use to compare a source and a target configuration file.
- 11 In the Associated Files section, you can use the default path to the source configuration file, or edit the path. You can click the plus button to add another path to a source configuration file you want to compare with a configuration file on the target.
- 12 When you are finished, click **OK**.
- 13 To finish configuring the audit, set the target servers, schedule, and notification for the audit.
- 14 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 15 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Creating an Audit Policy](#) on page 167.

Comparing Files in Audits with Configuration Templates

Another way you can audit files on a target server is to compare them against a source server file using application configuration (AppConfig) templates as the basis of comparison.

Configuration templates model the structure of a configuration file and determine its contents and organization. When you use configuration templates in an Audit's file rule to compare files, the audit uses the configuration template to filter both the source and target files' contents for the comparison. This ensures that you are comparing only the value sets defined in the template when you run the audit and compare the files.

For example, you might want to compare the `/etc/passwd` file on several target servers to make sure they contain only the values defined in the `/etc/passwd` file on a "golden" server that you know has acceptable values. Using the configuration file comparison feature, you select a configuration template that models the `/etc/passwd` file (`passwd.tpl`) and associate that configuration template with the actual `passwd` file on both the golden source server and the servers targeted by the audit.

You create the association by selecting the template, then entering the file pathname to where the file exists on the target servers. You can also compare multiple files using this feature. For example, you can select a directory that you know contains several configuration files to compare and you can associate configuration templates with directories you know contain the files you want to compare.

To use the configuration file comparison feature in an audit, perform the following steps:

- 1 Create the new audit using one of the methods in [Creating an Audit](#) on page 113.
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (If you select a snapshot, you will only be able to compare those files captured in the snapshot.)
- 3 In the Audit window, from the View pane, select Rules ► Files.
- 4 In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a file to compare, or a directory that contains the files you want to compare.
- 5 Click the right arrow button to move the folder or file into the Selected for Audit section.
- 6 In the Selected for Audit section, select the folder or file.
- 7 In the bottom section of the audit window, select the Configuration file comparison option and then click **Associations**.
- 8 In the AppConfig File Comparison window, in the top AppConfig Templates section, select the check box of the configuration template you want to use as a basis for comparison. For example, if you want to compare the /etc/hosts file of a source server against a target server, select the hosts.tpl configuration template. (Configuration templates use the TPL file extension.)
- 9 In the Associated Files section at the bottom of the window, enter the pathname to where the actual source and target configuration file exists on both the source and target servers. Note that the files you want to compare with the configuration template must exist in the same directory.
- 10 (Optional) If you want to make more than one association for a template, click the plus sign and enter another directory. Each directory you add applies to whatever template you have selected in the AppConfig Templates section in the top part of the window. You can make as many associations as you want in this window.
- 11 When you are finished, click **OK**.
- 12 To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 13 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 14 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Creating an Audit Policy](#) on page 167.

Configuring Hardware Rule

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

- **Interfaces:** Compares duplex mismatch and all network interfaces on a server.
- **CPU:** Compare CPU type and specification of target server.
- **Memory:** Compare memory of the target server.

- **Storage:** Compare storage capacity on the target server.
- **Interfaces:** Compare all network interfaces attached to the device.



If you are auditing or taking a snapshot of the Hardware rule on a server that just recently had the SA agent installed on it, it is possible that the hardware has not been fully registered with the Model Repository, and you won't be able to audit or snapshot accurate hardware information. (The SA Agent registers hardware usually within 24 hours after agent installation.)

If you are not sure, contact your SA Administrator or the person who installed the SA Agent on the server. For more information, see Appendix A: SA Agent Management for instructions on how to register a server's hardware manually.

To configure hardware rules, perform the following steps:

- 1 Create the new audit using one of the methods for creating an audit listed in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Hardware.
- 4 In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a hardware category to create a rule for.
- 5 Click the right arrow button to move the hardware item into the Selected for Audit section. All items that you select will be used to audit or snapshot the target server.
- 6 To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 8 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Creating an Audit Policy](#) on page 167.

Configuring IIS Metabase Rule

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

If you are checking ACLs for Metabase rule, and the user and group ACL does not exist, then after the audit is run and after remediation, if user and group does not exist on target a temporary user and group will be created as unknown name. The next time you run Audit it shows up as unknown, which shows name other than the source user.

Additionally, if you create an IIS Metabase rule from a source server and the metabase object selected for the rule inherits its values from a parent Metabase object, differences will show after an audit is run. For example, if you remediate once and then rerun the audit, if the source key was not inherited and the attribute has an IED when it gets created on target server, the object will be created based on parent key inheritance. When you rerun the audit,

the results will show the IED as a difference for the object's attribute.

For more information on remediation, see [Remediating Audit Results](#) on page 176.



If you want to audit Microsoft IIS 7.0 on a Windows 2008 server, create and configure the IIS 7.0 rule in your audit. For more information, see [Configuring IIS 7.0 Rule](#) on page 141.

To configure IIS Metabase rules, perform the following steps:


- 1 Create the new audit using one of the methods for creating an audit listed at [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► IIS Metabase.
- 4 In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select an IIS Metabase folder or object to create a rule for. (You can select any metabase folder or object for the rules, but you cannot select the root folder to use as a rule.)
- 5 Click the right arrow button to move the IIS Metabase folder or object into the Selected for Audit section. All items you select will be used to audit or snapshot the target server.
- 6 To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 8 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Creating an Audit Policy](#) on page 167.

Configuring Internet Information Server Rule

The Microsoft Internet Information Server rule allow you to use real time information about IIS for your audit, such as a Windows server, such as server name, server type, server state, log file path, document file path, and so on.

To configure the Internet Information Server rule, perform the following steps:

- 1 Create the new audit using one of the methods in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Internet Information Server.
- 4 In the content pane of the Audit window, expand the top level node in the Available for Audit section and select an Internet Information Server rule that you want to create a rule from.

- 5 Click the right arrow button to move the rule object into the Selected for Audit section. All Internet Information Server rules that you configure will be audited on the target servers or snapshot specification.
- 6 For each rule, select one of the following check types:
 - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.
 - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
 - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.
- 7 You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object . When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.

For example, you could enter simply * which would match everything on the target, P* would match all objects that begin with a capital P, while *P would match all elements ending with uppercase character 'P'.

After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.

It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.
- 8 To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 10 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Creating an Audit Policy](#) on page 167.

Configuring IIS 7.0 Rule

In SA 9.0, you can create audit and snapshot specification rules for Microsoft IIS 7.0 running on Windows 2008. You can expand and browse IIS 7.0 Application Pools, Web Sites, and features and add them to your audits or snapshot specifications to determine if they meet your organization's compliance standards. After your audit or snapshot has run, you can view the results and remediate any discrepancies found (with some exceptions).

For example, you might want to audit several Windows 2008 servers running IIS 7.0 to make sure that Anonymous Authentication is enabled on each server.

To perform this compliance check, select a Windows 2008 server that has Anonymous Authentication enabled to be the *source* server of the audit. Then, configure the audit rule to check that Anonymous Authentication is enabled on all servers targeted by the audit.

When you run the audit (which you can schedule on a recurring basis), the rule will check the target servers and discover if any do not have Anonymous Authentication enabled. If the audit finds any discrepancies, you can remediate those servers to enable their IIS 7.0 Anonymous Authentication.



You cannot remediate ISAPI filters for the IIS 7.0 audit rule in this release.

To configure the IIS 7.0 rule, perform the following steps:

- 1 Create a new audit using one of the methods in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source.

Some audit rule types, such as Application Configuration and Windows User's and Groups, must have a source server upon which to base the rule. Some specific rules and criteria, such as checking IIS 7.0 Anonymous Authentication, also require that you select a source server. If you do not select a source server, you will be limited on the specificity of the rule.
- 3 In the Audit window, from the View pane, select Rules ► IIS 7.0.
- 4 In the content pane of the Audit window, in the Available for Audit section expand one of the IIS 7.0 elements you want to create a rule for, such as Application Pools, Sites, or Features. This may take a few moments to load if this is the first time you are loading one of the elements.
- 5 Select an element from the list and then click the right arrow button to move the rule object into the Selected for Audit section, which enables you to create a rule for the element. For example, you could expand the Authentication folder and select Anonymous Authentication, then click the right arrow button to add the selection to your audit.
- 6 For each rule, in the lower section of the audit window you can select one of the following rule criteria types:
 - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.
 - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.


Remediation of the IIS 7.0 rule is possible only when an audit is setup with the Equivalent to source check.

- **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.

For example, if you wanted to check that a target server (or group of servers) running IIS 7.0 has Anonymous Authentication enabled, in the bottom of the audit window, you would select:

- Property Values
- Status
- =
- Enabled

This tells the audit to find out if each target server's IIS 7.0 Anonymous Authentication is enabled.

- 7 You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object  *. When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.

For example, you could enter simply * which would match everything on the target, P* would match all objects that begin with a capital P, while *P would match all elements ending with uppercase character 'P'.

After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.

It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.


- 8 To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy, which enables other users to access the rule set you create in the audit. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 10 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Configuring Local Security Settings Rule

The Local Security Settings rule allows you to use real time information about security settings, such as password policy, audit policy, user rights, and security options in your rule.

To configure the Local Security Settings rule, perform the following steps:

- 1 Create the new audit using one of the methods in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Local Security Settings.
- 4 In the content pane of the Audit window, expand the top level node in the Available for Audit section and select an Internet Information Server rule that you want to create a rule from.

- 5 Click the right arrow button to move the rule object into the Selected for Audit section. All Internet Information Server rules that you configure will be audited on the target servers or snapshot specification.
- 6 For each rule, select one of the following check types:
 - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.
 - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
 - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.
- 7 You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object . When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.

For example, you could enter simply * which would match everything on the target, P* would match all objects that begin with a capital P, while *P would match all elements ending with uppercase character 'P'.


After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.

It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.
- 8 To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 10 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Configuring Registered Software Rule

The Registered Software rule allows you to audit use all installed packages or patches actually installed on a source server to build your rule, whether or not the patches or packaged have been registered by the SA model repository.

To configure the Registered Software rule, perform the following steps:

- 1 Create the new audit using one of the methods in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Registered Software.
- 4 In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a patch or a package that you want to create a rule from.
- 5 Click the right arrow button to move the rule object into the Selected for Audit section. All rules that you configure will be audited on the target servers or snapshot specification.
- 6 For each rule, select one of the following check types:
 - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.
 - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
 - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.
- 7 You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object  *. When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.

For example, you could enter simply * which would match everything on the target, P* would match all objects that begin with a capital P, while *P would match all elements ending with uppercase character 'P'.

After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.

It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.
- 8 To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.

- 9 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 10 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Configuring Storage Rule

The storage rule allows you audit servers for storage devices and SAN devices and connections in your data center, if your core is configured to connect to SE.



In order to audit and snapshot SAN objects, Storage Essentials (SE) version 6.1.1 or later is required and the Server Automation SE Connector component must be installed and configured on your SA core. For more information, see your SA administrator or the Storage Visibility and Automation documentation.

To configure the storage rule, perform the following steps:

- 1 Create the new audit using one of the methods in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Storage.
- 4 In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Storage rule that you want to create a rule from. Each storage audit rules check for the acceptable values for each category. You can configure the rule to check for minimum, maximum, or exact numbers.
 - **Unmounted Volume Capacity:** Acceptable total capacity of unmounted volumes in bytes.
 - **Unmounted Volume Count:** Acceptable number of unmounted volumes.
 - **Fabrics:** Acceptable number of fabrics.
 - **FCA:** Acceptable number of Fibre Channel Adapters (FCAs).
 - **Initiator Ports:** Acceptable number of initiator ports
 - **Switches:** Acceptable number of SAN switches.
 - **Target Ports:** Acceptable number of target ports.
 - **RAID Types:** Acceptable RAID types on the target storage array. (**Note:** The audit will fail if this rule is selected and no RAID type is specified.)



The compliance rules that involve ports, switches, or fabrics, check active ports only. These types of compliance rules do not check for physical port connectivity.

- 5 Click the right arrow button to move the rule object into the Selected for Audit section. All storage rules that you configure will be audited on the target servers or snapshot specification.
- 6 For each rule, select one of the following check property:


- An operator, such as equal to (=), less than (<), less than or equals to (<=), and so on.
 - A value, depending upon the rule type, such as a number.
- 7 To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
 - 8 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
 - 9 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Configuring Windows .NET Framework Configurations Rule

The Windows .NET Framework Configuration rule allows you to use time information about Assembly Cache and Configured Assembly List, such as assembly name, version, locale, public key token, cache file (GAC or ZAP), processor architecture, custom, and file name in your audits.

To configure the Windows .NET Framework Configuration rule, perform the following steps:

- 1 Create the new audit using one of the methods in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Windows .NET Framework Configuration.
- 4 In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Windows .NET Framework Configuration rule that you want to create a rule from.
- 5 Click the right arrow button to move the rule object into the Selected for Audit section. All Windows .NET Framework Configuration rules that you configure will be audited on the target servers or snapshot specification.
- 6 For each rule, select one of the following check types:
 - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.
 - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
 - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.

- 7 You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object  *. When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.

For example, you could enter simply * which would match everything on the target, P* would match all objects that begin with a capital P, while *P would match all elements ending with uppercase character 'P'.

After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.

It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.
- 8 To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 10 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Configuring Windows Registry Rule

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

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

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

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

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

You can also choose to compare Access Control Levels (ACLs) for registry rule.



If you are checking ACLs for Registry rule, and the user and group ACL does not exist, then after the audit is run and after remediation, if user and group does not exist on target a temporary user and group will be created as unknown name. The next time you run Audit it shows up as unknown, which shows name other than the source user. For more information on remediation, see [Remediating Audit Results](#) on page 176.

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

- 1 Create the new audit using one of the methods for creating an audit listed in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Windows Registry.
- 4 In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Windows Registry folder or key to create a rule for.
- 5 Click the right arrow button to move the Windows Registry folder or key into the Selected for Audit section. All items that you select will be used to audit or snapshot the target server.
- 6 For each registry entry key rule you create, you can set two options to include when the audit checks the target:
 - Also Compare Contents of Sub-Keys: Will evaluate all subkeys belonging to the selected registry key.
 - Also Compare ACLs: Will also compare ACLs of the selected registry key.
 - Use case-insensitive compare:
- 7 To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 8 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 9 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Configuring Windows Services Rule

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

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

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

To configure windows services rules, perform the following steps:


- 1 Create the new audit using one of the methods for creating an audit listed in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Windows Services.
- 4 In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a windows services to create a rule for. (You can select any available service for the rule, but you cannot select the root folder for all services.)
- 5 Click the right arrow button to move the selected windows services into the Selected for Audit section. All items that you select will be used to audit or snapshot on the target server.
- 6 To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7 Save the audit.
- 8 To run the audit, from the **Actions** menu select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Configuring Windows/UNIX Users and Groups Rule

The Windows or Unix Users and Groups rule allows you to access local users and groups information from Windows and Unix servers.

To configure the Users and Groups rule, perform the following steps:

- 1 Create the new audit using one of the methods in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3 In the Audit window, from the View pane, select Rules ► Windows/Unix Users and Groups.
- 4 In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Users and Groups rule that you want to create a rule from.
- 5 Click the right arrow button to move the rule object into the Selected for Audit section. All Users and Groups rules that you configure will be audited on the target servers or snapshot specification.
- 6 For each rule, select one of the following check types:
 - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array. For some property types you can select the values from the 'value selector box'.

- **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
 - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object does exist on the target server, then the rule is out of compliance. The rule is considered compliant if no objects are found.
- 7 You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object . When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.
- For example, you could enter simply * which would match everything on the target, P* would match all objects that begin with a capital P, while *P would match all users with name ending with uppercase character 'P'.
- After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.
- It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.
- 8 To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 10 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Configuring Compliance Checks

If you subscribe to the BSA Essentials Subscription Services, you have access to dozens of compliance rules and their components, known by content developers as *compliance checks*.

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





If you do not subscribe to BSA Essentials Subscription Services, you will not see any compliance checks in your audits, audit policies, snapshots, or the Compliance Check Editor. If you would like more information on content subscriptions and obtaining compliance checks, contact your BSA Essentials Subscription Services sales representative.

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

For more information on managing your core's compliance checks, such as editing check property data or creating compliance check groupings, see [Managing Compliance Checks](#) on page 154.

To configure compliance checks in audits or snapshot specifications, perform the following steps:

- 1 Create an audit or snapshot using one of the methods described in [Creating an Audit](#) on page 113. (If you want to create this rule for a snapshot specification, see [Creating a Snapshot Specification](#) on page 193.)
- 2 Select an Audit Source: Server, Snapshot, Snapshot Specification, or No Source.
- 3 In the Audit window, from the View pane expand the Rules object.
- 4 Select the Compliance Checks  rule.
- 5 In the content pane of the Audit window (right side of window), click the Add  button.
- 6 In the Select Check window, from the Browse tab you can browse for the compliance checks categories and select a check for the audit or snapshot.

Alternately, you can select the Search tab and search for check by name. The check search tool searches on the name of a check and any words in a check's description. For example, if you wanted to find all rules that check for maximum password length, you could enter `max password` in the Keywords field.

The Advanced search option allows you to set more specific parameters to find checks.


- 7 When you select a check (or multiple checks using CTRL or SHIFT + click), click **OK** to add the checks to your audit.
- 8 Select the check and then define or set the following parameters:

Input Value

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

Target Value

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

- **Operator:** To build an expression from the output of the script, choose an Operator, such as equals (=), not equals (<>), less than (<), greater than (>), and so on.
- **Reference:** Choose the source of the script output.
- **Source:** This will use the value from the source server and compare that value to with the value found on the target server or servers.
- **Value:** Enter your own value. This option uses the value you enter and compares it with the value returned on the target server. You can get the value from the source server if you click the eyedropper  icon. The returned value is displayed in the text box, which you can accept as is or edit to your liking.
- **Server Attribute:** Select to compare a server attribute located on the source server.
- **Custom Attribute:** Select to compare a custom attribute found on the target server.

Remediation Value

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

- 9 To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 10 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see [Saving an Audit as Audit Policy](#) on page 115.
- 11 To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see [Running an Audit](#) on page 172.

Renaming Compliance Checks

You can easily rename instances of compliance checks in an audit, audit policy, or snapshot specification using the right-click menu.

For information on renaming compliance checks and editing their properties, see [Managing Compliance Checks](#) on page 154.

To rename a compliance check name, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Audit and Remediation ► and open an audit, audit policy, or snapshot specification.
- 2 From inside the audit (or audit policy or snapshot specification) window, from the Views pane select a specific rule that contains custom checks. For example, Users and Groups.
- 3 In the Contents pane of the window (right side), in the Available for Audit section, select a custom rule check, right-click and select **Rename Rule** to rename the rule.






You cannot rename a rule check if the audit or snapshot specification is linked to an audit policy.

Searching for Compliance Checks from the Audit/Snapshot Specification Window

Since your SA core can potentially contain dozens, if not hundreds, of compliance checks, you can use the search tool inside the Audit or Snapshot Specification window to find the checks you want.

To search for compliance checks from inside an audit or snapshot specification, perform the following steps:

- 1 In the Audit or Snapshot Specification window, from the View pane expand the Rules object.
- 2 Select the Compliance Checks  rule.
- 3 In the content pane of the window (right side), click **Add** .
- 4 In the Select Check widow, from the Browse tab you can browse for the compliance checks categories and select a check for the audit or snapshot.
- 5 Select the Search tab to search for checks by name. The check search tool searches on the name of a check and any words in a check's description. For example, if you wanted to find all rules that check for maximum password length, you could enter `max password` in the Keywords field.

- 6 Click the Advanced Search link to build more specific search criteria. Advanced search allows to look for a text string plus restrict the query to values in the check's properties, such as Security Level, External ID, Platforms, and Test ID. Click  to add additional advanced search parameters.

For information on how to add a Test ID, Security Level, or External ID to your compliance check properties, see [Editing Compliance Check Properties](#) on page 154.
- 7 To execute the search, click **Search**.
- 8 In the search results, you can select the checks you want to add to the audit or snapshot specification and then click **OK**.

Managing Compliance Checks

The Compliance Check Editor allows you to browse, regroup, and edit property information (metadata) about your core's BSA Essentials Subscription Services compliance checks.

For example, your organization might require that an external numbering system be associated with all the compliance checks that you run against servers in your data center.

Using the Compliance Check Editor, you can add an external ID to those checks. You can also create custom groupings for checks you modify with this external ID, so that anytime you need to access those checks you can easily find them in the custom folder. This external ID can also be used as a search criteria so you can search all checks with the ID number, or string.

You can also edit information about custom check, such as changing a check's name, adding a custom security level, or modifying descriptive information about a check. For example, you can add to a check's remediation description to more clearly explain what happens during remediation, so that someone else who wants to use the check can better understand its behavior.



You must have permissions to access the Compliance Check Editor. To obtain these permissions, contact your SA administrator. Or, see the *SA Policy Setter's Guide* for more information.

Editing Compliance Check Properties


The Compliance Check Editor allows you to modify a compliance check's properties, such as renaming it, adding a description, modifying its property information, adding an external ID to it, and so on.


To edit compliance check property information, perform the following steps:

- 1 Inside the SA Client, from the **Tools** menu, select **Compliance Check Editor**. (If you do not see the menu item, contact your SA administrator to obtain permissions to access the Compliance Check Editor.)
- 2 In the Compliance Check Editor window, in the Browse tab you can expand the different Custom Checks categories to find the check you want to edit. You can narrow the list by selecting the Platforms filter.
- 3 Select the Search tab if you want to search for a check by a name or a keyword in its name and Description fields.

For example, if you wanted to find all rules that check for security logs, you could enter `security` and `log` in the **Keywords** field. If you wanted to narrow the search further, you could add the keyword `size` to find all checks that audit for security log file size.

The Advanced search option allows you to set more specific parameters to find checks. Using advanced search, you can filter by other properties such as security level, external ID, platforms, or test ID.

To add additional search parameters, click the plus  icon.

- 4 To edit a check's property information, select the check from the Browse tab or Search tab results.
- 5 On the right side of the Compliance Check Editor, in the Properties tab you can edit the following check information:
 - **Name:** Double-click inside the Name's value field to modify the check's name.
 - **Categories:** Click the Click to Edit link to add the check to a custom folder. For example, click the link and in the Categories window, press ENTER on your keyboard and then type a name to create a new compliance check category. Click **Apply**. To create the custom grouping folder, click **Apply Changes** at the bottom of the Compliance Check Editor window. For information on creating custom grouping for your checks, see [Creating Custom Compliance Check Categories](#) on page 155.
 - **External ID:** Double-click inside the value field to add or modify an External ID.
 - **Security Level:** Double-click inside the value field to enter or modify security level for the check.
- 6 Click **Apply Changes** at the bottom of the Compliance Check Editor window to apply the modifications to the checks.
- 7 To edit a check's descriptions, in the right-hand lower side of the Compliance Check Editor, select the Description, Remediation Description, or Technical Descriptions tabs to edit the descriptive text for each.
- 8 To access the HTML editor for the descriptions, click the Edit  icon.
- 9 In the HTML editor, click the HTML Edit icon at the bottom left of the window.
- 10 Edit the HTML description.
- 11 Click **Apply**. (If you want to undo any changed, from the **File** menu, select **Revert**.)
- 12 Click **Apply Changes** at the bottom of the Compliance Check Editor window to apply the description modifications to the checks.

Creating Custom Compliance Check Categories

The Compliance Check Editor allows you to create your own custom categories to contain the compliance checks installed on your core. For example, you could create a custom category to contain all the checks that audit user and group settings on your Windows servers. Or, you might only be interesting in accessing specific Linux services-related checks and could create a category to contain them.

To create custom compliance check categories, perform the following steps:

- 1 Inside the SA Client, from the **Tools** menu, select **Compliance Check Editor**. (If you do not see the menu item, contact your SA administrator to obtain permissions to access the Compliance Check Editor.)

- 2 In the Compliance Check Editor window, in the Browse tab you can expand the different Custom Checks categories to find the check you want to edit. You can narrow the list by selecting the Platforms filter.
- 3 Select a compliance check.
- 4 In the upper right side of the Compliance Check Editor window, Properties tab, Categories row, click Click to Edit.
- 5 In the Categories window, place your mouse point at the end of the main check category name and then press ENTER on your keyboard.
- 6 Type a name to create a new compliance check category. This creates a new compliance check category in the Compliance Check Editor. To add more categories, press ENTER again to start a new line and then type the name of the category. The selected check will be added to each new category.
- 7 Click **Apply**.
- 8 To create the custom grouping folder, click **Apply Changes** at the bottom of the Compliance Check Editor window.
- 9 To delete the custom category, repeat the process and delete the name of the category in the Categories window.

Restoring Compliance Checks to Defaults

If you would like to restore all of your compliance checks to their default state — their original state when they were first downloaded from the BSA Essentials Subscription Services portal — use the Restore Defaults operation.

Restore Defaults delete any customizations made to your compliance checks, reverts them to their original released state.

To restore compliance checks to their default state, perform the following steps:

- 1 Inside the SA Client, from the **Tools** menu, select **Compliance Check Editor**. (If you do not see the menu item, contact your SA administrator to obtain permissions to access the Compliance Check Editor.)
- 2 In the Compliance Check Editor window, from **the** Edit menu select **Restore defaults**.



Restoring defaults only applies to select compliance checks.

Showing Deprecated Checks

For those compliance checks that have been deprecated, you can choose to show them in the Compliance Check Editor.

To show deprecated checks in the Compliance Check Editor, perform the following steps:

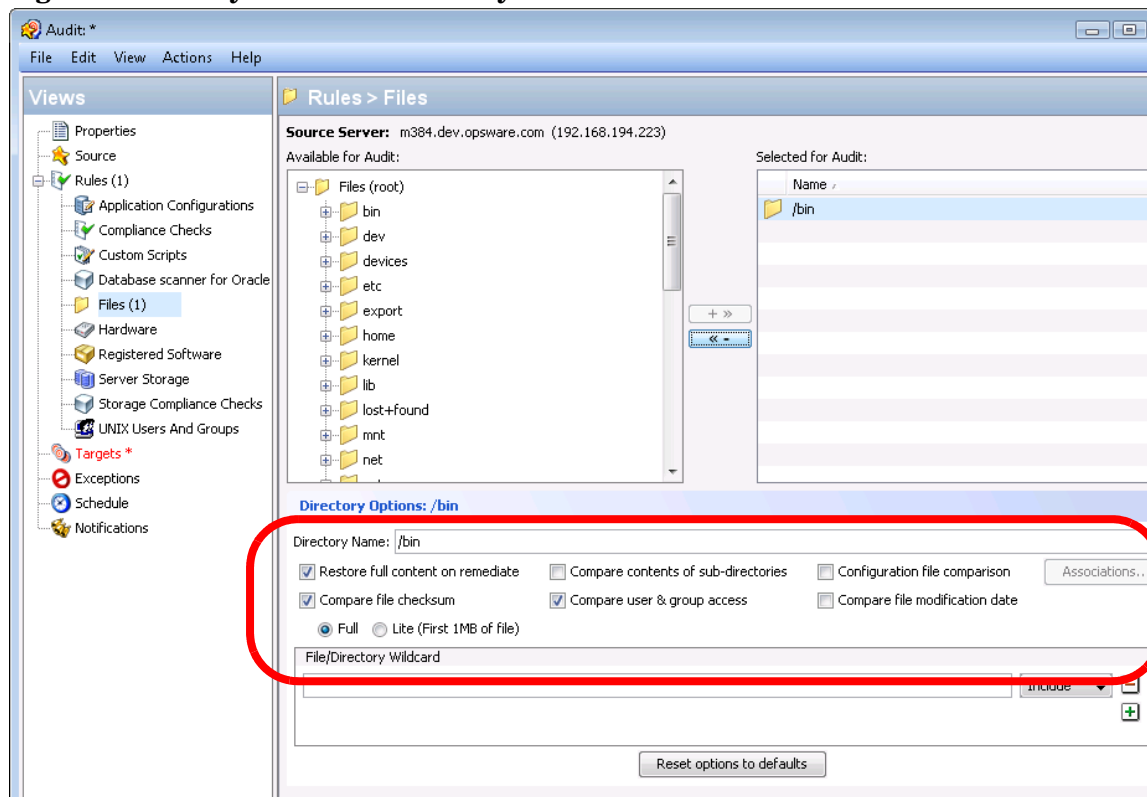
- 1 Inside the SA Client, from the **Tools** menu, select **Compliance Check Editor**. (If you do not see the menu item, contact your SA administrator to obtain permissions to access the Compliance Check Editor.)
- 2 From the **View** menu, select **Show Deprecated Checks**.
- 3 Expand any of the check categories to view any deprecated checks, which will appear in grayed out, italic font.

File Inclusion and Exclusion Rules

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

Inclusions and exclusion rules inside of an audit's file rule are found at the bottom of the audit or snapshot specification window, as shown in [Figure 23](#).

Figure 23 File System File/Directory Wildcard Inclusion and Exclusion Rules



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

For information on how to create and configure file system rules for an audit or snapshot specification, see [Configuring the File Rule](#) on page 135.

Inclusion and Exclusion Rule Types

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

- A file-type rule applies to the file name path and contains neither a “/” or a “\”.
- A relative-type rule applies to the relative path and can contain a “/” for Unix and a “\” for Windows, and is not fully qualified.

- An absolute-type rule applies to the absolute path. In Unix, an absolute path begins with a “/”. In Windows, an absolute path begins with a volume letter that is followed by “:\” and is fully qualified, such as “C:\”, “d:\”, “f:\”, and so on. If you use a “/” (forward slash) for Windows paths, Audit and Remediation will convert it to a “\” (backslash) to use it as a valid path.
- Environment variable and custom attribute parameterization for filenames and path. For more information, see [Parameterizing Filenames for SA/Custom Attributes](#) on page 162.

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

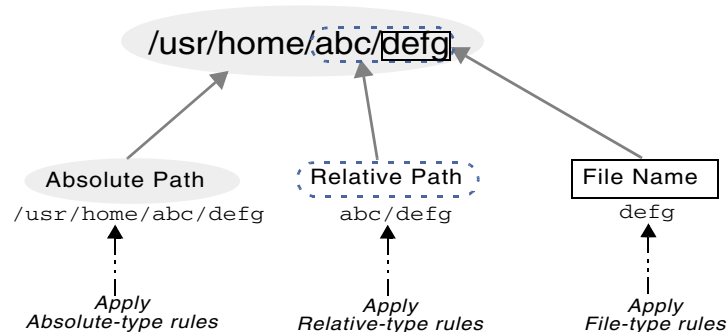
You can also use the asterisk (*) and the question mark (?) as valid wildcards in inclusion and exclusion rules. The wildcard character is a placeholder for matching a path, or one or more characters.

Depending on the type of inclusion and exclusion rule, the rule is applied only to the relevant subset of the absolute path of the file. In Audit and Remediation, there is one top level for each snapshot or audit. Each file that you compare against the inclusion and exclusion rules has an absolute path. In [Figure 24](#), the absolute path is `/usr/home/abc/defg`. A snapshot or an audit looks down the `/usr/home/abc/defg` absolute path and sees `abc/defg` as the relative path and `defg` as the file name. In this example, the inclusion and exclusion rules apply in the following manner:

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

See [Figure 24](#) for an illustration of how Audit and Remediation applies the inclusion and exclusion rules to a relative subset of the path of the file.

Figure 24 How Inclusion and Exclusion Rules Apply



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

A sample file system structure used in [Example: Including all .txt Files in a Snapshot or Audit](#) on page 159 and [Example: Including last temp.txt file and exclude all else](#) on page 160 is as follows:

```

/dir1/dir2/a
/dir1/dir2/b
/dir1/dir2/names.txt
/dir1/dir2/temp.txt
/dir1/dir2/version1.exe

```

/dir1/dir2/subdir/version2.exe

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

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

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

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

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

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

Example: Including Only File a in a Snapshot or Audit

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

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

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

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

- f Compare b to *, which is a match; compare b to a which is not a match. The file is excluded.

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

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

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

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

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

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

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

File Rule Overlap

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

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

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

Example A

If you create a snapshot using the following two rules:

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

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

The snapshot will record the following file system information:

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

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

Example B

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

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

The snapshot will record the following file system information:

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

Example C

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

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

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

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

The snapshot will record the following file system information:

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

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

Parameterizing Filenames for SA/Custom Attributes

When you create a file rule in an audit or snapshot specification, you can also reference environment variables and custom attributes in the file name. In the File/Directory Wildcard area of the rule window, you can edit the file name to add these references.

To add a reference to a Windows environment variable, the syntax is `%envVarName%` and for Unix, the syntax is `${varName}`.

The syntax for specifying custom attributes is `@varName@`. For example:

```
@/customattribute/custAttributeNAME@\rest\of\the\path
@/customattribute/FacilityCustomAttributeNAME@\rest\of\the\path
@/customattribute/CustomerCustomAttributeNAME@\rest\of\the\path
@/customattribute/ServerAttributeNAME@\rest\of\the\path
@/customattribute/GrpAttributeNAME@\rest\of\the\path
```

This allows for auditing relative paths on both source and target servers using a parameterized environment variable or custom attribute in the filename.

Examples of Parameterizing Filenames

For example, on the servers you want to audit you know the relative path to an application, but not necessarily the absolute path for all servers. You can parameterize the path in your audit's File rule so the relative pathname is eliminated and the Audit checks the relative path anywhere it exists on the target server.

For example, you want to Audit a target servers against a golden source server where `'%ProgramFiles%'` is `:\Program Files` against target servers where `%ProgramFiles%` is `D:\Program Files`.

In the File/Directory Wildcard section of the File rule, you can specify the root of the directory rule in the Audit to be `%ProgramFiles%\Company\MyApp`. The audit will remove `%ProgramFiles%` from the paths of the servers it targets when you run the Audit. In other words, `C:\Program Files\Company\MyApp\file1.txt` on the source server will be compared with `D:\Program Files\Company\MyApp\file1.txt` on the target servers.

In another example, you may want to audit an application that is installed into two completely different subdirectories on two different servers.

For example, in your audit you choose from a golden source server configuration the installation path of the following:

```
/usr/local/app-version-1232/prog
```

And, your target servers have the application installed anywhere under this path:

```
/usr/local/app
```

In order to audit the target server, you can defines a custom attribute `APP_INSTALL_LOC` with a value of `/usr/local/app-version-1232/prog` for the golden server and `/usr/local/app` for the production servers. The File rule in the Audit would look something like this:

```
@/customattribute/APP_INSTALL_LOC@/prog
```

This would cause the audit to treat `@/opsware/customattribute/APP_INSTALL_LOC@` as if it were an environment variable on the target server and do a path replacement.

If you wanted to reference a server attribute, the path would be entered like this:

```
@/server/APP_INSTALL_LOC@/prog
```

Using Environment Variables in Pathnames

If you want to use environment variables in file name PATH on Unix (also known as “parameterized checks”), it is best to define those environment variables under following file and directory: `etc/opt/opsware/snapshot/env`. Be sure that you do not use `/etc/profile` to source environment variables on Unix.

To define environment variables that can be sourced for the File rule configuration, you can create a file with the variables on the managed server you want to audit or snapshot. For example:

- 1 ssh to the managed server that you want to audit or snapshot.
- 2 Create a new directory in the following location:

```
mkdir /etc/opt/opsware/snapshot
```
- 3 Create a new empty file. For example:

```
touch /etc/opt/opsware/snapshot/env
```
- 4 Define the environment variables you want to source from the File rule by entering them in the new file. For example:

```
TEST1='/tmp/test1'  
TEST2='/home/test2'  
export TEST1 TEST2
```
- 5 When you have finished editing, save the file.

Audit Rule Exceptions

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

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

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

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

Rules That Cannot Have Exceptions

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

Considerations When Applying Exceptions to Device Groups

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

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


Adding a Rule Exception to an Audit

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

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

Some audit rules and audit rule collections cannot be excepted. For more information, see [Rules That Cannot Have Exceptions](#) on page 164.

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

- 1 First, create an audit. For information see [Creating an Audit](#) on page 113.
- 2 Configure audit rules for the audit. For information on configuring audit rules, see [Audit and Remediation Rules](#) on page 123.
- 3 From the audit view pane on the left, select the Exception  object.
- 4 Next, from the content pane, click **Add**.



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

- 5 In the Add Exception window, from the Select Target Server section, select a server, multiple servers, or device groups to which you want to apply the rule exception.
- 6 Next, from the Select Rule section, select one or more rules you want to associated with the servers you selected in the previous step.
- 7 (Optional) In the Reason for Exception section, add an explanation.
- 8 (Optional) In the Ticket ID section, add the ticket ID associated with this exception.


- 9 In the Expires section, either enter a date to indicate when the exception expires, or select a date from the drop down list.
- 10 When you are finished configuring the exception, click **Add**.
- 11 You now see a list of rule exceptions that will be applied when you run the audit.

Editing or Deleting a Rule Exception


You can edit an exception in one of two ways:

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

To edit an exception, perform the following steps:

- 1 Open an audit window.
- 2 From the audit's View pane on the left, select the Exception  icon.
- 3 From Contents pane, double-click an exception.
- 4 In the Edit Exception window, you can edit any of the exceptions and servers or device groups they are assigned to. When you have edited the exception, click **Add**.
- 5 If you want to completely change and the rule, click the **Add** button and then in the Add Exception window, change the rule by selecting target server and one or more rules. When you are finished, click **Add** to change the exception.

To delete an exception, perform the following steps:

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

Audit Policies

Audit policies allow you to define and store a centralized and reusable sets of server configuration compliance rules and allowing you to link them to audits, snapshot specifications, and other audit policies.

Typically, audit policies are created by *policy setters* who understand the exact compliance standards that a company wants its servers to meet. Another set of users, whose job it is to manage and audit actual servers, can utilized predefined audit policies by linking them to their audits or snapshot specifications. If any changes are made to the audit policy, the audit or snapshot specification that links to it will reference the audit policy's updated rules. This way, those users who audit servers can be sure their audits always reflect the latest compliance standards in their organization.

Another useful feature of audit policies is that they can link to other audit policies. For example, you could combine several different discrete audit policies together one master policy that defines how Windows services should be configured. After you run the audit, if any discrepancies are discovered you can remediate them from the audit results.

You can create an audit policy from scratch, or you can save the rules of an audit, snapshot specification (or another audit policy) as an audit policy. All audit policies are stored in the SA Client Library.

For information on creating rules for an audit policy, see [Audit and Remediation Rules](#) on page 123.

Linking vs. Importing Audit Policies

An audit policy can be used inside audits and snapshot specifications, or other audit policies, through *linking*. Audits and snapshot specifications and also use audit policies through *importing*.

Linking an audit policy to an audit or snapshot specification enables the audit or snapshot specification to use the exact same rule set of the audit policy. If any of the rules in the audit policy change, the same changes are reflected in the audit and snapshot specification's rules the next time they are run, since they link to the rule set defined in the audit policy. (You can break this link by selecting the Enable Unlinked rules option.)

Audit policies can be also be linked to other audit policies, and you can link as many audit policies as you want into an audit policy. When you link one or more audit policies to an audit policy, the linked audit policies become children of the parent audit policy. Thus, if you create an audit that links to the parent audit policy, when you run the audit on a target server, the rules from all linked policies are run against on the target server.

Importing an audit policy into an audit or snapshot specification imports all the rules from the audit policy, and once imported the rules are editable. When you import an audit policy into an audit, you can choose to replace any current values in the audit or merge rules from the audit policy with those in the audit or snapshot specification. Audit policies cannot import rules from another audit policy, but they can link to other audit policies.

For more information about linking and importing audit policies, see [Linking and Importing Audit Policies](#) on page 168.

Rule Overlap with Multiple Linked Audit Policies

Because you can link your audit or snapshot specifications to an audit policy that may references other audit policies, it is possible that some of the linked policies might contain the same rules but with different configuration options.

Rules become merged in audit results when you identify the same object for a rule and the only way to customize the rule is by setting options. The options may or may not be different, but they still get merged into one rule before running and there is only one result. If the options are different, the options are OR'ed together into the single rule. Examples include file rules, registry rules, metabase rules (legacy comparison type), Windows Service rules, etc.

Rules that take parameters or you specify the compliance criteria are merged if and only if the parameters and the criteria are exactly the same. Otherwise they are executed as separated rules. Examples include compliance (pluggable) rules, custom script rules, and server module based rules.

Creating an Audit Policy



When you create an audit policy, you have the option of creating its rules using a live server as a source to pick and choose for rules, by creating your own custom rules, or linking to the rules of another audit policy.

Using a source server for building audit policy rules allows you to base the audit policy's rules on the actual configurations of a managed server. The source server used to build rules, however, is not used once the audit policy is linked to an audit or snapshot specification.



All audit policies must be saved to a folder in the SA Client Library. To save an audit policy to a folder, you must have permissions to write to that folder. For more information about folder permissions, see the *SA Administration Guide*.

To create an audit policy, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Audit and Remediation ► Audit Policies, and then select Windows or Unix.
- 2 From the **Actions** menu, select **New**.
- 3 In the Content pane, for the audit policy's Properties, enter a name and description.
- 4 In the Properties, you also need to specify a location in the Library where you want to save the audit policy. Next to the Location field, click **Select**.
- 5 In the Select Folder window, select a location to save the audit policy. You must have permissions to write to the folder where you save the policy. Click **Select** when you have chosen a location.
- 6 Next, from the Views pane on the left side of the Audit Policy window, select Source if you would like use a managed server to base the audit policy's rules on.
- 7 From the Content pane, click a **Select** to choose a source server for the audit policy.
- 8 In the Select a Source window, select a server and then click **OK**.
- 9 From the Views pane, select the Rules view.
- 10 If you want to link other audit policies into this audit, click the Add  icon to select an audit policy.
- 11 If you would like to edit any of the linked audit policies, from the Rules list, select an audit policy and click the View  icon.
- 12 In the Select a Policy window, select one or more audit policies to link to the audit policy, and click **OK**. If you link one or more audit policies to an audit policy, you can still configure individual rules in the audit policy. All rules from externally revered audit policy will be combined with any invalidity rules you create to build one single rule set.
- 13 From the Rules list in the audit policy View list (left side), create any other rules you want to include in the audit policy. For more information on how to configure specific audit and remediation rules, see [Configuring Specific Audit and Snapshot Rules](#) on page 126.
- 14 When you are finished configuring the audit, from the **File** menu select **Save**. Once saved, the audit policy is ready to be linked to an audit, snapshot specification, or other audit policy.

Linking and Importing Audit Policies

You can import or save an audit policy into an audit, snapshot specification, or other audit policy:

- [Linking an Audit Policy to Audits or Snapshot Specifications](#)
- [Linking Audit Policies to an Audit Policy](#)
- [Importing Audit Policy Rules](#) (replace or merge)
- [Saving as Audit Policy](#)

Linking an Audit Policy to Audits or Snapshot Specifications

Linking an audit policy into an audit or snapshot specification creates a link that uses the rules from the audit policy for the audit or snapshot specification. Linking to an audit policy is useful if a policy setter wants to define a server configuration policy for servers and have others users link their audits and snapshot specifications to the audit policy. If the policy setter makes any changes to the audit policy, the changes will be reflected audit or snapshot specification that is linked to the policy.

When an audit policy is linked into an audit or snapshot specification, the rules cannot be modified in the context of the audit or snapshot specification. However, you can access the audit policy and edit its rules (given proper user permissions).

If the audit or snapshot specification you are linking the audit policy to already has some rules defined, then those preexisting rules in the audit or snapshot specification will be overwritten when you link to an external audit policy.

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

- 1 Open an existing audit from the Library using one of the following methods:
 - From the Navigation pane, select **Library** ► **Audit and Remediation** ► **Snapshot Specification**, and then open the audit.
 - From the Navigation pane, open an existing snapshot specification from **Library** ► **Audit and Remediation** ► **Snapshot Specification**.
- 2 From the **Actions** menu, select **Link to Policy**.
- 3 In the Select a Policy window, select an audit policy to link to the audit or snapshot. You can only link to one audit policy per audit or snapshot specification. However, you can link multiple audit policies to an audit policy. For more information, see [Creating an Audit Policy](#) on page 167 or [Linking Audit Policies to an Audit Policy](#) on page 169.
- 4 When you have selected an audit policy, click **OK**.
- 5 If you are linking an audit policy into an audit or snapshot specification that already has rules defined, a message window asks if you want overwrite any existing rule definitions. Click **Yes** to import the audit policy.
- 6 To save the audit or snapshot specification, from the **File** menu, select **Save**.



Linking Audit Policies to an Audit Policy

Linking an audit policy to other audit policies enables you to combine multiple audit policies into a single, “master” audit policy. Because you can link as many audit policies as you want into an audit policy, you can build and reuse existing audit policies into a single audit policy that meets a specific auditing need.

When you link one or more audit policies to an audit policy, the linked audit policies become children of the parent (or master) audit policy. If you create an audit that links to the parent audit policy, when you run the audit on a target server, the rules from all linked policies are run against on the target server.

For example, your SA Client library might contain several individual audit policies that define compliance standards for a set of Unix servers. One policy might contain a set of rules that checks to make sure the FTP services are enabled, while another policy makes sure that cron logging is always enabled. You can create a single audit policy that links to these other policies, and so on, and then this master policy can be referenced by other audits.

To link audit policy to other audit policies, perform the following steps:

- 1 From the Navigation pane, select **Library > By Type > Audit and Remediation > Audit Policies**, and then select **Windows or Unix**.
- 2 Select an existing audit policy, or, to create a new audit policy, from the **Actions** menu, select **New**.
- 3 If this is a new audit policy, in the Content pane, for the audit policy’s Properties, enter a name and description. If you are linking audit policies to an existing audit policy, skip to [step 9](#) on page 169.
- 4 In the Properties, you also need to specify a location in the Library where you want to save the audit policy. Next to the Location field, click **Select**.
- 5 In the Select Folder window, select a location to save the audit policy. You must have permissions to write to the folder where you save the policy. Click **Select** when you have chosen a location.
- 6 Next, from the Views pane on the left side of the Audit Policy window, select **Source** if you would like use a managed server to base the audit policy’s rules on.
- 7 From the Content pane, click a **Select** to choose a source server for the audit policy.
- 8 In the Select a Source window, select a server and then click **OK**.
- 9 From the Views pane, select the **Rules** view.
- 10 Click the Add  icon to select an audit policy.
- 11 In the Select a Policy window, select one or more audit policies to link to the audit policy, and click **OK**. If you link one or more audit policies to an audit policy, you can still configure individual rules in the audit policy. All rules from externally referenced audit policy will be combined with any invalidity rules you create inside the audit policy.
- 12 From the Rules list in the audit policy (left side), create any other rules you want to include in the audit policy. For more information on how to configure specific audit and remediation rules, see [Configuring Specific Audit and Snapshot Rules](#) on page 126.
- 13 If you would like to edit any of the linked audit policies, from the Rules list, select an audit policy and click the View  icon.
- 14 When you are finished configuring the audit policy, from the **File** menu select **Save**.

Importing Audit Policy Rules

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

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

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

- 1 Open an existing audit or snapshot specification from the Library using one of the following methods:
 - From the Navigation pane, select **Library** ► **Audit and Remediation** ► **Audits**, and then open the audit.
 - From the Navigation pane, open an existing snapshot specification from **Library** ► **Audit and Remediation** ► **Snapshot Specification**.
- 2 From the **Actions** menu, select **Link to Policy**.
- 3 If the audit or snapshot specification already has rules defined, choose to either to overwrite the existing rules, or merge the audit policy rules with the existing rules:
 - If you click **Yes**, then the audit policy will overwrite any existing rules in the audit or snapshot specification.
 - If you click **No**, then the audit policy will merge the audit policy rules with any existing rules. If any conflicts are found, then the audit policy rules will overwrite any existing rules.
- 4 To save the audit or snapshot specification, from the **File** menu, select **Save**.

Saving as Audit Policy

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



All audit policies must be saved to a folder in the SA Client Library. To save an audit policy to a folder, you must have permissions to write to that folder. For more information about folder permissions, see the *SA Administration Guide*.

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

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

- 6 In the Select Folder window, choose a folder where you want to save the audit policy, and then click **OK**. The audit policy is saved and can be accessed at **Library ► Audit and Remediation ► Audit Policies**.

Locating an Audit Policy in the Folder Library

Once you create and save an audit policy to the folder library, you can easily find the audit policy in the library by using the Locate in Folders feature.

To locate an audit policy in folder, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Audit and Remediation ► Audit Policies**, and then select **Windows** or **Unix**.
- 2 Select an audit, right-click, and select **Locate in Folders**. The location where the audit policy is saved is now visible.

Exporting an Audit Policy to HTML or CSV

If you want to get a list of all the rules contained and configured in an audit policy, you can export the policy to either HTML or CSV.

To export an audit policy to HTML or CSV, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Audit and Remediation ► Audit Policies**, and then select **Windows** or **Unix**.
- 2 Open an audit policy by double-clicking it, or, right-clicking and selecting **Open**.
- 3 From the Actions menu, select **Export ► HTML** or **CSV**.
- 4 Select a path and filename for the file, and then click **Export**.
- 5 Open the file to view the exported information.

Running an Audit

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

- [Running an Audit from the Library](#)
- [Running an Audit on a Server from All Managed Servers](#)
- [Re-running an Audit from Audit Results](#)

Running an Audit from the Library

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

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

- 1 From the Navigation pane, select **Library ► By Type ► Audit and Remediation**.
- 2 Select Audits, and then select either Windows or Unix.
- 3 Select the audit you want to run, right-click, and select **Run Audit**.
- 4 In the Run Audit window, step one shows you the name of the audit, the source server or snapshot being used in the Audit, the total number of rules defined in the audit, and all targets of the audit (servers and snapshot). Click **View Rule Details** to view the rule definitions.

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

Running an Audit on a Server from All Managed Servers

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

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

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

(If you would like to run the audit immediately, click the Start Job button at any point in the process.)
- 6 Click Next.

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

Re-running an Audit from Audit Results

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

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

To rerun an audit, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Audit and Remediation**.
- 2 Select Audits, and then select either Windows or Unix.
- 3 Select an audit, and in the Details pane, select an audit result for the audit. (Each time the audit is run, its results are accumulated in the Details pane.)
- 4 Double-click the audit result to open it.
- 5 From the **Actions** menu, select **Re-Run audit**.
- 6 In the Run Audit window, step one shows you the name of the audit, the source server or snapshot being used in the Audit, the total number of rules defined in the audit, and all targets of the audit (servers and snapshot). Click **View Rule Details** to view the rule definitions.

(If you would like to run the audit immediately, click the Start Job button at any point in the process.)
- 7 Click Next.
- 8 In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 9 Click Next.
- 10 In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click Add Notifier and enter an email address.

- 11 (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.
- 12 (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.
- 13 Click Next.
- 14 In the Job Status page, click Start Job to run the audit. When the audit has run, click View Results to view the results of the audit.

Scheduling an Audit

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



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

Scheduling a Recurring Audit

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

To schedule a recurring audit, perform the following steps:

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

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

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.

- 5 In the Time and Duration section, for each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose a date to end the audit schedule, select **End** and then choose a date. The Time Zone is set according to the time zone set in your user profile.
- 6 (Optional) Deselect the **End** option, if you want the audit schedule to run indefinitely.
- 7 To save the audit schedule, from the **File** menu, select **Save**. The audit will now run according to the defined schedule.

Editing an Audit Schedule

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

To edit a scheduled audit, perform the following steps:

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

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

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.

- 7 In the Time and Duration section, for each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose a date to end the audit schedule, select **End** and then choose a date. The Time Zone is set according to the time zone set in your user profile.
- 8 (Optional) Deselect the **End** option, if you want the audit schedule to run indefinitely.
- 9 To save the audit schedule, from the **File** menu, select **Save**. The audit will now run according to the defined schedule.

Viewing a Completed Audit Job

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

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

Remediating Audit Results

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

Whether or not you can remediate a rule depends upon the rule type. The rule must support remediation and the source of the audit rule for that server must contain enough data to support the remediation.

For example, some rules do not support remediation, such as a Hardware rule. You cannot “remediate” a server's physical memory or hardware. If your audit is using a snapshot as a source, and the snapshot was unable to gather sufficient information from a rule, then that rule will not be remediable as well.

For those audits that link to audit policies, the results will show all of the rules in the audit, but the results do not show the audit policy or policies where the rules were originally defined.

For more information, see the following topics:

- [Remediating Rules with Inherited Values](#)
- [Accessing Audit Results](#)
- [Audit Results Window](#)
- [Remediation Methods: All, By Server, or By Rules](#)
- [Viewing and Remediating Audit Results Differences](#)
- [Viewing Audit Results with Exceptions](#)

Remediating Rules with Inherited Values

If you create an audit rule based on an object that inherits properties from a parent object, be aware that if you remediate the rule, the target server object will not inherit the parent object's properties.

For example, if you created a rule for a Registry entry, and that registry entry inherited some values from a parent, when you remediate the rule on to a target server, none of the values inherited from its parent will be remediated, and the rule will show in the audit results as a difference.

Additionally, if your audit checks ACLs for the File, Registry, or IIS Metabase rules, and the user and group ACL does not exist, then after the audit is run and after remediation, if user and group does not exist on target a temporary user and group will be created as unknown name. The next time you run Audit it shows up as unknown, which shows name other than the source user.

Additionally, if you create an IIS Metabase rule from a source server and the metabase object selected for the rule inherits its values from a parent Metabase object, differences will show after an audit is run. For example, if you remediate once and then rerun the audit, if the source key was not inherited and the attribute has an IED when it gets created on target server, the object will be created based on parent key inheritance. When you rerun the audit, the results will show the IED as a difference for the object's attribute.

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

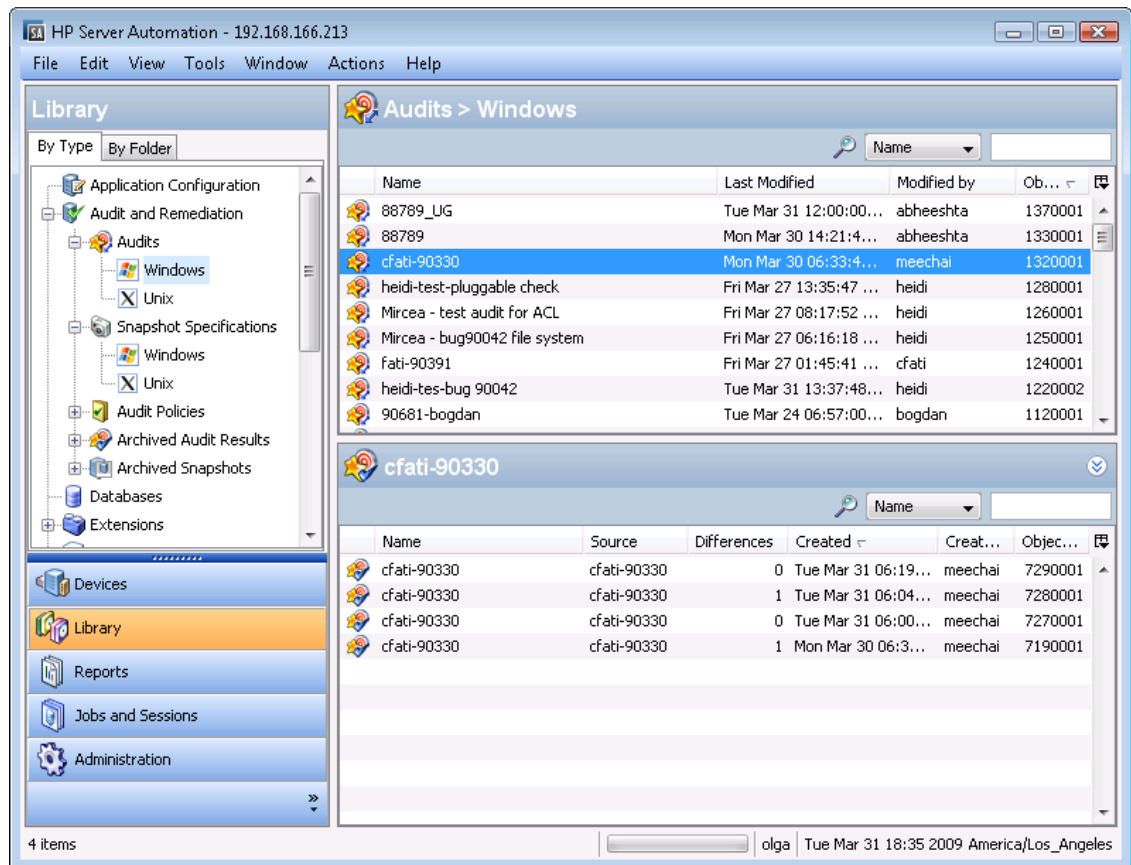
- SA Audit and Remediation does not support the remediation of the following two values on Windows 2000 servers for the Windows Local Security Settings rule, under Security Options: Rename AdministratorAccount and Rename Guest Account.

- You cannot remediate ISAPI filters for the IIS 7.0 audit rule in this release.

Accessing Audit Results

In the SA Client, you can view audit results for any audit, as shown in [Figure 25](#).

Figure 25 Audit Results



When you select an audit in the library, all results associated with the audit appear in the Details pane below.

Audit Results Window

The Audit Results window provides information regarding the audit job, most importantly, any differences between servers targeted by the audit and the rules defined in the audit. From this window, you can determine the extent to which the servers that have been audited are in compliance with the standards set for your data center.



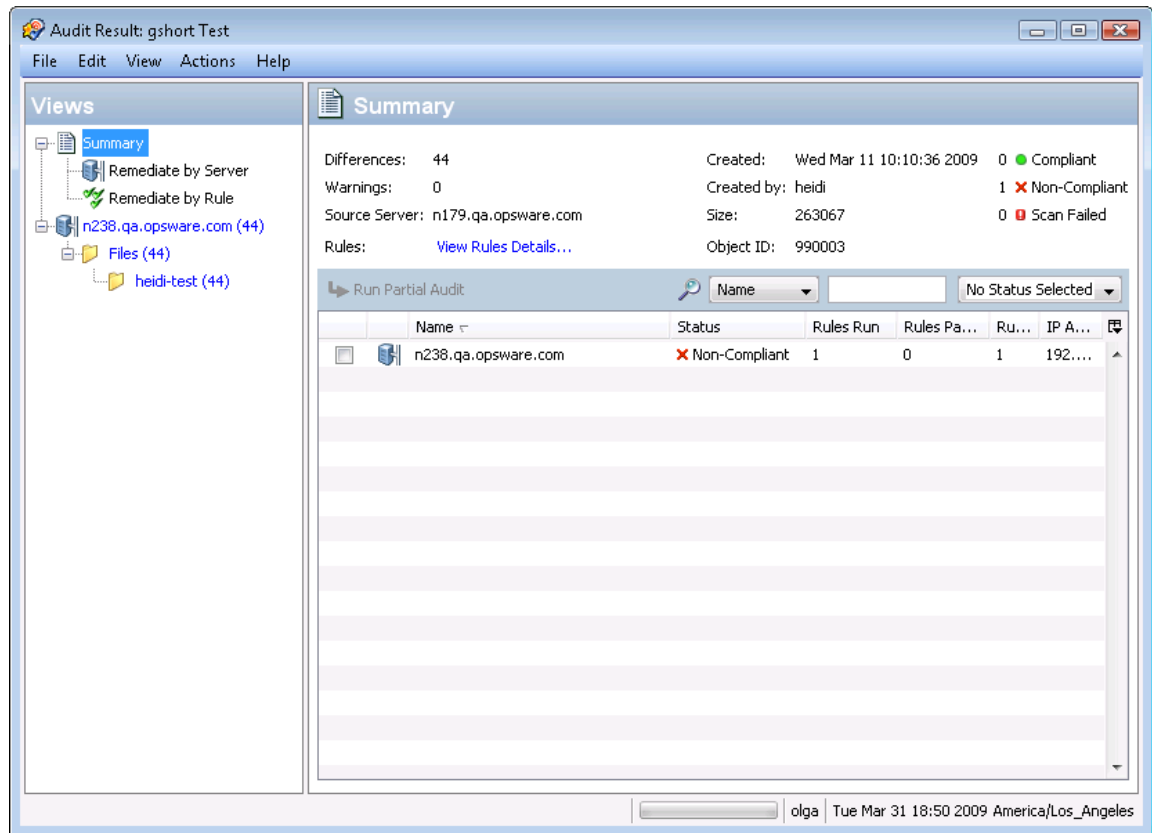
For those audits that link to an audit policy, the results show all of the rules in the audit, but the results do not show the audit policy or policies where the rules were originally defined.

Each Audit Results window provides two main elements:

- **Summary:** A an overview of the audit job that provides a compliance status for the audit and remediation options that allow you to remediate by server, by rule, or remediate all rules on all servers. Remediation is only available in instances where the target server configurations do not match the rule definitions in the audit.

- **Server List:** A list of all servers that the audit was run against, and a list of all rules from the audit, showing which rules if any are out of compliance.

Figure 26 Audit Results Window



Audit Results Summary

The audit results summary view describes the following information about an audit job:

- The number of differences discovered during the audit, which is the sum total of all configuration differences found on all target servers and the defined in the audit rules.
- Compliance information related to the audit, detailing which rules in the audit are
 - Compliant: ● Number of rules for which target server configurations matched the rules in audit.
 - Non-compliant: ✗ Number of target server configurations that did not match the rules in the audit.
 - Scan Failed: ❗ Number of rules for which the audit was unable to determine the target server configuration.
- The user who created the audit, when it was created, the source server used in the audit the results are based upon.
- The **View Rule Details** link launches the audit window so you can view the audit's rule.
- The **Run Partial Audit** link allows you to to select servers and re-run the audit on only those rules that have a Non-Compliant or Scan Failed compliance status.


For more information about audit remediation types, such as remediate by server or remediate by rule, see [Remediation Methods: All, By Server, or By Rules](#) on page 180.

Remediation Methods: All, By Server, or By Rules

From the audit results window, there are three different ways to remediate non-compliant rules in audit results:

- **Remediate All:** From the Audit Results window, in the Actions menu select Remediate all to remediate differences found in the audit results.
- **Remediate by Server:** Remediate by servers targeted by the audit results.
- **Remediate By Rule:** Remediate specific, individual audit rules.

Remediate All

You can select to remediate all the differences found in an audit result for all rules that are remediable. This option remediates all remediable rules on all servers targeted by the audit. Rules that have a status of Compliant  are not remediated when the audit is run.

To remediate all differences found in an audit results, perform the following steps:

- 1 From the Navigation pane select **Library > By Type > Audit and Remediation > Audits**.
- 2 Select an audit. From the Details pane below the audit list, you see all audit results associated with the audit.
- 3 Select an audit result, right-click, and select **Open**.
- 4 In the Audit Result window, from the Actions menu select **Remediate All**.
- 5 In the Remediate Audit window, step one shows you the name of the audit, the target of the Audit, the total number of rules defined in the audit. If you wish to bypass all of the audit task steps, click **Start Job** to immediately run audit job.
- 6 Click **Next**.
- 7 In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 8 Click **Next**.
- 9 In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 10 (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.
- 11 (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Professional Services has integrated SA with your change control systems. It should be left blank otherwise.
- 12 Click **Next**.
- 13 In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

Remediate By Rule

You can remediate specific differences found in rules in audit results by selecting individual rules that are out of compliance, and then re-running the audit to remediate only the rules you select. You can select to remediate by individual rule for all servers targeted by the audit, or choose only selected servers to have rules remediated.

To remediate specific differences found in an audit results, perform the following steps:

- 1 From the Navigation pane select **Library ► By Type ► Audit and Remediation ► Audits**.
- 2 Select an audit.
- 3 From the Details pane below the audit list, you see all audit results associated with the audit.
- 4 Select an audit result, right-click, and select **Open**.
- 5 In the Audit Result window, expand the Summary list, and then select Remediate By Rule. You see all differences discovered by rule in the audit results.
- 6 For each rule you want to remediate, select the check mark in the list in the All Servers column, which means that when you remediate the audit results, the rule will be remediated on all servers targeted by the audit that the rule is applied to.

If you want to globally select all rules, right-click and select **Select All**. To deselect all rules, right-click and select **Deselect All**.
- 7 When you have selected the rules you want to remediate, from the **Actions** menu, select **Remediate**.
- 8 In the Remediate Audit window, step one shows you the name of the audit, the target of the Audit, the total number of rules defined in the audit.
- 9 Click **Next**.
- 10 In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 11 Click Next.
- 12 In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click Add Notifier and enter an email address.
- 13 (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.
- 14 (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Professional Services has integrated SA with your change control systems. It should be left blank otherwise.
- 15 Click Next.
- 16 In the Job Status page, click Start Job to run the audit. When the audit has run, click View Results to view the results of the audit.

Remediate by Server

You can remediate specific differences found in rules in audit results by the server that the audit targets. You can select to remediate all rules on all servers, or, for all rules on selected servers.

To remediate specific differences found in an audit results by server, perform the following steps:

- 1 From the Navigation pane select **Library ► By Type ► Audit and Remediation ► Audits**.
- 2 Select an audit.
- 3 From the Details pane below the audit list, you see all audit results associated with the audit.
- 4 Select an audit result, right-click, and select **Open**.
- 5 In the Audit Result window, expand the Summary list.
- 6 In the Contents pane, you can see the list of servers targeted by the audit. For each server you want to audit, select the check box next to the server, and then click Run Partial Audit

Or

You can and then expand the list of servers in the Views pane, and for each server, you see all differences discovered on all servers targeted by the audit.

For each server you want to remediate, select the check mark in the list in the All Rules column, which means that when you remediate the audit results, all rules will be remediated on the selected servers.

Or

If you want to globally select all servers in the audit results, right-click and select **Select All**. To deselect all servers, right-click and select **Deselect All**.

You can also

- 7 When you have selected the servers you want to remediate, from the **Actions** menu, select **Remediate**.
- 8 In the Remediate Audit window, step one shows you the name of the audit, the target of the Audit, the total number of rules defined in the audit.
- 9 Click **Next**.
- 10 In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 11 Click Next.
- 12 In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click Add Notifier and enter an email address.
- 13 (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.
- 14 (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Professional Services has integrated SA with your change control systems. It should be left blank otherwise.
- 15 Click Next.

- 16 In the Job Status page, click Start Job to run the audit. When the audit has run, click View Results to view the results of the audit.

Remediating Comparison-Based Audit Results

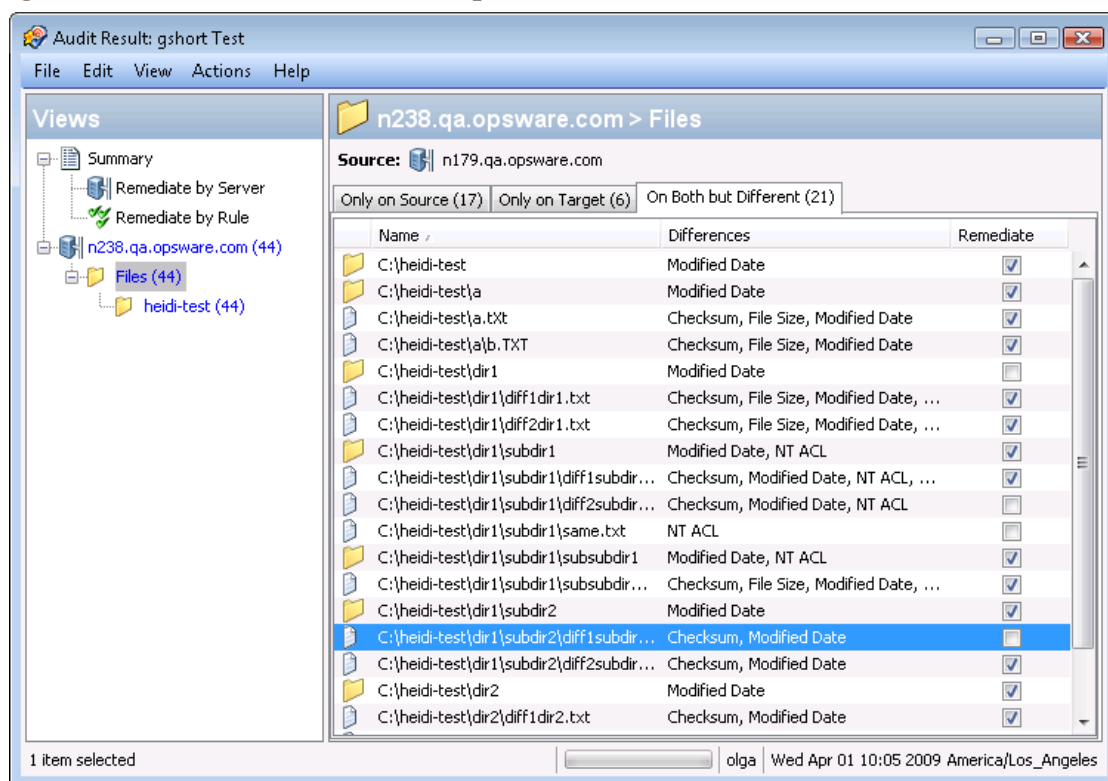
Audit results based on a comparison-based audit allow you to view differences between the source server (or snapshot) and target servers or snapshot. If the audit results fails — that is, it finds differences between source and target — you can remediate the differences (for most rule types). You can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

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

For example, [Figure 27](#) shows audit results for a windows file system rule, where the selected file and path exist on both the source (audit rule source server) and the target, but are different, located under the Only Both But Different tab of the Audit Results Window.

From the Audit Results Window, you can select the Files rule, and from the **Actions** menu select **Remediate**.

Figure 27 Audit Results For a Comparison-Based Audit Rule



In this example where file difference were found between the source and the target, you can double click the rule to view those differences in a separate window, to make sure you want to perform the remediation. Then, you can select **Remediate** from the **Actions** menu and remediate the out of compliance rule — or, schedule the audit to run at a later time. When you remediate, the values from the audit (derived from the source) will replaces those on the target server.



When remediating COM+ objects from snapshot or audit results, the SA Client does not check the version of the COM+ object, and thus will always remediate the object, whether or not there is any difference between them.

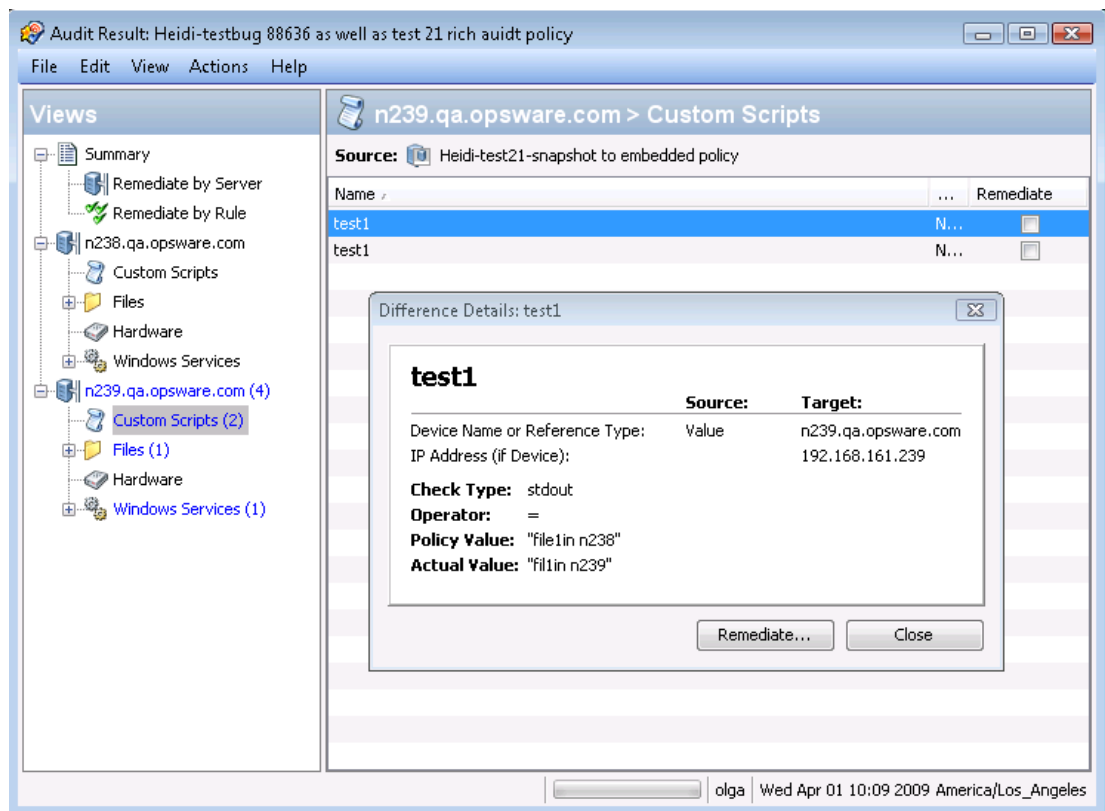
Viewing Value-Based Audit Results – Audit Rule Remediation

Value-based audit results indicates if the server configuration matches the values defined in the audit rule. You can view the differences between what was defined as the expected value in the rule and the actual value found on the target server. Depending on the rule, you can remediate the difference found on the target server by replacing it with the value specified in the rule.

Some value-based rules are not remediable. For example, Windows/Unix users and groups, the Property value check is not remediable.

Figure 28 shows a value-based audit rule in the form of a custom script where the output of the script was different than the results of the same script run on the source server. The Status column for the rule indicates Non-Compliant, which means the output of the script rule is different between the source and the target. To fix the discrepancy, select the Remediate option and select **Remediate** from the **Actions** menu. Or, double-click the rule and click the **Remediate** button.

Figure 28 Audit Results for a Value-Based Audit Rule



Viewing and Remediating Audit Results Differences

For some objects in an audit result, you can view those differences between object that exist on both the target and the source and that have differences between them. You can also see what is different about them and remediate them, if necessary.

For some audit rules, you can view general differences, such as a service's status, the release number for a patch, a registry key's value, and so on. For other server objects, such as files, you can view the differences of the file's contents.

Viewing and Remediating File Differences

For some rules, such as file system, you can view differences between files side by side and line by line. You can see lines that were added, deleted, or modified.

To view and remediate contents of two files that differ in an audit, perform the following steps:

- 1 From the Navigation pane select **Library ► By Type ► Audit and Remediation ► Audits**.
- 2 Select an audit.
- 3 From the Details pane below the audit list, you see all audit results associated with the selected audit.
- 4 Select an audit result, right-click, and select **Open**.
- 5 In the Views pane of the Audit Result window, expand one of the target servers and select a result.
- 6 In the Content pane, expand a target server and select one of the results.
- 7 Next, in the Content pane, select the On Both but Different tab.
- 8 Select a file, right-click, and select View Differences.
- 9 In the Comparison window, select an item from the Encoding drop-down list to specify the character encoding of the data displayed.



If the file in question exceeds 2MB in file size, Audit and Remediation cannot display the file differences.

- 10 Click the arrows to find the first, next, previous, or last lines that were added, deleted, or modified. Differences are highlighted according to the following color scheme:
 - **Green:** This content was added.
 - **Blue:** This content was modified.
 - **Red:** This content was deleted.
 - **Black:** No changes were made to this content.
- 11 Click **Close** to close this window.
- 12 To remediate file differences, from inside the Audit Results window, select either the the Only On Source tab or On Both But Different tab, select a file, right-click and select **Remediate**.
- 13 In the Select Server window, select a server you want to copy the file from the source to, and then click **OK**.

Viewing and Remediating Object Differences

For many server objects, such as Users and Groups, IIS Metabase, Windows Registry, and so on, when there are differences between the source object and the target object, you can view differences in object properties side by side. Each server object will show different windows, depending on the object and if the audit rule set was comparison-based (comparison between source and target) or value-based (comparison between user-defined audit rule and target).

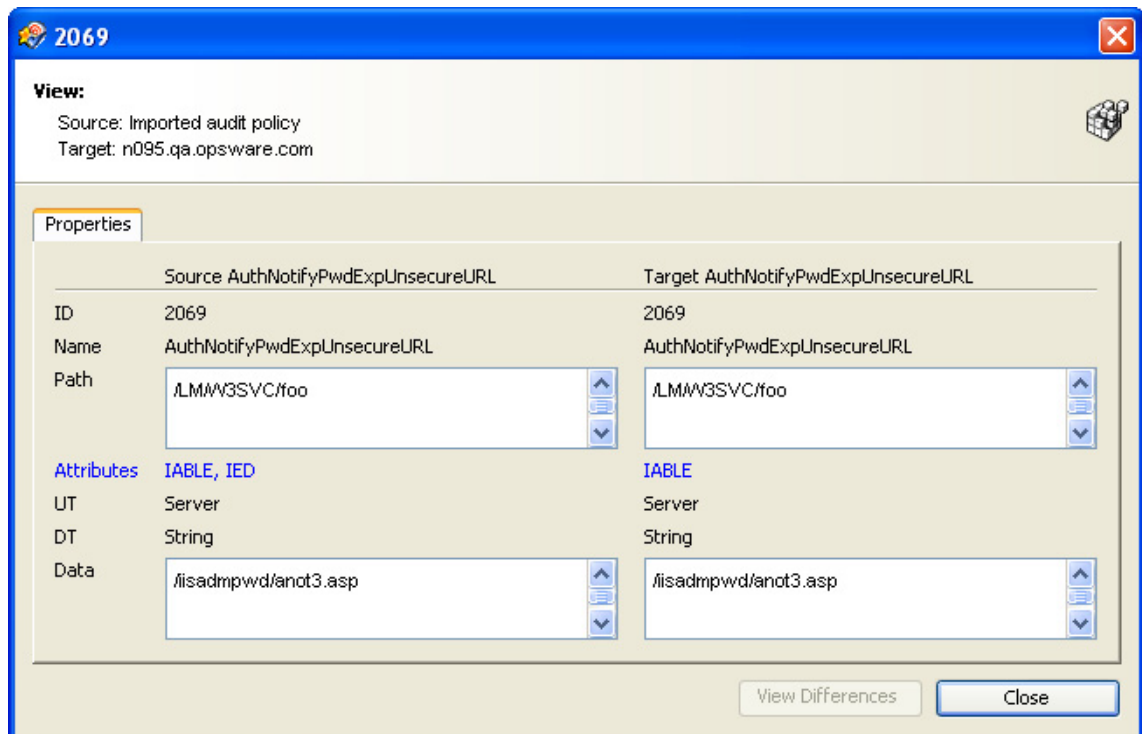
For some value-based audit rules, you can remediate the values on the target server.

To view the contents of two objects that differ, perform the following steps:

- 1 From the Navigation pane select **Library ► By Type ► Audit and Remediation ► Audits**.
- 2 Select an audit.
- 3 From the Details pane below the audit list, you see all audit results associated with the selected audit.
- 4 Select an audit result, right-click, and select **Open**.
- 5 In the Views pane, expand one of the target servers and select a result.
- 6 In the Views pane, select an object.
- 7 In the Content pane, select the On Both but Different tab.
- 8 In the Content pane, select an object, right-click, and select **Open**. You will see a window that shows the differences between the object as defined the audit and the object on the target server.

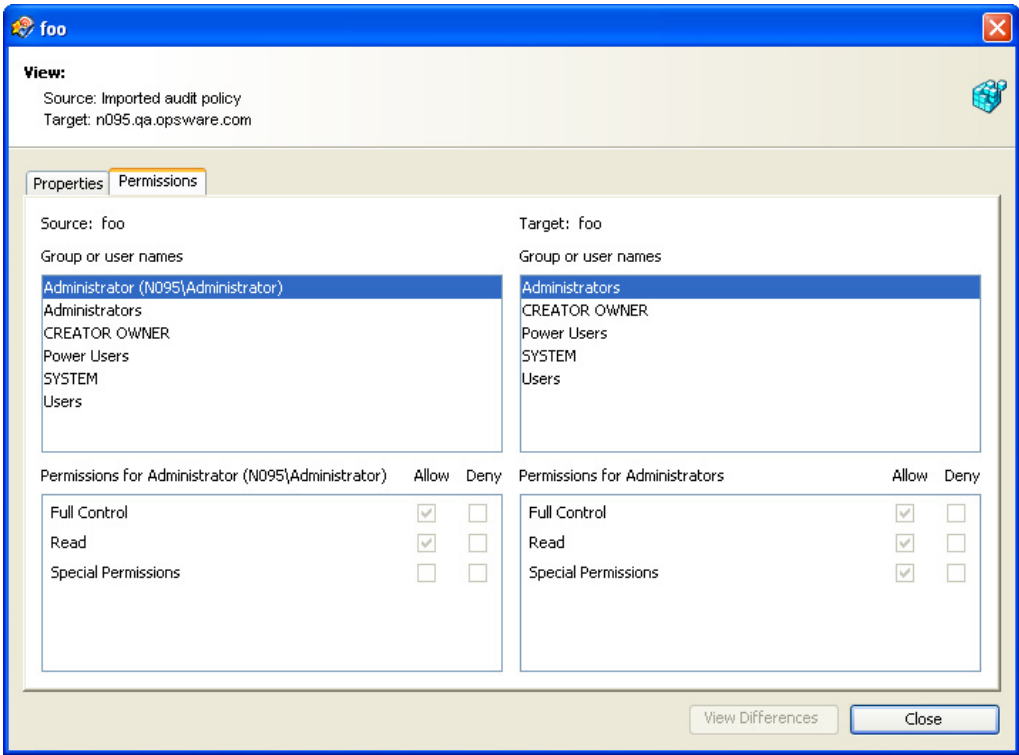
The example in [Figure 29](#) displays the audit Result differences for two IIS Metabase objects, showing an attribute of the object that exists on the server but does not exist on the source server, displayed in blue font.

Figure 29 Comparison-Based Audit Results Difference: IIS Metabase Objects



For a value-based rule, the difference window will be slightly different and will also include a Remediate option, if remediation is possible. This difference window displays the audit rule, including the policy value and the actual value found on the target server. The example in [Figure 30](#) shows the permissions differences for a value-based Windows Registry rule.

Figure 30 Rule-Based Audit Results Difference: Windows Registry Permissions Differences




- 9 To remediate the differences, select the Remediate check mark next to each rule.
- 10 From the **Actions** menu select **Remediate**.
- 11 In the Remediate window, follow the steps to run or schedule the remediation. For more information on remediating audit results, see [Viewing and Remediating Audit Results Differences](#) on page 185.

Viewing Audit Results with Exceptions

If an audit contains rule exceptions, then the excepted rules are not checked on the target servers when the audit is run. However, your audit results will show which of the rules in the audits are exceptions, including details about the rule exceptions.

The manner in which rule exceptions are displayed in audit results depends on the type of rule that has been excepted:

- Custom script and custom or pluggable check rule exceptions (such as those created by developers or provided by a EP Content Subscription) appear in the Contents pane of the Audit Results window. You can double-click the rule exception for details on the exception.

- All other rule exceptions, such as file system, registry settings, services, IIS Metabase, and COM+ rules, the Audit Results window will display an Exceptions icon  in the Views pane, which you can select and see the details of the exception in the Contents pane.

Searching for Audits

You can use the SA Client Search tool to find audits in your facility. You can search for audits by name, by the operating system, and many other criteria.

To search for audits, perform the following steps:

- 1 From inside the SA Client, ensure that the search pane is activated by selecting View ► Search pane.
- 2 From the top drop-down list, select Audit.
- 3 Click the green arrow button or ENTER to execute the search.
- 4 The results appear in the Content pane.

If you want to extend your search criteria, add new criteria in the search parameters section at the top of the Content pane. You can also save the search by clicking **Save**, or export the Search results to .html or .csv.

Deleting Audits

To conserve disk space, you can delete audits that you no longer need. You can choose to archive all audit results generated from the audit, if you would like to keep a record of the results.

To delete an audit, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Audit and Remediation ► Audits**.
- 2 Choosing either Windows or Unix, select one or more audits and then select **Actions ► Delete**.
- 3 In the Confirmation Dialog, click **Yes** to delete this audit, or click **No** if you do not want to delete it. You can also select the Archive Audits option, which will archive all audit results generated from the audit. If you do not select the Archive option, all audit results from the selected audit will be deleted.



When you delete an audit, all schedules associated with it will be also deleted. See [Scheduling an Audit](#) on page 174 in this chapter for more information.

Deleting Audit Results

As a best practice, you should delete audit results that you no longer need. If you would like to save audit results, you can choose to keep them.



You must have read permissions for the snapshot to be able to delete it. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

To delete a snapshot, perform the following steps:

- 1 Select a snapshot or select multiple snapshots and then select **Actions** ► **Delete**.
- 2 In the Confirmation Dialog, click **Yes** to delete this snapshot or click **No** if you do not want to delete it.
- 3 If you want to archive the snapshot instead of delete it, select the snapshot, right-click, and select **Archive**.



When you delete a snapshot, you do not delete the snapshot specification that was used to create it. See [Deleting a Snapshot Specification](#) on page 203 in this chapter for more information.

Archiving Audit Results

Some audits yield numerous results, especially those audits scheduled to run on a recurring basis. You can archive all audit results to keep a record of all audit results run from an audit. When you archive an audit result, it removes its connection to the original audit, but the results and targets of the audit are kept intact.

To archive audit results, perform the following steps:

- 1 From the Navigation pane select **Library** ► **By Type** ► **Audit and Remediation** ► **Audits**.
- 2 Select an audit.
- 3 From the Details pane below the audit list, you see all audit results associated with the selected audit.
- 4 To archive an audit result, select it, right-click, and select **Archive**.
- 5 You are asked to confirm if you want to archive the audit result, since doing so will remove the link between the result and the audit. Click **Yes** to archive the audit result.
- 6 To view all archived audit results, From the Navigation pane select **Library** ► **By Type** ► **Audit and Remediation** ► **Archived Audit Results**.

Snapshots

A snapshot captures the configuration of a managed server at a particular point in time, and provides a means of capturing the current state of a known working (or, not working) server. A snapshot is useful for capturing a server configuration that you know represents a desired state of configuration. You can also compare the snapshot with other servers in your facility by using the snapshot in an audit.

A snapshot is also a useful way to back up a managed server, especially if you plan to make changes to the server and want to keep a record of it before you change anything.

In addition to recording information about objects on managed servers, a snapshot can contain the content of some objects. A server snapshot also identifies attributes of other objects on specific types of operating systems, such as the Windows Registry and Windows Services, application configurations, COM+ objects, hardware information, installed patches, and more. You can even create custom scripts that gather data from the target managed servers.



VMware ESXi servers cannot be the source or the target of a snapshot.

Snapshot Specification and Snapshot

Snapshots are configured in similar way as you configure an audit. First you create a *snapshot specification*, which is like a template that defines exactly what you want to capture of a server's configuration. Then, you configure the snapshot specification's rules, and then run it. The results are a snapshot — a picture of a server's configuration. The main difference between a snapshot and an audit is that a snapshot takes a picture of a server's configuration, whereas an audit compares a server configuration with the rule values that you define.

You can schedule when you want a snapshot to be created (either once or as a recurring job) and who you want to receive email notification about the status of the job.

Snapshot Used in an Audit

You can use a snapshot in an audit to compare managed servers, groups of servers, and snapshots. By using a snapshot in an audit, you can compare a problematic server (target of the audit) with a known working server (snapshot as source for the audit). To further extend the audit definition, you can also define rules for server objects.

When a snapshot is used as the source for an audit, all server configuration values captured in the snapshot results are available to use as rules for the audit. For more information about using a snapshot in an audit, see [Configuring an Audit](#) on page 117.

Snapshot Specification Used in an Audit

You can use a snapshot specification as the source of an audit if you want to keep track of a server's configuration over time and monitor any changes that occur. For example, you might want to keep track of a specific application to make sure that its configuration remains correct over a period of time. If this application runs on several servers, you can create a snapshot specification that defines a desired state of server configuration, and then run the snapshot.

Next, you can create an audit and use the original snapshot specification as the source for your audit. Each server that was targeted by the snapshot are now also included as targets of the audit. Next, when you run the audit (either on-demand or on a scheduled basis), each server's current configuration will be compared with the state originally captured when you took the initial snapshot. Any changes are displayed in the audit results window.

For more information, [Configuring an Audit](#) on page 117.

Audit Policies and Snapshot Specifications

An audit policy is collection of rules that defines a desired state of a server's configuration. An audit policy can be used inside a snapshot specification, either through linking or importing. An audit policy is useful because it allows a policy setter to define server configuration compliance values, which then can be used by others in their snapshot specifications.

Because an audit policy can be linked to an audit or snapshot specification, whenever a change is made to the policy, the audit or snapshot specification using the policy will also reflect the latest changes. Or, an audit policy can be imported into a snapshot specification, without keeping the link to the source audit policy. When you import an audit policy into a snapshot specification, you can choose to replace any current values in the audit or merge values from the audit policy with those in the snapshot specification.

For more information on importing or linking an audit policy to a snapshot specification, see [Linking and Importing Audit Policies](#) on page 168.

Snapshot Specification Elements

An snapshot specification consists of the following elements:

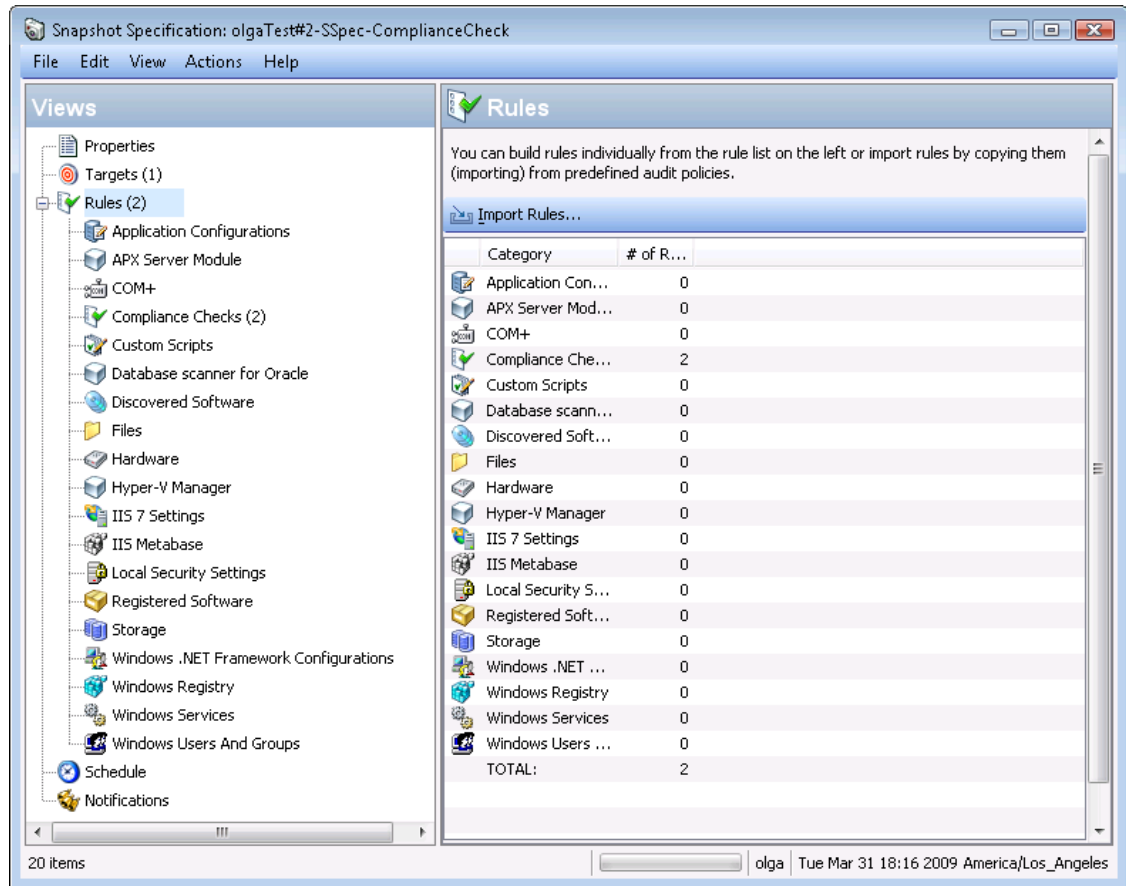
- **Properties:** The name and description of the snapshot specification. If you want to create an inventory of some snapshot specification rules, you can select the Perform Inventory and the snapshot result will collect all information about the specific rules from the target servers. This option applies to the following rules: Discovered Software, Internet Information Server, Local Security Settings, Registered Software, Runtime State, Windows and Unix Users and Groups.
- **Targets:** The servers that you want to take a snapshot of — that is, capture the specific server configuration as defined in the snapshot specification's rules. You can choose as many servers and groups of servers as you want.
- **Source:** The source of a snapshot specification. If you choose a server then you can select server objects from that server as the basis of your snapshot. The source of a snapshot specification can be a server, or no source at all. (Some rules require a source server. Other rules can be defined by your own custom values without a source.)
- Note that the value of a source parameter is not used when taking a snapshot. It only has meaning when defining a snapshot specification.
- **Rules:** A check on a particular server object with a desired value and an optional remediation value. For example, you might check if a server contains a specific Windows Service, and if found, determine if the service is turned off. For a description of server objects that you can define rules for in a snapshot specification, see [Audit and Remediation Rules](#) on page 123.
- **Schedule:** The time the snapshot will run. You can run the snapshot specification as a job on a onetime basis, or on a recurring schedule.

- **Notifications:** The email notification send after the snapshot has run. You can base the notification on success, failure, or simply the completion of the snapshot specification job.

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

Figure 31 shows a snapshot specification that has three rules that will capture configuration information about the target server for event logging, operating system, and windows services.

Figure 31 Snapshot Specification Elements



The Snapshot Process

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

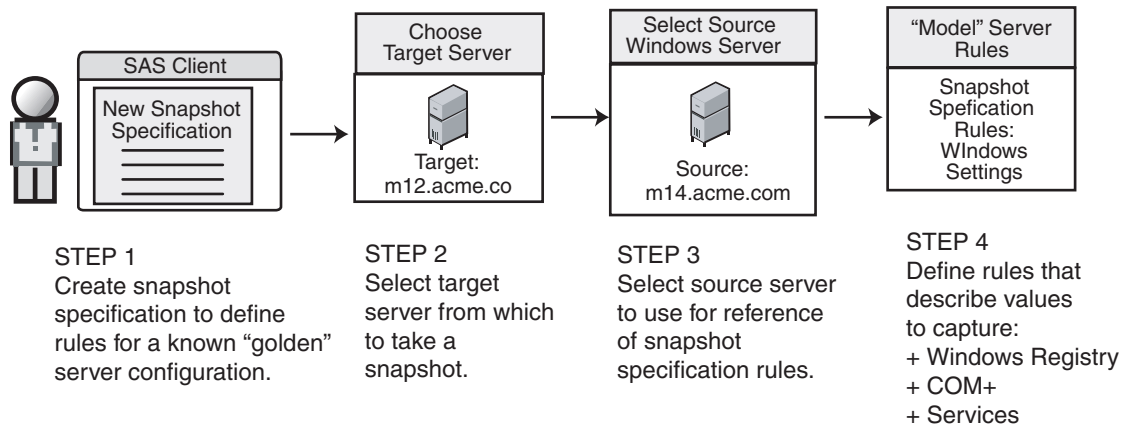
- Creating a snapshot specification, which is a template that defines the configuration parameters captured on a target server.
- Running the snapshot specification job that results in a snapshot.

Figure 32 illustrates an example of the snapshot process.

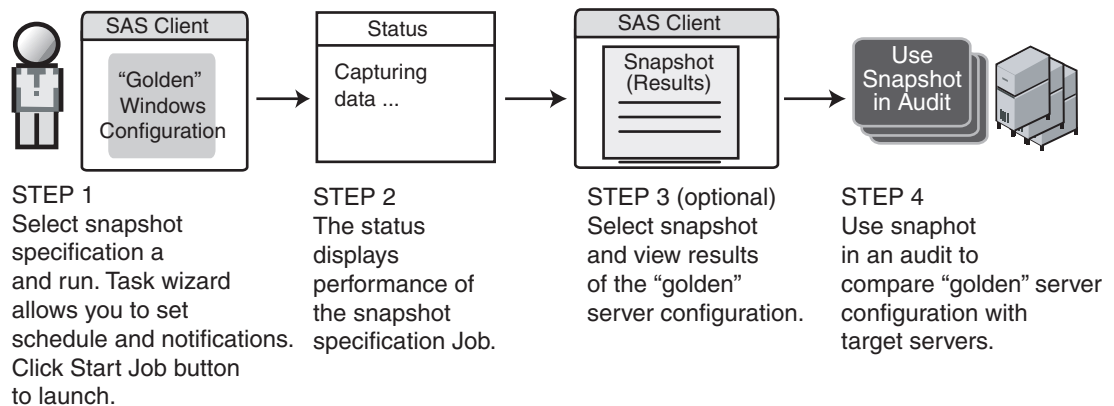
Figure 32 Snapshot Process

SNAPSHOT PROCESS - Windows Server Snapshot

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



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



Creating a Snapshot Specification

You can create a snapshot specification from two different locations inside the SA Client, depending on your purpose. You can create a snapshot specification from the following locations inside the SA Client:

- [Creating a Snapshot Specification from a Server](#)
- [Creating a Snapshot Specification from the Library](#)



You must have a set of permissions to create and modify snapshot specifications. To obtain these permissions, contact your SA administrator. See the *SA Policy Setter’s Guide* for more information.

Creating a Snapshot Specification from a Server

When you create a new snapshot specification from a managed server, the snapshot specification will use the selected server as its source. You can choose several different server sources for the snapshot specification as you define the rules, or choose no source at all and define your own custom rules. Some rules, however, require a source.



To take a snapshot of a managed server, the server must be reachable and you must have access to the server.

To create a snapshot specification from a server, perform the following steps:

- 1 From the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2 Select a server, then select **Actions > Create Snapshot Specification**.

Creating a Snapshot Specification from the Library

If you want to create a new snapshot specification and set all your own rules, create the audit from the SA Client Library by performing the following steps:

- 1 From the Navigation pane, select **Library > By Type > Audit and Remediation**.
- 2 In the Navigation pane, select snapshot specifications, then Windows or Unix.

Configuring a Snapshot Specification

To configure a snapshot specification, performing the following tasks:

- Name and describe the snapshot specification, and decide if you want to perform an inventory.
- Choose target servers you want to take a snapshot of. You can choose to snapshot multiple servers or groups of servers.
- Configure your own custom rules, or choose settings from a source server to serve as the basis for the snapshot specification rules.
- Schedule the snapshot specification job to run once or on a recurring schedule.
- Set up email notifications to notify users when the snapshot specification job finishes successfully, if the job fails, or on both conditions.
- Save the snapshot specification.



If you take a snapshot of COM+ objects from a 32 bit Windows server, and you attempt to remediate the results using copy to onto a Windows 64 bit server, it may not work

Configuring a Snapshot Specification

To configure a snapshot specification, perform the following steps:



VMware ESXi servers cannot be the target of an audit or snapshot.

- 1 From the Navigation pane, select **Library ► By Type ► Audit and Remediation**.
- 2 In the Navigation pane, select Snapshot Specifications, then either Windows or Unix.
- 3 From the **Actions** menu, select **New**.
- 4 In the Snapshot Specification window, enter the following information:
 - **Properties:** Enter a name and description for the snapshot specification. Also, for certain snapshot specification rules (Discovered Software, Internet Information Server, Local Security Settings, Packages and Patches, Runtime State, Windows and Unix Users and Groups), you can select the Perform Inventory option, which will capture all resources associated with the rule.
 - **Source:** Select a source for the snapshot specification. By default, the source server for the snapshot specification will be the managed server that you chose as the source for the snapshot specification. Browse the source server for values to populate the snapshot specification's rules. You can also choose a different source server as the basis of the snapshot specification for each rule category, or no source at all. If you choose no source, you must define your own rules, or choose to link to an audit policy in the rules section.
 - **Rules:** Choose a rule category from the list to begin configuring your snapshot specification's rules. Since each rule is unique and requires its own instructions, to configure specific rules, see [Audit and Remediation Rules](#) on page 123.

If you want to use an audit policy to define the rules of your snapshot specification, click either **Link Policy** or **Import Policy**. When you link an audit policy, the snapshot specification maintains a direct connection with the audit policy, so if any changes are made to the policy, the snapshot specification will update it with the new changes. If you import an audit policy, the snapshot specification will use all the rules defined in the policy but will not maintain a link to the audit policy. For information on how to import or link to a snapshot specification, see [Linking and Importing Audit Policies](#) on page 168.

- **Targets:** Choose the Targets of the snapshot specification. These are servers or groups of servers that you want the configured snapshot specification rules to capture. To add a server or group of servers, click **Add**. To choose a source server to use to create the snapshot specification rules, click **Select**.
- **Schedule:** Choose to run the snapshot specification immediately, or on a recurring schedule. Choose whether you want to run it once, daily, weekly, monthly, or on a custom schedule. Parameters include:
 - **None:** No schedule will be set. To run the snapshot specification, select the snapshot specification, right-click, and select **Run snapshot specification**.
 - **Daily:** Choose this option to run the snapshot specification on a daily basis.
 - **Weekly:** Choose a day of the week to run the snapshot specification.
 - **Monthly:** Choose the months to run the snapshot specification.

- **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

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

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.

- **Time and Duration:** For each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the snapshot specification will keep running indefinitely. To choose an end date to end the snapshot specification schedule, select **End**, and from the calendar selector, choose a date. The Time Zone is set according to the time zone set in your user profile.
- **Notifications:** Enter the email addresses (separated by a comma or a space) of those you want to receive an email when the snapshot specification Job finishes running. You can choose to send the email notification on both success and the failure of the snapshot specification job (not the success of the audit rules). To add an email address, click **Add Notification Rule**.

- 5 When you have finished configuring the snapshot specification, from the **File** menu, select **Save**.



To prevent runaway processes, the snapshot process will time-out if it exceeds 60 minutes or if the data that is collected from a managed server exceeds one gigabyte (GB). If you specify that you want to collect the full contents of files in the selection criteria, the data collected might exceed the maximum size that can be successfully recorded in a snapshot.

Configuring Snapshot Specification Rules

For information on how to configure specific snapshot specification rules, see [Audit and Remediation Rules](#) on page 123.

Saving a Snapshot Specification as an Audit Policy

You can save selection criteria used in a snapshot and save it as an audit policy. This can be useful if you would like to use the rules configured in a snapshot specification for other snapshot specifications or audits.



In order to save a snapshot specification or an audit as a policy, the policy must be saved to a SA Client Library folder, and your user must have write permissions to the folder you want to save to. For more information on permissions, see the *SA Administration Guide*.

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



- 1 Launch the SA Client. From the Navigation pane, select **Library > By Type > Audit and Remediation**.
- 2 Select Snapshots Specification, and then double click a snapshot specification you want to save as an audit policy.
- 3 In the Snapshot Specification window, select **File > Save As**.

- 4 In the Save As window, enter a name and (Optional) description.
- 5 From the Type drop-down list, select Audit Policy.
- 6 Click **Save**. The selected snapshot specification has been saved as an audit policy. To view the audit policy, from the Navigation pane, select **Library ► By Type ► Audit and Remediation ► Audit Policies**. For more information about using audit policies, see [Audit Policies](#) on page 166.

Running a Snapshot Specification

When you run a snapshot specification, it captures from the target servers all configuration parameters configured in the rules. After you run a snapshot specification, the results of the snapshot job become a snapshot and can be viewed inside the snapshot.

To run a snapshot specification, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Audit and Remediation**.
- 2 In the Navigation pane, select Snapshot Specifications then either Windows or Unix.
- 3 Select a snapshot specification, right-click, and select **Run**.
- 4 In the Run Snapshot Specification window, step one shows you the name of the snapshot, the total number of rules defined, and all targets). Click **View Rules Details** to view the rule definitions.
- 5 Click **Next**.
- 6 In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select the second option and choose a day and time.
- 7 Click **Next**.
- 8 In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 9 (Optional) You can specify if you want the email to be sent on success of the audit job () or failure of the audit job ()
- 10 (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when SA Professional Services has integrated SA with your change control systems. It should be left blank otherwise.
- 11 Click **Next**.
- 12 In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

Scheduling Snapshot Jobs

A snapshot specification job enables you to specify when you want the SA Client to create a snapshot (either once or on a recurring basis) and who you want to receive email notification about the status of the job. You can also view, edit, and delete existing snapshot specification schedules. When you delete a snapshot specification, all schedules associated with that snapshot specification will be deleted.

This section discusses the following topics:

- [Scheduling a Recurring Snapshot Job](#)
- [Viewing and Editing a Snapshot Job Schedule](#)
- [Viewing and Editing a Snapshot Job Schedule](#)
- [Deleting a Snapshot Job Schedule](#)

Scheduling a Recurring Snapshot Job

After you have created, configured, and saved an snapshot specification, you can schedule snapshot specification a recurring snapshot job. After the schedule is set, you can edit the schedule according to your needs.

To schedule a recurring snapshot specification, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Audit and Remediation ► Snapshot Specifications**.
- 2 Select either Windows or Unix, and then double-click a snapshot specification to open it.
- 3 From the Snapshot Specification window Views pane, select Schedule.
- 4 In the Schedule section, choose to run the snapshot job immediately or on a recurring schedule. Choose to run it once, daily, weekly, monthly, or on a custom schedule:
 - **None:** No schedule will be set. To run the snapshot job, select the snapshot specification, right-click, and select **Run Audit**.
 - **Daily:** Choose to run the snapshot job on a daily basis.
 - **Weekly:** Choose a day of the week to run the snapshot specification job.
 - **Monthly:** Choose the months to run the snapshot specification job.
 - **Custom:** In the Custom Crontab string field, enter a string the indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

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

An asterisk (*) in any of these fields represent all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.

- In the Time and Duration section, for each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the snapshot specification job will keep running indefinitely. To choose an end date to end the audit schedule, select End, and then choose an end date. The Time Zone is set according to the time zone set in your user profile.

- (Optional) Deselect the End option if you want the snapshot specification job to run indefinitely.
- 5 To save the snapshot specification job schedule, from the **File** menu select **Save**. The snapshot specification will now run according to the defined schedule.

Viewing and Editing a Snapshot Job Schedule

You can edit a snapshot specification schedule after you have created (or edited) and saved it.

To edit a scheduled snapshot specification, perform the following steps:

- 1 From the Navigation pane, select Jobs and Sessions.
- 2 Select Recurring Schedules.
- 3 From the drop-down list at the top of the Contents pane, select Create Snapshot. The list shows all scheduled snapshot specification jobs.
- 4 To view a scheduled snapshot specification, double-click one.
- 5 Select the Schedule object in the Views pane.
- 6 To edit the snapshot specification job schedule, modify the following parameters:
 - **Schedule:** Choose to run the snapshot specification immediately, or on a recurring schedule. Choose to run it once, daily, weekly, monthly, or on a custom schedule. Parameters include:
 - **None:** No schedule will be set. To run the snapshot specification, select the snapshot specification, right-click, and select **Run snapshot specification**.
 - **Daily:** Choose to run the snapshot job on a daily basis.
 - **Weekly:** Choose the day of the week you want the snapshot job to run.
 - **Monthly:** Choose the months to run snapshot specification job.
 - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

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

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For more information about crontab entry formats, consult the Unix man pages.

- **Time and Duration:** For each type of schedule, specify the hour and minute, the day of the week (and month) you want the daily schedule to start. Unless you specify an end time, the snapshot specification job will keep running indefinitely. To choose a date to end the snapshot specification job schedule, select End and then choose a date. The Time Zone is set according to the time zone set in your user profile.
 - (Optional) Deselect the End option if you want the snapshot specification schedule to run indefinitely.
- 7 To save the snapshot specification schedule, from the **File** menu select **Save**. The snapshot job will now run according to the defined schedule.

Deleting a Snapshot Job Schedule

To delete a snapshot job schedule, perform the following steps:

- 1 From the Navigation pane, select **Jobs and Sessions**.
- 2 Select **Recurring Schedules**.
- 3 From the drop-down list at the top of the Contents pane, select **Create Snapshot**.
- 4 The Content pane displays all snapshot specification jobs that have been run on this SA core. To display only snapshot specification jobs, from the drop-down list at the top of the Content pane, select **Run Snapshot Task**. If you want to see only those snapshot specifications that you have scheduled or run, enter your user ID in the User ID field at the top of the Content pane.
- 5 To delete the schedule, select it, right-click, and select **Delete Schedule**.

Locating Snapshots

After you have created a snapshot, you can find it in several locations inside the SA Client.

Locating Snapshots In the Library

- 1 From the Navigation pane, select **Library > By Type > Audit and Remediation > Snapshot Specifications**.
- 2 Select either **Windows** or **Unix**.
- 3 From the list, select a snapshot specification. The Details pane at the bottom of the application window displays all snapshots run from the selected snapshot specification.

Locating Snapshots in the Device Explorer

To locate snapshots associated with a specific server, you can view them in server's Device Explorer by performing the following steps,

- 1 From the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2 Select a server from the list, right-click, and select **Open**.
- 3 In the Device Explorer window, select **Inventory > Snapshot Specification**.
- 4 In the Content pane, select a snapshot specification and all associated snapshots appear in the Details pane at the bottom of the window.
- 5 To view a snapshot, double-click it to open.

Searching for Snapshots

You can use the SA Client Search tool to find snapshots in your facility. You can search for snapshots by name, by the operating system, and many other criteria.

To search for snapshots, perform the following steps:

- 1 From inside the SA Client, ensure that the search pane is activated by selecting **View > Search Pane**.

- 2 From the top drop down list, select Snapshot.
- 3 Click the green arrow button or ENTER to execute the search. The results appear in the Content pane. If you want to extend your search criteria, you can add new criteria in the search parameters section at the top of the Content pane. You can also save the search by clicking **Save**, or export the Search results to .html or .csv.

Viewing Snapshot Results

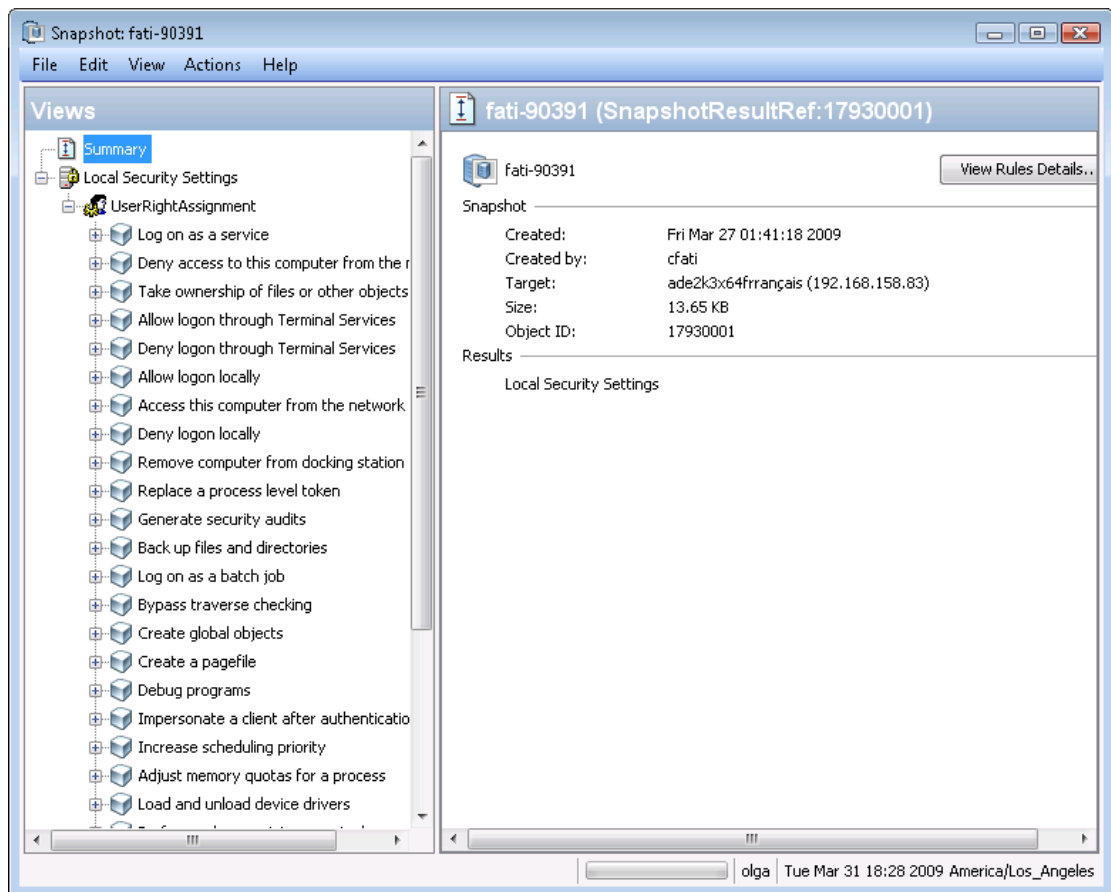
You can view the contents of a snapshot and view detailed information about the server configurations that were recorded.

For information about remediating snapshot results, see [Copying Objects from a Snapshot to a Server](#) on page 204.

To view the contents of a snapshot, perform the following steps:

- 1 From one of the starting points described in [Locating Snapshots](#) on page 200, open a snapshot.

Figure 33 Sample Snapshot of a Windows Server



- 2 In the snapshot window, you can select:
 - **Summary:** Displays general information about a snapshot, such as the date and time the snapshot was created and by whom, the snapshot source (name of the managed server), the size of the snapshot file, and a snapshot ID number.

You can also click **View Rules Details** to see the snapshot specification which this snapshot is based on.

- **Compliance Library:** Information relevant to the specific compliance checks configured in the snapshot specification. For more information about the types of BSA Essentials Subscription Services compliance checks available and how to configure them, see [Configuring Compliance Checks](#) on page 151.
- **Installed Hardware:** Information about the type of CPU processor and speed, cache size, memory size for SWAP and RAM, and storage devices that were recorded in the snapshot.
- **Installed Patches:** Displays information about the installed patches that were recorded in the snapshot, such as the patch type.
- **Installed Packages:** Displays information about the installed packages that were recorded in the snapshot, such as package type, package version, and release number.
- For .zip packages, the Snapshots do not show a version number, but instead displays the install path of the package on the server.
- **Event Logging:** Displays security, application, and system log files recorded in the snapshot.
- **File System:** Displays the directories, file properties, attributes, and contents of the files recorded in the snapshot.



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

- **Windows Services:** Displays information about the running services recorded in a snapshot, such as the name, description, startup state, startup type, and log on account.
- **Windows Registry:** Displays information about Windows Registry entries in the snapshot, such as the registry key, registry value, and subkey. A registry key is a directory that contains registry values, where registry values are similar to files within a directory. A subkey is similar to a subdirectory. The content area in this window excludes subkeys. Audit and Remediation supports the following Windows Registry keys: HKEY_CLASSES_ROOT, HKEY_CURRENT_CONFIG, HKEY_LOCAL_MACHINE, and HKEY_USERS.
- **COM+:** Displays information about Windows COM (Component Object Model) objects in the snapshot, such as the name and GUID (Globally Unique Identifier) of the object, and the path to the in-process server DLL.
- SA provides warning messages that explain how Windows COM folders were processed. The following scenarios apply:
 - When you create a snapshot and select a Windows COM folder that does not contain any objects, the snapshot window displays a summary. SA displays a warning that the GUID (Globally Unique Identifier) for that folder is invalid, which means that the Windows COM folder does not contain any objects.
 - When you create a snapshot specification and select a Windows COM+ object that does not exist on a target, SA displays a warning that the folder is invalid.
 - When you create a snapshot and select a Windows COM+ folder that does not contain any objects, SA displays a warning that the folder is empty.
- **Metabase:** Displays information about IIS Metabase objects in the snapshot, such as the ID, name, path, attributes, and data of the object.
- **Custom Scripts:** Displays information about the custom script rule recorded in the snapshot.

- **Users and Groups:** Displays information about users and groups on servers, such as user name for last login, whether or not CTRL + ALT + DELETE is enabled, and so on.
- 3 Click **Close** to close the object browser.

Archiving Snapshots

Some snapshot specification yield numerous snapshots, especially those scheduled to run on a recurring basis. You can archive all snapshots to keep a record of all snapshots run for a server or group of servers.

When you archive a snapshot, it detaches the snapshot from the server and removes its connection to the original snapshot specification.

To archive audit results, perform the following steps:

- 1 From the Navigation pane select **Library > By Type > Audit and Remediation > Audits**.
- 2 Select an audit.
- 3 From the Details pane below the audit list, you see all audit results associated with the selected audit.
- 4 To archive an audit result, select it, right-click, and select **Archive**.
- 5 You are asked to confirm if you want to archive the audit result, since doing so will remove the link between the result and the audit. Click **Yes** to archive the audit result.
- 6 To view all archived audit results, From the Navigation pane select **Library > By Type > Audit and Remediation > Archived Snapshots**.

Deleting a Snapshot Specification

To conserve disk space, you can delete snapshot specifications that you no longer need. You can choose to archive all snapshots generated from the snapshot specification, if you would like to keep a record of the results. Or, you can choose to delete the snapshot specification and all snapshots associated with it.

To delete an snapshot specification, perform the following steps:

- 1 From the Navigation pane, select **Library > By Type > Audit and Remediation > Snapshot Specifications**.
- 2 Choosing either Windows or Unix, select one or more Snapshot Specification and then select **Actions > Delete**.
- 3 In the Confirmation Dialog, click **Yes** to delete this snapshot specification, or click **No** if you do not want to delete it. You can also select the Archive Snapshots option, which will archive all snapshots generated from the snapshot. If you do not select the Archive option, all snapshots generated from the selected snapshot specification will be deleted.



When you delete a snapshot specification, all schedules associated with it will be also deleted. See [Scheduling Snapshot Jobs](#) on page 198 in this chapter for more information.

Deleting a Snapshot

As a best practice, you should delete snapshots that you no longer need from the Software Repository to conserve disk space.



You must have read permissions for the snapshot to be able to delete it. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

To delete a snapshot, perform the following steps:

- 1 Select a snapshot or select multiple snapshots and then select **Actions ► Delete**.
- 2 In the Confirmation Dialog, click **Yes** to delete this snapshot or click **No** if you do not want to delete it.
- 3 If you want to archive the snapshot instead of delete it, select the snapshot, right-click, and select **Archive**.



When you delete a snapshot, you do not delete the snapshot specification that was used to create it. See [Deleting a Snapshot Specification](#) on page 203 in this chapter for more information.

Copying Objects from a Snapshot to a Server

After viewing snapshot contents, you can copy certain objects to a target server. Audit and Remediation allows you to copy directories, files, windows services (state only), IIS Metabase objects, COM+ objects and categories, and Windows Registry keys to a managed server.



In order to copy COM+ rule snapshot results from a snapshot to a server, you must have selected the Archive all associated files option when you configured the COM+ rule. Also the COM+ object being copied must not be in use by any application in order for the copy to remediation to work. For more information, see [Configuring COM+ Rule](#) on page 131.

Before you copy these objects over to a managed server, it is important to understand what actually gets copied to or created on the destination server:

- When you select a directory, only the directory will be copied to the destination server, excluding any files in that directory. For example, if dir1 contains file1 and file2, and you select dir1, Audit and Remediation copies only dir1 (not file1 and file2) to the destination server.
- When you select a file and its parent directory does not exist on the destination server, Audit and Remediation will create the directory on and copy the files to the destination server. For example, if you select file1 and dir1 does not exist on the destination server, Audit and Remediation will create dir1 on and copy file1 to the destination server.
- When you copy a Windows Services object, you copy the state of the service, such as started, stopped, paused, and so on. You can select one or more Windows Services objects for a single copy process.
- When you copy a Windows Registry object, you can select one or more registry keys and subkeys for a single copy process.

- ACLs are not copied along with COM+ objects or Microsoft IIS objects to the target server.
- When remediating COM+ objects from snapshot results using copy to, the SA Client does not check the version of the COM+ object, and thus will always copy the object, whether or not there is any difference between them.



You must have write permission on the destination server to be able to copy an object to it. To obtain these permissions, contact your SA administrator. See the *SA Policy Setter's Guide* for more information.

Copying Objects to a Server from a Snapshot

To copy an object from a snapshot to a managed server, perform the following tasks:

- 1 From one of the starting points described in [Locating Snapshots](#) on page 200, open a snapshot.
- 2 In the Views pane, select a file system, Windows Services, or Windows Registry object.
- 3 In the Content pane, select one or more objects that you want to copy.
- 4 Select **Actions** ► **Copy To**.
- 5 In the Select Server window, select a destination server.



Use the search tool to dynamically filter this list by entering a server name, IP address, or operating system.

- 6 Click **Select** to copy the object to that managed server or click **Cancel** to close this window without saving your changes.

3 Server Compliance

Overview of Server Compliance

The SA Client Server Compliance Dashboard feature allows you to view overall compliance levels for all servers and groups of servers in your facility and enables you to remediate servers that are out of compliance. You can view compliance for an individual servers, multiple servers, groups of servers, or for all servers under SA management.

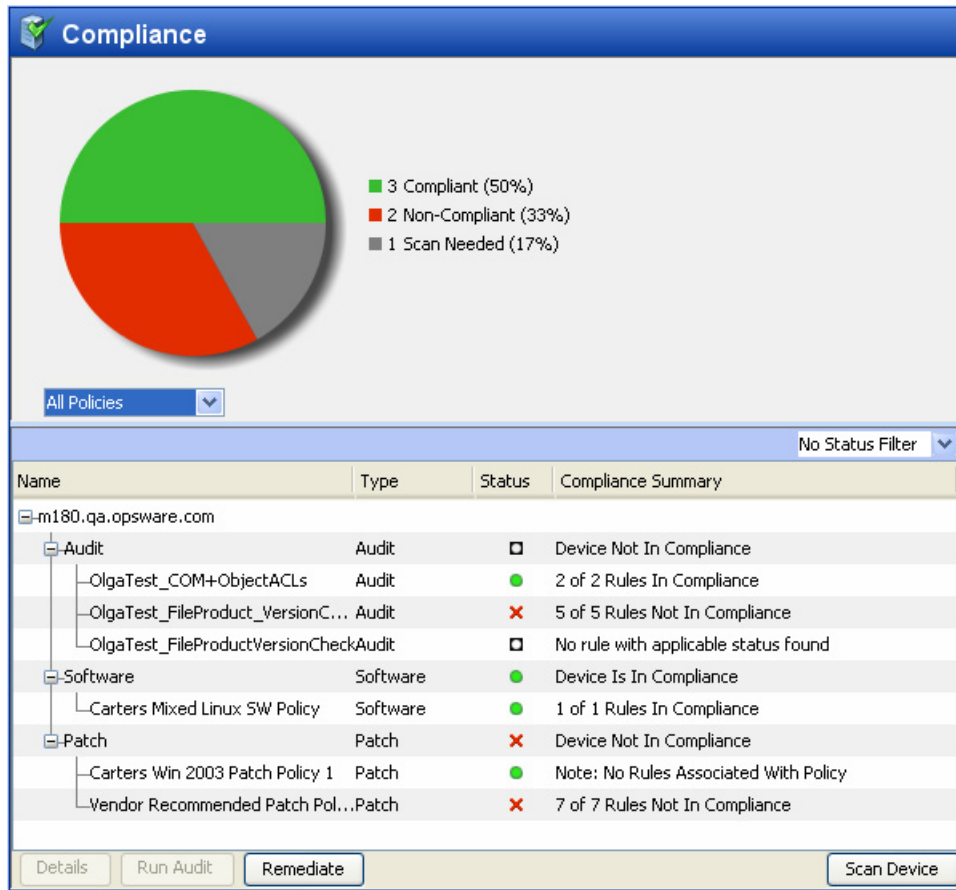
The Compliance View displays the results of all compliance statuses on servers (or groups of servers) for audits, software policies, patch policies, and application configurations. A server's compliance status is based upon a compliance *policy*, which defines unique server configuration settings or values to ensure that your IT environment is configured as it should be.

Compliance policies are created and defined by a user often known as a policy setter, though sometimes an ad-hoc policy might be created by a systems administrator. The policy setter creates compliance policies and then attaches them to servers in order to ensure that servers are compliant with the organization's standards and policies.

For example, a policy setter can create a software policy that defines a standard set of patches and packages that should be installed on a server, or the policy setter could define the manner in which certain application files should be configured on a server. A server or group of servers is considered *compliant* if its configuration matches the rules defined by the policy setter in the compliance policy.

The Compliance View in the Device Explorer allows you to determine if the server's actual installed software, packages, patches, and configuration files settings match the configuration defined in the software policy. The Compliance View in the Device Group Explorer allows you to view compliance for groups of servers, showing a compliance status rollup for all members and (sub-group members) of a group. From the Compliance View, you can discover servers and groups of servers that are out of compliance and remediate any problems.

Figure 34 Compliance View in the SA Client



The information displayed in the Compliance View is as up to date as the last time the SA Client requested compliance information from the core. By default, the SA Client checks for new compliance information every five minutes, but this time interval can be changed. For information on how to change this time interval, see [Setting Automatic Compliance Check Frequency](#) on page 228

To immediately get the latest compliance information, press Control + F5.

Compliance Dashboard Usage: Proactive and Reactive

You can proactively use the Compliance View by viewing it on a regular basis to assess server compliance levels, and to take the necessary action to fix problems.

For example, you can use the Compliance View to determine the status of an individually scheduled audit that makes sure a Web application's configuration (such as Apache's http.conf file) meets the standards set by your group. You want to ensure that no one has changed the application's configuration. To verify that no unwanted changes have been made, you should regularly check the Compliance View on this server's Device Explorer to see if this scheduled audit's compliance status has changed to red (non-compliant), and if so, view the audit results and remediate the problem.

In other situations, you can reactively use the Compliance View to answer a specific question or diagnose a specific problem. For example, you can create a scheduled audit that defines security standards for a group of servers in your facility. The audit will ensure that all

Windows 2003 servers contain a specific patch. When Microsoft releases a new security patch, you want identify the Windows 2003 servers that contain the new patch of and those that do not. You can update the audit to contain the new patch and then browse the Windows 2003 servers in the Device Group's Compliance View. When you rerun the audit, you can discover the servers that need the patch and remediate them by installing the new patch.

Compliance Terms and Concepts

- **Compliance:** The degree to which a server's actual configuration conforms to the rules defined in a compliance policy.
- **Compliance Category:** The Compliance View displays compliance statuses for four different compliance categories, including Audit, Software, Patch, and App Config (Application Configuration).
- **Compliance Policy:** The user-defined configuration that expresses the desired state for a server or device configuration or setting. For example, a patch policy defines the specific patches that should be installed on a computer. An audit policy might define that a certain Windows service should be disabled at all times. An application configuration policy defines the way in which a configuration file should be configured.
- **Compliance Rule:** The content or setting inside of a policy that defines an ideal configuration for a server, such as a patch or package, a file configuration, software installation order, user and group membership and privileges, and so on.
- **Compliance Statuses:** Indicates the compliance status for a compliance category, reporting the difference between what should be (compliance policy) and what actually is (server configuration). For example, software compliance category in the Compliance View displays a status of Compliant if all configurations defined in the policy match the server configuration. Compliance calculation for groups is slightly different than individual servers. For more information on compliance statuses, see [Compliance Dashboard Statuses](#) on page 211.
- **Compliance Scan Results:** The results of a compliance scan. These results report the compliance status, details, and can also include remediate options.
- **Compliance Scan:** The mechanism that checks servers targeted by a compliance policy (audit, software, patch, and application configuration) and returns the results to the SA Client. A compliance scan could check to see what patches are installed on a computer targeted by a patch policy or software policy and return the results, or, it can check a configuration file's contents and determine if it matches the rules defined in an application configuration. In the Compliance View, you can perform a compliance scan for the Software, Patch, and App Config compliance categories. Audits do not have a scan feature, but running an audit achieves the same results. Running an audit checks the servers targeted by the audit to determine if they are in compliance with an audit's rule definitions.
- **Compliance View:** Displays overall and individual compliance levels for all managed servers or groups of servers in your facility.

For more information on compliance scans, see [Scanning for Compliance](#) on page 229.

Server Compliance Dashboard Categories

The Compliance View for servers and groups of servers displays compliance for the following categories:

- **Audit:** Audit compliance represents an aggregate of all audits that run on a recurring schedule and indicates whether or not the rules defined in a scheduled audit match what is installed and configured on a the target server or servers.
- For more information, see [Audit Compliance](#) on page 234.

- For more information about creating and running audits, see [Audit and Remediation](#) on page 153.
- **Software:** Software compliance is determined by whether or not a software policy definition matches what is installed on a server. A software policy defines patches, packages, and application configurations, scripts as well as a host of other server objects such as services, Windows registry, COM+, IIS Metabase, and so on. A software policy can also contain other software policies.

For more information, see [Software Compliance](#) on page 237.

For more information on creating software policies, see “Software Management Setup” on page 51.

- **Patch:** Patch compliance is determined by whether or not the patch policy definition matches the patches are installed on a server or group of servers. The Compliance View displays compliance information for Windows patches only.

For more information, see [Patch Compliance](#) on page 241.

For more information on creating Windows patch policies, see [Patch Management for Windows](#) on page 287.


- **Application Configuration:** An application configuration’s compliance is determined by whether or not the application configuration definition matches the configurations on a server or group of servers. An application configuration defines the configuration settings and values for application configuration files.


For more information, see [Application Configuration Compliance](#) on page 244.

For more information on creating, configuring, and using application configurations, see the *SA Application Configuration User Guide*.

For information on how to remediate servers and groups of servers are out of compliance for each of these compliance categories, see [Compliance Dashboard Remediation](#) on page 231.

Compliance Dashboard Statuses

In general, a server or group of servers can be *Compliant* or *Non-Compliant*. A server is considered Compliant if the rules defined in the policy match the actual configuration on the server that the policy is attached to. When a server is in compliance with the policy attached to it, the Compliance View displays a Compliant icon .

If the server’s actual configuration does not match the rules configured in a policy, then the Compliance View displays a status of Non-Compliant .

For example, you can configure an audit to make sure that a Windows 2003 server has the Windows CIS recommended minimum password length of at least eight characters. When the audit runs and checks the server’s user password and discovers a user password that is only four characters, then the Device Explorer’s Compliance View shows the server’s audit policy as Non-Compliant.

If the server has more than one audit attached to it, then the Compliance View shows an aggregate, or roll up, compliance status for all audits attached to the server. If at least one of the audits targeting the server is Non-Compliant, then the overall Audit compliance status for the server is Non-Compliant.

If this server belongs to a device group of multiple servers, you can access the Compliance View for the group to see compliance status levels for all audits that run on all servers in the group, as well all servers in any sub-groups. The method used for determining compliance

statuses for groups is based upon a default calculation. The group of servers is considered Compliant if at least ninety five percent of the servers that belong to the group have a status of Compliant. If less than ninety five percent of the servers have a status of Compliant, then the status of the group is partially compliant

The default compliance status threshold for groups of servers can be customized to fit your needs. For more information, see [Changing Device Group Compliance Settings](#) on page 214.

Compliance Status Definitions

Table 13 lists default compliance statuses for a policies, servers, and device groups.

Table 13 Compliance Dashboard Compliance Status Statuses







Icon	Compliance Status Description
	Compliant <ul style="list-style-type: none"> Policy: All rules or items defined in the policy match the actual server configuration. Servers: Compliance scan ran successfully and the server configuration matches <i>all</i> of the rules defined in <i>all</i> of the policies attached to the server. Device Groups: Compliance scan ran successfully and the percentage of compliant servers is greater than the minimum threshold set in the Compliance Status section in the Opsware Administration pane. By default, this threshold for a Compliant status is ninety five percent of servers in the group, but this value can be modified.
	Partial <ul style="list-style-type: none"> Policy: One or more rules or items defined in the policy does not match the actual server configuration, due to an exception applied to one of the rules. (Windows Patch policies only.) Servers: Compliance scan ran successfully and the server configuration did not match at least one of the rules defined in any of the policies attached to the server, due to an exception applied to one of the rules. (Windows Patch policies only.) Device Groups: Compliance scan ran successfully, and a number of servers in the group meet the threshold for Non-Compliance set in the Compliance Status section in the Opsware Administration pane, while the rest of the servers in the group are Compliant. The compliance threshold definitions for Partial Compliance can be modified.
	Non-Compliant <ul style="list-style-type: none"> Policy: One or more rules or items defined in the policy does not match the actual server configuration. Servers: Compliance scan ran and the actual server configuration does not match at least one or more of the rules defined in the policy. Device Groups: Compliance scan ran and enough servers in the group meet the criteria for Non-Compliance set in the Compliance Status section in the Opsware Administration pane to indicate the group is Non-Compliant. The compliance threshold definitions for Non-Compliance can be modified

Table 13 Compliance Dashboard Compliance Status Statuses

Icon	Compliance Status Description
	Scan Failure Compliance scan was unable to run.
	Scan Needed Results undefined, perhaps because a compliance scan was never run (for example, on a new installation), or the configuration on the server (or servers in the device group) changed since the last time information was reported to the SA Client.
	Scanning: Compliance scan currently running.
—	No Tests Defined No compliance policies of this type are attached to the server or all servers in the device groups (including all servers in any sub-groups).



It is possible that actual server configurations as well as policy information might have changed from the last time you viewed compliance for a server or group in the Compliance View. To get the latest compliance data from the SA core, select **Refresh** from the **View** menu. (Or, press Control + F5.) Or, you can run a compliance scan on the server or group to determine compliance status. For more information, see [Scanning for Compliance](#) on page 229.

Compliance Status Thresholds — Policy, Server, and Group

Compliance status for a policy — an audit, a software policy, a patch policy, an application configuration — is based upon all the rules in the policy. All it takes is one of the rules in a policy to be Non-Compliant (does not match the actual configuration on the server) and the entire policy is also considered Non-Compliant for a server.

Compliance status for a server is based upon all the policies attached to the server or that define the server as a target. If any one of the compliance categories has a compliance status of Non-Compliant, then the server's overall compliance status is also considered Non-Compliant. Stated another way, all of the policies in all of the compliance categories must be Compliant for the server's overall compliance status to be Compliant.

For information on how compliance status for a server is displayed in the Device Explorer, see [Device Explorer Compliance Summary Pie Chart and Details](#) on page 217.

Device Groups Compliance Status Thresholds

Whether or not a server is considered Compliant or Non-Compliant is important when viewing device group compliance in the Compliance View, which is based upon a default threshold calculation — and which you can configure and customize.

In the Device Group Compliance View, in order for a compliance category (Audit, Software, Patch or App Config) to display a status of Non-Compliant, more than five percent of all servers in a group must have the status of Non-Compliant for that category. Another way to state Non-Compliance for a group is when less than ninety five percent of the servers are Compliant.

In the Device Group Compliance View, in order for a compliance category (Audit, Software, Patch or App Config) to display a status of Partial-Compliant, more than two percent but less than or equal to five percent of all servers in a group must have the status of Non-Compliant for that category. Another way to state Partial-Compliance for a group is when less than ninety eight percent but at least ninety five percent of the servers are Compliant.

In the Device Group Compliance View, in order for a compliance category (Audit, Software, Patch or App Config) to display a status of Compliant, less than two percent of all servers in a group must have the status of Non-Compliant for that category. Another way to state Compliance for a group is that at least ninety eight percent of the servers are Compliant.

Group status is calculated based on all policies (in all compliance categories) attached to all servers that belong to the group. This includes servers in all sub-groups that are children to the selected group.

You can change the default thresholds used to calculate compliance status. For example, you could configure that group compliance status be calculated non-recursively, which would exclude all sub-group server members from the compliance calculation.

For more information on how to change default compliance settings for device groups, see [Changing Device Group Compliance Settings](#) on page 214.

Changing Device Group Compliance Settings

By default, the SA Client allows you to configure the manner in which compliance for a device group is determined.



In order to change device group compliance settings, your user must be a member of a group that is assigned permission to the SA feature Model: Opsware. For more information on what type of permissions your user has been granted, contact your SA Administrator.

To change the settings for device group compliance, perform the following steps:

- 1 From the Navigation panel ► **Opsware Administration** ► **Compliance Settings** ► **Device Group Compliance**.
- 2 Click **Edit Settings**.
- 3 In the Device Group Settings window, you can configure the following settings:
 - **Display Device Group Rollup Compliance:** This option allows you to show or hide the icon that indicates compliance status of the parent group shown at the top of each compliance category column. This icon indicates a compliance status rollup for all members of a selected group.

For example, if this option is selected, when you select a group and from View drop-down list select Compliance, the top column heading for each compliance category column (Audit, Software, Patch, App Config) shows an icon that indicates the compliance status for all servers in the selected group. You can mouse-over this column heading to view compliance status counts all devices in this category.

- **Member Calculations:** This option allows you to choose whether or not you want to include servers that belong to sub-groups when calculating overall group compliance level for a compliance category. For example:
 - **Server and group members are considered:** This means that the compliance status for a device group will recursively check compliance for all servers in a group, and all servers in all sub-groups that belong to the selected device group.
 - **Only server members are considered:** This means that the compliance status for the selected device group will only check compliance for servers at the top level of the group, and will exclude any servers that belong to any sub-group members.
- **Thresholds:** Allows you to change the compliance threshold calculation used to determine device group compliance status for all compliance categories.

By default, a group will display a status of Non-Compliant if greater than five percent of its members are Non-Compliant; a status of Partial Compliant if greater than two percent but less than five percent of its members are Non-Compliant; and, a status of Compliant if two percent or less of its members are Non-Compliant. You can set your own default compliance status thresholds here.

Viewing Compliance Dashboard in the SA Client

In the SA Client, you can view compliance for individual servers, servers and groups together, and for groups of servers:

- [Viewing Individual Server Compliance](#)
- [Viewing Compliance for Multiple Servers](#)
- [Viewing Group Compliance in the Device Group Explorer](#)



When viewing compliance status for groups, it is possible that there are servers in the group that your user does not have permission to see. In addition, your user account might not have permissions to view some of the policies (audit, software, patch) used to calculate the compliance status for a group of servers.

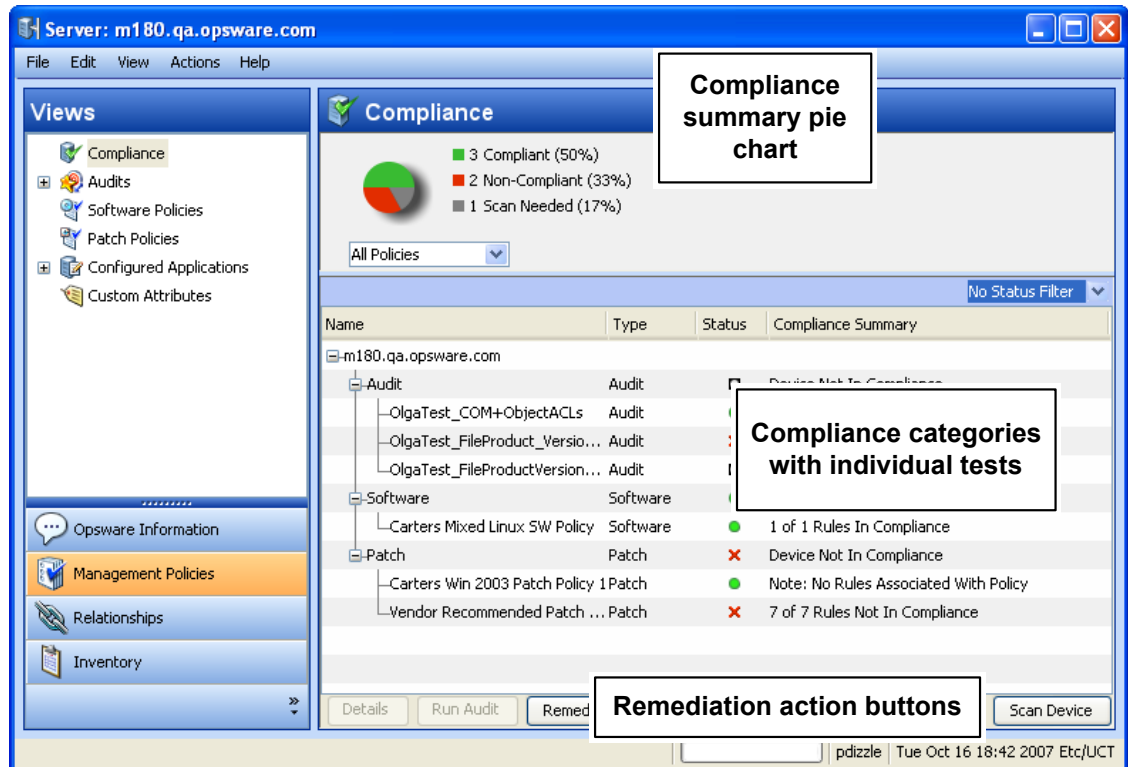
In these cases, even though you cannot see some servers and some policies, you will still be able to see overall compliance status for groups your user has access to view, and you will still be able to see compliance category roll ups, even though some of the policies may be hidden from your view.

Viewing Individual Server Compliance

To view compliance information for an individual server, perform the following steps:

- 1 From the Navigation pane, select **Devices ► All Managed Servers** (or **Virtual Servers**).
- 2 From the Content pane, select a server, right-click, and select **Open**. (Or, you can double-click the server).
- 3 In the Device Explorer, from the Views pane, select **Management Policies**.
- 4 Select **Compliance** from the Views pane. On the right side, the Content pane displays a compliance summary pie chart of compliance statuses for each compliance category, as well as detailed status information for individual policies, as shown in [Figure 35](#).

Figure 35 Compliance View for an Individual Server



- 5 To perform an action on one of the compliance categories, or an individual policy in the categories, make a selection in the Details pane and select Run Audit (for audits only), Remediate, or Scan Device. For more information on types of remediation you can perform for each policy category, see [Compliance Dashboard Remediation](#) on page 231.

► The ability to both view policies and perform remediation operations on them is determined by your user's permissions. If you are not able to view a policy or perform an action on one, consult your SA Administrator.

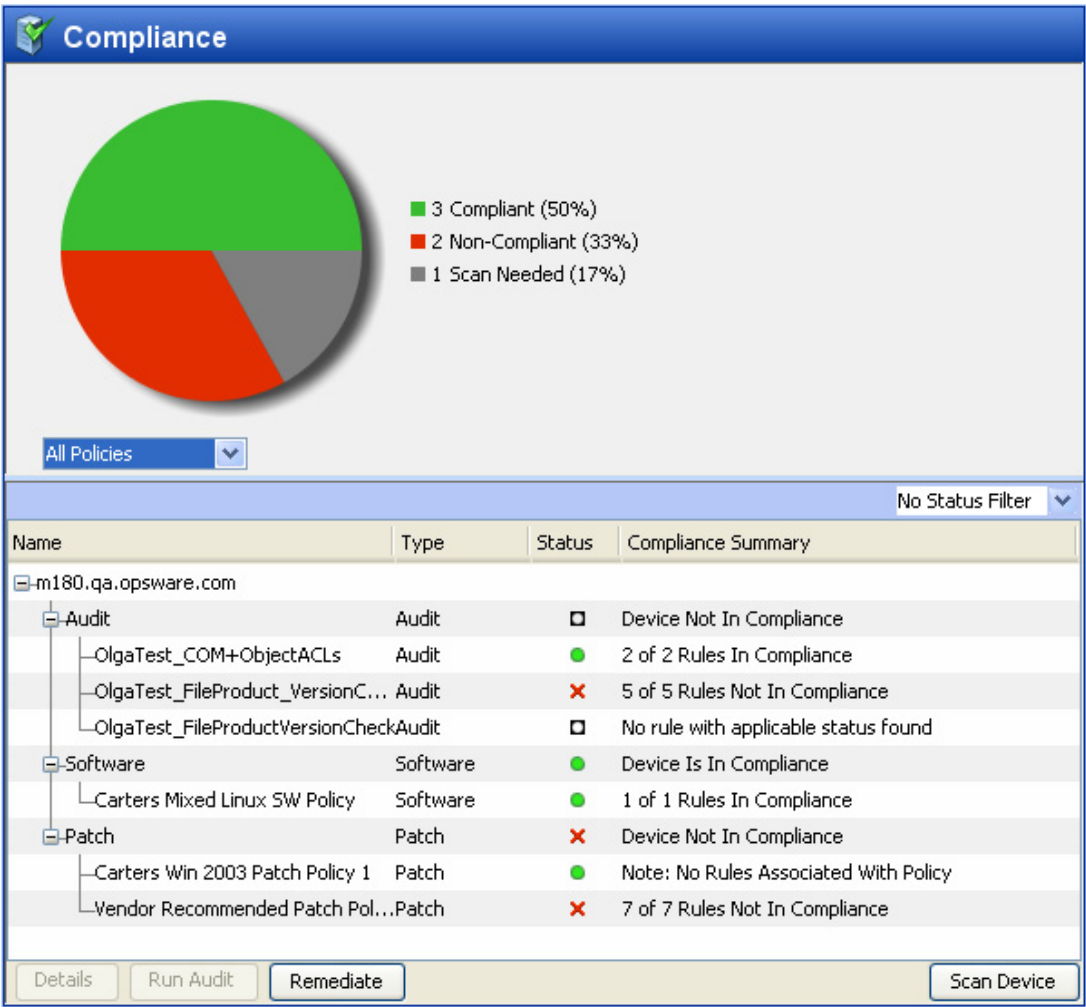
Device Explorer Compliance Summary Pie Chart and Details

The Device Explorer's Compliance View contains two main sections: the compliance summary pie chart and the compliance summary details list.

- Compliance summary pie chart (upper pane), which provides a graphical display of the overall compliance status for all policies attached to the selected server, and breaks down the percentage of each status level by category, such as Audit, Software, and AppConfig. This pie chart can also be filtered to show status only for a specific compliance category.
- Compliance summary details (lower pane) chart, which allows you to drill down in each category to see overall compliance status for each category, the individual policies contained in each category and the compliance status for each policy (Compliant, Non-Compliant, and so on), and a summary description for each. Depending upon your


selection, you can launch actions to remediate Non-Compliant policies, such as scanning the device for compliance, running an audit, viewing details of a policy, as shown in [Figure 35](#).

Figure 36 Compliance Summary Pie Chart for an Individual Server



You can select the drop-down list beneath the pie chart to view the pie chart filtered by each compliance test category, such as selecting Audits Policies, as shown in [Figure 37](#).

Figure 37 Compliance Summary Pie Chart Showing Compliance Levels for Audits Only



You can also choose to filter the compliance policy breakdowns in the details pane below the pie chart to see all compliance policies that contain a certain compliance status. For example, in [Figure 38](#), the compliance view has been filtered to show only all compliance policies that are non-compliant.

Figure 38 Server Compliance Filtered to Show Only Non-Compliant Policies

Name	Type	Status	Compliance Summary
m180.qa.opsware.com			
Audit	Audit	<input type="checkbox"/>	Device Not In Compliance
OlgaTest_FileProduct_VersionCheck	Audit	✗	5 of 5 Rules Not In Compliance
Patch	Patch	✗	Device Not In Compliance
Vendor Recommended Patch Policy f...	Patch	✗	7 of 7 Rules Not In Compliance

Details Run Audit Remediate Scan Device

In the above example, the Compliance View details pane shows all Non-Compliant policies attached to the server. A policy is considered Non-Compliant if at least one of the rules configured in the policy does not match the configuration on the server.

For a list of the actions you can take for a compliance test, such as remediate or scan device, see [Compliance Dashboard Remediation](#) on page 231

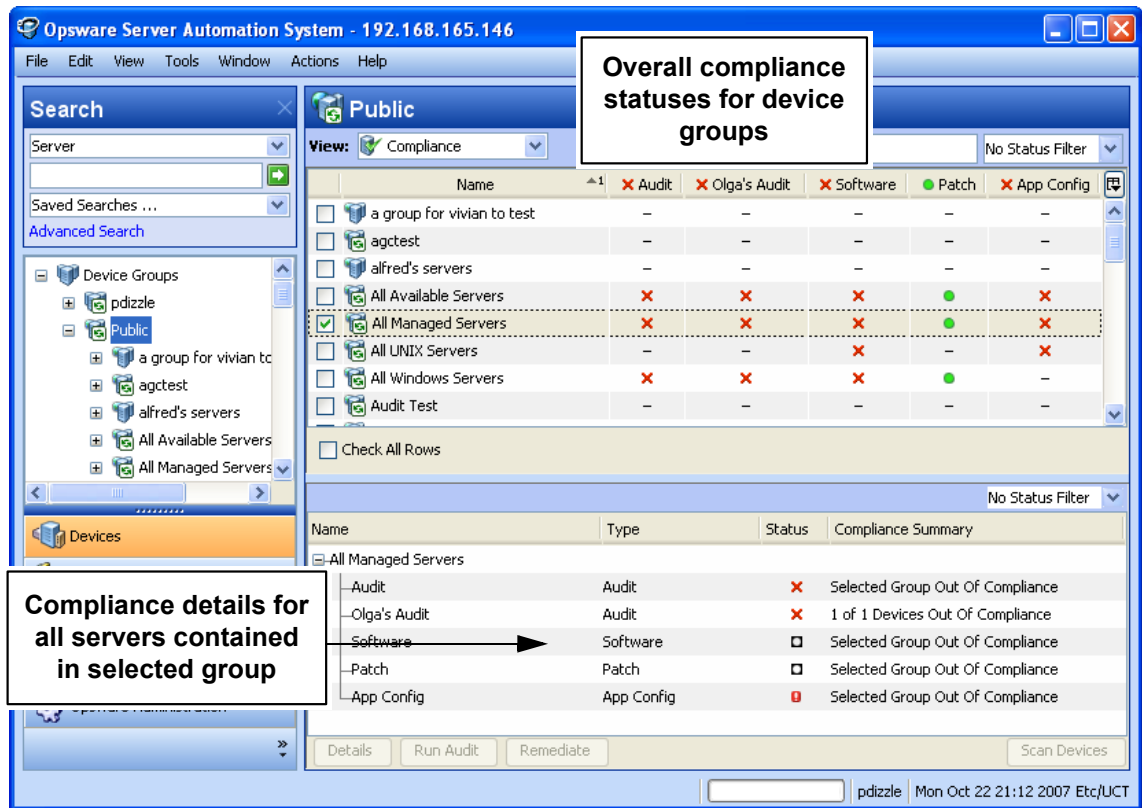
Viewing Compliance for Multiple Servers

To view compliance information for groups of servers, perform the following steps:

- 1 From the Navigation pane, select **Devices ► Device Groups**.

- 2 In the Device Groups tree, select Public. Or, select your own user's group list. The Content pane on the right side displays the contents of all the groups in the list, either all public groups or all groups your user created.
- 3 From the Views drop-down list above the Content pane, select Compliance.
- 4 For one or more of the device groups (or, any servers), select the check mark next to it to include it in the Compliance View's Details pane.
- 5 The Contents pane displays compliance summary and detail information for the selected group, as shown in [Table 39](#).

Figure 39 Compliance View for Device Groups



- 6 You can use the Status Filter drop-down list to filter the view by status; for example, you could choose to view only those groups that show a compliance status of Non-Compliant.

In the Details pane, you can select one of the categories, and depending upon on the category (and your user's permissions), click one of the action buttons at the bottom of the pane for more details, launch an audit, remediate a software or patch policy, or run a compliance scan on all members of the group.

For a list of the actions you can take for a compliance test, such as remediate or scan device, see [Compliance Dashboard Remediation](#) on page 231

Device Group Compliance — Content Pane

The Device Group Content pane displays a summary of compliance status roll ups for all the group members (and contents of groups) that you have selected from the Navigation pane **Devices ► Device Groups**.

Compliance status (Compliant, Non-Compliant, Partial, and so on) icons in the column heading at the top of the list indicate the rollup status for all groups in the list. If you move your mouse pointer over the top of the column for a category, pop-up text displays the overall compliance for the compliance category for all the visible groups.

In each row of the list, this view displays compliance status for each group in all four compliance categories for each group in the list, which include Audit, Software, Patch, and App Config, as well as any individually scheduled audits that you choose to display in this view, as shown in [Figure 36](#).

Figure 40 Compliance Roll Ups for Groups

The screenshot shows a web interface titled 'Public' with a 'View: Compliance' dropdown. Below the header, there is a table with columns for 'Name', 'Audit', 'Olga's Audit', 'Software', 'Patch', and 'App Config'. The table lists several server groups, including 'a group for vivian to test', 'agctest', 'alfred's servers', 'All Available Servers', 'All Managed Servers', 'All UNIX Servers', 'All Windows Servers', and 'Audit Test'. The 'All Managed Servers' row is highlighted with a green checkmark in the 'Name' column and red 'X' marks in the 'Audit', 'Olga's Audit', 'Software', and 'App Config' columns, and a green dot in the 'Patch' column. A 'Check All Rows' checkbox is at the bottom left.

Name	Audit	Olga's Audit	Software	Patch	App Config
a group for vivian to test	-	-	-	-	-
agctest	-	-	-	-	-
alfred's servers	-	-	-	-	-
All Available Servers	X	X	X	●	X
All Managed Servers	X	X	X	●	X
All UNIX Servers	-	-	X	-	X
All Windows Servers	X	X	X	●	-
Audit Test	-	-	-	-	-

In [Figure 36](#), each compliance category (Audit, Software, and so on) displays a compliance status for all policies of each type that are attached to servers in the group. The group named All Managed Servers, for example, displays all categories as Non-Compliant X except for the Patch category. This means that other than Patch, more than five percent of the servers in the group have a status of Non-Compliant for Audits, Software, and App Config (as well as the custom column named Olga's Audits).

The Patch category, however, shows a Compliant ● status, which means that at least 95 percent of patch policies attached to servers in this group have a Compliant status. (For information on how to change the compliance status thresholds for device groups, see [Changing Device Group Compliance Settings](#) on page 214.)

In addition, the scheduled audit named Olga's Audits" has been added to the list, which shows the status of all the servers targeted by that specific audit. For information on how to add or remove compliance categories, see [Adding and Removing Compliance View Columns](#) on page 225.

Device Group Compliance — Details Pane

When you select one or more groups from the Content pane (or all of them), the Details pane displays device compliance aggregate rollups in each column of the summary pane for all members of the group, as displayed in [Figure 41](#).

Figure 41 Device Group Members Compliance Status Rollup in the Details Pane

Name	Type	Status	Compliance Summary
All Available Servers			
Audit	Audit	✗	Selected Group Out Of Compliance
Olga's Audit	Audit	✗	1 of 1 Devices Out Of Compliance
Software	Software	☐	Selected Group Out Of Compliance
Patch	Patch	●	Selected Group Is In Compliance
App Config	App Config	⚠	Selected Group Out Of Compliance

Details

Run Audit

Remediate

Scan Devices

You can use the Status Filter drop-down list to filter the view by status; for example, you could choose to view only those groups that show a compliance status of Non-Compliant.

You can also select one of the compliance aggregate columns, and depending upon on the columns (and your user's permissions), click one of the action buttons at the bottom of the pane for more details, launch an audit, remediate a software or patch policy, or run a compliance scan on all members of the group. For a list of the actions you can take for a compliance test, such as remediate or scan device, see [Compliance Dashboard Remediation](#) on page 231

Viewing Group Compliance in the Device Group Explorer

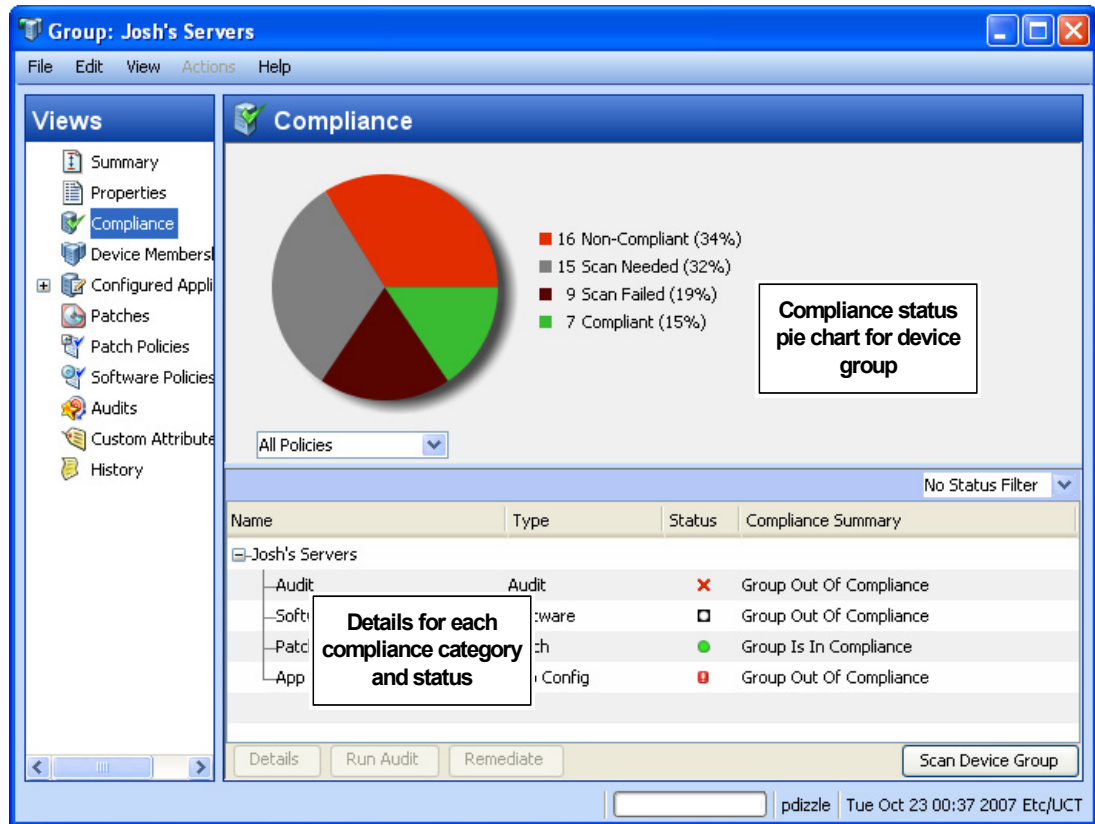
To view detailed compliance information about a group of servers, open a group's Device Group Explorer and select its Compliance View, which shows a rollup of compliance policy aggregates for each policy type for all members of the group as a whole, as opposed to compliance status for individual servers. This gives you a sense of whether or not the group is compliant for each policy type, and for all servers in the group (and any sub-groups).

You can use the Status Filter drop-down list to filter the view by status; for example, you could choose to view only those groups that show a compliance status of Non-Compliant. You can also select one of the categories, and depending upon on the type of category (and permissions to view), click one of the action buttons at the bottom of the pane for more details, launch an audit, remediate a software or patch policy, or run a compliance scan on all members of the group.

To view a group of servers in the Device Group Explorer, perform the following steps:

- 1 From the Navigation pane, select **Devices ► Device Groups**.
- 2 In the Device Groups tree, navigate and select Public (or, select your own user's group list), and then select a group specific group.
- 3 From the **Actions** menu, select **Open**.
- 4 From the View pane of the Device Group Explorer, select Compliance. The Compliance View displays summary and rollup compliance status information about all servers in the group. shows the Compliance View for a device group in the Device Group Explorer, as shown in [Figure 38](#).

Figure 42 Device Group Compliance View

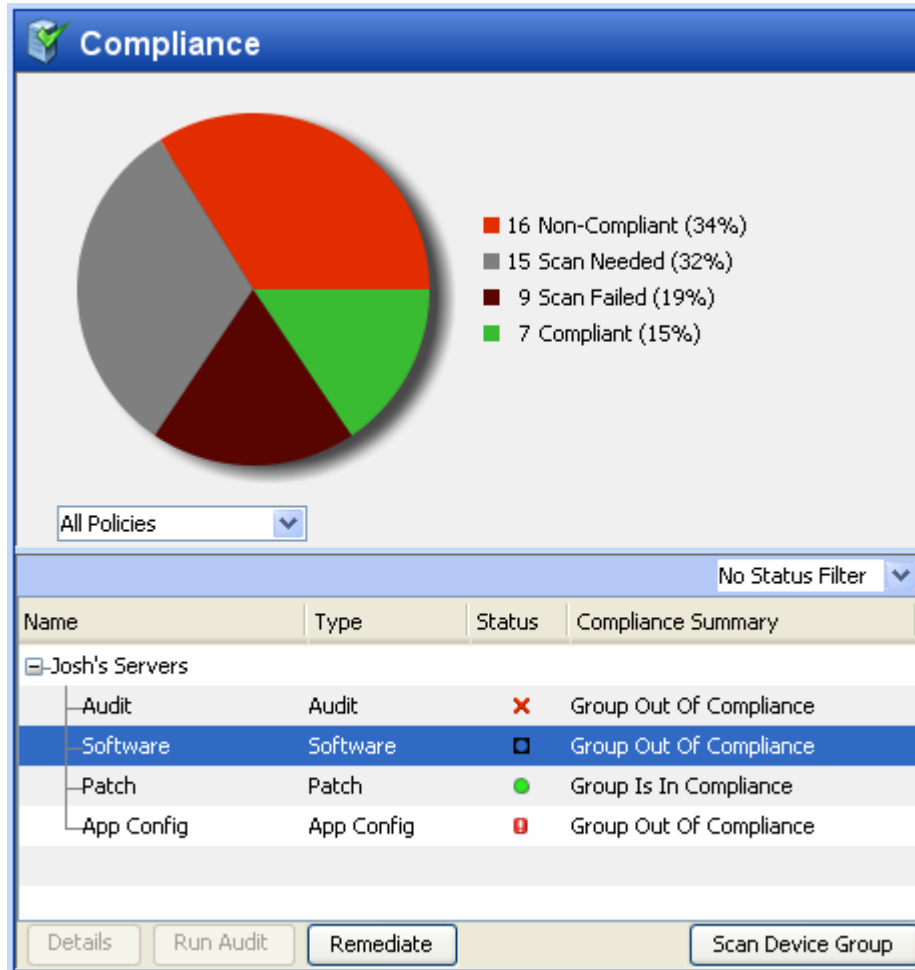


Device Group Compliance Summary Pie Chart and Details Pane

The Device Group Explorer's Compliance View contains two main sections:

- Compliance summary pie chart (upper pane), which provides a graphical display of the overall compliance status for all policies aggregates for all associated servers in the group, and breaks down the percentage of each status level by category, such as Audit, Software, and AppConfig. This pie chart can also be filtered to show status only for a specific compliance category.
- Compliance summary details (lower pane) chart, which allows you see device group compliance status for each category (Compliant, Non-Compliant, and so on) and view a summary for each. Depending upon your selection (and your user's permissions), you can launch actions to remediate non-compliant policies, scan device for compliance, run an audit, view details of a policy, as shown in [Figure 43](#).

Figure 43 Device Group Browser Compliance Summary Pie Chart and Details Pane



By default, a device group is Non-Compliant if more than five percent of the servers in the group have a status of Non-Compliant. You can filter the pie chart to show only those servers that have a specific compliance, such as, show all servers in the group that have a status of Compliant.

Figure 43, shows that sixteen of the policies (thirty four percent) attached to servers in the selected device group meet have a status of Non-Compliant. You can select a compliance category from the Details pane below the pie chart (for example, Software), and perform a remediation action on the group, depending upon your user's permissions. For example, you can select Software and then click Remediate to remediate the Software Policy on to the servers in the group.

Adding and Removing Compliance View Columns

When you view compliance for device groups in the Compliance View, by default, all four compliance categories are displayed as columns in the Content pane — Audit, Software, Patch, and App Config. You can, however, add or remove any of these categories, as well as add or remove any individual policy in each category.

To add or remove device group compliance categories in the Compliance View, perform the following steps:


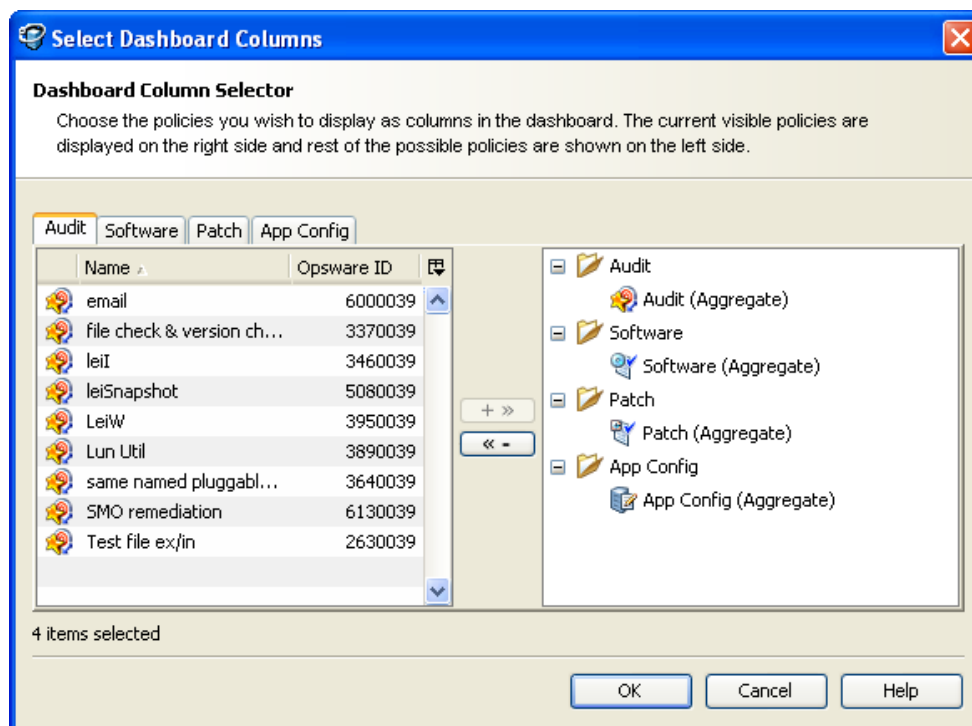
- 1 From the Navigation pane, select **Devices ► Device Groups**.
- 2 In the Device Groups tree, navigate and select Public (or, select your own user's group list), and then select a specific group. The Content pane shows a list of all four compliance categories and statuses for each member of the group.
- 3 To add or remove a category, click the Column Selector  button on the upper right corner of the Content pane.
- 4 In the Select Dashboard Columns window, the left side of the window displays four tabs, one for each compliance category and all compliance policies in those categories your user has permissions to see. The right side of the window displays the currently visible policies in each category in the Compliance View. By default, the Compliance View displays the Aggregate (rollup) of all policies in the category, as shown in [Figure 44](#).

Figure 44 Select Compliance View Columns



- 5 To add an individual policy as a column in the Compliance View, from the left side, select a compliance category tab and then a policy and click the right arrow button.
- 6 To remove an individual policy or an aggregate column from the Compliance View, select one from the right-side of the window and then click the left arrow button.
- 7 When you are finished, click **OK**. You can now view your changes in the Compliance View.

Filtering By Compliance Status


When you view compliance for groups of servers (and individual servers) in the Compliance View, you can filter the view to show only groups and servers that have at least one server that matches a specific compliance status for any of the displayed compliance categories.

For example, when you select a group and then select Compliance View, you can use the status filter to only show members of the selected group (individual servers and those in any sub-groups) that have a Non-Compliant status for each of the compliance categories, such as Audit, Software, and so on.

To filter the Compliance View by compliance status, perform the following steps:

- 1 From the Navigation pane, select **Devices ► Device Groups**.
- 2 In the Device Groups tree, navigate and select Public (or, select your own user's group list), and then select a group specific group. The Content pane shows the Compliance View statuses for all members of the selected group.
- 3 To filter this view by compliance status, select one from the compliance status drop-down list, as shown in [Figure 45](#).

Figure 45 Compliance Status Drop-Down List



- 4 The Compliance View displays only those members of the group (individual servers and those in any sub-groups) have a status of Non-Compliant.
- 5 You can select any of the servers or sub-groups in the group listed, and Details pane below will show the compliance status information for those servers. You can further filter the Details pain by using the status filter on the upper right corner of the Details pane.

Refreshing For Latest Compliance Information

When you first select the Compliance View, the information displayed shows the latest information reported from the SA core for each compliance category. It is possible, however, that a server's configurations has changed since you last looked at the Compliance View. It is also possible that a policy has changed since you last viewed server and groups in the Compliance View.

If this is the case, you might want to scan for compliance, or rerun an audit in order to generate new data for the Compliance View to display.

As a best practice, it's useful to refresh the Compliance View to ensure that you are looking at the latest compliance information in your core. To get the latest compliance information from the core, from the **View** menu, select **Refresh**, click **Refresh**, or press Control + F5.

Setting Automatic Compliance Check Frequency

By default, the SA Client will check the core for new or changed compliance information every five minutes. However, you can change this time interval using the Set Options window.



If you want the SA Client to immediately check for new compliance information from the core, press Control + F5.

To set the automatic compliance check frequency, perform the following steps:

- 1 From inside the SA Client, from the **Tools** menu select **Options**.
- 2 In the Set Options window, select General from the left pane.
- 3 In the General section on the right pane, in the Cache — Check for updates section, enter a time interval for how often you want the SA Client to check the core for new compliance information.



Note that this check applies to all information accessed from the core by the SA Client, not just compliance information. A longer interval increases the likelihood that the information you are viewing is out of date, while a shorter interval increases network traffic flowing to and from your core.

- 4 When you are finished, click **Save**.

Scanning for Compliance

When you scan for compliance, you are scanning the servers targeted by a compliance policy — Software, Patch, or App Config — to determine if the target server configurations match the policy's rule definitions. For an audit, when you run an audit it checks the target server configuration to determine the extent to which it matches the audit's rule definitions.

For example, a compliance scan can check to see what patches are installed on a computer, compare that with a patch or software policy, and then return the results to the Compliance View. Or, a compliance scan can check the contents of a configuration file on a server in order to determine if it matches the rules defined in an application configuration.

In the Compliance View, you can perform a compliance scan for the Software, Patch, and App Config compliance categories. Audits do not have a scan feature, but running an audit achieves the same results.

Specifically, each different feature category performs the following actions when scanning for compliance:

- **Software Compliance Scan:** Compares configuration files on a server to determine if they match the values stored in the software policies attached to the server or group of servers. The results of this scan show you the servers that are in compliance (have all required software policy items installed) and the servers that are out of compliance (do not have all required software policy items installed). For more information on scanning software policies, see [The Software Policy Compliance Scan](#) on page 282.
- **Patch Compliance Scan:** Compares patches that are installed on a server with patch policies and patch policy exceptions that are attached to that server. The results of this scan show you the servers that are in compliance (have all required patches installed) and the servers that are out of compliance (do not have all required patches installed). Scanning for compliance relates only to Windows patching; Unix patching is encompassed within software policies.

For more information on patch compliance see [Patch Management for Windows](#) on page 287.

- **App Config Compliance Scan:** Compares configuration files on a server with the template definitions defined application configurations that are attached to that server. The results of this scan show you the servers that are in compliance (configuration file definitions match the configuration templates) and the servers that are out of compliance (configuration file definitions do not match the configuration templates).

For more information on App Config compliance, see [Application Configuration Compliance](#) on page 244.

Exporting Compliance View Information

If you want to view all the information displayed in the Compliance View to a file, you can export the view to either .html or .csv.

To export Compliance View information to a file, perform the following steps:

- 1 To view the Compliance Dashboard, from the Navigation pane, select **Devices ► Device Groups**.
- 2 Select a group that you want to view compliance for, and from the **View** menu, select **Compliance**.
- 3 Right-click inside the Contents pane and select **Export**.
- 4 In the Export Compliance View window, enter a name for the file, and choose if you want to export to .html or .csv. You can also change the encoding if you want the saved file to use a specific encoding scheme.
- 5 Click **Export**.

Compliance Dashboard Remediation

In addition to providing compliance status information for servers and groups, the Compliance View enables you to remediate server configurations that are not in compliance with your organization's standards, as defined by your audit, software, patch, and application configuration compliance policies.

Generally speaking, the act of remediating a server or group of servers means finding how and where a server or group is out of compliance (Non-Compliant), and then making sure that a server's actual configuration conforms to your compliance policies.

From the Compliance View for a server or group of servers, you can perform the following actions:

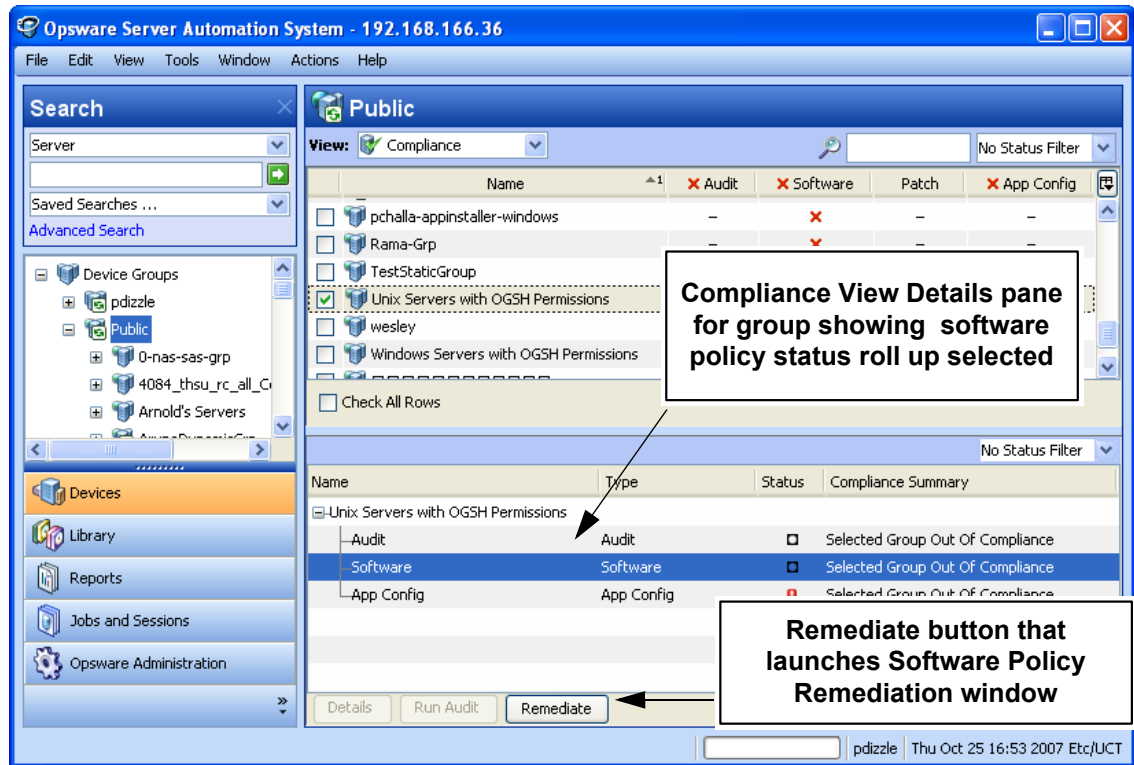
- Remediating a software or patch policy
- Running, viewing, and remediation audit results
- Pushing an application configuration on to a server
- Running a compliance scan for patches, software, or application configurations to get the latest compliance information for your servers.

When you select a server or group of servers from in the Compliance View, or view them in the Device or Device Group Explorer, the Details pane provides action buttons for actions that help you discover and remediate out of compliance policies. The type of actions available depending upon which varies depending upon the type of policy, whether you select a single server or group of servers, and whether or not you select an individual policy, multiple policies, or the roll up of a compliance category (such as Audit, Software, Patch, or App Config) in the Details pane.

Group Compliance Remediation

[Figure 46](#) displays how the Compliance View enables a group's compliance remediation options as located in the Details pane.

Figure 46 Details Pane Displaying Remediation Actions Available for a Group of Servers



In Figure 46, the Details pane for the selected group shows a summary of all policies attached to all servers in the group (and all servers in any sub-groups) arranged by compliance category — Audit, Software, Patch, and App Config. When you select a group, you can only remediate an entire category of policies, such as all software or patch policies attached to all servers in the group that are out of compliance. If you select the Software category in the Details pane, the **Remediate** button activates and when clicked launches the Software Policy Remediation window, allowing you to remediate any out of compliance policy configurations for all servers in the group.

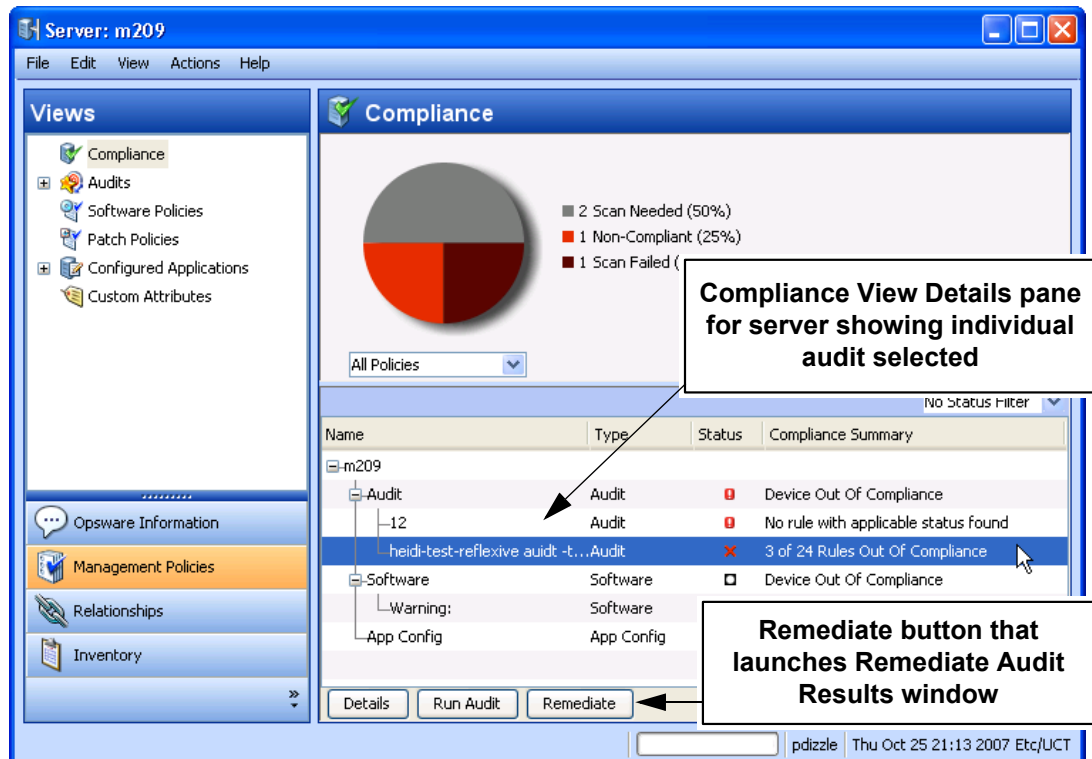
You can view this same information and access these option by selecting the group, and from the Actions menu select **Open**. Doing this launches the Device Group Explorer and displays the same Details pane for the group, along with the action buttons at the bottom of the pane.

Compliance Remediation for Servers

With groups, your remediation options always apply to all members of the group. For an individual server, however, you can launch the server's Device Explorer and you can remediate either all or specific policies that are attached to the server. For example, you can launch a server, and from the server's Device Explorer select **Management Policies** ► **Compliance**, and view all the compliance policies attached to the server.

From the Details pane, you can select an audit or a software policy, and view the audit, run the audit, remediate the software policy, scan the device for compliance, and so on, as shown in Figure 47.

Figure 47 Device Explorer Showing Compliance and Remediation Options



Audit Compliance

The Audit and Remediation feature allows you to define server configuration policies in an *audit*, which helps ensure that the servers in your facilities meet your audit policy standards. An audit consists of a collection of rules that you can define to model those standards.

For example, an audit might consist of Windows COM+ configurations, registry settings, services, file system settings, hardware configuration, user and group password settings, software installation, packages, storage settings, and so on, that define an ideal server configuration. Or, the audit might represent a negative server configuration that enables you to determine the way a server should *not* be configured.

Audit compliance determines if the rules defined in a recurring audit match the actual server configuration for all servers targeted by the audit. The Compliance View allows you to see both the aggregate and individual compliance status of all audits that run on a recurring schedule on a server or group of servers. If any of the audits are Non-Compliant, you can remediate any differences found between the audit and the audit's target server or servers.

The Compliance View derives audit compliance servers and groups of servers from regularly scheduled audits. For more information on audit remediation, and how to create and schedule recurring audits, see [Audit Compliance Remediation](#) on page 235.

Audit Compliance Status

Audit compliance status is determined by the following criteria:

- **Audit Compliance — Single Server:** If a single rule inside an audit does not match the target server's configuration, then the server's audit compliance status is Non-Compliant. The Details pane of a server's Device Explorer window shows the Audit category as Non-Compliant, and the summary column indicates how many rules are Non-Compliant out of the total number of rules.

For example, if an audit has ten rules, and four of the rules are Non-Compliant, then the audit's status is listed as Non-Compliant and the summary description reads: "4 of 10 Rules Out of Compliance."

If more than one audit targets the server, and if at least one of those audits is Non-Compliant, then the aggregate compliance status for audits is displayed as Non-Compliant as well. You can expand the Audit category of the Details pane to see which of the audits are not in compliance, as well as a breakdown of how many rules in each audit are in compliance or out of compliance.

- **Audit Compliance — Device Groups:** An audit that targets a group of servers (and all the servers in all sub-groups) is considered Compliant if at least 95 percent of the servers in the group that are targeted by the audit have a compliance status of Compliant.

If more than five percent of the servers in the group targeted by an audit have an status of Non-Compliant, then the aggregate compliance for audits will display as Non-Compliant. Another way to state Non-Compliance for a group is when less than ninety five percent of the servers are Compliant.

However, if more than two percent but less than or equal to five percent of all servers in a group have the status of Non-Compliant for that category, then the status is Partial-Compliance. Another way to state Partial-Compliance for a group is when less than ninety eight percent but at least ninety five percent of the servers are Compliant.

If less than two percent of all servers in a group have an Audit status of Non-Compliant for that category then the overall status is Compliant. Another way to state Compliance for a group is that at least ninety eight percent of the servers are Compliant for a given category.

The Details pane for a group of servers in the Compliance View shows whether or not all of the audits are compliant or not, but does not expand to show a breakdown of individual servers and audits.

You can modify the thresholds used to determine compliance for groups of servers. For more information, see [Changing Device Group Compliance Settings](#) on page 214.

Audit Compliance Remediation

The Compliance View allows you to view all audits that target a server or group of servers and to remediate those results that are out of compliance, to ensure that a server's configuration complies with the rules defined in an audit.

For each audit rule that is out of compliance on the target server (the server's configuration either mismatched the rule definition or simply did not exist), remediation copies the rule object onto the target server so it matches the rule. Or, in the case of a value-based audit rule, changes the target server's configuration to match the rule.

For example, you have an audit that checks a group of Windows servers to make sure that they contain specific registry keys and ACLs. After the audit runs against an actual Windows server, it is possible that several of the rules are out of compliance — which means the Registry keys specified in the audit rules were not found on the target servers.

When you remediate, the audit feature copies the Registry keys specified in the Audit rule on to the target servers, ensuring that the servers have the specific keys and associated ACLs. For a group of servers, remediation has the same results, only the remediation operation applies to all servers in the group, including all the server contained in any sub-groups.

Remediating Audits for One or More Servers

You can remediate an audit that is attached to a single server or one attached to several servers that you select in the Device Groups list. You can only remediate individual audits, but not aggregate audits at the top of level. For any Group that is selected, all direct server children in that group are the subject of the remediation.



If the remediate button is not enabled, even though a single policy is selected in the detail pane (and one or more servers are selected in the summary pane), this likely means that there is no audit result for that policy to remediate.



You cannot run an audit on a group of servers from the Compliance View. However, you can create an audit that runs against a group of servers and remediate those audit results for a group of servers from the Audit Results window. For more information, see [Configuring an Audit](#) on page 168 and [Remediating Audit Results](#) on page 240.

To remediate an individual audit on one or more servers, perform the following steps

- 1 To remediate an individual audit on a single server in the Device Explorer, from the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2 From the list, select a server.

- 3 From from the **Actions** menu, select **Open**.
- 4 From inside the Device Explorer's View pane, select **Management Policies ► Compliance**.
- 5 In the Details pane of the Compliance View, expand the Audit category and select an individual policy.
Or
- 6 Select multiple servers by selecting the check box next to the server.
- 7 To remediate an individual audit for several servers, from the Navigation pane, select **Devices ► Device Groups** and select a group.
- 8 From the View drop-down list, select Compliance.
- 9 In the Details pane of the Compliance View, expand the Audit category and select an individual audit that is targeting all of the selected servers.
- 10 The following list describes the types of remediation you can perform for an individual audit on individual or multiple servers selected in the summary pane:
 - **Details:** Launches the Launch Audit Results window, which allows you to view all differences found between the audit and the target, and to remediate the differences by rule or by server. For more information, see [Remediating Audit Results](#) on page 240.
 - **Run Audit:** Launches the Run Audit task window and allows you to run the audit immediately or schedule to run the audit at a later time. The audit will run against all servers targeted by the audit.
 - **Remediate:** Launches the audit Remediate Audit Results window, which allows you to remediate target server configurations that are out of compliance with the audit rules. You can remediate differences by rule or by server. For more information, see [Remediating Audit Results](#) on page 240.
 - **Scan Device:** Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. This does not have any effect on the audits that target this server. For more information, see [Scanning for Compliance](#) on page 229.

Software Compliance

The Software Management feature allows you to create *software policies* that enable you to install software and configure applications simultaneously. A software policy can contain several different kinds of items, such as packages, RPM packages, patches, application configurations, and other software policies. After creating a software policy, you can attach it to servers or groups of servers.

Software compliance indicates whether or not the items in a software policy are compliant with the actual server configuration. If the actual server configuration does not match the software policy definitions, then the server's software policies are Non-Compliant.

The Compliance View derives software compliance information for software policies when you scan a server or group for software compliance. For more information on how to scan for software compliance, see [Scanning for Compliance](#) on page 229.

For more information software management and how to create and manage software policies, see [Software Management Setup](#) on page 13.

Software Compliance Status

Software compliance status is determined by the following criteria:

- **Software Compliance — Single Server:** If at least one item in a software policy does not match what is discovered (or does not exist) on the server the policy is attached to, the server's software compliance status is Non-Compliant. The Details pane of a server's Device Explorer window shows the Software category as Non-Compliant and the summary column indicates how many rules (software policy items) are Non-Compliant out of the total number of rules.

For example, if a software policy contains ten items, and six of the items are Non-Compliant, then the software policy's status is listed as Non-Compliant and the summary description reads: "6 of 10 Rules Out of Compliance."

If more than one software policy targets a single server, and if at least one of those policies is Non-Compliant, then the aggregate compliance status for Software is displayed as Non-Compliant as well. You can expand the Software category of the Details pane to see which of the policies are not in compliance, as well as a breakdown of how many rules in each policy are either in or out of compliance.

- **Software Compliance — Device Groups:** A software policy attached to a group of servers is considered Compliant if more than five percent of the servers in the group attached to the policy have a status of Non-Compliant. If this is the case, the aggregate compliance for software policy will display as Non-Compliant. Another way to state Non-Compliance for a group is when less than ninety five percent of the servers are Compliant.

However, if more than two percent but less than or equal to five percent of all servers in a group have the status of Non-Compliant for that category, then the status is Partial-Compliance. Another way to state Partial-Compliance for a group is when less than ninety eight percent but at least ninety five percent of the servers are Compliant.

If less than two percent of all servers in a group have a Software Policy status of Non-Compliant for that category then the overall status is Compliant. Another way to state Compliance for a group is that at least ninety eight percent of the servers are Compliant for a given category.

The Details pane for a group of servers in the Compliance View shows whether or not all of the software policies are compliant or not, but does not expand to show a breakdown of individual servers and policies.

You can modify the thresholds used to determine compliance for groups of servers. For more information, see [Changing Device Group Compliance Settings](#) on page 214.

Software Compliance Remediation

The Compliance View allows you to view all software policies attached to a server or groups of servers and to remediate those servers that are out of compliance, and in the process ensure that a server's software configuration complies with the software policy definition.

For each software policy item — such as software, packages, patches, scripts, application configurations — software remediation installs (or for a script, executes) those items on the target server. If the items do not exist on the server, then they get installed. If the items existed but did not match the policy, they get updated with the correct version.

For example, you have a software policy that consists of several packages, patches, a few scripts, and an application configuration all organized in the order in which they are to be installed and executed. First, you remediate the software the policy onto a servers to make sure the server is in compliance with your company's software installation standards. Over time, some of the items in the software policy get updated — such as a new set of packages gets added— and for whatever reason, a software item on the server was uninstalled.

When you perform a software compliance scan, the scan determines the server's compliance status by comparing the software policy contents with the actual software installed on the server. Even if only one software item attached to one of the servers is not in compliance with the policy, the server will have a software compliance status of Non-Compliant.

When you remediate a server or group of servers, the patches, packages, and application configurations specified in the policy are installed and applied in the order specified in the policy. For a group of servers, remediation has the same results, only the remediation operation applies to all servers in the group, including all the server contained in any sub-groups. (For information on how to change this setting, see [Changing Device Group Compliance Settings](#) on page 214.)

For more information on software policy remediation, see [Software Management](#) on page 257.

Remediating Software Compliance — Single or Multiple Servers

When you remediate software compliance for a single server or multiple servers, you can choose to remediate all of the policies attached to the servers or select to remediate individual policies.

You can select the Software Aggregate policy, which remediates all software policies for all servers selected. If a group is selected, it remediates against all direct server children in that group. If a single software policy is selected in the details pane, then the entities selected in the summary pane have that policy remediated.

To remediate software policies on single or multiple servers, perform the following steps:

- 1 To remediate software policies for a single server in the Device Explorer, from the Navigation pane, select **Devices ► Servers ► All Managed Servers**.
- 2 Select a server from the list.
- 3 From from the **Actions** menu, select **Open**.

- 4 From inside the Device Explorer's View pane, select **Management Policies ► Compliance**.
- 5 In the Details pane of the Compliance View, expand the Software category and select an individual software policy or the top level Software category, which will enable you to remediate all of the policies attached to the server.
Or
- 6 In the Content pane that shows a list of servers that belong to the group, select multiple servers by selecting the check box next to the server.
- 7 To remediate software policies for multiple servers, from the Navigation pane, select **Devices ► Device Groups** and select a group.
- 8 From the View drop-down list, select Compliance.
- 9 In the Details pane of the Compliance View, expand the Software category and select a software policy that is attached to the selected servers. Or, select the top level Software category if you want to remediate all of the software policies attached to the selected servers.
- 10 From the bottom of the Details pane, you have the following options:
 - **Remediate:** Remediate the selected software policy or policies against the selected server or servers.
 - **Scan Device:** Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. For more information, see [Scanning for Compliance](#) on page 229.

Remediating Software Compliance — Groups

When you remediate software policies for a group or multiple groups of servers, you can remediate all the policies attached to all servers in the group or multiple groups. However, when you select a group or multiple groups, you can only remediate *all* of the software policies attaches to all the servers in the group and any sub-groups.

To remediate software policies for groups or multiple groups of servers, perform the following steps:

- 1 To remediate software policies for a single server in the Device Explorer, from the Navigation pane, select **Devices ► Servers ► All Managed Servers**.
- 2 Select a server from the list.
- 3 From from the **Actions** menu, select **Open**.
- 4 From inside the Device Explorer's View pane, select **Management Policies ► Compliance**.
- 5 In the Details pane of the Compliance View, expand the Software category and select an individual software policy or the top level Software category, which will enable you to remediate all of the policies attached to the server.
Or
- 6 In the Content pane that shows a list of servers that belong to the group, select multiple servers by selecting the check box next to the server.
- 7 To remediate software policies for multiple servers, from the Navigation pane, select **Devices ► Device Groups** and select a group.
- 8 From the View drop-down list, select Compliance.

- 9 In the Details pane of the Compliance View, expand the Software category and select a software policy that is attached to the selected servers. Or, select the top level Software category if you want to remediate all of the software policies attached to the selected servers.
- 10 From the bottom of the Details pane, you have the following options:
 - **Remediate:** Remediate the selected software policy or policies against the selected server or servers.
 - **Scan Device:** Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. For more information, see [Scanning for Compliance](#) on page 229.

Patch Compliance

The Patch Management (Windows and Unix) feature enables you to identify, install, and remove patches on managed servers and groups of servers. With Windows Patch Management you can identify and install patches for the Windows 2000, Windows 2003, and Windows NT 4.0 operating systems, include Service Packs, Update Rollups, and hotfixes.

In the Compliance View, you can view compliance status for patch policies in order to see whether or not your servers have the correct patches installed on them. During a Patch compliance scan, Patch Management checks managed servers and public device groups to determine whether all patches in a policy and a policy exception were installed successfully. If the patches installed (or not installed) on the server does not match the patch policy definitions, then the Compliance View displays the server's patch policies as Non-Compliant.

Compliance scans can be run on a one time basis, or can be scheduled on a recurring basis. You can remediate a patch policy to a server in order to ensure a server's or groups patch compliance.

For more information on patch compliance, see:

- [Patch Compliance](#) on page 323
- [Verifying Patch Policy Compliance](#) on page 319
- [Scanning for Compliance](#) on page 229
- [Scheduling a Patch Compliance Scan](#) on page 329

Patch Compliance Status

Patch compliance status is determined by the following criteria:

- **Patch Compliance — Single Server:** If at least one item in a patch policy does not match what is discovered (or does not exist) on the server the policy is attached to, the server's patch compliance status is Non-Compliant. The Details pane of a server's Device Explorer window shows the Patch category as Non-Compliant and the summary column indicates how many rules (patch policy items) are Non-Compliant out of the total number of rules.

For example, if a patch policy contains ten items, and six of the items are Non-Compliant, then the patch policy's status is Non-Compliant and the summary description reads: "6 of 10 Rules Out of Compliance."

If more than one patch policy targets a single server, and if at least one of those policies is Non-Compliant, then the aggregate compliance status for Patch is displayed as Non-Compliant as well. You can expand the Patch category of the Details pane to see which of the policies are not in compliance, as well as a breakdown of how many rules in each policy are either in or out of compliance.

- **Patch Policy — Rule Exception:** If a rule exception is applied to one of the patch policy items, then the server's Patch compliance will display a compliance status of Partial. Patch is the only compliance category that allows rule exceptions at the policy level, and thus is the only category that can have a Partial compliance status.
- **Patch Compliance — Device Groups:** A patch policy attached to a group of servers is considered Compliant if more than five percent of the servers in the group attached to the policy have a status of Non-Compliant. If this is the case, the aggregate compliance for Patch Policy will display as Non-Compliant. Another way to state Non-Compliance for a group is when less than ninety five percent of the servers are Compliant.

However, if more than two percent but less than or equal to five percent of all servers in a group have the status of Non-Compliant for that category, then the status is Partial-Compliance. Another way to state Partial-Compliance for a group is when less than ninety eight percent but at least ninety five percent of the servers are Compliant.

If less than two percent of all servers in a group have a Patch Policy status of Non-Compliant for that category, then the overall status is Compliant. Another way to state Compliance for a group is that at least ninety eight percent of the servers are Compliant for a given category.

The Details pane for a group of servers in the Compliance View shows whether or not all of the patch policies are compliant or not, but does not expand to show a breakdown of individual servers and policies.

You can modify the thresholds used to determine compliance for groups of servers. For more information, see [Changing Device Group Compliance Settings](#) on page 214.

For more information on creating and using patches and patch policies, see [Patch Management for Windows](#) on page 287 or [Patch Management for Unix](#) on page 395

Remediating Patch Compliance — Single or Multiple Servers

When you remediate patch compliance for a single server or multiple servers, you can choose to remediate either all of the policies attached to the servers or only remediate individual policies.

You can remediate patch policies for a single server by viewing the server's Device Explorer, or you can remediate patch policies for multiple servers by selecting the policies in the Device Groups list.

To remediate patch policies on single or multiple servers, perform the following steps:

- 1 To remediate patch policies for a single server in the Device Explorer, from the Navigation pane, select **Devices ► Servers ► All Managed Servers**.
- 2 Select a server from the list.
- 3 From the **Actions** menu, select **Open**.
- 4 From inside the Device Explorer's View pane, select **Management Policies ► Compliance**.
- 5 In the Details pane of the Compliance View, expand the Patch category and select an individual policy or the top level Patch category, which will enable you to remediate all of the patch policies attached to the server.

Or

- 6 In the Content pane that shows a list of servers that belong to the group, select multiple servers by selecting the check box next to the server.
- 7 To remediate patch policies for multiple servers, from the Navigation pane, select **Devices ► Device Groups** and select a group.
- 8 From the View drop-down list, select Compliance.
- 9 In the Details pane of the Compliance View, expand the Patch category and select a software policy that is attached to the selected servers. Or, select the top level Patch category if you want to remediate all of the policies attached to the selected servers.
- 10 From the bottom of the Details pane, you have the following options:
 - **Remediate:** Remediate the selected patch policy or policies against the selected server or servers.

- **Scan Device:** Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. For more information, see [Scanning for Compliance](#) on page 229.

Remediating Patch Compliance — Groups

When you remediate patch policies for a group or multiple groups of servers, you can remediate all the policies attached to all servers in the group or multiple groups. However, when you select a group or multiple groups, you can only remediate all of the patch policies attaches to all the servers in the group and any sub-groups.

To remediate patch policies for groups or multiple groups of servers, perform the following steps:

- 1 To remediate patch policies for a single server in the Device Explorer, from the Navigation pane, select **Devices ► Servers ► All Managed Servers**.
- 2 From from the **Actions** menu, select **Open**.
- 3 From inside the Device Explorer's View pane, select **Management Policies ► Compliance**.
- 4 In the Details pane of the Compliance View, expand the Patch category and select an individual patch policy or the top level Patch category, which will enable you to remediate all of the patch policies attached to the server.

Or

- 5 To remediate patch policies for multiple servers, from the Navigation pane, select **Devices ► Device Groups** and select a group.
- 6 In the Content pane that shows a list of servers that belong to the group, select multiple servers by selecting the check box next to the server.
- 7 From the View drop-down list, select Compliance.
- 8 In the Details pane of the Compliance View, expand the Patch category and select a policy that is attached to the selected servers. Or, select the top level Patch category if you want to remediate all of the policies attached to the selected servers.
- 9 From the bottom of the Details pane, you have the following options:
 - **Remediate:** Remediate the selected patch policy or policies against the selected server or servers.
 - **Scan Device:** Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. For more information, see [Scanning for Compliance](#) on page 229.

Application Configuration Compliance

An application configuration (App Config) manages configuration files on a managed server. An application configuration can manage one or several configuration files for a server or group of servers. Each application configuration is made up of one or more templates which model an ideal configuration state for the fields and are targeted to manage configuration values (key-value pairs) for specific files on a server.

For example, you can create an application configuration that manages the hosts file for servers in your data center. You can define the IP address-hostname key-value pairs for a standard Unix hosts file, and then attach the application configuration to a several servers or a group of servers that contain the file. The application configuration serves as the policy that helps ensure that the hosts files on the target servers have the correct IP address-hostname definitions.

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. In the case of the hosts file example, if the information inside the hosts file on actual server configuration does not match the values defined in the application configuration, then the server's App Config is Non-Compliant. If more than one application configuration is attached to a server, and any one of the actual configuration files targeted by the application configuration is different, then the entire server is Non-Compliant in the Compliance View.

Conversely, if there are no differences found between the application configuration and the files on server, then the App Config compliance status is Compliant. All application configurations must be 100 percent compliant for the server's App compliance status to be considered Compliant in the Compliance View.

To check the latest state of a configuration file targeted by an application configuration, you can you perform an application configuration compliance scan to determine if there are any differences between the application configuration and the actual configuration files on the server.

For more information on running a compliance scan, see [Scanning for Compliance](#) on page 229.

For more information on creating and using Application Configurations, see the *SA Application Configuration User Guide*.

Application Configuration (App Config) Compliance Status

Application Configuration (App Config) compliance status is determined by the following criteria:

- **App Config Compliance — Single Server:** If any differences are discovered between the application configuration and the actual configuration file on the target server, the server's App Config compliance status is Non-Compliant. The Details pane of a server's Device Explorer window shows the App Config category as Non-Compliant. If the server has several application configurations attached to it, and any one of the actual configuration files targeted by the application configuration is different than the application configuration, then the entire server is considered Non-Compliant in the Compliance View.
- **App Config Compliance — Device Groups:** An Application Configuration attached to a group of servers is considered Compliant if more than five percent of the servers in the group attached to the Application Configuration have a status of Non-Compliant. If this is

the case, the aggregate compliance for App Config will display as Non-Compliant. Another way to state Non-Compliance for a group is when less than ninety five percent of the servers are Compliant.

However, if more than two percent but less than or equal to five percent of all servers in a group have the status of Non-Compliant for that category, then the status is Partial-Compliance. Another way to state Partial-Compliance for a group is when less than ninety eight percent but at least ninety five percent of the servers are Compliant.

If less than two percent of all servers in a group have a App Config status of Non-Compliant for that category, then the overall status is Compliant. Another way to state Compliance for a group is that at least ninety eight percent of the servers are Compliant for a given category.

The Details pane for a group of servers in the Compliance View shows whether or not all of the application configurations are compliant or not, but does not expand to show a breakdown of individual servers and application configurations.

You can modify the thresholds used to determine compliance for groups of servers. For more information, see [Changing Device Group Compliance Settings](#) on page 214.

Remediating App Config Compliance — Servers and Groups

Remediation for an application configuration is slightly different than the other compliance category types. Rather than remediating a policy onto a server, as you can with Software or Patch or Audits, to remediate an application configuration you select an application configuration from inside either the Device Explorer or Device Group Explorer and use the Push function to push the values defined in the application to onto the actual configuration files on the server (or group of servers).

When you push an application configuration, all values defined in the application configuration templates add or replace those on the target configuration files.



The manner in which some value in an application configuration get pushed — for example, sequences (of lists and scalars) — depends upon how those values have been set in the application configuration inheritance hierarchy and what sequence merge modes have been configured in the configuration template. For more information about sequence merging, see the *SA Application Configuration User Guide*.

To remediate application configurations server or group of servers, perform the following steps:

- 1 To remediate application configuration for a single server in the Device Explorer, from the Navigation pane, select **Devices ► Servers ► All Managed Servers**, then select a server.
Or
- 2 To remediate patch policies for a group of servers, from the Navigation pane, select **Devices ► Device Groups** and select a group.
- 3 From from the **Actions** menu, select **Open**.
- 4 From inside the Device Explorer's View pane, select **Management Policies ► Configurations**. Or, from a Device Group Explorer, select Configured Applications. For more information on how to configure and push application configuration values, see the *SA Application Configuration User Guide*.

4 SA Client Reports

Overview of SA Client 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 parameterized reports are presented in graphical and tabular format, and are actionable—which means that you can perform appropriate actions on objects, such as a policy or an audit, within the report. These reports are also exportable to your local file system (as .html, .pdf, or .xls files) to facilitate use within your organization.

This section contains information about the types of SA Client reports, how to modify report parameters, how to run the reports, and how to perform actions in the report results.

Reports Features

SA Client Reports enable you to perform enterprise health assessments by providing the following features:

- Actionable reports that enable you to take the appropriate action on objects within the reports. For example, in the list view of a compliance report, you can select a server and open a Remote Terminal or Server Explorer to browse it, perform an audit, create a snapshot, create a package, and so on.
- A single entry point in the SA Client Dashboard for all reports.
- Reports that are data-secured—controlled by the user's permissions. You can view all objects that you have read permissions for. You can perform actions on objects that you have write permissions for.
- Reports that are exportable to .html, .pdf, and .xls formats. You can export reports to your local file system for use within your organization.

HP Server Automation Client Reports

Table 14 lists the SA Client Reports by report folders.

Table 14 SA Client Reports

Report Folder	Report Title
Server Reports	Servers by Customer
	Servers by Facility
	Servers by Manufacturer
	Servers by Model
	Servers by Operating System
	Servers by Use
Virtualization Reports	Virtualization by Virtual Technology
	All Virtual Servers
	Solaris 10
	Virtual Servers by Hypervisors (zones only)
	Resource Allocation by Hypervisors (zones only)
	VMware ESX 3
	Virtual Servers by Hypervisors (VMs only)
	Resource Allocation by Hypervisors (VMS only)
User and Security Reports	Client and Feature Permissions
	Customer/Facility Permissions and Device Group Permission Overrides
	User Groups Memberships
	User Login
	Administrator Actions
	Users and Authorizations, By User Group
	Users and Authorizations, By Individual User Group
	Administrator Customer Groups
	Server Permissions, By User
	Server Permissions, By Server
	OGFS Permissions, By User
	OGFS Permissions, By Server

Table 14 SA Client Reports (cont'd)

Report Folder	Report Title
Network Reports	Connections by Network Device
	Connections by Server
	Duplex Compliance (All Servers)
	Duplex Compliance by Customer
	Duplex Compliance by Facility

See the following documentation for more information about the SA Client features that support information in these reports:

- [Software Management](#) on page 257
- [Audit and Remediation](#) on page 153
- [Patch Management for Windows](#) on page 287
- [Patch Management for Unix](#) on page 395
- *The SA Integration Guide*
- “Server Management in SAS Web Client” in the *SA Users Guide: Server Automation*

User Permissions

Reports are controlled by the user’s permissions. You can view all objects that you have read permissions for, and you can perform actions on objects that you have write permissions for.

To view or run a network report, SA/NA integration must be configured. See the *SA Integration Guide*.

To view or run a user and security report, system administrator permissions are required.

Launching the Reports Feature

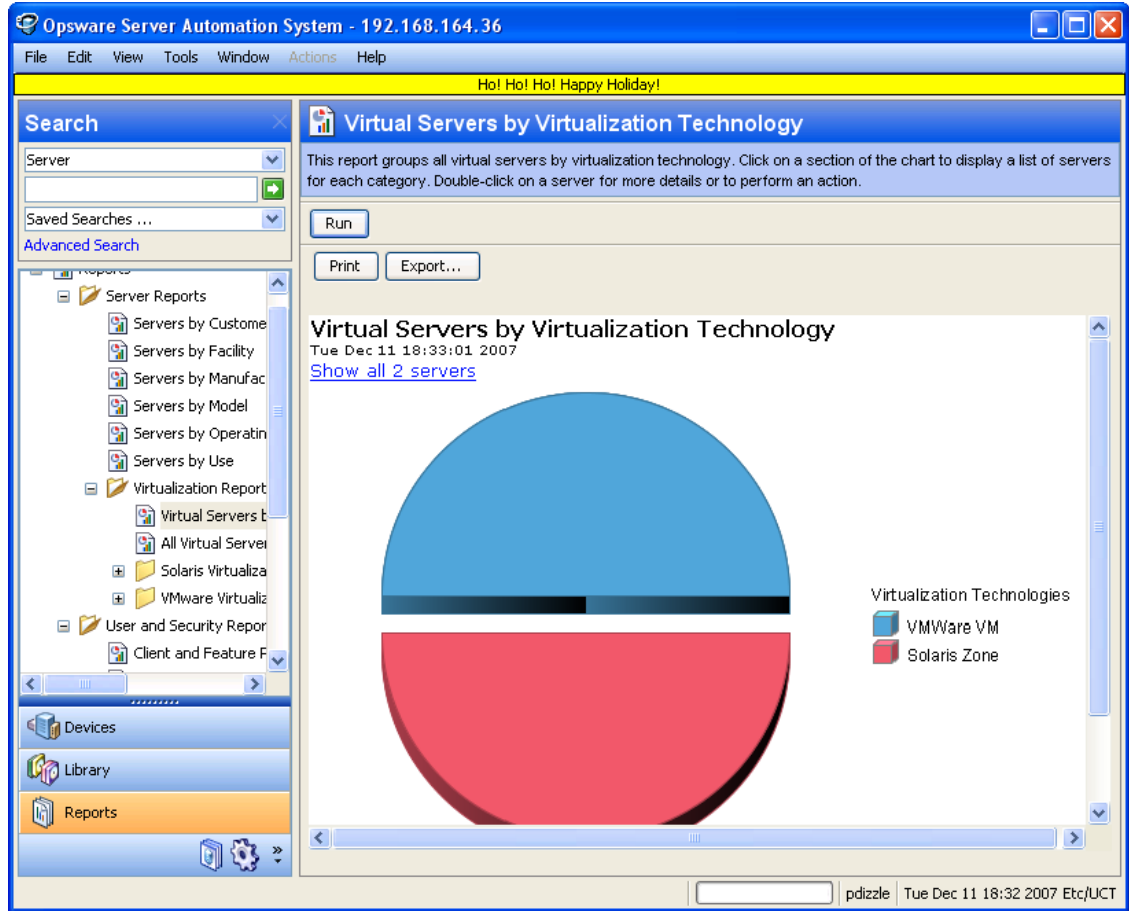
To launch the Reports feature, perform one of the following steps:

- From the **View** menu, select **Reports ► Dashboard**.
- From the **View** menu, select **Reports ► Reports**.
- From the Navigation pane, select Reports.

Reports Display

The Reports feature display consists of a Search pane, report parameters, report folders, and other filtering tools.

Figure 48 The Reports Feature Display



Search Pane

In the Reports feature, you can use the SA Client Search feature to find reports by defining specific filter criteria. See “SA Client Search” in the *SA Users Guide: Server Automation*.

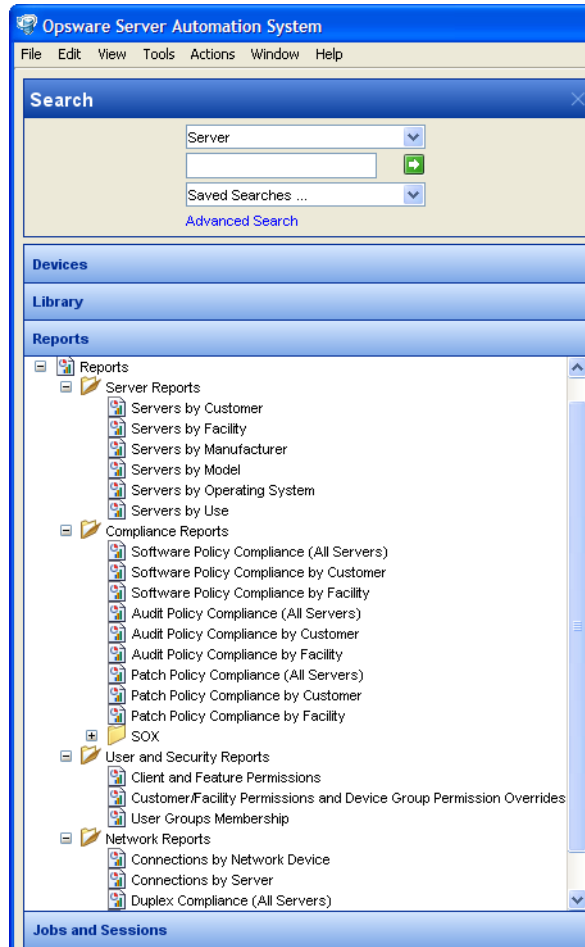
Report Folders

Reports are organized in the following folders according to regulatory or IT best practice standards:

- **Server Reports:** This folder contains reports about servers by customer, facility, manufacturer, model, operating system, and server usage.
- **Compliance Reports:** This folder contains reports about compliance for software policies, audit policies, and patch policies by servers, customer, and facility.
- **SOX Reports:** This folder contains reports about compliance standards based on Sarbanes-Oxley, including the COSO process model and the CobiT control model.
- **Network Reports:** This folder contains reports about connections and duplex compliance for network devices and servers. You must have NA installed to see this folder in the Navigation pane.

- **User and Security Reports:** This folder contains reports about client and feature permissions; customer, facility, and device group permissions; and user group memberships. You must have system administrator permissions to see this folder in the Navigation pane.
- **Custom Reports:** This folder contains any custom reports you have created.
- [Figure 49](#) illustrates the Report folders in the Navigation pane, including the reports you will find in each folder.

Figure 49 Report Folders



Report Parameters

Many reports require input parameters in order to be run. For reports that require parameters, you can run the report with its default parameter values or modify the parameter values. If you want to run a report that includes or excludes certain servers, customers, or hardware models, you need to specify this criteria in the report parameters. See [Running a Report](#) on page 252.

Running a Report

To run a report, perform the following steps:

- 1 From the Navigation pane, select Reports.
- 2 Expand the Reports folder and then expand the Server Reports and Virtualization Reports.
- 3 Select one of the virtualization report listed in the folder.
- 4 If there are no report parameters in the Content pane, click **Run**.
- 5 If there are report parameters in the Content pane, you can either use the default parameters or change them:
 - To use the default report parameters, click **Run** to run the report.
 - To change the report parameters, see [Modifying Report Parameters](#) on page 252.

Modifying Report Parameters

To modify the default parameters and run a report that includes certain servers, customers, hardware models, and so on, perform the following steps:

- 1 In the drop-down list for (the Server, Customer, Model, and so on), select Contains, Equals, Begins With, or Ends With.
- 2 (Optional) Select the ellipsis button to open the Select Values window.
- 3 In the Select Values window, select a value in the Available or Selected pane and then use the directional buttons to include it in or exclude it from your search criteria.
- 4 Click **OK** to save your changes.
- 5 Click **Run** to run the report.



If data cannot be found to run the report, a “No records to display!” error displays.

Report Results Restriction

The following reports have a limit of 2000 “items” that can be displayed in their results:

- Server Permissions By Server
- Server Permissions By User
- OGFS Permissions By Server
- OGFS Permissions By User

In these reports, if the results reach 2000, the report will stop, because depending on the specified search parameters, they can yield thousands of results and slow performance of the Opware core.

For example, the Server Report by User will run successfully if you specify 10 users and 200 servers in the search parameters, but will not run if you specify 10 users and 201 servers.

To avoid this problem, either modify your search parameters to yield less results, or break the report query into smaller searches and run as many smaller reports as you need to achieve your results.

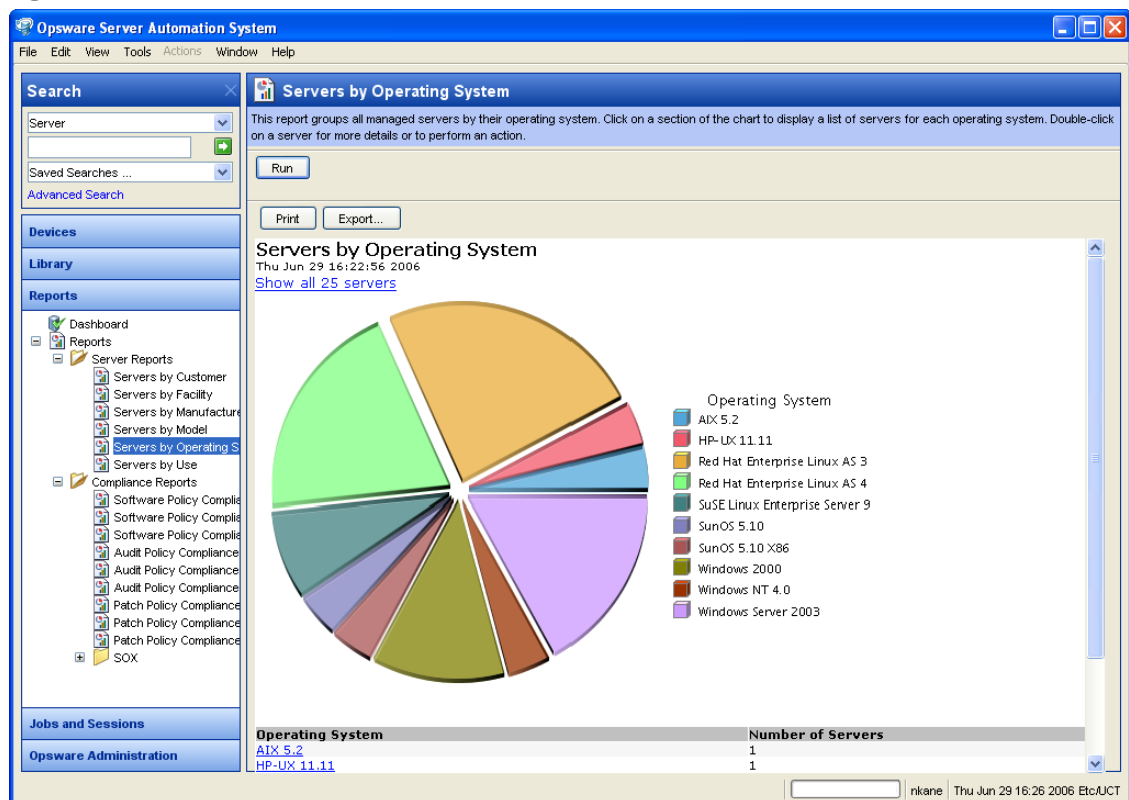
Report Results

Report results initially appear in a graphical or list view. The graphical report is an overview of available data for this report displayed in a pie chart or in a bar graph. You can drill down for more detail in the chart or graph by clicking on any of the sections or bars. For example, you can drill down to individual servers that appear in a report and get detailed information about them.

Graphical Report

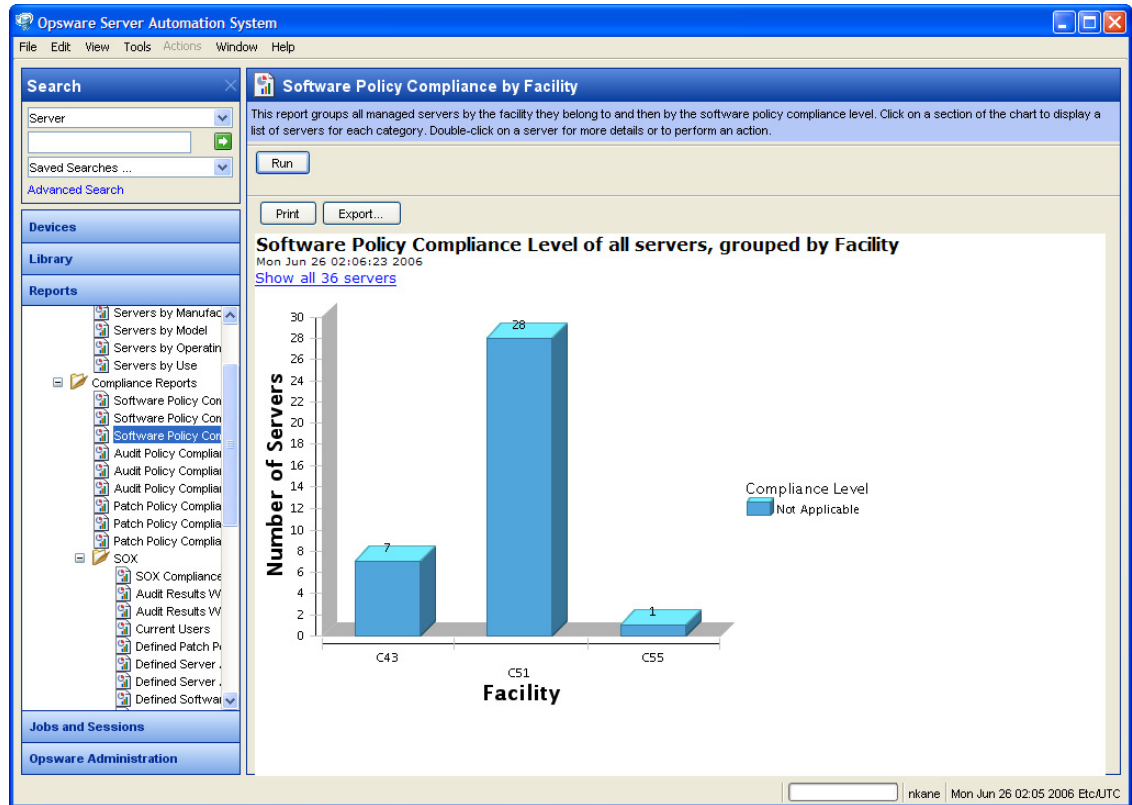
A graphical report is a pie chart or a bar graph. Click on a section of the chart or graph to drill down for more details or to perform an action. You can also click on the “Show all <number> servers” link to display a list of servers. See [Figure 50](#) and [Figure 51](#) for examples.

Figure 50 Pie Chart



To display the corresponding list view of a bar graph, click on the front part of the bar. Do not click on the top or shaded part of the bar.

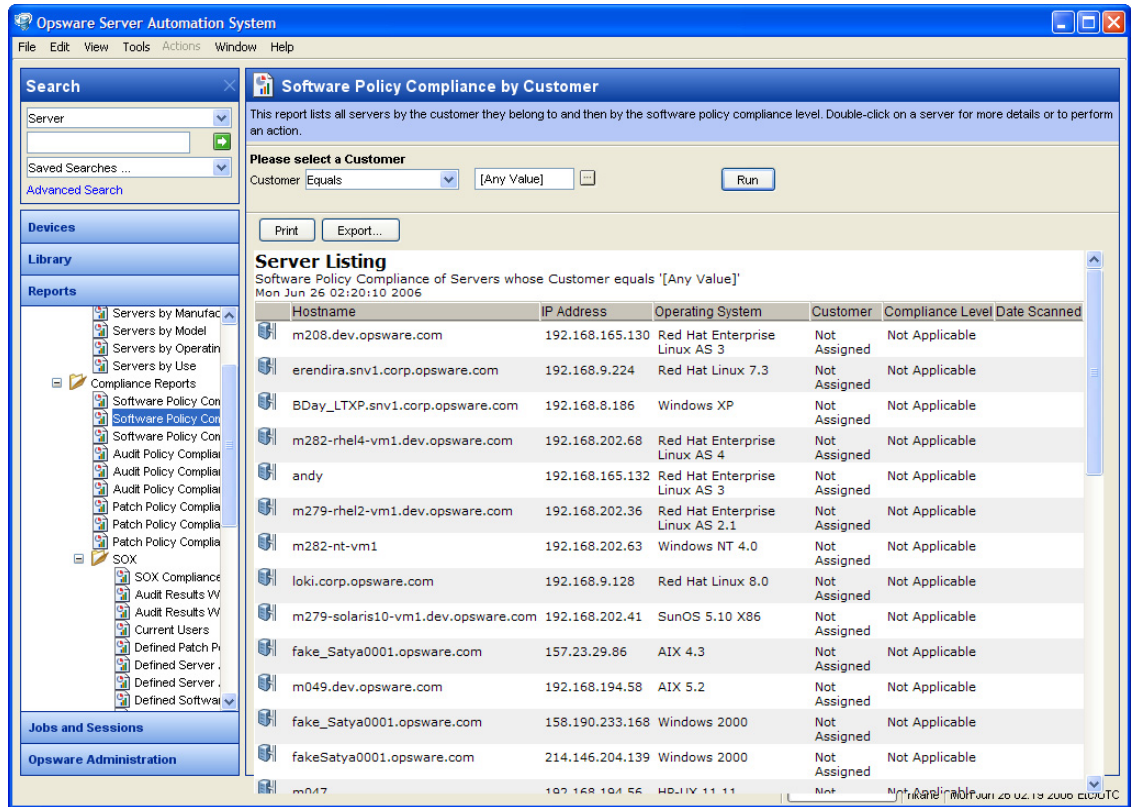
Figure 51 Bar Graph



List Report

A list report is a tabular display of information. Double-click on a row in the list, such as a server, audit, or policy, for more detail or to perform an action. See [Figure 52](#) for an example.

Figure 52 List Report



Report Results Restriction

The following reports have a limit of 2000 “items” that can be displayed in their results:

- Server Permissions By Server
- Server Permissions By User
- OGFS Permissions By Server
- OGFS Permissions By User

In these reports, if the results reach 2000, the report will stop, because depending on the specified search parameters, they can yield thousands of results and slow performance of the SA core.

For example, the Server Report by User will run successfully if you specify 10 users and 200 servers in the search parameters, but will not run if you specify 10 users and 201 servers.

To avoid this problem, either modify your search parameters to yield less results, or break the report query into smaller searches and run as many smaller reports as you need to achieve your results.

Exporting a Report

You can export a report for use in other applications in your environment and attach a report for email distribution. Depending on the report format, you can export a report to your local file system in either .html, .pdf, or .xls file formats. You can export a graphical report to .html or .pdf only. You can export a list report to either .html, .pdf, or.xls file formats.



When you export a report in the SA Client, the time that you will see marked on the exported report will be the time when the report was exported, not the time when the report was generated.

To export a report, perform the following steps:

- 1 From the report, click **Export** to open the Save window.
- 2 In the Save in field, enter a location that identifies where you want to save the file to, or select from the drop-down list.
- 3 Enter a file name.
- 4 Select the file type.
- 5 Click **Save**.

Printing a Report

To print a report, perform the following steps:

- 1 From the report, click **Print** to open the Print window.
- 2 Use the default print options or modify them, and then click **OK**.

5 Software Management


SA Software Management is a powerful feature that allows you to model enterprise software using *Software Policies* which automate the process of installing software and configuring applications on SA Managed Servers. You can create software policies and enforce their application against SA Managed Servers. You can identify SA Managed Servers that are non-compliant with the software policies you have created and remediate those non-compliant servers to bring them into compliance.

SA Software Management further provides an organizational structure that allows you to distribute your software resources in folders and use folder permissions to protect and control access to those resources.

Policy-Based Software Management

One of the most difficult problems in the IT organizations today is the deployment of software to servers running in the operational environment. Many IT organizations define policies for creating systems in their operational environment and set standards for server configuration. Often, policies are written by subject-matter experts, like the application developer or security administrator in an IT organization, or policies are leveraged from industry standards organizations as shown in [Figure 53](#).

Figure 53 Defining Policies for Software Deployment



Clearly defining policies is the correct approach to managing software deployment. However, these policies are typically static and unenforceable because the policy creators have no means to enforce their policies. SA solves the software deployment problem by introducing the Software Management feature in the SA Client.

SA Software Management allows IT policy setters to enforce policies, standardize operations, and ensure compliance against defined software policies. To do so, policy setters use the SA Client to define software policies that create a model of their IT environment. Defining a software policy is similar to defining a server *baseline* that ensures that all servers are provisioned in a standardized way with approved content.

A software policy specifies the software to be installed, the configurations to be applied, the scripts to be run during installation, and more to the Managed Servers in an IT environment. System administrators can then manage the servers in their environment by attaching and remediating software policies to the servers. *Attaching a software policy* to a Managed Server associates that policy with the server. A Managed Server must be in compliance with its attached software policies. *Remediation* applies software policy specifications to the Managed Server.

Policy setters can modify a software policy simply by changing the baseline defined in the policy. When an existing software policy that is attached to Managed Servers is modified, those servers will become non-compliant. A Software Policy Compliance Scan identifies the non-compliant servers which can then be remediated against the modified software policy to bring them back into compliance.

This automated and systematic method for keeping large numbers of servers in compliance eliminates the need to store and manage hundreds of rigid images and provides the means to easily restore or rollback servers to a previous working state.

SA Software Management Features

SA Software Management provides the following:

- **Organizational structure for software**

The SA Software Management folder hierarchy provides a way to organize your software resources. Folders act as containers for packages, scripts, software policies, OS sequences, and server objects.

- **Security boundaries for folders**

Folders allow you to define security permissions to control access to their contents across users and user groups. You can set folder permissions to determine which user groups can view, use, and modify items within a folder.

- **Model-based approach to manage the IT environment**

SA Software Management enables a policy setter to create a model of their IT environment by defining software policies. A software policy specifies the packages, patches, scripts, and server objects to be installed, and the configurations to be applied to the managed servers in their IT environment. A system administrator can then bring the Managed Servers in their environment into compliance by applying the software policy to the servers (remediation).

- **Sharing of software resources among user groups**

A software policy can specify various software resources that may be managed by different user groups and located in different folders, thus allowing software resources to be shared across different groups.

- **Simultaneous installation and configuration of applications**

Software policies can be applied to multiple Managed Servers in a single step.

- **Deployment of multiple application instances on a single server**

SA Software Management can install multiple instances of an application on a server by using relocatable ZIP packages.

- **Deployment of application installation media on Managed Servers**

SA Software Management can import a software application provided by a software vendor into SA and deploy that application to a Managed Server.

- **Precise control of the software installation process**

SA Software Management allows you to separate the different stages (analysis, download, and installation) of the software installation process. You can independently schedule the various stages of the software deployment process. You can be notified of job status upon successful completion of a stage by email and associate a Ticket ID with each job.

- **Flexible software installation/uninstallation**

SA Software Management allows you to install and uninstall software with or without a software policy.

- **Software policy compliance scan**

The Compliance View (for individual servers or groups of servers) allows you to view the compliance state of Managed Servers to determine compliance or non-compliance with attached software policies.

- **Reporting**

The Reporting feature allows you to generate reports that provide a summary of software policy compliance across servers. You can generate reports that provide information about software policies on a specified server. These reports can be exported to .html and .xls formats.

- **Comprehensive search for software resources and servers**

The SA Client search functionality allows you to search for Managed Servers, software policies, folders, application configurations, patches, and software, and perform actions on the search results.

Overview of Software Policies

A software policy defines the software configuration that must be applied to any Managed Server that has been attached to the policy. It can specify packages, RPMs, patches, application configurations, scripts, and server objects. Software policies can also specify how software installation/uninstallation is ordered, when reboots are allowed, and more.

Software policies can also be associated with OS Provisioning *OS Sequences* which allows you added control over the manner in which a particular OS should be installed, including the proper OS Installation Profile to be used, Application and Patch policies, and how these policies should be remediated either before or after the OS is installed.

Attaching a Software Policy to Managed Servers or Device Groups

After creating a software policy, you can attach it to Managed Servers or Device Groups. When you remediate a server or Device Group, the patches, packages, RPMs, scripts, server objects, and application configurations specified in the attached software policy are automatically installed and applied.

When you attach a software policy to a Managed Servers or Device Group, the policy has a *persistent* association to those servers. Therefore, whenever the software policy is updated, you receive a notification indicating which servers or groups of servers are affected by the

updated software policy. You can then choose to remediate the servers or groups of servers to reflect the changes to the software policy. See the [Attach a Software Policy to a Server](#) on page 268 for more information.

Remediation

Remediation compares the software that is actually installed on a Managed Server to the software that should be installed on the server as specified in a software policy. SA then determines what operations are required to make the server compliant with the software policy. See [Remediating Policies](#) on page 271 for more information about the remediation process.

Software Templates

SA Software Management also provides the *Software Template* which is a software policy that specifies only other software policies (sub-policies). A software template gives you the option of installing software and configuring applications simultaneously without *persistently* associating a software policy with a Managed Server(s) or Device Group. Note that if a sub-policy specified in a software template is updated, the changes are not reflected on the Managed Servers or Device Groups that template is attached to. See the [Overview of Software Templates](#) on page 277 for more information about installing software using a software template.

A software template can be associated with either a single operating system family or multiple operating system families. When you add software resources to a software template, the software resources must belong to the same operating system family as the software template. For example, if you define the operating system for a software template as HP-UX, you can only add software resources applicable to versions of HP-UX.

Managing Software Resources Using a Software Policy

A software policy can specify packages, RPMs, patches, application configurations, scripts, and server objects (Windows COM+, Windows Services, windows Registry, Unix Users and Groups, Windows Users and Groups, Local Security Settings, .NET Framework Configurations). See the *SA Users Guide: Server Automation* for more information about server objects.

A software policy can also specify other software policies (sub-policies). The specified child sub-policies and the parent software policy must belong to the same operating system family. This is not to be confused with a *software template* which specifies *only* sub-policies.

After you specify software resources in a software policy, you can specify the order in which you want them to be installed. When you attach a software policy to a Managed Server and remediate the server, SA installs the software resources in the software policy in the specified order.

When a software policy specifies sub-policies, all the software resources from the sub-policies are grouped together and then installed as a unit. Thus, policy inclusion provides a way to organize your software and manage dependencies between the software resources across sub-policies.



You must have permissions to set the installation order of software resources in a software policy. To obtain these permissions, contact your SA administrator. See Appendix A of the *SA Administration Guide* for more information.



If a software policy specifies sub-policies then during remediation, SA does not consider the install order specified in the sub-policies, it only considers its own install order.

The Software Management Process

SA software management can install/uninstall software in two ways:

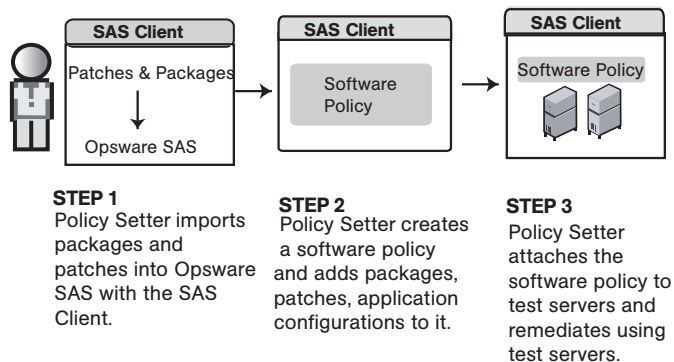
- Directly on a Managed Server using the SA Client Install or Uninstall Windows
- Using a Software policy that is attached to Managed Server(s) or Device Group(s) and then remediated against the server(s) or group(s)

Figure 54 shows the software policy process. This process can include such tasks as creating and attaching software policies, remediating software policies, running software compliance scans to identify and remediate non-compliant servers, and generating software compliance reports.

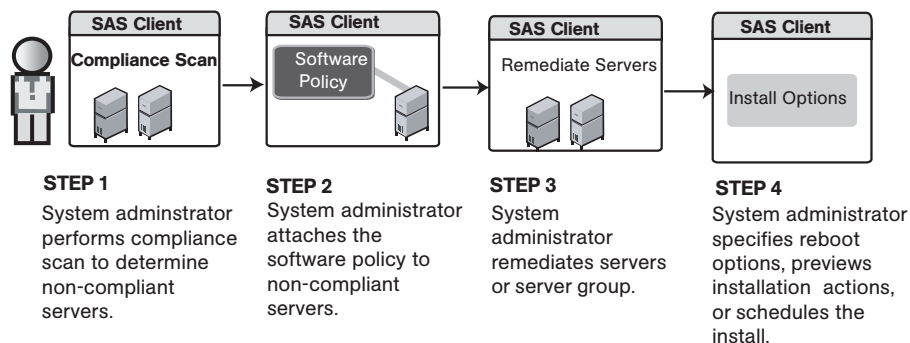
Figure 54 Software Management Process – Software Policy Remediation

SOFTWARE MANAGEMENT PROCESS (REMEDiate)

Part A: Set Up Software Policies



Part B: Attach Software Policies to Servers and Remediate



Installing/Uninstalling Software without a Software Policy

This section describes installing/uninstalling software without the use of a software policy. For information about using software policies, see [Install Software Using a Software Policy](#) on page 267



You must have the *Allow Install / Uninstall Software* permission. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide* for more information.

The Install or Uninstall Software Windows

In the SA Client, you can access the *Install Software* or *Uninstall Software* windows in several ways.

From the Server List:

From the SA Client Navigation pane:

- 1 select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.
or
select **Devices** ► **Device Groups**. The device group list appears in the Content pane.
- 2 From the Content pane, select a server or device group.
- 3 From the **Actions** menu:
select **Install** ► **Software**. The *Install Software* window appears.
or
select **Uninstall** ► **Software**. The *Uninstall Software* window appears

From the By Type View in the Library:

- 1 From the SA Client Navigation pane, select **Library** ► **By Type** ► **Type of Software**. The Software List appears in the Content pane.
- 2 From the Content pane, select the software you intend to install.
- 3 From the **Actions** menu, select **Install Software**. The *Install Software* window appears.

From the SA Client Server Explorer

- 1 From the SA Client Navigation pane:
select **Devices** ► **Servers** ► **All Managed Servers**. The Server List appears in the Content pane.
or
select **Devices** ► **Device Groups**. The Device Group List appears in the Content pane.
- 2 From the Content pane, select a server from the Server List or from the Device Group list.
- 3 From the **Actions** menu, select **Open**. The *Server Explorer* window appears.

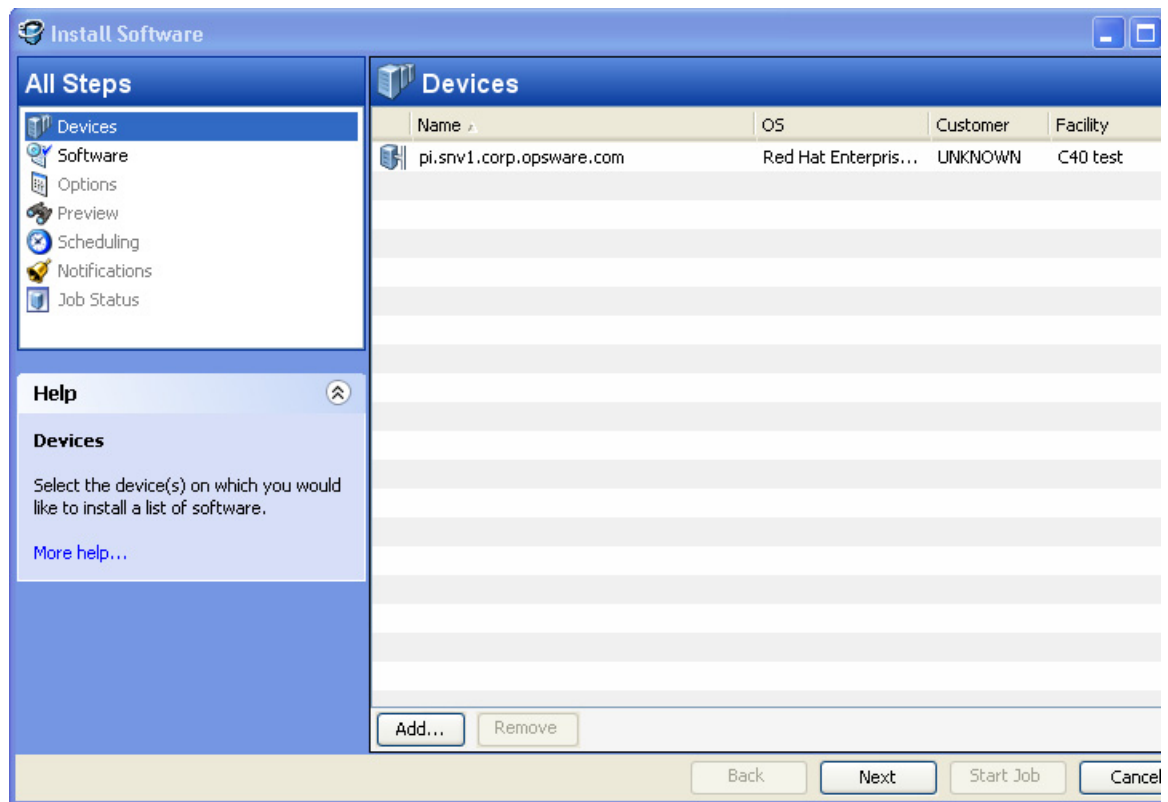
- 4 From the Views pane, select **Inventory** and then select the type of software.
- 5 From the **Actions** menu:
select **Install Software**. The *Install Software* window appears.
or
select **Uninstall Software**. The *Uninstall Software* window appears.

Install or Uninstall Software

The *Install Software* window as shown in [Figure 55](#) allows to you install software directly on an SA Managed Server and provides the following options:

- [Select Devices](#)
- [Select Software](#)
- [Specify Options](#)
- [Preview](#)
- [Schedule Stages](#)
- [Setting Email Notifications](#)
- [Viewing Job Status](#)



Figure 55 Install Software Window



Select Devices

You can specify the servers on which to install or uninstall software.





Perform the following tasks to select the servers:

- 1 Open the *Install Software* window or *Uninstall Software* window using one of the methods described in [The Install or Uninstall Software Windows](#) on page 262
- 2 In left panel of the *Install Software Policy* window, select **Devices**.
- 3 (Optional) Click  to add additional servers to the list. Or, to remove a server, select the server in the list and then click .
- 4 In the Select Servers and Device Groups window, select and add servers.
- 5 Click **Next** to proceed to the Select Software step.

Select Software

Specify the software (packages, RPMs, patches, etc.) to install or uninstall. You can also specify the order in which you want to install or uninstall the software.

Perform the following tasks to select the software:

- 1 In the *Install Software* window or *Uninstall Software* window, click . The **Select Library** window appears.
- 2 In the **Select Library** window, highlight the software to be installed or uninstalled and click **Select**.
- 3 (Optional) If you need to remove any of the software you have added, highlight the software and click  to remove it.
- 4 (Optional) Click  or  to reorder the software in the list.
- 5 Click **Next** to proceed to the Specify Options Step.

Specify Options

The following installation or uninstallation options are available:

- The reboot actions required for the installation or uninstallation process. You can control when to reboot servers during installation or uninstallations to minimize the downtime caused by server reboots.
- The option to continue installation or uninstallation when an error occurs.
- The scripts to run on a server before or after installation or uninstallation. The scripts include:
 - **Pre-Download:** (*Installation Only*) A script that runs *before* software or patches are downloaded from the Software Repository to the Managed Server.
 - **Post-Download:** (*Installation Only*) A script that runs *after* software or patches are downloaded from the Software Repository to the Managed Server but before the software or patch is installed

- **Pre-Install/ Pre-Uninstall:** A script that runs *before* software or patches are installed or uninstalled.
- **Post-Install/ Post-UnInstall:** A script that runs *after* software or patches are installed or uninstalled.

Reboot Options

Perform the following tasks to specify the options for installation or uninstallation:

- 1 Select one of the following Reboot options:

- **Reboot servers as specified by individual software items**

This option allows you to reboot servers depending on the reboot option specified in the software resources properties window.

- **Reboot servers after each installation or uninstallation**

This option allows you to reboot servers after installing or uninstalling software.

- **Hold all server reboots until all actions are complete**

If the *reboot option is selected* in the software resources properties, this option *reboots the servers after* all the software resources are installed and uninstalled. If the *reboot option is not selected* in the software resources properties, this option *does not reboot the server after* all the software resources are installed and uninstalled.

- **Suppress all reboots**

This option allows you to suppress the reboots even if the reboot option is selected in the software resources properties.

Continue on Error

- 2 Select **Attempt to continue running if an error occurs** if you want the installation or uninstallation process to continue even when an error occurs with any of the software, patches or scripts.

Run Scripts

- 3 In the Scripts section, select the **Pre-Download**, **Post-Download** or **Pre-Install**, **Post-Install** tab. You can specify different scripts and options on each of the tabs. You must have certain *Script permissions* in order to enable these options. See the *SA Administration Guide* for more information about the permissions required.
 - a On each tab, you can select **Enable Script**. Selecting Enable Script enables the remainder of the fields on the tab. Enable Script must be selected for a script to run.
 - b Select **Saved Script** or **Ad-Hoc Script** from the drop-down list. A Saved script is stored for future use after you upload the script to SA. An Ad-Hoc script must be entered manually and is intended only for a single operation and is not stored in SA.
 - If you select Saved Script from the drop-down list, click **Select** to specify the script. The **Select Script** window appears. Select the script(s) to run and click **Select**.
 - If you selected Ad-Hoc Script from the drop-down list, select the type from the Type drop-down list and then enter the script content in the Script field.
 - c Enter any *command-line flags* in the **Command field** as required.
 - d Enter a script *time-out value* in minutes in the **Script Timeout** field.
 - e In the **User section**, select *Root* to execute the script as root. To execute the script as a specified user, select *Name* and enter the user name and password.

To enter a *Windows Domain Name* in the pre-download, post-download, pre-install, post-install scripts, use the following format in the **Name** field:

DomainName\UserName

and enter the password in the password field.

f Select *Stop job if script returns an error* to stop the installation if the script returns an error.

4 Click **Next** to proceed to the Preview step.

Preview

You have the option to preview the installation or uninstallation process.

The Preview option allows you to view a detailed list of actions that will be performed on a server as a result of software installation or uninstallation. It displays information for each server that is selected for installation or uninstallation.

Preview displays:

- the software resources that will be installed on or uninstalled from a Managed Server
- the application configurations that will be applied to the server
- the dependency information required for the software or patches
- any reboots required during the installation or uninstallation process
- the scripts that will be executed.

If you select an object that has other software dependencies, during the remediation preview you may see other objects listed (such as packages, ZIP files, and so on).

Perform the following steps to preview the installation or uninstallation process:

- 1 Click **Preview** to view the actions that will be performed during the installation or uninstallation process. To view the details of each of the actions, select a row in the table. The details for each action appear.
- 2 Select **Output** to preview the job output or select **Errors** to preview possible errors.
- 3 Click **Next** to proceed to the Scheduling step.

Schedule Stages

Schedule the analysis, download and installation or uninstallation stage to be run immediately or at a specified date and time.

Perform the following tasks to schedule the installation or uninstallation process:



- 1 In the *Schedule Analysis* section, select one of the following options:
 - **Run at Job Start:** If enables, run the job immediately.
 - **Start time:** Specify a later date and time to schedule the job.
 - **Use Preview Results:** Use Preview results. This option is available only if you select to run a Preview.
- 2 In the *Schedule Download* section, select one of the following options: (*Installation Only*)
 - **Run Immediately After Analysis:** Download software immediately.
 - **Start time:** Specify a later date and time to download software.

- **Use Preview Results:** Use preview results. This option is available only if you select to run a Preview.
- 3 In the *Schedule Install* or *Schedule Uninstall* section, select one of the following options:
 - **Run Immediately:** Install or Uninstall software immediately.
 - **Start time:** Specify the date and time to install or uninstall software.
 - **Use Preview Results:** Use the preview results if you have run a preview. This option is available only if you select to run a Preview.
 - 4 Click **Next** to proceed to the Email Notifications step.

Setting Email Notifications

Set email notifications to alert users on the success or failure of the installation or uninstallation process. You can associate a **Ticket ID** with the installation or uninstallation process.

Perform the following tasks to specify email notifications:

- 1 Click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 2 To trigger notification on the *success* of a Job, select the  icon.
 To trigger notification on the *failure* of a Job, select the  icon.
- 3 Enter a *Ticket ID* to be associated with a Job in the **Ticket ID** field.
- 4 Click **Next** to go to the Job Status display.

Viewing Job Status

View the summary information about the progress of a job. You can also view the status of each action required complete the job. If you chose to run the job immediately in the Scheduling step, the job begins immediately and will appear in the Job Status view. If you scheduled the job for a later time, the job will run at the scheduled time and only then will appear in the Job Status view.

Perform the following tasks to view the Job Status:

- 1 The job progress appears in the *Install Software* window.
- 2 To view the details of each action, select a row in the table. The details for each action appear.
- 3 Select **Output** to view the job output or select **Errors** to view the error details.
- 4 You have the option to click **End Job** to stop the job or click **Close** to close the Install Software window.

Install Software Using a Software Policy

This section describes installing software on a Managed Server using a software policy. For information about installing software without a software policy, see [Installing/Uninstalling Software without a Software Policy](#) on page 262. Software policies are typically created by an SA Administrator. While users can attach software policies to Managed Servers and Device

Groups and Managed Servers and Device Groups to policies, and remediate those policies, they usually do not have the ability to modify the policy. See the *SA Policy Setter's Guide* for information about creating, removing, and modifying software policies.

Using a software policy to install software has two phases:

- Attach a software policy to a server or attach a server to a software policy
- Remediate a server against a software policy

Attach a Software Policy to a Server

When you attach a software policy to a Managed Server or Device Group, the software policy is only associated with that server or group. Simply attaching a policy to a server does not install the software specified in a software policy, you must *remediate* the server with the software policy in order to install the software. See [Remediating Policies](#) on page 271.

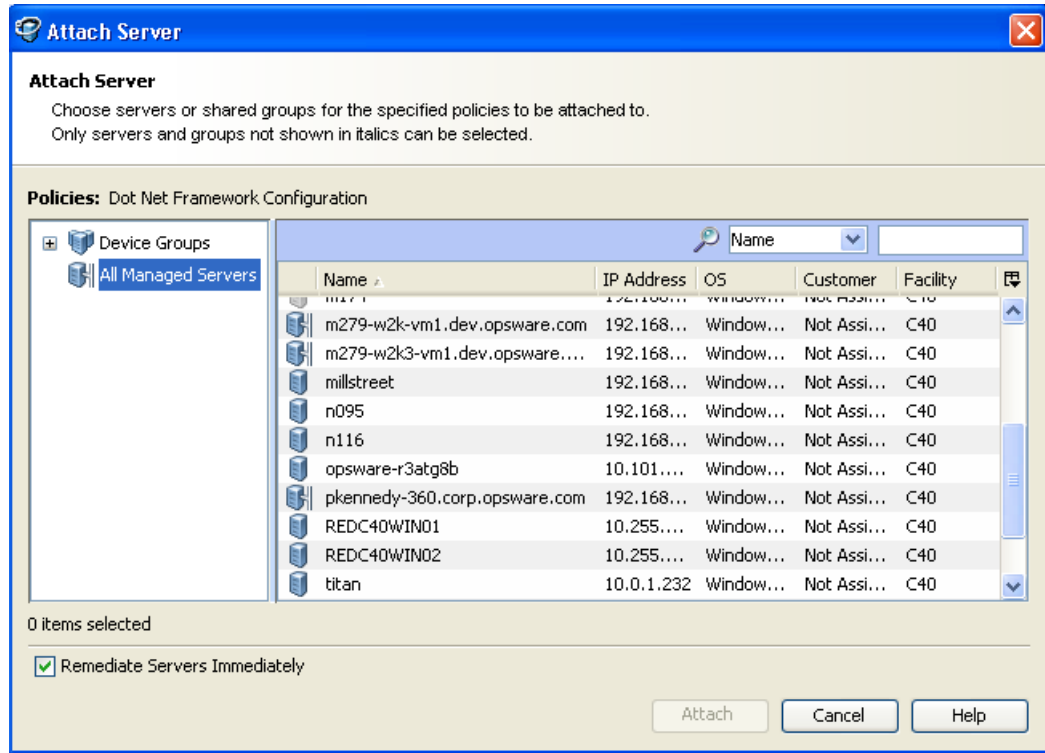


You must have the required permissions to attach a software policy to a server. To obtain these permissions, contact your SA Administrator. See Appendix A of the *SA Administration Guide* for more information about permissions.

Perform the following tasks to attach a software policy to a server:

- 1 From the SA Client Navigation pane, select **Library ► By Type ► Software Policies**. A list of available software policies appears in the Content pane.
- 2 From the Content pane, select a software policy.
 - a Open the software policy. The Software Policy Window appears.
 - b From the View pane, select **Server Usage**.
 - c From the Content pane, select a Managed Server from the Server List.or
 - a From the **View** drop-down list in the Content pane, select **Server Usage**.
 - b Select a Managed Server from the Server List.
- 3 From the **Actions** menu, select **Attach Server**. The **Attach Server** window appears as shown in [Figure 56](#):

Figure 56 The Attach Server Window in the SA Client



- 4 (Optional) Enable **Remediate Servers Immediately** to remediate the servers against the software policy. This option is only available if you have the Remediate Servers permission. See [Remediating Policies](#) on page 271 in this chapter for more information.
- 5 From the list in the left panel, select *All Managed Servers* or *Device Groups*, select a Managed Server and click the **Attach** button. Selecting this option displays the *Remediate* window. (You can only select servers that are not in *italics*. Servers in italics indicate that you do not have the necessary permissions to attach a software policy to that server.)
- 6 In the Remediate window, follow the steps to immediately run remediation or schedule it for a later time.

Attach a Server to a Software Policy

When you attach a Managed Server or Device Group to a software policy, the software policy is associated with that server or Device Group. This action does not install the software specified in the software policy. To install the software, you must remediate the Managed Server or Device Group with the software policy. See [Remediating Policies](#) on page 271.



You must have the required permissions to attach a server to a software policy. To obtain these permissions, contact your SA Administrator. See Appendix A of the *SA Administration Guide* for more information.

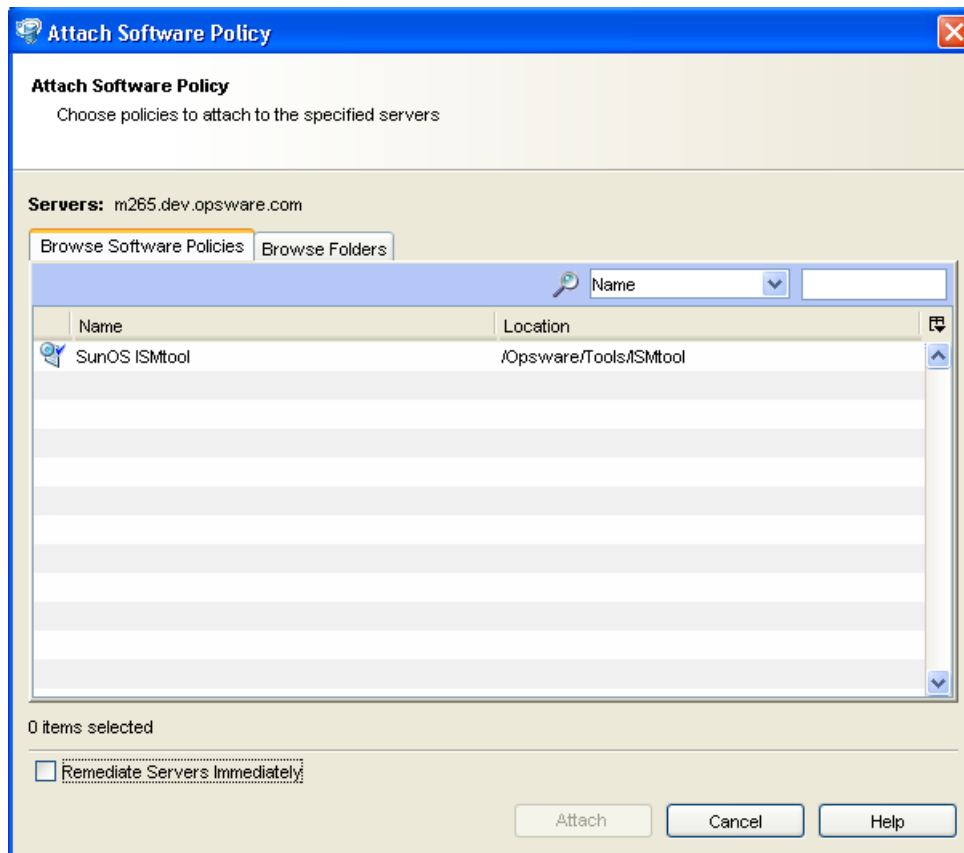
- 1 From the SA Client Navigation pane:
select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane

or

select **Devices** ► **Device Groups**. The device group list displays in the Content pane.

- 2 From the Content pane, select a Managed Server or a Device Group.
- 3 From the **Actions** menu, select **Attach** ► **Software Policy**. The Attach Software Policy window appears as shown in [Figure 57](#).

Figure 57 The Attach Policy Window in the SA Client



- 4 (Optional) Select **Remediate Servers Immediately** to remediate the servers against the software policy. This option is only available if you have the Remediate Servers permission. See [Remediating Policies](#) on page 271.
- 5 Select **Browse Software Policies** and select the software policies from the list.
or
select **Browse Folders** and then select the software policies from the folder hierarchy.
- 6 Click **Attach**.
- 7 In the Remediate window, follow the steps to immediately run remediation or schedule it for a later time.

The Remediation Window

The *Remediate Window* is where you remediate the servers against software policies and define the conditions for remediation.

Open the Remediate Window

There are several ways to open the Remediate Window:

From the Server List:

- 1 From the SA Client Navigation pane:
select **Devices** ► **Servers** ► **All Managed Servers**. The Server List appears in the Content pane.
or
select **Devices** ► **Device Groups**. The Device Group list appears in the Content pane.
- 2 From the Content pane, select a Managed Server or Device Group.
- 3 From the **Actions** menu, select **Remediate**. The Remediate window appears.

From the Software Policies List:

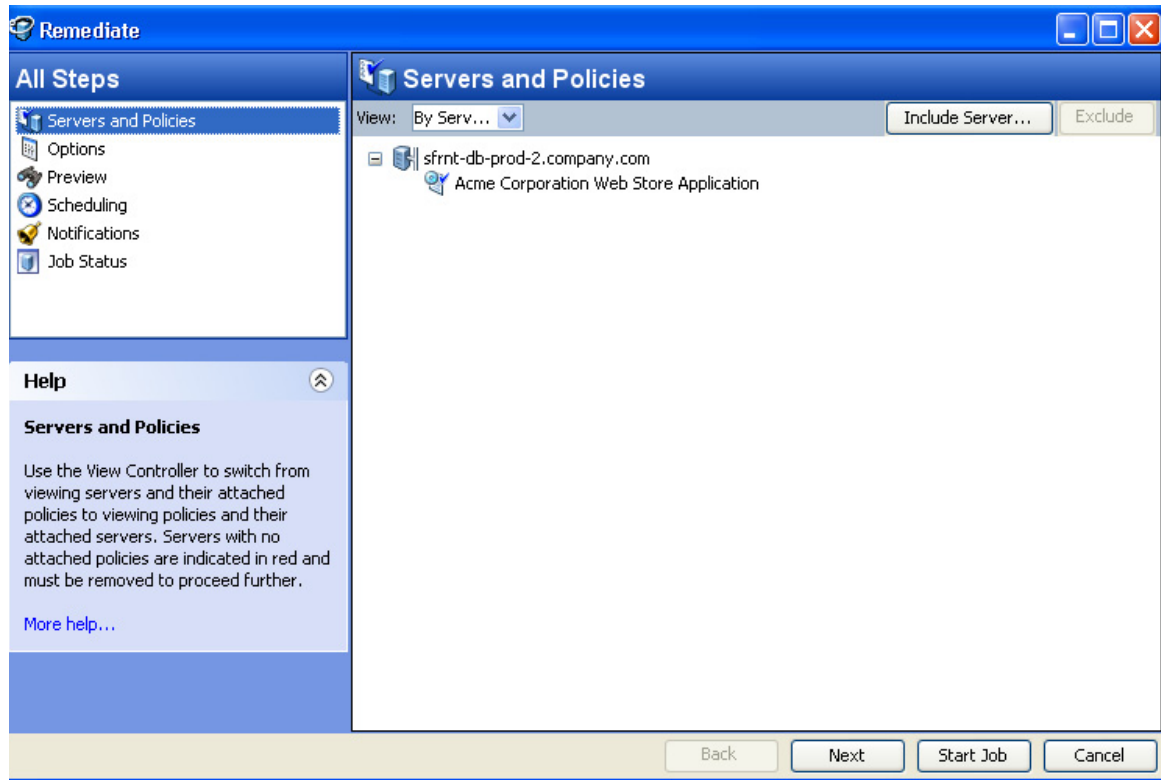
- 1 From the SA Client Navigation pane, select **Library** ► **By Type** ► **Software Policies**. The Software Policy List appears in the Content pane.
- 2 From the Content pane, select a software policy:
 - a From the **View** drop-down list, select **Server Usage**.
 - b Select a server(s) then select **Remediate** from the **Actions** menu. The Remediate window appears.or
 - a From the Content pane, open a software policy. The **Software Policy Window** appears.
 - b From the **View** pane, select **Server Usage**.
 - c Select a server(s) then select **Remediate** from the **Actions** menu. The Remediate window appears.

Remediating Policies

The Remediate window as shown in [Figure 58](#), allows to you remediate the servers against the policies and consists of the following steps:

- [Selecting Servers and Policies for Remediation](#)
- [Specifying Options for Remediation](#)
- [Preview Policy Remediation](#)
- [Schedule Software Policy Remediation](#)
- [Setting Email Notifications for Remediation](#)
- [Viewing Job Status](#)

Figure 58 The Remediate Window in the SA Client



Selecting Servers and Policies for Remediation

This step allows you to specify the servers (with software policies attached) for remediation. In this step, you can add and remove servers from the list, view all the policies attached to a server, and remove policies attached to servers.

Perform the following steps to select servers and policies for remediation:

- 1 Open the Remediate window from one of the methods described in [Open the Remediate Window](#) on page 271.
- 2 In the Remediate window, select **Servers and Policies** in the panel on the left. Managed Servers with attached software policies and patch policies are displayed.

A *software policy* is represented in the list by the icon .

A *patch policy* is represented by the icon .



You can view a list of policies with attached servers by selecting **By Policies** from the View drop-down list.

- 3 (Optional) Click **Include Server** to add additional servers to the list or select a Managed Server and click **Exclude** to remove that server from the list.
- 4 Select a Managed Server(s) with its attached software policies.
- 5 Click **Next** to proceed to Options setup.

Specifying Options for Remediation

The following remediation options are available:

- The reboot actions required for the installation or uninstallation process. You can control when to reboot servers during installation or uninstallations to minimize the downtime caused by server reboots.
- The option to continue installation or uninstallation when an error occurs.
- The scripts to run on a server before or after installation or uninstallation. The scripts include:
 - **Pre-Download:** (*Installation Only*) A script that runs *before* software or patches are downloaded from the Software Repository to the Managed Server.
 - **Post-Download:** (*Installation Only*) A script that runs *after* software or patches are downloaded from the Software Repository to the Managed Server but *before* the software or patch is remediated
 - **Pre-Remediate:** A script that runs *before* packages or patches are installed on the server.
 - **Post-Remediate:** A script that runs *after* packages or patches are installed on the server.

Reboot Options

Perform the following tasks to specify the options for remediation:

- 1 Select one of the following Reboot options:

- Reboot servers as specified by individual software items

This option allows you to reboot servers depending on the reboot option specified in the software resources properties window.

- Reboot servers after each installation or uninstallation

This option allows you to reboot servers after remediating software.

- Hold all server reboots until all actions are complete

If the reboot option is selected in the software resources properties, this option allows you to reboot the servers after all the software resources are installed and uninstalled. If the reboot option is not selected in the software resources properties, this option does not reboot the server after all the software resources are remediated.

- Suppress all reboots

This option allows you to suppress the reboots even if the reboot option is selected in the software resources properties.



If a software policy specifies multiple *non-RPM type* packages with the option `reboot=yes` selected for every package in the **Package Properties** window, and the option *Reboot as dictated by package properties* selected in the Remediate window, then remediating a sever with the software policy will reboot the server *every time* a package is installed.

If a software policy specifies multiple *RPM type* packages with the option `reboot=yes` selected for every RPM package in the Package Properties window, and the option *Reboot as dictated by package properties* selected in the Remediate window, then remediating a server with the software policy will reboot the server *only once* after all the RPM packages are installed.

- 2 Select *Attempt to continue running if an error occurs*, if you want the remediate process to continue even when an error occurs with any of the packages, patches or scripts.
- 3 In the **Scripts** section, select the *Pre-Download*, or *Post-Download*, or *Pre-Remediate*, or *Post-Remediate* tab. You can specify different scripts and options on each of the tabs. You must have certain *Script permissions* to select these options. See Appendix A of the *SA Administration Guide* for more information about the required permissions.
 - a Select **Enable Script**: enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
 - b Select **Saved Script** or **Ad-Hoc Script** from the drop-down list. A Saved script is stored for future use after you upload the script to SA. An Ad-Hoc script must be entered manually and is intended only for a single operation and is not stored in SA.
 - If you select Saved Script from the drop-down list, click **Select** to specify the script. The **Select Script** window appears. Select the script(s) to run and click **Select**.
 - If you selected Ad-Hoc Script from the drop-down list, select the type from the Type drop-down list and then enter the script content in the Script field.
 - c Enter the *command-line flags* in the **Command field**, if required.
 - d Enter a *script time-out* value in minutes in the **Script Timeout** field.
 - e In the **User** section, select Root to execute the script as root. To execute the script as a specified user, select **Name** and enter the user name and the password.

To enter a *Windows Domain Name* in the pre-download, post-download, pre-install, post-install scripts, use the following format in the **Name** field:

DomainName\UserName

and enter the password in the password field.

 - f Select *Stop job if script returns an error* to stop the installation if the script returns an error.
- 4 Click **Next** to proceed to the Preview option.

Preview Policy Remediation

The Preview option displays a detailed list of actions performed on a server as a result of software remediation. It displays information for each server that is selected for remediation. Preview displays:

- the software resources that will be installed on or uninstalled from a server
- the application configurations to be applied to a server
- the dependency information for the packages or patches
- the reboots required during the remediation process
- the scripts that will be executed during the remediation process

Perform the following tasks to preview the remediation process:

- 1 Click **Preview** to view the actions that will be performed during the remediation process. To view the details of each action, select a row in the table. The details for the action appear.
- 2 Select **Output** to preview the job output or select **Errors** to preview possible errors.
- 3 Click **Next** to proceed to the Scheduling setup.

Schedule Software Policy Remediation

Schedule the analysis, download and installation stage to be run immediately or later at a specified date and time.


Perform the following steps to schedule the remediation process:


- 1 From the Remediate window, click **Next** to advance to the **Scheduling** display.
- 2 In the **Schedule Analysis** section, select one of the following options:
 - **Run at Job Start**: run the job immediately.
 - **Start time**: schedule a later date and time for the job.
- 3 In the **Schedule Download** section, select one of the following options:
 - **Run Immediately After Analysis**: download software immediately after analysis.
 - **Start time**: schedule a later date and time to download software.
- 4 In the **Schedule Remediate** section, select one of the following options:
 - **Run Immediately After Remediate**: install software immediately.
 - **Start time**: schedule a later date and time to install software.
- 5 Click **Next** to proceed to Email Notifications setup.

Setting Email Notifications for Remediation

Specify email notifications to alert users on the success or failure of the remediation process. You can associate a *Ticket ID* with the remediation process.

Perform the following steps to set email notifications:

- 1 To add email addresses, click **Add Notifier** and enter the email addresses in the **Notification Email Address** field.
- 2 To trigger the notification on the *success* of a Job, select the  icon.

To trigger the notification on the *failure* of a Job, select the  icon.
- 3 Enter a *Ticket ID* to be associated with a Job in the **Ticket ID** field.
- 4 Click **Next** to go to the Job Status display.

Viewing Job Status

View the summary information about the progress of the remediation job and the individual status of each action required to complete the job.

Perform the following tasks to view the job status:

- 1 If you chose to run the remediate job immediately in the Scheduling setup, the job begins immediately. If you scheduled the job for a later time, the job status will be available when the job runs at the scheduled time. The job progress displayed in the Remediate window.
- 2 To view details about each action, select a row in the table. The details for each action appears.
- 3 Select **Output** to view the job output or select **Errors** to view the error details.
- 4 Click **End Job** to stop the job or click **Close** to close the Remediate window.



You can also view all your jobs in the SA Client job logs. See the *SA Users Guide: Server Automation* for information about job logs.

Uninstall Software Using a Software Policy

You can uninstall software installed using a software policy by detaching the policy from a Managed Server or Device group and remediating. Remediation causes the software to be uninstalled when the software policy has been detached.

Uninstalling software by detaching a software policy has two phases:

- [Detach the Software Policy from the Managed Server or Device Group](#)
- [Remediate the Software Policy to Remove Software](#)

Detach the Software Policy from the Managed Server or Device Group

Simply detaching a software policy from a server does not delete the software policy itself nor does it uninstall the software from the Managed Server or Device Group. To uninstall the software, you must detach the software policy from the server or group and then *remediate*.

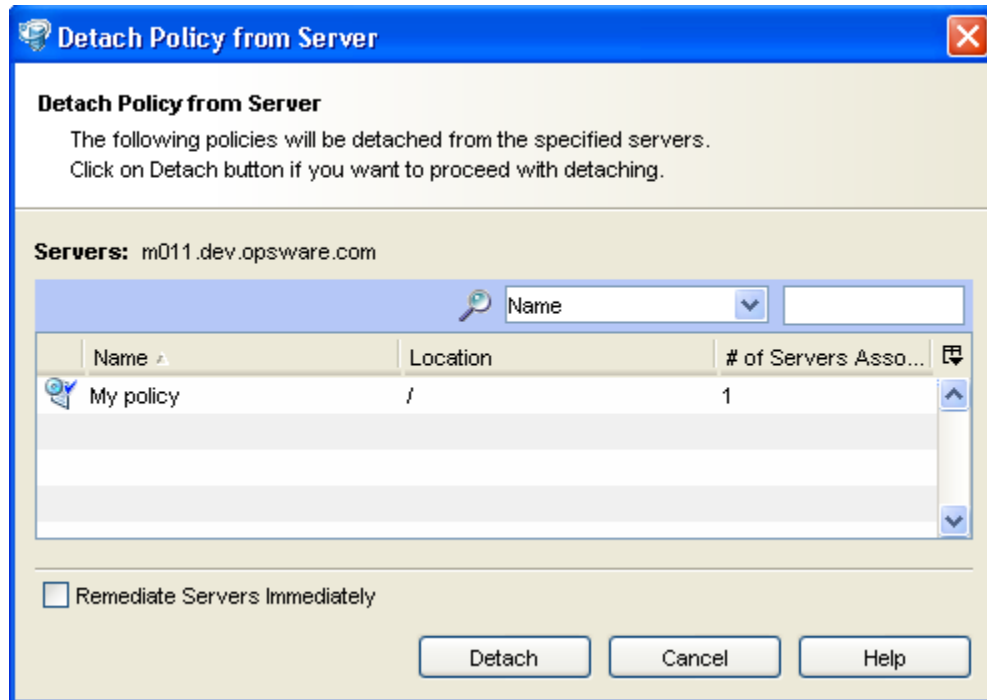


You must have a set of permissions to detach a software policy from a server. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide* for more information.

Perform the following steps to detach a software policy from a server:

- 1 From the SA Client Navigation pane:
select **Devices** ► **Servers** ► **All Managed Servers**. The Server List appears in the Content pane.
or
select **Devices** ► **Device Groups**. The Device Group List appears in the Content pane.
- 2 From the Content pane, select a Managed Server or a Device Group.
- 3 From the **View** drop-down list, select **Software Policies**.
- 4 From the **Actions** menu, select **Detach**. The **Detach Software Policy** window appears as shown in [Figure 59](#).

Figure 59 The Detach Software Policy Window in the SA Client



- 5 Highlight a software policy and click **Detach**.
- 6 (Optional) Select “Remediate Servers Immediately” to remediate the servers against the software policy. Selecting this option will display the Remediate window.

Remediate the Software Policy to Remove Software

Perform the tasks described in [Remediating Policies](#) on page 271. The software specified in the detached software policy will be removed from the Managed Server.

Overview of Software Templates

SA allows you to install software by using a software template. You use the same procedure to create a software template as you use to create a software policy, however, you specify that the policy is to be a software template.

A software template is a software policy that can only specify other software policies. A software template is not persistently associated with a server or group of servers. When you remediate a software template against a Managed Server or Device Group, the software policies specified in the software template (and the software specifications they specify) are installed.



If you update a software template, Managed Servers that already had that software template applied *will not* be automatically updated with the software template changes. You must again remediate the software template against any Managed Server it has been applied to in order to apply the changes.

Although a software template is created in the same way as a software policy, a software template differs from a software policy in the following ways:

- A software template is *not* associated with a Managed Server or Device Group.
- A software template specifies software policies. Software policies contain software specifications.
- A software template is associated with an *operating system family*.
- Software templates are *located in folders*.
- *Custom attributes* can be specified for software templates.

If you need information about Creating a software template (part of Software Policy creation) or adding software policies to a software template, see the *SA Policy Setter's Guide*. These tasks are typically performed by SA administrators.

Overview of Running ISM Controls

The **Run ISM Control** window in the SA Client allows you to run the control scripts in an ISM (Intelligent Software Module).

To run the control scripts in an ISM, you must first add the ISM package to a software policy and then attach the software policy to a Managed Server.

See the *SA Policy Setter's Guide* for information about adding an ISM package to a software policy. See also [Attach a Software Policy to a Server](#) on page 268.



You must have the required permissions to run an ISM control script. To obtain these permissions, contact your SA Administrator. See Appendix A of the SA Administration Guide for more information.

Open the Run ISM Control Window

From the Server List:

- 1 From the SA Client Navigation pane:
select **Devices** ► **Servers** ► **All Managed Servers**. The Server List appears in the Content pane.
or
select **Devices** ► **Device Groups**. The Device Group List appears in the Content pane.
- 2 From the Content pane, select a Managed Server or Device Group.
- 3 From the **Actions** menu, select **Run** ► **ISM Control**. The Run ISM Control window appears.

From the Software Policies List:

- 1 From the SA Client Navigation pane, select **Library** ► **By Type** ► **Software Policies**. The Software Policy List appears in the Content pane.

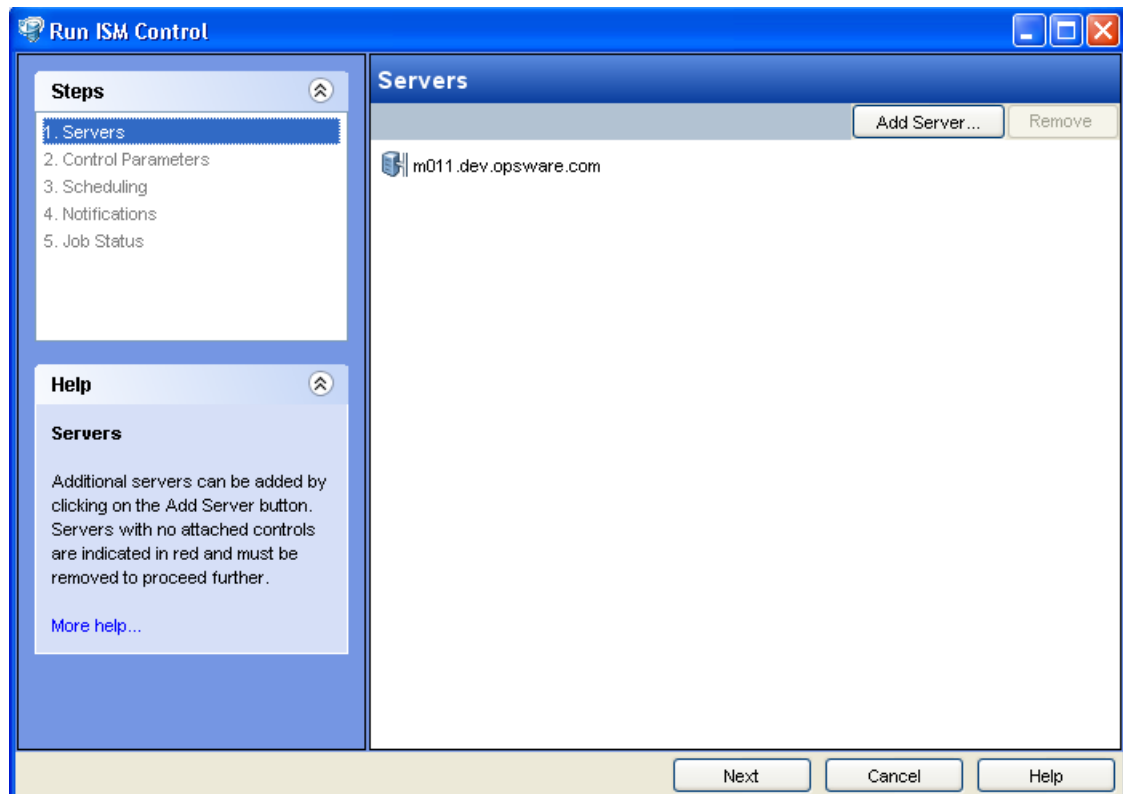
- 2 From the Content pane, select a software policy that specifies an ISM.
 - a From the **View** drop-down list, select **Server Usage**.
 - b Select a Managed Server(s) and then select **ISM Control** from the **Actions** menu. The **Run ISM Control** window appears.or
 - a From the Content pane, open a software policy that specifies an ISM package. The **Software Policy** Window appears.
 - b From the **View** pane, select **Server Usage**.
 - c Select a Managed Server(s) then select **Run ISM Control** from the **Actions** menu. The **Run ISM Control** window appears.

Running ISM Controls

The **Run ISM Control** window, as shown in [Figure 60](#), allows you to run an ISM Control on a server which consists of the following steps:

- [Selecting Managed Servers](#)
- [Specifying Control Parameters](#)
- [Schedule ISM Control Script Execution](#)
- [Set Email Notifications](#)
- [View Job Status](#)

Figure 60 The Run ISM Control Window in the SA Client



Selecting Managed Servers

Specify the Managed Server(s) on which to run an ISM Control script:

- 1 In the left panel of the **Run ISM Control** window, select **Servers**. A list of Managed Servers is displayed.
- 2 (*Optional*) Highlight a Managed Server and click **Include Server** to add additional servers to the list or click **Exclude** to remove a Managed Servers from the list.
- 3 Click **Next**.

Specifying Control Parameters

Specify the control parameters:

- 1 From the Run ISM Control window, click **Next** to advance to the Control Parameters setup.
- 2 From the **Software Policy** drop-down list, select an ISM package.
- 3 From the **Control Script** drop-down list, select a control script. The drop-down list contains only the control scripts assigned to the ISM package selected in the previous step.
- 4 In the **Parameters** section, the name of a parameter matches the name of its corresponding custom attribute name. The value of a custom attribute determines the value of the parameter.
- 5 Click **Next** to proceed to the Scheduling setup.

Schedule ISM Control Script Execution



Schedule an ISM Control script to be run immediately or at a specified date and time:

- 1 From the **Scheduling** window, select one of the following options:
 - **Run Task Immediately**: run the ISM control script immediately.
 - **Run Task At**: specify a later date and time to run the ISM control script.
- 2 Click **Next** to proceed to the Email Notification setup.

Set Email Notifications

Set email notifications to alert users on the success or failure of ISM control script execution. You can associate a *Ticket ID* with the ISM Control script execution job.

Perform the following steps to set email notifications:

- 1 To add email addresses, click **Add Notifier** and enter the email addresses in the **Notification Email Address** field.
- 2 To trigger the notification on the *success* of a job, select the  icon.
To trigger the notification on the *failure* of a job, select the  icon.
- 3 Enter a *Ticket ID* to be associated with the job in the **Ticket ID** field.
- 4 Click **Next** to go to the **Job Status** display.

View Job Status

View summary information about the progress of the ISM Control script execution job and the status of each action required for the job to be completed:

- 1 If you selected Run Task Immediately in the Scheduling setup, the job begins immediately. If you scheduled the job for a later time, the job will run at the scheduled time. The job progress appears in the Run ISM Control window.
- 2 To view the details of each action, select a row in the table. The details for each action will appear.
- 3 Click **End Job** to stop the Job or click **close** to close the Run ISM Control window.



You can also view all your jobs in the SA Client job logs. See the *SA Users Guide: Server Automation* for information about job logs.

Software Policy Compliance





A software policy compliance scan determines whether a Managed Server's software configuration is compliant with the specifications in the software policies attached to that server. If the Managed Server's configuration does not match its attached software policies' requirements, the server is considered to be *non-compliant*.



Scripts specified in a software policy are *not* used to calculate software compliance.

If a Managed Server is not compliant with even a single attached software policy, it is considered non-compliant. Non-compliant Managed Servers must be brought into compliance by remediating the software policy against the server.

The SA Client displays the following compliance information for Managed Servers:

- **Compliant:** If a server is compliant with all software policies attached to it, the server is considered compliant and displays this icon .
- **Non-compliant:** If a Managed Server is not compliant with one or more of the software policies attached to it, the server is considered non-compliant and is represented by the icon .
- **Scan Started:** When a software compliance scan is in progress and information is currently being calculated, the server is represented by the icon .
- **Scan Needed:** If a server's software compliance information must be (re)calculated or its compliance information could be inaccurate, it is represented by the icon .
- **Not Applicable:** Software compliance information is not applicable and the server is represented by a dash (—).

For example, if you detach a software policy from a Managed Server but do not remediate the server against the detached software policy to remove the installed software, the **server's** compliance status is displayed as Not Applicable.

Using the SA Client, you can scan for software compliance by selecting servers from the **Server List** or from a server's or Device Group's **Compliance View**.

See [The Software Policy Compliance Scan](#) on page 282 for information about performing a compliance scan from the server list.

See [Server Compliance](#) on page 207 in Chapter 3 for information about the Compliance View for a device or device group.

The Software Policy Compliance Scan



You must have the required permissions to perform a software policy compliance scan. To obtain these permissions, contact your SA Administrator. See Appendix A of the SA Administration Guide for more information.

Perform the following tasks to scan a Managed Server for software policy compliance:

- 1 From the SA Client Navigation pane, select **Devices ► Servers ► All Managed Servers**. The Server List appears in the Content pane.
- 2 From the Content pane, select a server to scan.
- 3 From the **Actions** menu, select **Scan ► Software Compliance**. During the scan, a dialogue shows the status of the scan. After the scan, the compliance status of the server appears in the Server List.

After you perform the software policy compliance scan, the results of this scan show you the servers that are in compliance (have all required software installed) and the servers that are out of compliance (do not have all required software installed). The software compliance scan is automatically updated after you install or uninstall software or remediate a software policy against a Managed Server.

For more information, see [Remediating Policies](#) on page 271 and [Install or Uninstall Software](#) on page 263.

Software Policy Reports

SA Reporting allows you to generate reports that provide a summary of software policy compliance across Managed Servers. You can also generate reports that provide information about software policies on a single Managed Server. After you generate reports, you can print them, export the reports to .html and .xls, and perform actions on the results.

SA provides the following software policy and policy compliance reports:

- **Software Policy Compliance:** This report groups all Managed Servers by their software policy compliance level to show compliant and non-compliant servers.
- **Software Policy Compliance By Customer:** This report lists all Managed Servers by the customer they are associated with and then by the software policy compliance level.
- **Software Policy Compliance By Facility:** This report displays a chart of all Managed Servers by the Facility they are associated with and then by the software policy compliance level.
- **Defined Software Policies:** This report lists all the software policies by name and their location in the folder hierarchy.

- **Servers With Attached Software Policies:** This report lists all Managed Servers that have one or more software policy attached.
- **Servers In Compliance With Their Software Policies:** This report lists all Managed Servers that are in compliance with all of their attached software policies.
- **Servers Not In Compliance with their Software Policies:** This report lists all Managed Servers that are not in compliance with all of their attached software policies.
- **Servers Without Attached Software Policies:** This report lists all servers that have no software policies attached.

See [SA Client Reports](#) on page 247 in Chapter 4 for information about how to run and view reports in the SA Client.

6 Patch Management for Windows

Overview of Patch Management for Windows

SA Windows Patch Management enables you to identify, install, and remove Microsoft® Windows patches and maintain a high level of security across managed servers in your organization.

Using the SA Client user interface, you can identify and install patches that protect against security vulnerabilities for the Windows Server 2000, Windows Server 2003, and Windows Server 2008 operating systems. These patches include Service Packs, Update Rollups, and hotfixes. This feature also supports patching of Windows Server 2003 x_64 and Windows Server 2008 x_64 operating systems and for Windows XP x_86 operating systems.

This section contains information about how to install Windows patches using patch policies and how to uninstall patches using a sequence of tasks. It also contains information about running patch compliance scans and generating patch policy compliance reports.

SA automates the key aspects of patch management, while offering a fine degree of control over how and under what conditions patches are installed.

Because patches are often released to address serious security threats, an organization must be able to roll out patches quickly, before systems are compromised. At the same time, however, patches themselves can cause serious problems; from performance degradation to server failures.

The SA Patch Management feature allows you to react quickly to newly discovered threats, but it also provides support for strict testing and standardization of patch installation. And, if patches cause problems, even after being tested and approved, the Patch Management feature also allows you to uninstall the patches in a safe and standardized way.

Patch Management is a fully integrated SA component. It leverages the SA features. SA, for example, maintains a central database (called the Model Repository) that has detailed information about every server under management, the patches and software installed on the servers, and the patches and software available for installation. You can use this data to determine the severity of your exposure to a newly discovered threat, and to help you assess the benefits of rolling out a patch versus the costs in downtime and testing requirements.

By automating the patching procedure, the Patch Management feature can reduce the amount of downtime required for patching. SA also allows you to schedule patch activity, so that patching occurs during off-peak hours.

Patch Management for Windows Features

SA automates patch management by providing the following features:

- A central repository where patches are stored and organized in their native formats
- A database that stores information about every patch that has been applied
- Customized scripts that can be run before and after a patch is installed

- Advanced search abilities that identify servers that require patching
- Auditing abilities for tracking the deployment of important patches

These features enable you to browse patches by a certain operating system, schedule patch downloads and installations, set up email notifications, preview a patch installation, use policies and remediation to install patches, and export patch information to a reusable file format.

Types of Patch Browsing

The SA Client interface organizes Microsoft patches by operating systems and displays detailed vendor security information about each patch, such as Microsoft Security Bulletins. You can browse patches by the date Microsoft released the patch, by the severity level, by the Security Bulletin ID, QNumber, and so on. You can also browse all patches that are installed on a server, and view and edit patch metadata.

Scheduling and Notifications

In Patch Management, you can separately schedule when you want patches to be imported from Microsoft into HP Server Automation (either automatically or on demand) and when you want these patches to be downloaded to managed servers. As a best practice, patch installations are typically scheduled for a time that causes minimal disruption to an organization's business operation. If you are installing one patch on one server, the installation operation will start only after the download operation has completed.

Patch Management also allows you to set up email notifications that alert you whether the download and installation operations completed, succeeded, or failed. When you schedule a patch installation, you can also specify reboot preferences to adopt, override, postpone, or suppress the vendor's reboot options.

Patch Policies and Exceptions

To provide flexibility in how you identify and distribute patches on managed servers or groups of servers, Patch Management allows you to create patch policies that define groups of patches that you need to install.

By creating a patch policy and attaching it to a server or a group of servers, you can effectively manage which patches get installed where in your organization. If you want to include or exclude a patch from a patch installation, Patch Management allows you to deviate from a patch policy by specifying that individual patch in a patch policy exception.

An additional patch is one that is not already specified in the patch policy and is one that you want to include in (add to) the patch installation. A patch that you want to exclude from a patch installation is one that is already specified in a patch policy and is identified in the patch policy exception as one you do not want installed.

In cases where it is already known that a certain Windows patch may cause a server or application to malfunction, you should create a patch policy exception to exclude it from being installed on that server or on all servers that have that application.

Patch Installation Preview

While Patch Management allows you to react quickly to newly discovered security vulnerabilities, it also provides support for strict testing and standardization of patch installation.

After you have identified patches to install, Patch Management allows you to simulate (preview) the installation before you actually install a patch. This preview process tells you whether the servers that you selected for the patch installation already have that patch installed. In some cases, a server could already have a patch installed if a system administrator had manually installed it.

After this type of patch installation, if a compliance scan has not been run or the installed patch has not been registered, HP Server Automation does not know about it. The preview process for an up-to-date report of the patch state of servers. The preview process also reports on patch dependency and supersedence information, such as patches that require certain Windows products, and patches that supersede other patches or are superseded by other patches.

Patch Policy Remediation

Patch Management also provides a solution for remediating servers that are not operating properly due to installed patches. If installed patches cause problems, even after being tested and approved, Patch Management allows you to uninstall the patches in a safe and standardized way. Patch Management allows you to specify uninstall options that control server reboots and the execution of uninstall commands, and pre-uninstall and post-uninstall scripts. Similar to previewing a patch installation, you can also preview a patch uninstallation.

Exporting Patch Data

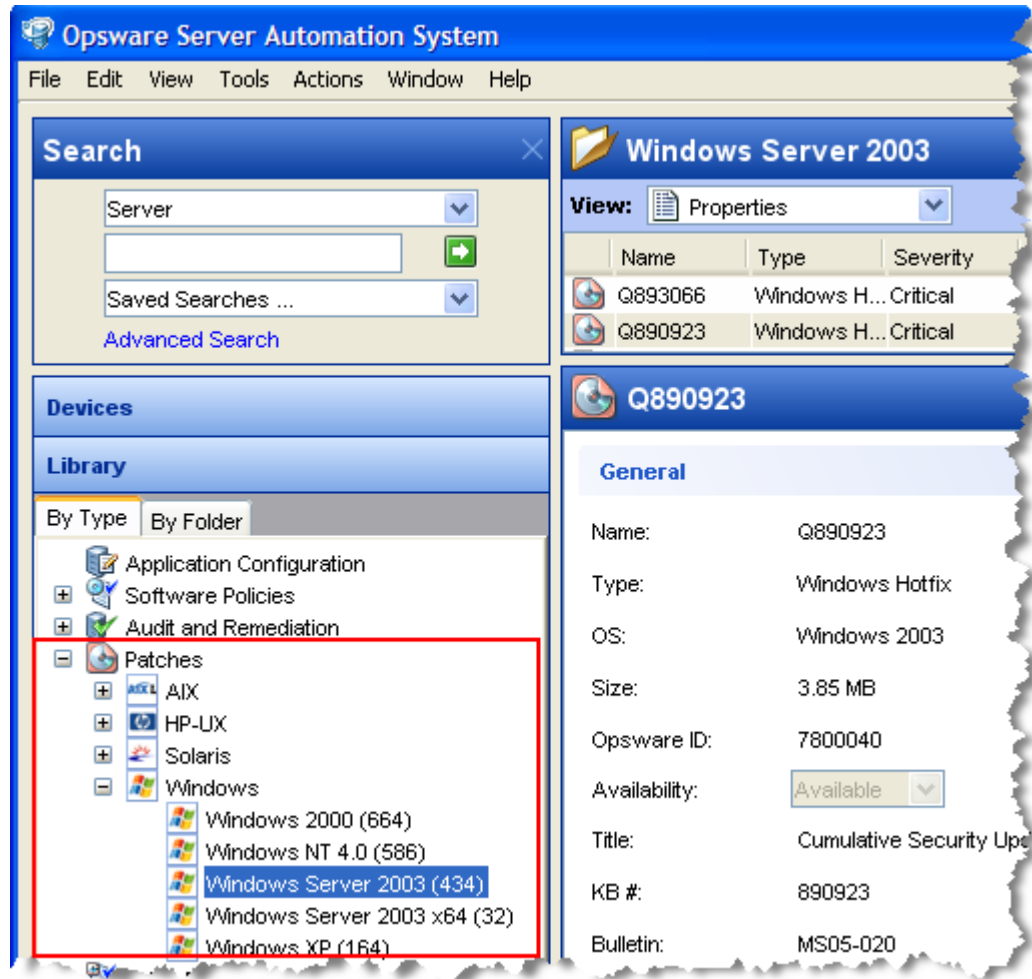
To help you track the patch state of servers or groups of servers, Patch Management allows you to export this information. This information can be exported in a comma-separated value (.csv) file and includes details about when a patch was last detected as being installed, when a patch was installed by HP Server Automation, the patch compliance level, what patch policy exceptions exist, and so on. You can then import this information into a spreadsheet or database to perform a variety of patch analysis tasks.

Library

The SA Client Library provides flexibility in searching for and displaying Microsoft patches by operating system, severity level, release date, bulletin ID, and so on. See [Figure 61](#). The number in parenthesis is the total number of patches (for that operating system version) that were uploaded from the Microsoft web site. In the Content pane, a dimmed patch icon indicates that the patch has not yet been uploaded to the Library. Use the column selector to control the columns of patch metadata data that you want to display.

Since the Library is integrated with Microsoft patch metadata, you can review vendor information (in real-time) in the Preview pane.

Figure 61 Windows Patches in the HP Server Automation Client Library



Patch Management for Windows Prerequisites

The managed servers that will be patched have the following requirements:

- Either Microsoft Core XML Services (MSXML) 3.0 (or later) or Internet Explorer (IE) 6.0 (or later) must be installed on the managed servers. These versions of MSXML and IE support the Microsoft XML parser and related DLL files. HP Server Automation uses MBSA version 2.1 for patch management. Vendor-recommended patches that are installed during the patch remediation process are based on MBSA 2.1. See the *SA Simple/Advanced Installation Guide* for instructions on installing MBSA 2.1 and associated files.
- Windows Installer 3.1 must be installed on the managed servers. This installer is available at the following URL:
<http://support.microsoft.com/kb/893803/>
- On the managed servers, the Automatic Update service must be set to either Automatic or Manual. To set a Windows service, from the Windows Control Panel select **Administration Tools ► Services**. This service setting is required because Patch Management relies on the MBSA 2.1 scanning engine (`mbsacli`) to detect installed and

recommended patches. If the Automatic Update service is disabled, `mbsaccli` will not work properly and the patching process will not continue after reboot. In this situation, the Agent is unable to report a complete set of installed and recommended patches.

- For Windows Server 2000 managed servers, SP4 must be installed. Servers with earlier service packs are not supported by Patch Management.
- For Windows Server 2003 and 2008 managed servers, MSXML 3+ and Windows Update Agent must be installed and the Windows (Automatic) Update service must not be disabled but must be set to never check for updates.
- For Windows Server 2008, the Add and Remove Programs dialogue must be closed when you run Windows patch management tasks.
- For Windows XP managed servers, SP2 must be installed.
- To use Patch Management on managed servers with SA Agent versions earlier than 6.1, the language (locale) of the managed server must be either English, Japanese, or Korean. To set the language, on the managed server, open the Control Panel, open the Regional and Language Options window, select the Regional Options tab, and select an item from the drop-down list at the top.
- Specific versions of the SA Agent are required to support the functions of Patch Management, as listed in [Table 15](#).

Table 15 SA Agent Requirements

Patch Functionality	SA Agent Version
Install Patch	5.5 or later
Uninstall Patch	5.5 or later
Remediate	5.5 or later

Windows Server 2008 Patch Management Support

As of HP Server Automation 8.0, Windows Server 2008 x86, Windows Server 2008 x64 and Windows Server 2008 R2 Patch Management is supported. SA Patch Management for Windows enables you to identify, install, and remove Microsoft® Windows patches and maintain a high level of security across managed servers in your organization.

SA Windows Server 2008 patch management support is compatible with a mixed-version Multimaster Mesh (where both patched and unpatched cores co-exist) with the following caveats:

- Install, uninstall and remediation of Windows Server 2008 patches work correctly only if invoked from an SA 8.0 patched core.
- Managed servers that register to a unpatched core can have an incorrect compliance status and server patch list (however, once the core is patched, the data is corrected after the next patch compliance scan).
- The MBSA patch database must be imported from a patched core.

SA patch management for Windows Server 2008 is identical to that for Windows Server 2003. In addition to the Windows platform patch management functions, after HP Server Automation 8.0 is installed the following will apply:

- Windows Server 2008 patches will appear under the Library after the MBSA patch database is imported.

- You can select `Windows Server 2008` under **Administration > Patch Settings > Patch Downloads > Patch Products** to specify whether to import Windows Server 2008 patch metadata.
- Windows Server 2008 patches can be managed just like patches for other Windows platforms.
 - Users can invoke a patch browser to edit patch properties, descriptions, and reboot/install/uninstall flags.
 - Users can see the following patch views when a Windows Server 2008 server is selected.
 - Patches Needed
 - Patches Recommended By Vendor
 - Patches with Policies or Exceptions
 - Patches Installed
 - Patches with Exceptions
 - All Patches
- You can import patch binaries from the vendor using the SA Client or from a file.
- You can attach Windows Server 2008 patch policies to servers and server groups.
- You can define patch policy exceptions for Windows Server 2008 patches on servers and server groups.

HP Server Automation 8.0 supports most features available For Windows Server patch management. Windows Server patch management features are supported for Windows Server 2008 if they are supported on another Windows server platforms, unless otherwise noted.

The `populate-opsware-update-library` Script

This script has been updated to include the following command line arguments:

- `no_w2k8`
specifies Windows Server 2008 x86 patch binaries should not be uploaded.
- `no_w2k8x64`
specifies Windows Server 2008 x64 patch binaries should not be uploaded.

Policies and Exceptions for Windows Server 2008 patches

HP Server Automation 8.0 provides a recommended patch policy for Windows Server 2008 x86 and patch policy for Windows Server 2008 x86_64. You can also define additional custom patch policies in the same way as described in the *SA User Guide: Application Automation*.

Remediate and ad-hoc Install/Uninstall

Like Windows Server 2003, you can remediate Windows Server 2008 patch policies and perform ad-hoc Windows Server 2008 patch installations and uninstallations. Windows Server 2008 patches can be remediated in software policies and ad-hoc installations using install/uninstall software. However software compliance does not account for applicability. This is consistent with other Windows Server platform patches contained in software policies.

Patch Compliance

You can perform Patch Compliance scans on Windows Server 2008 servers to determine compliance relative to attached 2008 policies and exceptions. Patch compliance is based on patch applicability on the selected server(s).

The Compliance dashboard will display compliance details for Windows Server 2008 servers just as it is for other Windows servers.

Known limitations

- The Install/Uninstall Patch task window typically allows you to specify install/uninstall flags when one or more patch is selected for installation/uninstallation if the patch is in .EXE format (Microsoft delivers Windows Server 2008 patches in both .EXE and .CAB format). In HP Server Automation 8.0, if a patch is in .CAB file format, you cannot specify install/uninstall flags in the Patch Browser, Install Patch, and/or Uninstall Patch task windows because command-line arguments are not supported for .CAB format patches.
- If you add install or uninstall flags to a Windows patch browser, any flags that SA would otherwise have used are overridden. Importantly, overriding the `-q` (if the patch supports `-q`) flag will cause patch installation to fail. That installation failure can take as long as an hour to time out.

Therefore, if you must use additional flags in a Windows patch browser, you must specify the `-q` flag with your additional flags. For example, if you want to log the install/uninstall process and do not want to override the default flags, specify the following:

```
/log:c:\mylog.txt /q /z
```

- Prior to the introduction of Windows Server 2008 CAB patches, each Windows patch identified specific applicable locales (such as `en` for English servers, `ja` for Japanese servers, and so on).
- The locale of CAB patches can be `ALL` meaning the patch can be applicable for servers of all locales.
- When you view the Vendor Documentation of a CAB patch, only the English version is displayed.

Microsoft Patch Database

The Microsoft patch database contains information about released patches and how they should be applied. Patch Management compares all Windows servers to this database to enable the policy setter to determine the patches that must be applied.

Microsoft posts patches on its web site on the second Tuesday of each month, unless a special circumstance requires an immediate release. Windows patches released on *patch Tuesday* are available immediately to import into HP Server Automation. Before Patch Management can install a patch on a managed server, the patch must be downloaded from the Microsoft web site and imported into the Software Repository. You can download and import patches with either the HP Server Automation Client or with a script.

Once every 24 hours, the SA Agent on a Windows server compares the server's current state against the Microsoft patch database (based on the latest version of the MBSA) that has been imported into HP Server Automation by the patch administrator. The SA Agent reports the results of that comparison and then stores the data in the Model Repository. When a user requests a patch compliance scan of a Windows server, the data is retrieved from the Model

Repository and displayed in the SA Client. By storing the data in the Model Repository, rather than performing an actual comparison on the server itself when a user requests an analysis, the data can be quickly retrieved and displayed.

If you perform a patch analysis of a Windows server immediately after importing a new version of the Microsoft patch database, the analysis does not yet include the data from the new patch database. Instead, HP Server Automation reports the data from the last time that the SA Agent recorded the results of its comparison. For example, the SA 5.5 Agent on a Windows server uses Microsoft's latest detection engine (MBSA 2.1) to identify installed patches. If you used a previous version of the SA Agent to create a package of installed patches (from a server snapshot), a previous version of Microsoft's detection engine (MBSA 1.2.1, 2.0.1) was used. Because different versions of MBSA were used to identify patches installed on a Windows server, you should expect to see a difference between the list of installed patches that the SA Client displays and the installed patches in the package that was created from a snapshot.



While MBSA 2.1 can include programs that are not patches in the Microsoft patch database, such as Malicious Software Removal Tool entries, these programs are excluded from Patch Management.

SA Integration

When a server is brought under management by SA, the SA Agent installed on the server registers the server's configuration, including installed patches, with SA. (The SA Agent repeats this registration every 24 hours.) This information, which includes data about the exact operating system version, hardware type, installed software and patches, is immediately recorded in the Model Repository. Also, when you first provision a server with SA, the same data is immediately recorded.

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

After a server is brought under management, you should install all Windows patches by using the Patch Management feature. If you install a patch manually, SA does not have data about that patch until the next software registration. If you install a patch manually, it can take as long as 24 hours until the data about that server in the Model Repository is up-to-date. However, whenever you install patches with SA, the SA Agent immediately updates the information about the server in the Model Repository.

You cannot use HP Server Automation to uninstall a patch that was not installed by using the Patch Management feature.

Support for Windows Patch Testing and Installation Standardization

HP Server Automation offers features to minimize the risk of rolling out patches. When a patch is initially imported into HP Server Automation, its status is marked as Limited and only administrators with the required permissions can install it.

The patch administrator then defines patch installation and uninstallation options and tests the patch. Only after the patch is tested and the patch administrator marks it as available for use (Available) can other administrators install the patch.

The Patch Management feature allows you to standardize the way that patches are installed and uninstalled, thereby preventing ad-hoc installation procedures. Patch administrators standardize patch installation by providing pre-install and post-install scripts, install and uninstall flags, reboot instructions, and how to handle error codes from the pre-install and post-install scripts

Supported Windows Patch Types

The following table lists the Windows patch types that Patch Management supports.

Table 16 Windows Patch Types

OS Versions	Patch Types
Windows Server2000	Windows Hotfix Windows OS Service Pack Update Rollup
Windows Server 2003	Windows Hotfix Windows OS Service Pack Update Rollup
Windows Server 2008	Windows Hotfix Windows OS Service Pack Update Rollup
Windows XP	Windows Hotfix Windows OS Service Pack Update Rollup

Supporting Technologies for Patch Management

Patch Management uses patching utilities and technologies for each supported Windows operating system. HP Server Automation uses these tools behind the scenes. This allows you to perform patch management through a single interface, without having to worry about invoking a number of different patching utilities.

The following patch management and installation tools are used for the supported Windows operating systems:

- **msiexec.exe**: Installs and uninstalls MSI packages.
- **qchain.exe**: Enables a single reboot when you are installing more than one hotfix.

- **unzip.exe:** Extracts info-zip compatible zip archives.
- **Windows Update Agent:** Enables access to the Microsoft framework for patch updates.

See [Importing Windows Patch Utilities](#) on page 330.

Windows Hotfixes

After a Microsoft Windows hotfix is imported into HP Server Automation, you can specify options to reboot the server when a hotfix is installed or uninstalled. A Windows hotfix typically requires a reboot if it updates system files. This reboot enables SA to use the newly updated system files.

When a hotfix is installed along with other hotfixes, this process is called hotfix chaining. If one or more hotfixes require that the server is rebooted, the reboot can sometimes be postponed until all hotfixes have been installed. The user performing the installation must first run qchain.exe before performing the reboot. This ensures that the Pending File Rename Queue is correctly ordered.

Postponing reboots is not always possible, due to a defect in qchain.exe that was resolved in December 2002. All Windows hotfixes created after May 2001 included the Pending File Rename Queue manipulation logic in qchain.exe. Therefore, all hotfixes created between May 2001 and December 2002 are vulnerable to the same qchain.exe defect. See the Microsoft Article for Q815062.

If a Windows Service Pack or Security Rollup Package is being installed in the same hotfix chaining process, a reboot is required. This reboot cannot be postponed. Before the reboot that is associated with this package occurs, qchain.exe must be run.

When multiple hotfixes are chained by HP Server Automation, the setting that specifies that a reboot on install is required for each hotfix is honored. HP Server Automation analyzes the set of hotfixes being installed to determine whether one or more reboots can be postponed until the end of the chaining operation.



If you are installing a Windows hotfix that does not support the -z flag, remember to use the /-z option to prevent the Patch Management feature from passing in the -z flag.

HP Server Automation examines the date each hotfix was created to determine whether any associated reboot can be safely postponed until the end of the chained installation.

HP Server Automation will *not* change the installation order of the chained hotfixes (as an attempt to further reduce the number of reboots), whether or not Service Pack or Security Rollup Packages are being installed in the chained operation.

When HP Server Automation installs a hotfix in isolation (not as part of a chained installation operation), HP Server Automation honors the value of the reboot on the installation operation.

HP Server Automation runs qchain.exe on the managed server after the installation of each Windows hotfix and before any associated reboot. This guards against problems associated with an incorrectly ordered Pending File Rename Queue. This problem could occur if another hotfix was installed on the managed server outside of HP Server Automation.

Searching for Patches and Policies

In the SA Client, you can search for information about your operational environment by using the SA Client Search feature. The Search feature enables you to search for patches, patch policies, servers, and so on. See “SA Client Search” in the *SA Users Guide: Server Automation*.

Roles for Windows Patch Management

HP Server Automation provides support for rigorous change management by assigning the functions of patch management to several types of users in an organization. These users include a policy setter, a patch administrator, and a system administrator.

- **Policy Setter:** The policy setter is a member of a security standards group that reviews patch releases and identifies the vendor patches that will be included in the organization’s patch policies. A policy setter is responsible for reviewing the latest security threats and the patches that vendors have released to address these problems. A policy setter is generally known as an expert in the operating systems and applications that they manage, and is able to assess the necessity of applying patches issued by vendors. A policy setter is also able to diagnose common problems that arise after patches are installed, allowing for a thorough test of the patch application process.
- **Patch Administrator:** The patch administrator has the authority to import, test, and edit patch options. The patch administrator is often referred to as the security administrator in an organization. A patch administrator is granted specific permissions to import patches into HP Server Automation to test the patches and then mark them as available for use. Basic users can import patches, but they cannot install them or mark them as available. Patch administrators are also able to edit patch options (such as installation scripts) through patch management. Other types of users are not allowed to import or edit patches. Typically, a patch administrator imports the Microsoft patch database and tests patches on non-production reference hardware. After testing the patches and determining that the patches are safe to apply to production systems, a patch administrator marks the patches available in the Library and then advises the system administrators that they must apply the approved patches.
- **System Administrator:** The system administrator installs patches (that have been approved for use) uniformly and automatically, according to the options that the patch administrator specifies. The system administrator is an SA user who is responsible for the day-to-day maintenance of the servers in a deployment. These users are not required to have the same level of expertise in low-level system details as the policy setter and patch administrator. Because the patch administrator has set up the patch installation, the system administrators can attach policies to servers, set an exception for a patch, and install patches on a large number of managed servers. They are responsible for searching for servers that require the approved patch, installing the patches, and verifying that the patches were successfully installed. The system administrator can import patches but cannot install a patch until the patch administrator has marked it as available. The system administrator can also uninstall patches.



These responsibilities are enforced by assigning permissions for managing patches in SA. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide*.

Patch Management Process

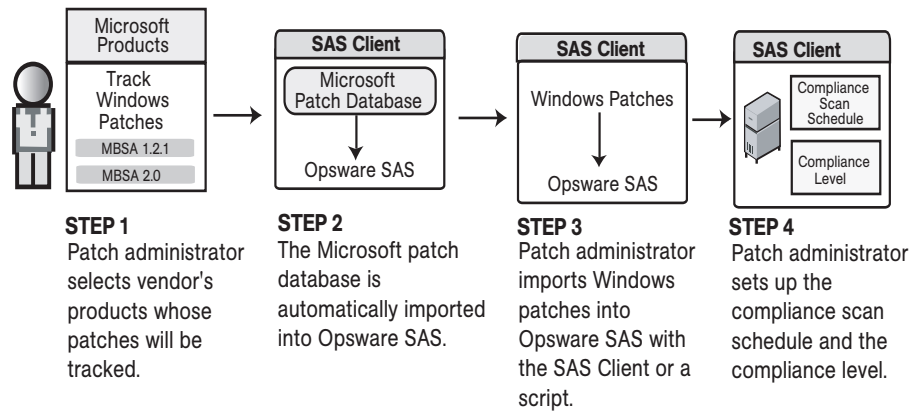
The Windows patching process consists of several key phases: setup, policy management, patch compliance, and deployment.

- Setup steps include getting the Microsoft database (patches and metadata) into HP Server Automation, identifying products you want to track patches for, and configuring patch compliance.
- Policy management steps include investigating released patches, creating and updating patch policies or exceptions, marking patches available to use, and attaching policies or exceptions to servers or groups of servers.
- Patch compliance steps include running compliance scans to determine whether a server is out of compliance, remediating policies, setting up installation options, and installing applicable patches.
- To deploy patches on demand, you can import the required patches, test them, update policies, create new policies, mark them as available to use, specify install options, and install the required patches. [Figure 62](#) and [Figure 63](#) illustrate these phases and steps.

Figure 62 Windows Patching Process: Part A and Part B

WINDOWS PATCHING PROCESS

Part A: Set Up Patch Management



Part B: Create and Attach Patch Policies to Servers

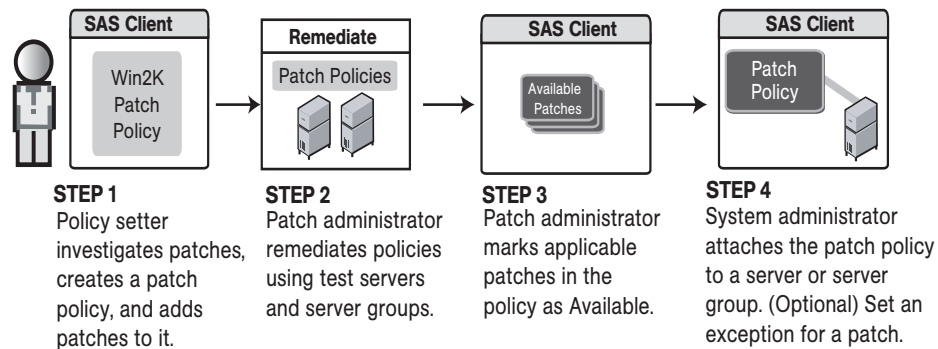
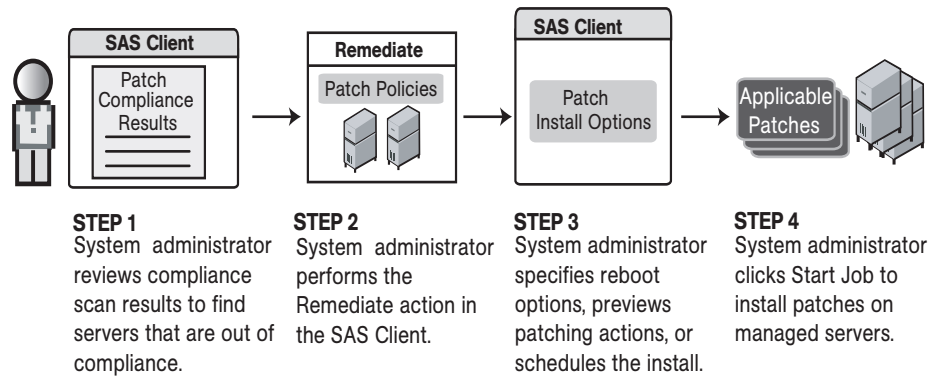


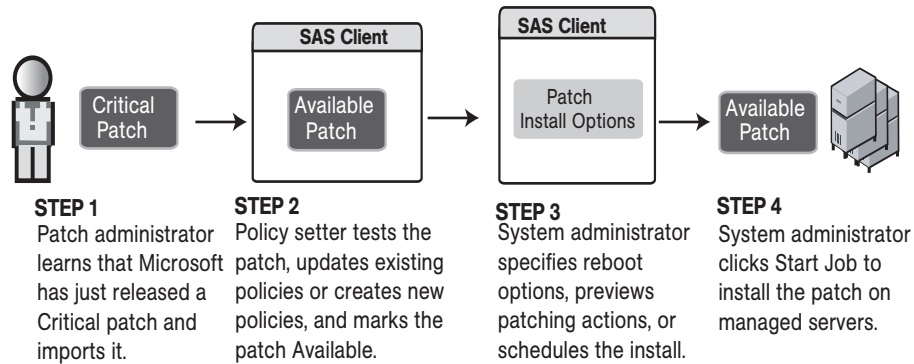
Figure 63 Windows Patching Process: Part C and Part D

WINDOWS PATCHING PROCESS

Part C: Install Patches By Remediating Policies



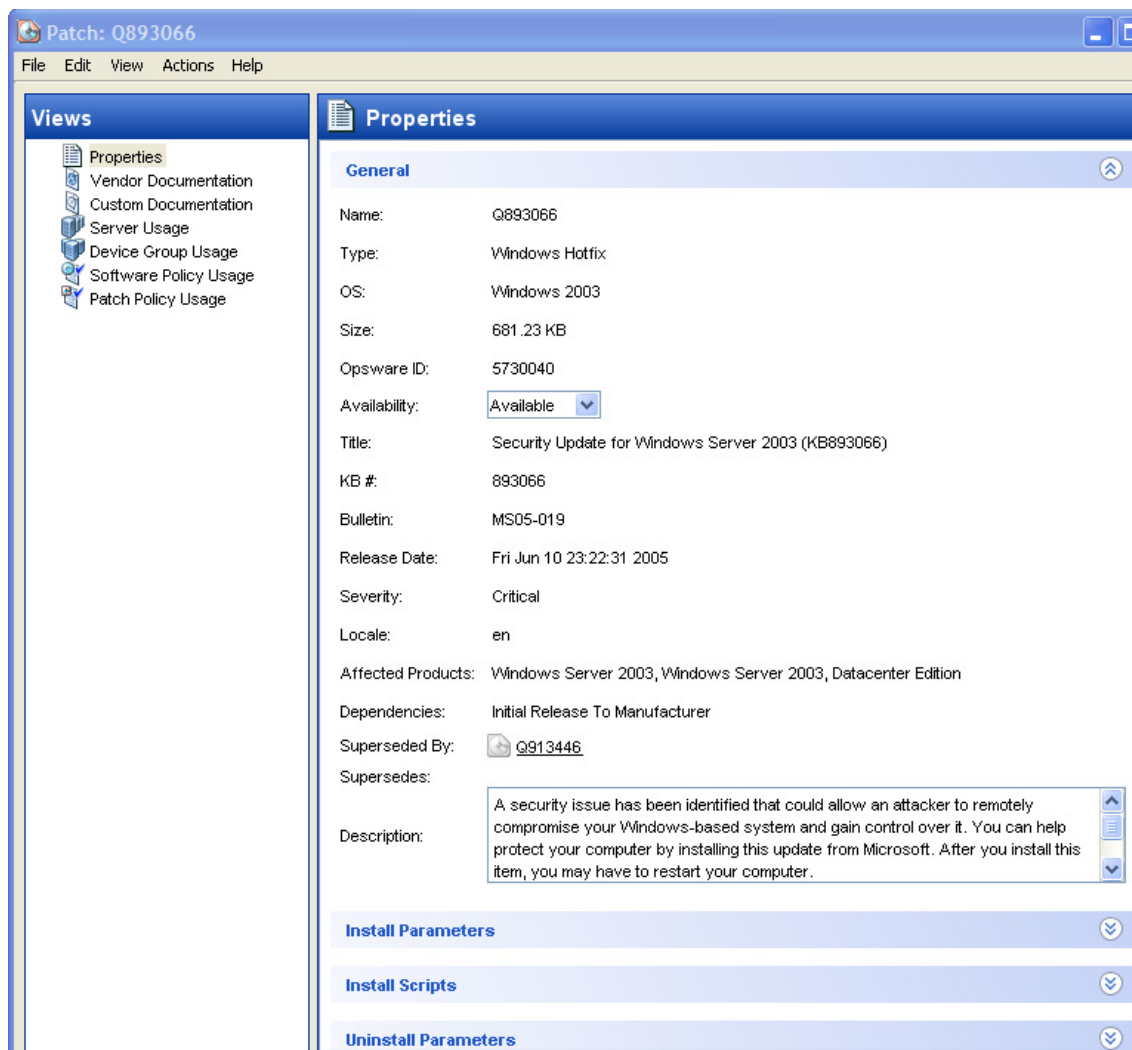
Part D: Install Patches on Demand



Patch Properties

Patch Management displays detailed information (properties) about a patch.

Figure 64 Windows Patch Properties



Patch properties include the following information:

- **Name:** The Microsoft name of the patch, such as QNumber, Windows 2000 Service Pack 4, and so on.
- **Type:** The type of patch, such as Windows Hotfix or Windows Update Rollup.
- **OS:** The Windows operating systems that are known to be affected by this patch.
- **Size:** The size of the patch file, in kilobytes (KB) or in megabytes (MB).
- **Opware ID:** The HP Server Automation unique ID for the patch.
- **Availability:** The status of a patch within HP Server Automation, which can be one of the following:
 - **Not Imported:** The patch is listed in the Microsoft Patch Database, but has not been imported (uploaded) into HP Server Automation.

- **Limited:** The patch has been imported into SA but requires additional permissions (Manage Patch: Read & Write) to be installed. This is the default patch availability. For more information on permissions, see the *SA Administration Guide*.
- **Available:** The patch has been imported into HP Server Automation, tested, and has been marked available to be installed on managed servers.
- **Deprecated:** The patch cannot be added to patch policies or set as a patch policy exception but can still be installed.
- **Title:** The title of the Microsoft Knowledge Base article for this patch.
- **KB #:** The Microsoft Knowledge Base article ID number for this patch.
- **Bulletin** (Optional): The Microsoft Security Bulletin ID number for this patch.
- **File Name:** The name of the .exe for this patch.
- **Release Date:** The date that Microsoft released this patch.
- **Severity** (Optional): The Microsoft severity rating for this patch, which can be one of the following:
 - **Critical:** A patch that if exploited could allow the propagation of an internet worm, without user action.
 - **Important:** A patch that if exploited could result in a compromise of the confidentiality, integrity, or availability of user data, or of the integrity or availability of processing resources.
 - **Moderate:** A patch that if exploited could result in minimal impact. Exploitability is mitigated to a significant degree by certain factors, such as default configuration, auditing, or difficulty of exploitation.
 - **Low:** A patch that is difficult to exploit or if exploited, could result in minimal impact.
- **Locale:** The locale this patch applies to.
- **Affected Products:** Information from MBSA that identifies other Microsoft software that is known to be affected by this patch.
- **Dependencies:** Microsoft products that this patch requires. The patch cannot be installed if these products do not already exist on the server.
- **Superseded By** (Optional): A list of patches that this patch is superseded by. This relationship does not apply to MBSA 1.2.1 patches.
- **Supersedes** (Optional): A list of patches that this patch supersedes. This relationship does not apply to MBSA 1.2.1 patches.

Patch Dependencies and Supersedence

Patch metadata identifies all known dependency and supersedence relationships between patches and Windows products, and between patches and other patches. Dependency relationships identify Windows products that must already exist on a server before you can install a certain patch. Supersedence relationships identify patches that supersede or are superseded by other patches. In Patch Management, *supersedes* means that one patch replaces another and *superseded by* means that the patch you are installing is replaced by another patch.

For all MBSA 2.1 patches, Patch Management analyzes this information to determine the viability of a patch installation. For example, if you are remediating patches and a superseding patch is already installed, the patch will not be installed. If you try to install a superseded patch and the superseding patch is available and included in a patch policy, the superseded patch will not be installed. Patch Management does not analyze this information for MBSA 1.2.1 patches.



Patch Management does not detect whether two patches are mutually exclusive, which is when either one can be installed but not both. Subsequently, Patch Management does not prevent you from installing both patches on a server. This means that you may be able to install both a superseded patch and a superseding patch on a server.

Viewing Windows Patches

The SA Client displays information about Microsoft Windows patches that have been imported into HP Server Automation.

To view information about a patch, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Windows operating system.

The Content pane will display all of the patches listed in the Microsoft Patch Database for the Windows operating system that you selected.

- 3 (Optional) Use the column selector to sort the patches according to Name, Type, Severity, Availability, Release Date, and Bulletin Number.
- 4 In the Content pane, open a patch to view its properties in the Patch window.

Editing Windows Patch Properties

You can edit a patch's Description, Availability, Install Parameters, and Uninstall parameters. Due to the nature of the type of patch, some properties are not editable. For example, you cannot turn the reboot-on-install option of a Windows Service Pack off.

The Availability property indicates the status of the patch in HP Server Automation. If the Availability is Not Imported, you cannot change this property.

You can set the install and uninstall parameters on either the patch properties page or in the Patch Actions only when you are installing or uninstalling one patch at a time. The parameters on the properties page are saved in the Model Repository, but the parameters in Patch Actions are used only for that action. The parameters in Patch Actions override those on the patch properties page.

To edit the patch properties, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 In the Content pane, open a patch to view its properties in the Patch window.

- 4 Edit any of the following fields: Description, Availability, and the Install and Uninstall parameters.
- 5 From the **File** menu, select **Save** to save your changes.

Importing Custom Documentation for a Patch

The Custom Documentation view of a patch displays text files that have been imported from the local file system. Non-plain text file types, such as .html or .doc, are not supported.

To import your own documentation for a patch, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 In the Content pane, open a patch to view its properties in the Patch window.
- 4 From the Views pane, select Custom Documentation.
- 5 From the **Actions** menu, select **Import Custom Documentation** or click **Import**.
- 6 In the Import Custom Documentation window, locate a text file and specify encoding.
- 7 Click **Import**.

Deleting Custom Documentation for a Patch

The Custom Documentation view of a patch displays text files that have been imported from the local file system. Non-plain text file types, such as .html or .doc, are not supported.

To delete custom documentation for a patch, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 In the Content pane, open a patch to view its properties in the Patch window.
- 4 From the Views pane, select Custom Documentation.
- 5 From the **Actions** menu, select **Delete Custom Documentation**.
- 6 In the Delete Custom Documentation window, click **Delete**.

Finding Vendor-Recommended Windows Patches

To find the patches that Microsoft recommends for a particular server (based on MBSA 2.1), perform the following steps:

- 1 From the Navigation pane, select Devices ► Servers ► All Managed Servers.
- 2 From the View drop-down list, select Patches.
- 3 From the Content pane, select a server that is running SA Agent 5.5 and Windows Server 2000 (Service Pack 3 or higher), Windows Server 2003, or Windows Server 2008.

- 4 From the Preview pane, select Patches Recommended By Vendor from the drop-down list. This displays the types of patches for the selected server.

Finding Servers That Have a Windows Patch Installed

To find the servers that have a particular patch installed, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list in the Content pane, select Server Usage.
- 5 From the Show drop-down list for the selected patch, select Servers with Patch Installed.

You can browse a server in this list to view a list of all installed patches. Please note that this list may display a more complete list of installed patches than the list you will find in the Windows Add or Remove Programs utility.

Finding Servers That Do Not Have a Windows Patch Installed

To find the servers that do not have a particular patch installed, perform the following steps:

- 1 From the Navigation pane, select Library ► Patches.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list, select Server Usage.
- 5 From the Show drop-down list, select Servers without Patch Installed.

Importing a Patch

Windows patches are downloaded from the Microsoft web site and then imported (uploaded) into HP Server Automation. To see if a patch has been imported, view the patch's Availability property. The Availability of an imported patch is either Limited, Available, or Deprecated. A patch can be imported with the SA Client or with a script. For information about the script, see [Automatically Importing Windows Patches](#) on page 306.

To import a patch with the SA Client, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Package Repository.
- 2 Expand the Package Repository and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.

- 4 To import a patch directly from the Microsoft web site, from the **Actions** menu, select **Import ► Import from Vendor**.

The Import from Vendor window displays the URL of the patch's location on the Microsoft web site. You can override this URL, as needed.

Or

To import a patch that has already been downloaded to your local file system, from the **Actions** menu, select **Import ► Import from File**.

In the file browser window, locate the patch.

- 5 Click **Import**.

Automatically Importing Windows Patches

The `populate-opsware-update-library` shell script downloads the Microsoft Patch Database and patches from the Microsoft site. The script also imports the database and patches into HP Server Automation. (To be imported, a patch must be in the Microsoft Patch Database that has been imported into the Software Repository.) Optionally, the script sets the initial status (Available or Limited) of newly imported patches. The script can also filter the patches imported according to operating system (such as Windows Server 2003/2008). The functionality of the script is also available in the SA Client, as described in [Importing the Microsoft Patch Database](#) on page 328.

To run the `populate-opsware-update-library` script, you need to log onto the Software Repository server as `root`. Typically, you schedule the script to run periodically as a `cron` job on the Software Repository server. To end users of the SA Client, the patches imported with the script appear to have been automatically imported. Do not run concurrent instances of the script.

The `populate-opsware-update-library` script is in the following directory:

`/opt/opsware/mm_wordbot/util/`

[Table 17](#) describes the script's options.

Table 17 Options of `populate-opsware-update-library`

Option	Description
<code>--spin hostname-or-IP</code>	Hostname or IP address of Data Access Engine (spin) host. Default value: spin
<code>--theword hostname-or-IP</code>	Hostname or IP address of Software Repository (theword) host. Default value: theword
<code>--cert_path file-path</code>	File specification of cert file to be used for Spin connection. Default value: <code>/var/opt/opsware/crypto/wordbot/wordbot.srv</code>
<code>--ca_path file-path</code>	File specification of CA file to be used for Spin connection. Default value: <code>/var/opt/opsware/crypto/wordbot/opsware-ca.crt</code>

Table 17 Options of populate-opsware-update-library (cont'd)

Option	Description
<code>--verbose</code>	Display copious output, including patches skipped during the upload.
<code>--no_w2k</code>	Do not process W2K patches.
<code>--no_w2k3</code>	Do not process W2K3 patches.
<code>--no_w2k3x64</code>	Do not process W2K3 (64-bit) patches.
<code>--no_w2k8</code>	Do not process W2K8 patches.
<code>--no_w2k8x64</code>	Do not process Windows 2008 (64-bit) patches.
<code>--no_xp</code>	Do not process Windows XP (32-bit) patches.
<code>--use_proxy_url url</code>	When downloading binaries, connect via this proxy URL.
<code>--proxy_userid userid</code>	Basic-auth userid to provide to proxy server.
<code>--proxy_passwd passwd</code>	Basic-auth passwd to provide to proxy server.
<code>--set_available</code>	Set availability status to Available when uploading patches. The <code>--set_available</code> and <code>--set_limited</code> options cannot be specified at the same time.
<code>--set_limited</code>	Set availability status to Limited when uploading patches.
<code>--no_hotfixes</code>	Do not upload hotfixes.
<code>--no_servicepacks</code>	Do not upload servicepacks.
<code>--no_updaterollups</code>	Do not upload updaterollups.
<code>--no_wsusscan_upload</code>	Do not upload the MBSA 2.1 patch database.
<code>--wsusscan_url_override url</code>	Download the MBSA 2.1 patch database from this URL.
<code>--update_all</code>	Refresh the patches already uploaded into SA.
<code>--download_only path</code>	Download files from the vendor's web site to the specified path (directory), but do not upload them into SA. The files are downloaded into the <i>platform_ver/locale</i> subdirectory beneath the specified path.
<code>--upload_from_update_root path</code>	Upload files from the specified path (directory), not from the vendor's web site. The script looks for patches in the <i>platform_ver/locale</i> subdirectory beneath the specified path. If it cannot find the patch in the that subdirectory, the script looks for the patch in the specified path. If a patch is not found, the script skips the patch and does not upload it. This option is ignored if <code>--download_only</code> is also specified.

Table 17 Options of populate-opsware-update-library (cont'd)

Option	Description
--help	Display the syntax of this script.

Exporting a Windows Patch

To export a patch from HP Server Automation to the local file system, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.
- 4 From the **Actions** menu, select **Export**.
- 5 In the Export Patch window, enter the folder name that will contain the patch file in the File Name field.
- 6 Click **Export**.

Exporting Windows Patch Information

You can export information about patches installed on a server and patches recommended by the vendor. You can also export information from patches recommended by the vendor along with model information on the selected server (such as patch policies or patch policy exceptions). The following information is exported into a .csv file:

- **Server Name:** The name of the managed server.
- **OS:** The operating system of the server.
- **Service Pack:** The service pack level of the server being reported, such as Service Pack 0, Service Pack 1, and so on.
- **KB#:** The Microsoft Knowledge Base Article number for the patch.
- **Bulletin:** The MSYY-XXX ID associated with a hotfix, such as MS05-012, MS06-012, and so on. If the MSYY-XXX ID is unknown, this column will be blank.
- **Description:** A brief description of the purpose of the patch.
- **Time Queried:** The last software registration by the Agent.
- **Time Installed:** The time that the patch was installed.
- **Type:** The patch type.
- **Compliance Level:** An integer that represents the compliance level.
- **Compliance:** Text that displays when you place your cursor over the Compliance column in the Patch Preview pane.
- **Exception Type:** The type of exception, such as Always Install or Never Install.
- **Exception Reason:** A description that explains the purpose of the exception.



Patch Management will display all of the text, including commas, from the Description field displayed in the Patch Properties window in the Description column in the .csv file. To preserve commas in the Description column and keep all text together in that column, double quotes will be converted to single quotes. This does not distort the semantics of the patch description.

To ensure that all of the text about a patch displays in the Description field in the .csv file, Patch Management surrounds the entire description (that you see in the Patch Properties window) with double quotes.

To export the patch information to a .csv file, perform the following steps:

- 1 From the Navigation pane, select **Devices ► All Managed Servers**.
- 2 From the Content pane, select one or more managed servers.
- 3 From the Show drop-down list, select an option.
- 4 From the **Actions** menu, select **Export Patch Info to CSV**.
- 5 In the Export to CSV window, navigate to a folder and enter the file name.
- 6 Verify that the file type is Comma Separated Value Files (.csv). If you did not include the .csv extension in the file name field, Patch Management will append it only if you have the .csv file type selected.
- 7 Click **Export** to save the patch information in a .csv file or click **Cancel** if you do not want to export the patch information.

Deleting a Patch

When you delete a patch, it is removed from HP Server Automation, but it is not uninstalled from managed servers. A patch cannot be deleted if it is attached to a policy or if an exception has been set for it.



Do not delete all of the patches from HP Server Automation. If you do so accidentally, contact your SA support representative for assistance in importing the patches back into SA.

To delete a patch, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**.
- 2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
- 3 From the Content pane, select a patch.
- 4 From the **Actions** menu, select **Delete Patch**.
- 5 In the Delete Patches windows, click **Delete**.

Policy Management

In Patch Management, patch policies and patch policy exceptions enable you to customize patch distribution in your environment. Policies and exceptions define the Windows patches that should be installed or not installed on certain managed servers.

You can choose to have patching in your server environment comply to the model that these policies and exceptions define or you can choose to deviate from this model. If you choose to deviate from the patch policies and exceptions and perform ad hoc patch installs, then you need to remediate. The remediation process ensures that the applicable patches get installed on servers.

Patch Policy

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

A patch policy provides broad flexibility for distributing patches. For example, you can create a patch policy that contains security patches that you want to distribute only to servers used by your sales force. You can also create a patch policy that contains security patches that are applicable to specific software that is already installed on a server, such as Exchange Server, Internet Information Services (IIS), SQL Server, and so on. Or, you can create a patch policy that includes all patches ranked critical (by Microsoft) and installs them on all servers that are used by everyone in your organization.



If you do not want to create a patch policy, you can use the vendor-recommended set of patches (by operating system) as a default patch policy, such as the patches provided by MBSA.

You can attach as many patch policies as you want to servers or groups of servers. If several policies are attached to one server, the installation logic is cumulative—all patches listed in all attached policies will be installed on the server. The Remediate window allows you to select an individual patch policy to remediate. You do not have to remediate all policies attached to a server. You cannot nest patch policies.

If a description of the patch policy is defined, it is recorded in the server's patched state (in the Model Repository). This information enables Patch Management to report on patch policies for patch compliance purposes. The patch compliance process compares patch policies with corresponding patch policy exceptions.

Patch Management supports the following types of patch policies:

- **User-defined patch policy:** This allows an HP Server Automation user to specify the patches that are included in a policy. User-defined patch policies can be edited or deleted by a user who has permissions.

A user-defined patch policy allows a policy setter to opt out of patches. The policy setter can create a (user-defined) patch policy that is a subset of all available patches (that are in a vendor-recommended patch policy). This enables the policy setter to apply only those patches that their environment needs.

- **Vendor-recommended patch policy:** Membership of patches is defined by MBSA recommendations on a server-by-server basis. Vendor-recommended patch policies are system defined and cannot be edited or deleted by a user.



You can only export user-defined patch policies. You cannot export vendor-recommended patch policies.

Patch policies have the following characteristics:

- All patches in a patch policy must apply to the same operating system, such as Windows.
- A patch policy is associated with an operating system version, such as Windows Server 2003/2008.
- A patch policy has a name and can (optionally) include a description that explains its purpose.
- A patch policy can be either user-defined or vendor-defined.
- A patch policy does not have sub-policies. There is no inheritance.
- A patch policy is Customer Independent, which means that patches in the policy can be installed on any managed server, no matter what customer is associated with it. See the *SA Users Guide: Server Automation*.
- A patch policy is always public.
- A patch policy can be attached to zero or more servers or public device groups.
- More than one patch policy can be attached to a server or public device group.
- Only user-defined patch policies can be created, edited, and deleted by a user who has permissions.

Patch Policy Exception

A patch policy exception identifies a single patch that you want to explicitly include or exclude from a specific managed server, along with an optional reason for why the exception exists. The patch in a patch policy exception must apply to the same Windows operating system that the established patch policy is attached to.

A patch policy exception allows you to deviate from an established patch policy (one that is already attached to a server or a group of servers). You can do this by deselecting or adding individual patches to a server. Since patch policy exceptions override all patch policies attached to a server, you can use them to intentionally deviate from a patch policy on a server-by-server basis.

If a reason for a patch policy exception is defined, the description is recorded in the server's patched state (in the Model Repository). This information enables Patch Management to report on patch policy exceptions for patch compliance purposes. The patch compliance results explain how patch policy exceptions compare with corresponding established patch policies. All users who have access to the managed server can view attached patch policy exceptions.

Patch Management supports the following types of patch policy exceptions:

- **Always Installed:** The patch should be installed on the server, even if the patch is not in the policy.
- **Never Installed:** The patch should not be installed on the server, even if the patch is in the policy.



If you ever need to override a patch policy exception, you can manually install a patch.

The following information summarizes characteristics of a patch policy exception:

- A patch policy exception can (optionally) include a description that explains its purpose.
- A patch policy exception can have a rule value of Never Installed or Always Installed.
- A patch policy exception can be set for one patch and one server of the same operating system version. If a patch policy exception is set for a public device group and a server in that group does *not* match the operating system version specified in the patch policy exception, the patch policy exception is *not* applied.
- A patch policy exception can be set, copied, and removed by users who have permissions.

Precedence Rules for Applying Policies

By creating multiple patch policies and patch policy exceptions (that are either directly attached to a server or attached to a group of servers), you control the patches that should be installed or not installed on a server. A precedence hierarchy in Patch Management delineates how a patch policy or a patch policy exception is applied to a patch installation. This hierarchy is based on whether the patch policy or patch policy exception is attached at the server or device group level.

The following precedence rules apply to policies and exceptions:

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

Remediation Process

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



The remediation process requires that the selected managed server is running SA Agent 5.5 and a Windows 2000 Service Pack 3 (or higher) operating system or a Windows Server 2003 or Windows Server 2008 operating system. You cannot use the remediation process if the selected managed server is running a Windows NT4.0 operating system, a Windows 2000 RTM (no service pack), Service Pack 1, or Service Pack 2 operating system, or if the server is not running SA Agent 5.5. Use the Install Patch window to install patches on servers that are running these operating systems or SA Agents 4.5 or earlier.

As a best practice, each time you review the latest Microsoft patch releases and subsequently update a patch policy (by adding new patches to a policy), you should perform remediation. In these situations, a remediation process provides demand forecasting information. This allows you to determine how patch policy changes will impact servers that this policy is attached to.

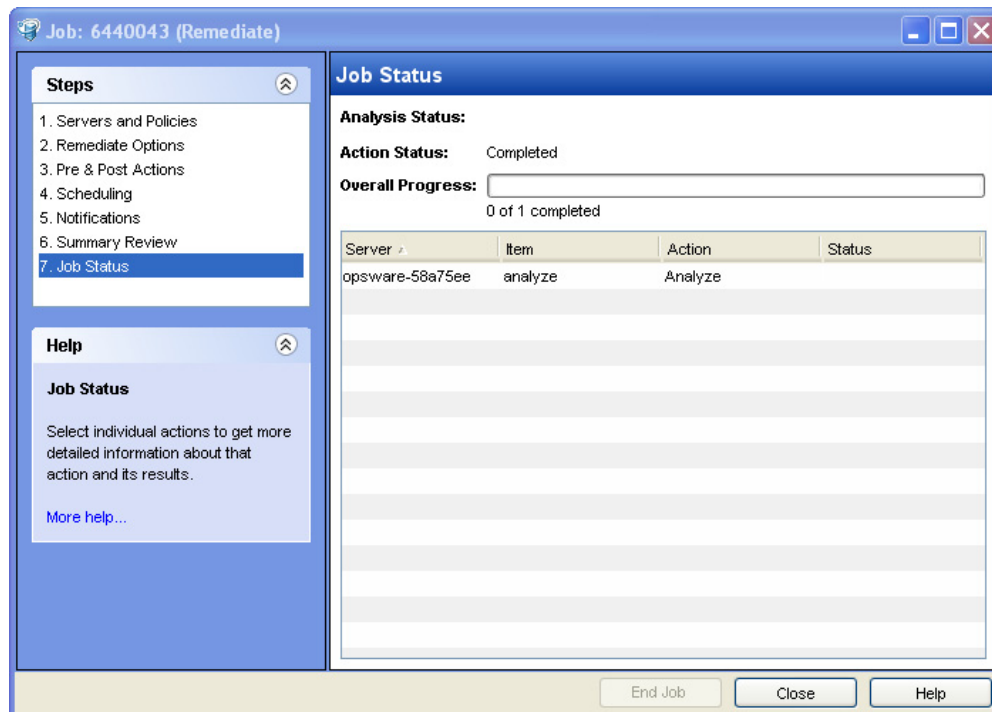
If the remediation process discovers any (applicable) missing patches, these patches will be installed on the servers.

If a patch was installed as part of a patch policy, the remediation process will not uninstall it. However, if a patch was installed as part of a software policy and it is no longer in the software policy, the remediation process will uninstall it.

After HP Server Automation determines the packages that need to be installed to complete the remediation process, remediation uses a set of standard system utilities to complete the operation. See [Supporting Technologies for Patch Management](#) on page 295.

To help you optimally manage the remediation conditions, Patch Management allows you to specify remediate options and pre and post actions, and set up ticket IDs and email notifications that alert you about the status of the remediate process. The Remediate window guides you through setting up these conditions.

Figure 65 Remediate Window



Remediating Patch Policies

This action installs the patches in a policy that has been attached to managed servers. (This action does not uninstall patches.) A patch policy can be overridden by an exception, which indicates that a patch is either always or never installed on a particular server.

When you invoke the Windows patch remediation process for a group of servers, patches will only be remediated for servers where:

- The SA Agent is from SA 5.5 or later; and

- The server is running Windows Server 2000 SP3 (or higher), Windows Server 2003, Windows Server 2008, Windows XP SP2 (or higher), Windows Server 2003 x_64, or Windows Server 2008 x_64.

To remediate a patch policy, perform the following steps:

- 1 From the Navigation pane, select **Library** ► **By Type** ► **Patch Policies**
- 2 Expand **Patch Policies** and select a specific Windows operating system. The Content pane will display all patch policies associated with that operating system.
- 3 From the Content pane, open a patch policy.
- 4 From the View drop-down list, select **Server Usage**.
- 5 From the Show drop-down list in the Content pane, select **Servers with Policy Attached**.
- 6 From the Preview pane, select one or more servers.
- 7 From the **Actions** menu, select **Remediate**. The first step of the Remediate window appears: **Servers and Device Groups**.

For instructions on each step, see the following sections:

- [Setting Remediate Options](#)
- [Setting Reboot Options for Remediation](#)
- [Specifying Pre and Post Install Scripts for Remediation](#)
- [Scheduling a Patch Installation for Remediation](#)
- [Setting Up Email Notifications for Remediation](#)
- [Previewing a Remediation](#)

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8 Click **Start Job** to launch the remediation job.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If you leave the Remediate window open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select **Refresh** from the **View** menu to update information in the Patch Preview pane.

Setting Remediate Options

You can specify the following remediate policy option:

“Do not interrupt the remediate process even when an error occurs with one of the policies.”

To set this option, perform the following steps:

- 1 From the Remediate window, click **Next** to advance to the Remediate Options step.
- 2 Select one of the following Staged Install Options:

Continuous: Run all phases as an uninterrupted operation.

Staged: Allow download and installation to be scheduled separately.

- 3 Select the Error Options check box if you want the remediation process to continue even when an error occurs with any of the patches or scripts. As a default, this check box is not selected.
- 4 Click **Next** to go to the next step or click **Cancel** to close the Remediate window.

Setting Reboot Options for Remediation

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

You can specify the reboot options in the following two places in the SA Client:

- Install Parameters tab of the patch properties window
- Pre & Post Actions step of the Remediate window



When you are selecting reboot options in the Remediate window, Hewlett Packard recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option. If it is not possible to use the Microsoft reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option. Failure to do this can result in the MBSA incorrectly reporting which patches are installed on the server until the next reboot occurs (outside of SA control).

The following options in the Remediate window determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Remediate window; they do not change the Reboot Required option, which is on the Install Parameters tab of the Patch Properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Hold all server reboots until after all packages are installed and/or uninstalled:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1 From the Remediate window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select one of the Reboot Options.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Remediate window.

Specifying Pre and Post Install Scripts for Remediation

For each patch remediation, you can specify a command or script to run before or after remediation. A pre-install script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-install script fails, the patches would not be

installed. A pre-install script could also be used to shut down a service or application before it is patched. A post-install script could be used to perform a certain cleanup process on the managed server.

You can specify the following types of scripts to run on the managed server before or after a remediation process:

- **Pre-Download:** A script that runs before patches are downloaded from SA to the managed server. This is available only if you select Staged in the Remediate Options step.
- **Post-Download:** A script that runs after patches are downloaded from SA to the managed server and before the patch is installed. This is available only if you select Staged in the Remediate Options step.
- **Pre-Install:** A script that runs before patches are installed on the managed server.
- **Post-Install:** A script that runs after patches are installed on the managed server.

To specify a pre-install script, perform the following steps:

- 1 From the Remediate window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select the Pre-Install tab.

You may specify different scripts and options on each of the tabs.

- 3 Select the Enable Script check box. This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
- 4 Select either Saved Script or Ad-Hoc Script from the drop-down list.

A Saved Script has been previously stored in HP Server Automation with the SAS Web Client. To specify the script, click **Select**.

An Ad-Hoc script runs only for this operation and is not saved in HP Server Automation. Select the Type, such as .bat. In the Script box, enter the contents of the script, including the drive letter of where the script is located, such as `echo dir>> C:\temp\preinstall1.log`. If you do not enter a drive letter, the default is %SYSTEMDRIVE%, which is where the system folder of Windows is installed.

- 5 If the script requires command-line flags, enter the flags in the Command text box.
- 6 In the User section, if the system is not Local System, select Name.
- 7 Enter the system name, your password, and the Domain name.
- 8 To stop the installation if the script returns an error, select the Error check box.
- 9 Click **Next** to go to the next step or click **Cancel** to close the Remediate window

Scheduling a Patch Installation for Remediation

You can schedule when you want patches installed and when you want patches downloaded.

To schedule a patch installation, perform the following steps:

- 1 From the Remediate window, select the Scheduling step. To reach this step, you must have completed the Pre & Post Actions step.

By default, the Scheduling step displays only the scheduling options for the installation phase. If you selected Staged in the Remediate Options step, the scheduling options for the download phase will also be displayed.



- 2 Select one of the following Install Phase options:

- **Run Task Immediately:** This enables you to perform the download or installation immediately.
 - **Run Task At:** This enables you to specify a date and time that you want the download or installation performed.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Remediate window.

Setting Up Email Notifications for Remediation

You can set up email notifications to alert users when the download and installation operations complete successfully or with errors.

To set up email notifications, perform the following steps:

- 1 From the Remediate window, click **Next** to advance to the Notifications step.
- 2 To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3 To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase. If you selected Staged in the Remediate Options step, the notification status for the download phase is also displayed.
- 4 Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5 Click **Next** to go to the next step or click **Cancel** to close the Remediate window.



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

Previewing a Remediation

The remediate preview process provides an up-to-date report about the patch state of servers. The remediate preview is an optional step that lets you see the patches that will be installed on managed servers. This preview process verifies whether the servers you selected for the patch installation already have that patch installed (based MBSA 2.1). In some cases, a server could already have the patch installed if a system administrator had manually installed it, which means that Patch Management does not know about it.

In the Preview, the servers, device groups, and patches that are listed in the Summary Step window will be submitted to remediation when you click **Start Job**. Patches that are not recommended by the vendor will be excluded from this list. If there are other patches in the policy with the same QNumber, only the vendor-recommended patch is displayed.

This list shows patches and their associated servers (regardless of any patch policy and server group membership changes that may have occurred). If you preview a remediation, this same list of servers, device groups, and patches will be used, even if changes have occurred to the patch policy or server group memberships.

If you modify parameters in the Remediate window after you have already clicked **Preview**, the preview process will produce an invalid summary of simulated patching actions. For example, if you have already clicked **Preview** and you add patches, patch policies, servers, or device groups, you must click **Preview** again for results that include your changes.



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

To preview a remediation, perform the following steps:

- 1 From the Remediate window, click **Next** to advance to the Summary Review step.
- 2 Verify the information displayed for the Servers, Device Groups, and Patches at the top of the window.
- 3 (Optional) Click **Preview** to see the separate actions that will be performed when the patch is installed. To view the details of a previewed action, select a row in the table.
- 4 To launch the installation job, click **Start Job**.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected a specific time, the job will run then.

- 5 The Job Progress displays in the Remediate window.

The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:

- **Analyze:** HP Server Automation examines the patches needed for the installation, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.
 - **Download:** The patch is downloaded from HP Server Automation to the managed server.
 - **Install:** After it is downloaded, the patch is installed.
 - **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
 - **Run Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the installation.
 - **Install & Reboot:** When a patch will be installed is also when the server will be rebooted.
 - **Verify:** Installed patches will be included in the software registration.
- 6 To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See the *SA Users Guide: Server Automation* for more information on browsing job logs.
 - 7 Click **Stop Job** to prevent the job from running or click **Close** to close the Remediate window. You can stop a job only if it is scheduled.

Verifying Patch Policy Compliance

To determine whether a managed server complies with patch policies and exceptions, perform the following steps:

- 1 From the Navigation pane, select **Devices ► All Managed Servers**.
- 2 From the Content pane, select **Patches** from the View drop-down list.
- 3 Examine the **Patch** column at the top of the pane. This column indicates the overall patch compliance for a server.
- 4 Select a server at the top of the Content pane and examine the **Compliance** column at the bottom. This column indicates the compliance status of each individual patch for the selected server.

Creating a Patch Policy

A patch policy is a set of patches that should be installed on a managed server. When it is first created, a patch policy contains no patches and is not attached to servers.

To create a patch policy, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies**.
- 2 Select a specific Windows operating system.
- 3 From the **Actions** menu, select **Create Patch Policy**.

The name of the policy you just created is **New Patch Policy n**, where **n** is a number based on the number of New Patch Policies already in existence.

- 4 From the Content pane, open the New Patch Policy.
- 5 (Optional) In the **Name** field of the Properties, enter a name that describes the purpose or contents of the policy.

Deleting a Patch Policy

This action removes a patch policy from HP Server Automation but does not remove or uninstall patches from managed servers. You cannot delete a patch policy if it is attached to servers or groups of servers. You must first detach the policy from the servers or groups of servers before removing it from HP Server Automation.

To delete a patch policy from HP Server Automation, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies**.
- 2 Select a specific Windows operating system.
- 3 From the Content pane of the main window, select a policy.
- 4 From the **Actions** menu, select **Delete Patch Policy**.

Adding a Patch to a Patch Policy

This action adds a patch to a patch policy, but does not install the patch on a managed server. The patch will be installed when the policy is remediated.

To add a patch to a patch policy, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies**.
- 2 Select a specific Windows operating system and view the list of Windows patches.
- 3 From the Content pane, select the patch.
- 4 From the View drop-down list, select **Patch Policies**.
- 5 From the Show drop-down list, select **Policies without Patch Added**.
- 6 Select a policy. From the **Actions** menu, select **Add to Patch Policy**.
- 7 In the Add to Patch Policy window, click **Add**.

Removing a Patch from a Patch Policy

This action only removes a patch from a patch policy. This action does not uninstall the patch from a managed server and does not remove the patch from HP Server Automation.

To remove a patch from a patch policy, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**.
- 2 Select a specific Windows operating system and view the list of Windows patches.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list, select **Patch Policies**.
- 5 From the Show drop-down list, select **Policies with Patch Added**.
- 6 Select a patch. From the **Actions** menu, select **Remove from Patch Policy**.
- 7 In the Remove Patch from Policy window, select the policy and click **Remove**.

Attaching a Patch Policy to a Server

This action associates a patch policy with a server (or group of servers). You must perform this action before you remediate a policy with a server (or group of servers).

To attach the policy, perform the following steps:

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

- 7 From the **Actions** menu, select **Attach Server**.
- 8 Click **Attach**.

Detaching a Patch Policy from a Server

This action does not delete the patch policy and does not uninstall patches from a managed server.

To detach the policy, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies**.
- 2 Select a specific Windows operating system and view the list of Windows patch policies.
- 3 From the Content pane, select a patch policy.
- 4 From the View drop-down list, select Server Usage (or Device Group Usage).
- 5 From the Show drop-down list, select Servers with Policy Attached (or Server Groups with Policy Attached).
- 6 From the Preview pane, select one or more servers.
- 7 From the **Actions** menu, select **Detach Server**.
- 8 Click **Detach**.

Setting a Patch Policy Exception

A patch policy exception indicates whether the patch is installed during the remediation process. (The Install Patch and Uninstall Patch actions ignore patch policy exceptions.) A patch policy exception overrides the policy. You can specify an exception for a particular patch and server (or group of servers), but not for a patch policy.

To set a patch policy exception, perform the following steps:

- 1 From the Navigation pane, select **Devices ► All Managed Servers**.
- 2 Select a server.
- 3 From the Content pane, select a server.
- 4 From the View drop-down list, select **Patches**.
- 5 From the Preview pane, select a patch.
- 6 From the **Actions** menu, select **Set Exception**.
- 7 In the Set Policy Exception window, select the Exception Type:
 - **Never Install**: The patch should not be installed on the server, even if the patch is in the policy.
 - **Always Install**: The patch should be installed on the server even if the patch is not in the policy.
- 8 (Optional) In the Reason field, enter an explanation. This explanation is displayed when you move the cursor over the Exception column in the Preview pane. The Patches with Exceptions option must be selected. When you are finished, click **OK**.

Finding an Existing Patch Policy Exception

You can search for managed servers that already have patch policy exceptions attached to them, and you can search for patches that have exceptions.

To find an existing patch policy exception, perform the following steps:

- 1 From the Navigation pane, select **Devices ► All Managed Servers**.
- 2 From the View drop-down list, select **Patches**.
- 3 From the Content pane, select a server.
- 4 From the Show drop-down list, select **Patches with Policies or Exceptions** or **Patches with Exceptions**.
- 5 In the Exception column, move the cursor over the icon to display the reason for this exception. The following icons indicate the type of patch policy exception:



An always install exception on a patch/server association.



An always install exception inherited to a server from a group of servers/patch association.



A never install exception on a patch/server association.



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

Copying a Patch Policy Exception

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

- 1 From the Navigation pane, select **Library ► By Type ► Patches**.
- 2 Expand the Patches and select a specific Windows operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list, select Server Usage (or Device Group Usage).
- 5 From the Show drop-down list, select Servers with Exception (or Server Groups with Exception).
- 6 From the Preview pane, select a server. This server is the source of the copied exception.
- 7 From the **Actions** menu, select **Copy Exception**.
- 8 In the Copy Policy Exception window, select the target servers or device groups.

These servers are the destinations of the copied exception. If this operation would result in replacing an existing exception, a message displays asking you to confirm whether this is the preferred action.

- 9 Click **Copy**.

Removing a Patch Policy Exception

To remove a patch policy exception, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**.
- 2 Expand the Patches and select a specific Windows operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list, select **Servers**.
- 5 From the Show drop-down list, select **Servers with Exception**.
- 6 From the Preview pane, select a server.
- 7 From the **Actions** menu, select **Remove Exception**.

Patch Compliance

Patch Management performs conformance tests (compliance checks) against managed servers and public device groups to determine whether all patches in a policy and a policy exception were installed successfully. To optimize patch compliance information for your organization, you can set the patch compliance levels and edit the rules of the customized patch compliance level.

Patch Compliance Scans

A patch compliance scan compares patches that are installed on a server with patch policies and patch policy exceptions that are attached to that server. The results of this scan show you the servers that are in compliance (have all required patches installed) and the servers that are out of compliance (do not have all required patches installed).

You should run or schedule patch compliance scans based on the dynamics of your patching environment. For example, if you updated a patch policy or installed a patch outside of (by not using) HP Server Automation, a compliance scan is required because the SA model has been changed and the compliance information must now be recalculated. Patch Management indicates these types of conditions by displaying Scan Needed. In this case, instead of waiting for the scan schedule to iterate, you can start compliance scan on one or more servers.

Ways to Start a Patch Compliance Scan

You can start a patch compliance scan in the following ways:

- Immediately, by selecting servers or groups and then selecting a menu item. See [Starting a Patch Compliance Scan Immediately](#) on page 324.
- Periodically, by setting up a schedule. See [Scheduling a Patch Compliance Scan](#) on page 329. By default, the scans are not scheduled.
- As a result of another task. HP Server Automation performs a patch compliance scan on a managed server at the end of the tasks described in the following sections:
 - [Installing a Windows Patch](#) on page 335

- [Uninstalling a Windows Patch](#) on page 343
- [Remediating Patch Policies](#) on page 313

Starting a Patch Compliance Scan Immediately

To start a scan on selected servers, perform the following steps:


- 1 From the Navigation pane, select **Devices**.
- 2 Select an entry from either the Managed Servers or Device Groups list.
- 3 Right-click and select **Scan ► Patch Compliance**.

Refreshing the Compliance Status of Selected Servers

You can refresh the compliance status of all Windows servers by selecting **View ► Refresh**. However, this global refresh operation can take a long time when scanning a large number of servers. To save time, you can refresh the compliance status of selected servers by performing the following steps:

- 1 From the Navigation pane, select **Devices**.
- 2 Drill down to the servers you want to check.
- 3 In the Contents pane, select one or more servers
- 4 Right-click and select **Refresh Server Status**.
- 5 Note any changed values in the Patch column.

Viewing Scan Failure Details





If the scan operation fails, you cannot determine whether a server is in compliance. A scan failure is indicated by the  icon. To find out why a patch compliance scan failed, perform the following steps:

- 1 From the Navigation pane, select **Devices**.
- 2 Drill down to the server you want to check.
- 3 In the Contents pane, select a server.
- 4 Right-click and select **Scan ► Show Patch Compliance Scan Failure Details**.
- 5 In the Patch Compliance Scan Failure Details window, select a server and examine the detailed error message that appears in the lower part of the window.

Patch Compliance Icons

Patch Management displays the following icons in [Table 18](#).

Table 18 Patch Compliance Status Icons

Status/Icon	Description
 Compliant	The server is compliant for all patches. Patches in policies attached to the server are all installed on the target server.
 Partial	The server is partially compliant for patches. An exception has been set for these patches.
 Non-Compliant	The installed patches on the server do not match the conditions defined in the patch policy.
 Scan Failed	The scan operation failed. Patch Management is unable to check the compliance of the server.

About Patch Non-Compliance

A Patch non-compliant status for a server or group of servers can be caused by different factors, such as the existence of applicable patches that need to be installed as defined in a patch policy attached to the servers. Or, there could be exceptions that affect a server patch compliance level.

For example, a server will be considered non-compliant if the patch policy has a patch marked as a “Never Install” exception but the target server does have that patch installed.

Also, if superseded patches are recommended and included in policies or exception, they are counted in the compliance calculations, and if they are missing on the target server, then the server's patch compliance status will be non-compliant.

Patch Compliance Levels

Patch compliance levels define your patch compliance rules. Results of a patch compliance scan can include only policies, both policies and exceptions, or your own customized level.

Patch Management supports the following compliance levels:

- **Policy Only:** Verifies whether the patches installed on a server comply with the patch policies.
- **Policy and Exception:** Verifies whether the patches installed on a server comply with the patch policies and any exceptions. The Partial (yellow) icon is displayed if the policy and exception do not agree and the exception does not have data in the Reason field.
- **Customized:** Verifies the rules that you edited for this compliance level.

Patch Compliance Rules

Patch compliance rules are the conditions that determine the compliance icons that are displayed in the Managed Server window.

Patch Management supports the following compliance rules:

- **Patch Added to Policy:** The patch has been added to the patch policy.
- **Patch Installed on Server:** The patch has been installed on the managed server.
- **Exception Type:** The Exception Type can have the following values:
 - **Always Installed:** The patch should be installed on the server, even if the patch is not in the policy.
 - **Never Installed:** The patch should not be installed on the server, even if the patch is in the policy.
 - **None:** An exception has not been specified for the patch and server.
- **Exception Reason:** A description entered in the Exception Reason of the Set Policy Exception window. In the Patch Compliance Rules window, the Exception Reason can have the following values.
 - **Yes:** The Exception Reason has data.
 - **No:** The Exception Reason is empty.
 - **N/A:** An exception has not been specified for the patch and server.
- 6 **Compliance Result:** The icon that indicates the result of the patch compliance scan. These icons are displayed in the Managed Server window.

Patch Compliance Reports

To help troubleshoot problems, you can run and examine several patch compliance reports that are based on Sarbanes-Oxley (SOX) standards. These reports identify whether all patches in a policy and a policy exception were installed successfully on managed servers. The Reports feature of the SA Client provides the following patch compliance reports.

- **Defined Patch Policies:** Lists patch policies by name, customer, and operating system, and includes the total number of patch policies.
- **Patch Policy Compliance (All Servers):** Groups all managed servers by their patch policy compliance level to show compliant and non-compliant servers.
- **Patch Policy Compliance by Customer:** Lists all servers by the customer they belong to and then by the patch policy compliance level.
- **Patch Policy Compliance by Facility:** Groups all managed servers by the facility they belong to and then by the patch software policy compliance level.
- **Servers in Compliance With Their Patch Policies:** Lists all managed servers that are in compliance with all of their attached patch policies.
- **Servers Not in Compliance With Their Patch Policies:** Lists all managed servers that are not in compliance with their attached patch policies.
- **Servers With Attached Patch Policies:** Lists all managed servers that have one or more patch policies attached, and includes the total number of servers with attached patch policies.
- **Servers Without Attached Patch Policies:** Lists all managed servers that do not have any patch policies attached, and includes the total number of servers without any attached patch policies.



See the *SA Users Guide: Server Automation* for information about how to run, export, and print these reports.

Patch Administration for Windows

You can customize patch administration for Windows to best support your environment in the following manner:

- You can specify whether you want patches immediately available for installation by using a command-line script or the SA Client.
- You can import the Microsoft patch database (on demand) by using a command-line script or the SA Client.
- You can track (and import) only patches that apply to certain Microsoft products or particular locales.
- You can import and export Windows patch utilities.
- You can manually launch (on demand) or schedule periodic policy compliance scans to determine the patch state of your managed servers.
- You can customize the icon display of policy compliance scan results.

Setting the Patch Availability

You can set the default patch availability with either the SA Client or a command-line script. The default used by the script overrides the default set by the SA Client. For information about the script, see [Automatically Importing Windows Patches](#) on page 306.

To set the default value for the Availability of a newly imported patch, perform the following steps:

- 1 From the Navigation pane, select Opsware Administration.
- 2 Select Patch Settings.
- 3 For the Patch Availability for Imported Patches, select either Available or Limited. The default is Limited.

If the patch is Available, it can be installed on managed servers. If the patch is Limited, it has been imported into HP Server Automation and can be installed only by a patch administrator who has the required permissions. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide* for an explanation of these permissions.

Importing the Microsoft Patch Database

You can import the Microsoft Patch Database by using a command-line script or the SA Client. For information about the script, see [Automatically Importing Windows Patches](#) on page 306.

To import the database with the SA Client, perform the following steps:

- 1 From the Navigation pane, select Opsware Administration.
- 2 Select Patch Settings.
- 3 To import the database from the Microsoft web site, click **Import from Vendor**.
A window appears with the default URL for the location of the database on the Microsoft web site. Click **Import**. To re-import a new version of the Microsoft database that is released monthly, you must use the default URL.
- 4 To import the database from the local file system, click **Import from File**.
A file browser window appears. Go to the folder containing the `wsusscan.cab` (MBSA 2.1) file and click **Import**. This file must have been previously downloaded from the Microsoft web site and copied to the local file system.



To be imported, a patch must be in the Microsoft Patch database that has already been imported into the Software Repository.

Selecting Windows Products to Track for Patching

This operation limits the patches tracked by HP Server Automation to specific Windows products. After performing this operation, the next time the Microsoft Patch Database is imported, any new patches listed by HP Server Automation are limited to the products that you select. Patches that were previously listed by HP Server Automation are still tracked. You can also track patches for all MBSA 2.1 products.

To limit the patches tracked to specific Windows operating systems, run the command-line script that automatically imports patches. For more information about the script, see [Automatically Importing Windows Patches](#) on page 306.

To select the Windows products to track for patching, perform the following steps:

- 1 From the Navigation pane, select Opsware Administration.
- 2 Select Patch Settings.
- 3 Select the Windows MBSA tab.
- 4 Click **Edit**.
- 5 In the Edit Patch Properties window, use the include and exclude arrows to select the products whose patches you want to track and then click **Select**.

Scheduling a Patch Compliance Scan

To schedule a patch compliance scan on all Windows managed servers, perform the following steps:

- 1 From the Navigation pane, select Opware Administration.
- 2 Select Patch Compliance Settings.
- 3 In the Patch Policy Compliance Scan Schedule section, click **Edit**.
- 4 In the Schedule Compliance Scan window, select Enable Compliance Scan.
- 5 In the Schedule drop-down list, select the frequency of the scans.

If you select Custom, specify the crontab string with the following values:

- Minute (0-59)
- Hour (0-23)
- Day of the month (1-31)
- Month of the year (1-12)
- Day of the week (0-6 with 0=Sunday)
- Any of these fields can contain an asterisk to indicate all possible values. For example, the following crontab string runs the job at midnight every weekday:

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

The crontab string can also handle serial (1,2,3,4) as well as range (1-5) values. For more information, consult the crontab man pages on a Unix computer.

- 6 In the Start Time field, specify the time you want the job to begin.
- 7 In the Time Zone drop-down list, select a default time zone for the job execution time or accept the default time zone. The default time shown converts the scheduled time to the time zone set in your user preferences. If you do not set a preferred time zone, the time zone is derived from the HP Server Automation core server, which is typically UTC.
- 8 In the Day(s) to Run field, select one or more days of the week that you want the scan to run.
- 9 Click **OK**.

Setting the Patch Policy Compliance Level

The patch policy compliance level defines your patch compliance rules. To view these rules or to set the patch policy compliance level, perform the following steps:

- 1 From the Navigation pane, select Opware Administration.
- 2 Select Patch Compliance Settings.
- 3 Select one of the following compliance levels: Policy and Exception, Policy Only, or Customized.

Importing Windows Patch Utilities

You can import the following Windows utilities from your local file system into HP Server Automation:

- parsembsaccli20.exe
- qchain.exe
- WindowsUpdateAgent-x86.exe
- WindowsUpdateAgent-x64.exe
- WindowsUpdateAgent-ia64.exe
- wusscan.dll

Initially, these files are imported into HP Server Automation during the installation of the core. To import a Windows patch utility, perform the following steps:

- 1 From the Navigation pane, select Opware Administration.
- 2 Select Patch Settings.
- 3 In the Patch Utilities section, select a utility and then click **Import Utility Update**.

Exporting Windows Utility Files

You can export the following Windows patch utilities from HP Server Automation to your local file system:

- parsembsaccli20.exe
- qchain.exe
- WindowsUpdateAgent-x86.exe
- WindowsUpdateAgent-x64.exe
- WindowsUpdateAgent-ia64.exe
- wusscan.dll

To export a Windows patch utility, perform the following steps:

- 1 From the Navigation pane, select Opware Administration.
- 2 Select Patch Settings.
- 3 In the Patch Utilities section, select one or more utilities and then click **Export Utility**.

Editing the Customized Patch Policy Compliance Level

Of the three compliance levels, only the Customized level can be edited. To edit this level, perform the following steps:

- 1 From the Navigation pane, select Opware Administration.
- 2 Select Patch Compliance Settings.
- 3 From the Compliance Level, select Customized.
- 4 In the Patch Policy Compliance Setting section, click **Edit**.

- 5 Select the Compliance Level icons that you want to change in the Compliance Result column: Non-Compliant, Compliant, No Indicator, or Partial.
- 6 Click **Apply** and then click **Close**.

Locales for Windows Patching

The locale of a patch identifies the language of the Windows servers that should receive the patch. A patch with the same name might be available for different locales. For example, a patch named Q123456 might be available for servers running the English and Japanese versions of Windows. Although they have the same name, the patches installed on the English and Japanese servers are different binaries.

Patch Management supports multiple locales in the same SA multimaster mesh. To install a patch on Windows servers with different locales, you specify the patch by name. During the installation (or policy remediation), SA matches the locale of the patch with the locale of each managed server. You do not need to repeat the installation for each locale.

Supported Locales

Patch Management supports Windows patches of the following locales:

- English (en)
- Japanese (ja)
- Korean (ko)

Overview of Locale Configuration Tasks

By default, Patch Management supports only the English locale. To set up Patch Management for non-English locales, step through the instructions in the following sections:

- [Configuring the SA Core for Non-English Locales](#) on page 331
- [Selecting the Locales of Patches to Import](#) on page 332
- [End User Requirements for Non-English Locales](#) on page 333

Configuring the SA Core for Non-English Locales

This task requires `root` access to core servers and a restart the OCC core component. To configure the core for non-English locales, perform the following steps on each core server running the OCC component:

- 1 Log onto the server as `root`.
- 2 With a text editor, in `/etc/opt/opsware/occ/psrvr.properties`, change the line for `pref.user.locales` to the following:
`pref.user.localesAllowed=en;ja;ko`

- 3 Restart the OCC component of the core:

```
/etc/init.d/opsware-sas restart occ.server
```

- 4 In a text editor, open the following file:

```
/opt/opsware/occclient/jnlp.tpl
```

- 5 For the Japanese language, In the <resources> section of the `jnlp.tpl` file, add the following XML element:

```
<property name="com.opsware.ngui.font.japanese" value="Arial Unicode MS"/>
```

- 6 For the Korean language, In the <resources> section of the `jnlp.tpl` file, add the following XML element:

```
<property name="com.opsware.ngui.font.korean" value="Arial Unicode MS"/>
```

- 7 In the `/opt/opsware/occclient` directory, if the following files exist, delete them:

```
$HOST_ja.jnlp  
$IP_ja.jnlp  
$HOST_ko.jnlp  
$IP_ko.jnlp
```

- 8 Follow the steps in [Selecting the Locales of Patches to Import](#) on page 332.

Selecting the Locales of Patches to Import

Follow the instructions in [Configuring the SA Core for Non-English Locales](#) on page 331 before performing the steps in this section.

This operation selects the locales of the Windows patches to import into HP Server Automation. The selections take effect the next time patches are imported into HP Server Automation. After the patches have been imported, they can be installed on managed servers. If you remove locales from the list with this operation, patches with those locales that have already been imported are not removed from HP Server Automation.

To select the locales of the Windows patches to import into SA, perform the following steps:

- 1 In the SA Client, from the Navigation pane, select Opsware Administration.
- 2 Select Patch Settings.
- 3 On the Windows MBSA tab, select Patch Locales.
- 4 Click **Edit**.
- 5 In the Edit Patch Locales window, use the include and exclude arrows to select the locales whose patches you want to import. If you want to select a locale that is not listed in [Supported Locales](#) on page 331, contact support.
- 6 Click **Select**.
- 7 Follow the instructions in [End User Requirements for Non-English Locales](#) on page 333.

End User Requirements for Non-English Locales

To view non-English fonts in the SA Client, end users must perform the following steps:

- 1 The end user verifies that the Windows desktop running the SA Client uses the Arial Unicode MS font.
- 2 After the SA Administrator performs the steps in [Configuring the SA Core for Non-English Locales](#) on page 331, the end user logs onto the SAS Web Client and goes to the My Profile page,
- 3 On the My Profile page, the end user updates the Locale field on the User Identification tab. For example, if the SA Administrator configured the core for Japanese, then the end user sets the Locale field to Japanese.

Patch Installation

Patch Management provides the following two phases in the patch installation process:

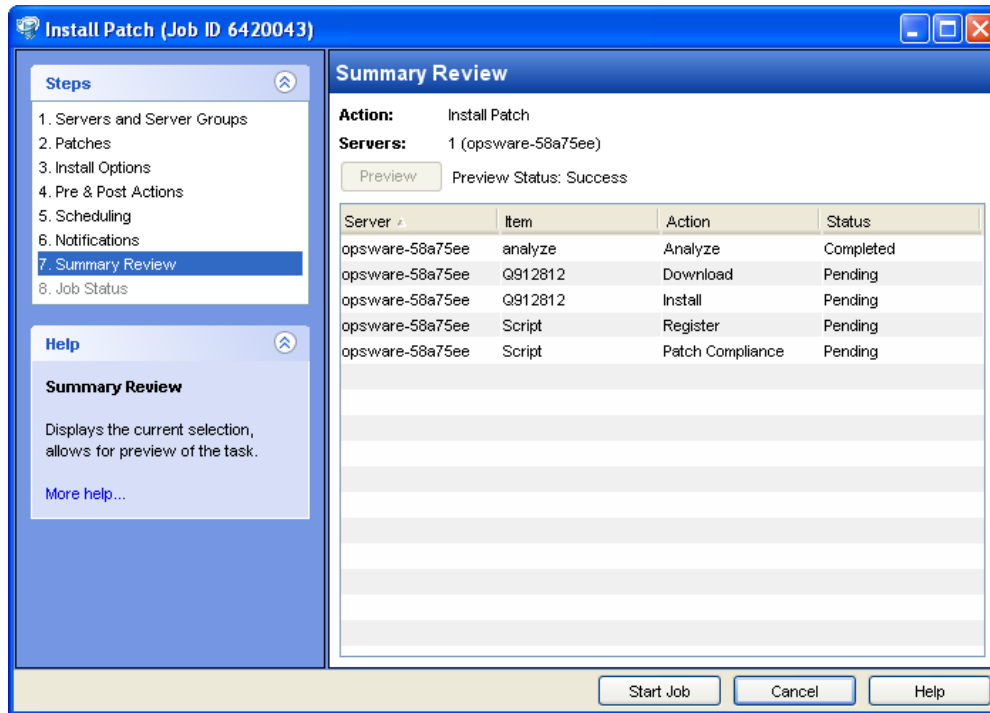
- **Download Phase:** This is when the patch is downloaded from HP Server Automation to the managed server. This phase is commonly referred to as the staging phase.
- **Installation Phase:** This is when the patch is installed on the managed server. This phase is commonly referred to as the deployment phase.

You can specify whether you want the installation to occur immediately after the patch is downloaded (staged) or you can schedule the installation to occur at a later date and time. Patch Management also supports the need for best-effort installations of multiple patches by allowing you to specify that the patch installation process will continue even when an error occurs with one of the patches.

Patch Management displays the name of the command (.exe file and any predefined command-line arguments) that the SA Agent runs on the managed server to install the patch. You can override these default command-line arguments.

To help you optimally manage Windows patch installation, Patch Management allows you to manage server reboot options, specify pre and post installation scripts, simulate (preview) a patch installation, and set up email notifications to alert you about the status of the installation process. The Install Patch window guides you through setting up these conditions.

Figure 66 Install Patch Window



Installation Flags

You can specify installation flags that are applied whenever a Windows patch is installed. However, HP Server Automation also uses default installation flags and requires that patches are installed with these flags. You must therefore be certain that you do not specify any installation flags that override or contradict the default flags passed by HP Server Automation. See [Setting Windows Install Options](#) on page 336 for information about how to specify commands and flags.



Some Windows hotfixes do not support the `-z` flag, some do not support the `-q` flag, and some do not support either. In such cases, you must use a special expression: `/-z` or `/-q` or `/-z -q` respectively. This prevents the Patch Management feature from passing in the `-z` or `-q` or `-z -q` flag. By default, HP Server Automation adds `/z /q` to the command line arguments when installing patches. To override this, specify `/-z /-q`. For example, if you prefer to not suppress the reboot, specify `/-z`.

The following table lists the default installation flags that HP Server Automation uses.

Table 19 Default Installation Flags

Windows Patch Type	Flags
Windows Hotfix	<code>-q -z</code>
Windows Security Rollup Package (treated identically to a Hotfix by the Patch Management feature)	<code>-q -z</code>
Windows OS Service Pack	<code>-u -n -o -q -z</code>

Application Patches

The Patch Management feature does not allow you to apply a patch to an operating system for which the patch is not intended. When you are installing an application patch, Patch Management does not automatically filter out servers that do not have the corresponding application installed. Although Patch Management does not prevent you from doing so, you should not attempt to apply application patches to servers that do not have the necessary applications installed. If a patch is for an application that is not installed on the server, the patch will not be applied and an error message will display, such as “There was an error with package <name of the package>”.

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

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

Service Packs, Update Rollups, and Hotfixes

When you try to install a Service Pack, Update Rollup, or a Hotfix, there is a known delay when a confirmation dialog displays. Since the SA Agent is installing or uninstalling the patch, it cannot respond to the confirmation dialog. The Agent will time out an installation or uninstallation process if you do not click **OK** in the confirmation dialog. For Hotfixes, the Agent will time out if five minutes have lapsed and you have not clicked **OK** in the confirmation dialog. For Service Packs and Update Rollups, the Agent will time out if 60 minutes have lapsed and you have not clicked **OK** in the confirmation dialog.

To prevent this from happening, patch install and uninstall commands should have arguments that invoke silent mode installs and uninstalls. By default, the -q flag is set.

Installing a Windows Patch

Before a patch can be installed on a managed server, it must be imported into HP Server Automation and its status must be Available. Administrators who have the required permissions can install patches that are marked Limited.



You must have a set of permissions to manage patches. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide*.

You can perform the installation by explicitly selecting patches and servers, and you can install a patch even if the patch policy exception is Never Install.

To install a patch on a managed server, perform the following steps:

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

- 2 Expand the Patches and select a specific Windows operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list, select Servers (or Device Groups).
- 5 From the Show drop-down list, select Servers without Patch Installed (or Device Groups without Patch Installed).
- 6 From the Preview pane, select one or more servers.
- 7 From the **Actions** menu, select **Install Patch**.

The first step of the Install Patch window appears: Servers and Device Groups. For instructions on each step, see the following sections:

- [Setting Windows Install Options](#)
- [Setting Reboot Options for a Windows Patch Installation](#)
- [Specifying Install Scripts for a Windows Patch Installation](#)
- [Scheduling a Windows Patch Installation](#)
- [Setting Up Email Notifications for a Windows Patch Installation](#)
- [Previewing a Windows Patch Installation](#)
- [Viewing Job Progress of a Windows Patch Installation](#)

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8 When you are ready to launch the installation job, click **Start Job**.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If the Install Patch window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select **Refresh** from the **View** menu to update information in the Patch Preview pane.

See [Remediating Patch Policies](#) on page 313 for another method of installing a patch.

Setting Windows Install Options

You can specify the following types of patch installation options:

- Perform the patch installation immediately after the patch is downloaded or at a later date and time.
- Do not interrupt the patch installation process even when an error occurs with one of the patches.
- Use different command-line options to perform the installation.

To set these options, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Install Options step.
- 2 Select one of the following Staged Install Options:
 - **Continuous:** This allows you to run all phases as an uninterrupted operation.

- **Staged:** This allows you to schedule the download and installation to run separately.
- 3 Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.
 - 4 In the Install Command text box, enter command-line arguments for the command (.exe file) that is displayed. By default, HP Server Automation adds /z /q. If you want to override these install flags, enter /-z /-q in the text box.
 - 5 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Setting Reboot Options for a Windows Patch Installation

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is installed on it, completely suppress all server reboots, or postpone reboots until all patches have been installed.



When you are selecting reboot options in the Install Patch window, Hewlett Packard recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option. If it is not possible to use the Microsoft reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option. Failure to do this can result in MBSA incorrectly reporting the patches that are installed on the server until the next reboot occurs (outside of SA control).

The following options determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Install Patch window; they do not change the Reboot Required option, which is on the Install Parameters tab of the patch properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install:** Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Hold all server reboots until after all packages are installed and/or uninstalled:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select one of the Rebooting Options.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Specifying Install Scripts for a Windows Patch Installation

For each patch, you can specify a command or script to run before installation or after installation. A pre-install script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-install script fails, the patch would not be installed. A pre-install script could also be used to shut down a service or application before it is patched. A post-install script could be used to perform a certain cleanup process on the managed server.

You can also specify the following types of scripts to run on the managed server before or after an installation or download phase:

- **Pre-Download:** A script that runs before patches are downloaded from SA to the managed server. This is available only if you select Staged in the Install Options step.
- **Post-Download:** A script that runs after patches are downloaded from SA to the managed server and before the patch is installed. This is available only if you select Staged in the Install Options step.
- **Pre-Install:** A script that runs before patches are installed on the managed server.
- **Post-Install:** A script that runs after patches are installed on the managed server.

To specify a pre-install script, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select the Pre-Install tab. You may specify different scripts and options on each of the tabs.
- 3 Select Enable Script. This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
- 4 Select either Saved Script or Ad-Hoc Script.

A Saved Script has been previously stored in HP Server Automation with the SAS Web Client. To specify the script, click **Select**.

An Ad-Hoc script runs only for this operation and is not saved in HP Server Automation. Select the Type, such as .bat. In the Script box, enter the contents of the script, including the drive letter of where the script is located, such as `echo dir>> C:\temp\preinstall1.log`. If you do not enter a drive letter, the default is %SYSTEMDRIVE%, which is where the system folder of Windows is installed.

- 5 If the script requires command-line flags, enter the flags in the Command text box.
- 6 Specify the information in the User section. If you choose a system other than Local System, enter the User Name, Password, and Domain. The script will be run by this user on the managed server.
- 7 To stop the installation if the script returns an error, select the Error check box.
- 8 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Scheduling a Windows Patch Installation

Since the two phases of patching can be decoupled, you can schedule that you want patches installed independently of when patches are downloaded.

To schedule a patch installation, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Scheduling step.
By default, the Scheduling step displays only the scheduling options for the installation phase. If you selected Staged in the Install Options step, the scheduling options for the download phase will also be displayed.
- 2 Select one of the following Install Phase options:
 - **Run Task Immediately:** This enables the system to perform a preview analysis in the Summary Review step. The scheduling option for the download phase is **Run Immediately Following Download**.
 - **Run Task At:** This enables you to specify a later date and time that you want the installation or download performed.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.





A scheduled patch installation can be cancelled (prior to its execution), even if the patch download has already completed.

Setting Up Email Notifications for a Windows Patch Installation

You can set up email notifications to alert users when the download and installation operations complete successfully or with errors.

To set up email notifications, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Notifications step.
- 2 To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3 To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase.
- 4 Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.



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

Previewing a Windows Patch Installation

The installation preview process provides an up-to-date report about the patch state of servers. The installation preview is an optional step that lets you see the patches that will be installed on managed servers and the type of server reboots that are required. This preview process verifies whether the servers that you selected for the patch installation already have

that patch installed (based on the MBSA). In some cases, a server could already have the patch installed if a system administrator had manually installed it, which means that Patch Management does not know about it.

The preview process also reports on dependency and supersedence information, such as patches that require certain Windows products, and patches that supersede other patches or are superseded by other patches. If a dependency is not met, Patch Management will display an error message indicating this condition.

For example, if a managed server is not running Windows Server 2000 Service Pack 3 (or higher), Windows Server 2003 x_86/x_64, or Windows Server 2008 x_86/x_64, and an HP Server Automation 5.5 Agent, Patch Management will report that a dependency has not been fulfilled. If you try to install a patch for Service Pack 4 and your server is using Service Pack 3, the remediate preview will display a “Will Not Install” error message to indicate this discrepancy. The Install Patch window allows superseded patches to be installed.

Cases in which a patch will not be installed, as displayed in the Preview step of the Install Patch or Remediate Patch Window:

- This patch has a Never Install patch policy exception, so it will not be installed.
- This patch is superseded by another patch in the same job, so it will not be installed. This means that another patch in the current job is more up to date than the marked patch.
- This patch is superseded by another patch, so it will not be installed. This means that the patch installed on the server is more recent than the patch in the policy, and thus will not be installed.
- This patch is not applicable because it is not recommended by MBSA, so it will not be installed.
- This patch is for a different locale, so it will not be installed.



This information is also displayed in the Job results window, and in an email if email notification has been configured for the patch install job.



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

To preview a patch installation, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Summary Review step.
- 2 (Optional) Click **Preview** to see the separate actions that will be performed when the patch is installed. To view the details of a previewed action, select a row in the table.
- 3 Click **Start Job** to launch the installation job or click **Cancel** to close the Install Patch window without launching the installation.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Windows Patch Installation

You can review progress information about a patch installation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Job Progress step. This will start the installation job.

The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:

- **Analyze:** HP Server Automation examines the patches needed for the installation, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.
 - **Download:** The patch is downloaded from HP Server Automation to the managed server.
 - **Install:** After it is downloaded, the patch is installed.
 - **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
 - **Pre/Post Install/Download Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstallation.
 - **Install & Reboot:** When a patch is installed, the server is also rebooted.
 - **Verify:** Installed patches will be included in the software registration.
- 2 To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See the *SA Users Guide: Server Automation* for more information about browsing job logs.
 - 3 Click **Stop Job** to prevent the job from running or click **Close** to close the Install Patch window.

Patch Uninstallation

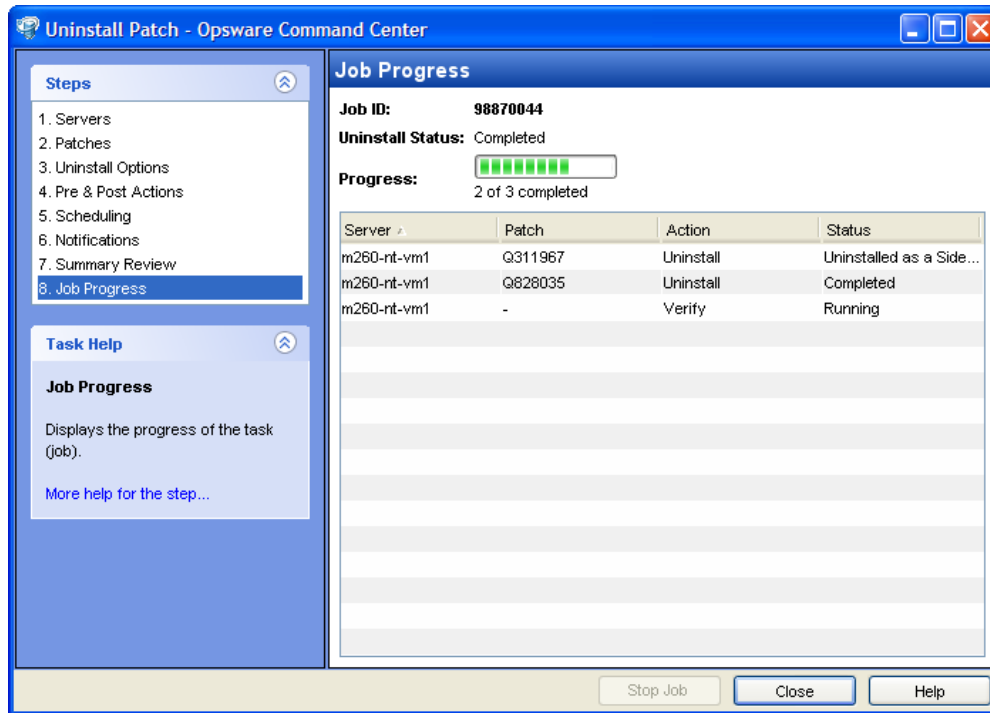
Patch Management provides granular control over how and under what conditions Windows patches are uninstalled (removed) from managed servers. To minimize problems, you can only uninstall one patch at a time. You cannot use HP Server Automation to uninstall a patch that was not installed by using the Patch Management feature.

To help you optimally manage these conditions, Patch Management allows you to do the following:

- Manage server reboot options, and pre and post installation scripts.
- Simulate (preview) a patch uninstallation.
- Set up email notifications to alert you about the status of the uninstallation process.

The Uninstall Patch window guides you through setting up these conditions.

Figure 67 Uninstall Patch Window



Uninstallation Flags

You can specify uninstallation flags that are applied whenever a Windows patch is uninstalled. However, SA also uses default uninstallation flags and requires that patches are uninstalled with these flags. You must therefore be certain that you do not specify any uninstallation flags that override or contradict the default flags passed by HP Server Automation.



Some Windows hotfixes do not support the `-z` flag, some do not support the `-q` flag, and some do not support either. In such cases, you must use a special expression: `/-z` or `/-q` or `/-z -q` respectively, to prevent the Patch Management feature from passing in the `-z` or `-q` or `-z -q` flag. By default, HP Server Automation adds `/z /q` to the command line arguments when uninstalling patches. To override this, specify `/-z /-q`. For example, if you prefer to not suppress the reboot, specify `/-z`.

Table 20 lists the default uninstallation flags used in SA.

Table 20 Default Uninstallation Flags

Windows Patch Types	Flags
Windows Hotfix	<code>-q -z</code>
Security Rollup Package	<code>-q -z</code>
Windows OS Service Pack	Not uninstallable

Uninstalling a Windows Patch

To remove a patch from a managed server, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand the Patches and select a specific Windows operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list, select Servers.
- 5 From the Show drop-down list, select Servers with Patch Installed.
- 6 From the Preview pane, select one or more servers.
- 7 From the **Actions** menu, select **Uninstall Patch**. The first step (Servers) in the Uninstall Patch window appears.
For instructions on each step, see the following sections:

- [Setting Uninstall Options](#)
- [Setting Uninstall OptionsSetting Reboot Options for a Windows Patch Uninstallation](#)
- [Specifying Install Scripts for a Windows Patch Uninstallation](#)
- [Scheduling a Windows Patch Uninstallation](#)
- [Setting Up Email Notifications for a Windows Patch Uninstallation](#)
- [Viewing Job Progress of a Patch Uninstallation](#)

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8 When you are ready to launch the uninstallation job, click **Start Job**.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If the Uninstall Patch window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select Refresh from the View menu to update information in the Patch Preview pane.

Setting Uninstall Options

You can specify the following types of patch uninstallation options:

- Do not interrupt the patch uninstallation process even when an error occurs with one of the patches.
- Use different command-line options to perform the uninstallation.

To set these options, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Uninstall Options step.
- 2 Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.

- 3 In the Uninstall Command text box, enter command-line arguments for the command (.exe file) that is displayed. By default, HP Server Automation adds /z /q. If you want to override these uninstall flags, enter /-z /-q in the text box.
- 4 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Setting Reboot Options for a Windows Patch Uninstallation

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is removed from it, completely suppress all server reboots, or postpone reboots until all patches have been uninstalled.



When you are selecting reboot options in the Uninstall Patch window, Hewlett Packard recommends that you use Microsoft's reboot recommendation. This is the "Reboot servers as specified by patch properties" option. If it is not possible to use the Microsoft reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option. Failure to do this can result in MBSA incorrectly reporting which patches are installed on the server until the next reboot occurs (outside of SA control).

The following options determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Uninstall Patch window; they do not change the Reboot Required option, which is on the Uninstall Parameters tab of the patch Properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install:** Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Hold all server reboots until after all packages are installed and/or uninstalled:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select one of the Rebooting Options.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Specifying Install Scripts for a Windows Patch Uninstallation

For each patch, you can specify a command or script to run before uninstallation or after uninstallation. A pre-uninstall script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-uninstall script fails, the patch would not be removed from a server. A pre-uninstall script could also be used to shut down a service or application before it is removed from a server. A post-uninstall script could be used to perform a certain cleanup process on the managed server.

You can specify the following types of scripts to run on the managed server before or after a patch uninstallation:

- **Pre-Uninstall:** A script that runs before the patch is removed from a managed server.
- **Post-Uninstall:** A script that runs after the patch is removed from a managed server.

To specify a script, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select the Pre-Uninstall or Post-Uninstall tab.
You may specify different scripts and options on each of the tabs.
- 3 Select Enable Script.
This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
- 4 Select either Saved Script or Ad-Hoc Script.
A Saved Script has been previously stored in HP Server Automation with the SAS Web Client. To specify the script, click **Select**.
An Ad-Hoc script runs only for this operation and is not saved in HP Server Automation. Select the Type, such as .bat. In the Script box, enter the contents of the script, including the drive letter of where the script is located, such as `echo dir>> C:\temp\preinstall1.log`. If you do not enter a drive letter, the default is %SYSTEMDRIVE%, which is where the system folder of Windows is installed.
- 5 If the script requires command-line flags, enter the flags in Commands.
- 6 Specify the information in the User section. The script will be run by this user on the managed server.
- 7 To stop the uninstallation if the script returns an error, select Error.

Scheduling a Windows Patch Uninstallation

You can remove a patch from a server immediately, or at a later date and time.

To schedule a patch uninstallation, perform the following steps:



- 1 From the Uninstall Patch window, click **Next** to advance to the Scheduling step.
- 2 Select one of the following Install Phase options:
 - **Run Task Immediately:** This enables you to perform the uninstallation in the Summary Review step.
 - **Run Task At:** This enables you to specify a later date and time that you want the uninstallation performed.

- 3 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Setting Up Email Notifications for a Windows Patch Uninstallation

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

To set up email notifications, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Notifications step.
- 2 To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3 To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the uninstallation phase.
- 4 Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Previewing a Windows Patch Uninstallation

The uninstallation preview process provides an up-to-date report about the patch state of servers. The uninstallation preview is an optional step that lets you see the patches that will be removed from managed servers. This preview process verifies whether the servers you selected for the patch uninstallation have that patch installed (based on the MBSA).



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

To preview a patch uninstallation, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Summary Review step.
- 2 Verify the information displayed for the Servers, Device Groups, and Patches at the top of the window.
- 3 (Optional) Click **Preview** to see the separate actions that will be performed when the patch is uninstalled. To view the details of a previewed action, select a row in the table.
- 4 Click **Start Job** to launch the job or click **Cancel** to close the Uninstall Patch window without launching the uninstallation.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Patch Uninstallation

You can review progress information about a patch uninstallation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Job Progress step. The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:
 - **Analyze:** HP Server Automation examines the patches needed for the uninstallation, checks the managed servers for the most recent patches installed, and determines other actions it must perform.
 - **Uninstall:** The patch is uninstalled.
 - **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
 - **Pre/Post Uninstall Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstallation.
 - **Uninstall & Reboot:** When a patch is installed, the server is also rebooted.
 - **Verify:** Installed patches will be included in the software registration.
- 2 To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select *Jobs and Sessions* to review detailed information about the job. See the *SA Users Guide: Server Automation* for more information on browsing job logs.
- 3 Click **Stop Job** to prevent the job from running or click **Close** to close the Uninstall Patch window.

7 Patch Management for Solaris

Overview of Solaris Patching

HP Server Automation automates the process of keeping your Sun Solaris servers running with current patches. With SA you can:

- Determine which patches your managed servers need.
- Download Solaris patches and patch clusters and store them in the SA library.
- Create Solaris patch policies from downloaded Solaris patches and patch clusters.
- Resolve all the dependencies for a set of patches including required patches, obsolete patches, superseding patches, incompatible and withdrawn patches.
- Download information associated with each patch including:
 - Platform settings including support for multiplatform patches
 - Patch dependencies
 - Reboot settings
 - Vendor documentation
- Install patches and patch clusters by remediating patch policies on managed Solaris servers. Remediation automatically handles various patch reboot settings including single-user mode, reconfiguration reboot and reboot immediate.
- Perform compliance scans including server platform, patch supersedence and package applicability checks.
- Patch Solaris zones (virtual machines).

Policy Based Patch Management

With Solaris patch policies you can ensure your Solaris servers have the right patches by creating a patch policy, which is a model of your desired IT environment. A Solaris patch policy defines a server baseline to ensure that all servers are provisioned with standard contents. Using SA, you can automatically download Solaris patches, organize them into policies, define installation order among patches in the policy, automatically resolve all dependencies for the patches and set reboot settings for all patches in the policy.

System administrators can then manage the servers in their environment by applying the Solaris patch policy to the servers. SA applies the changes to the managed servers when you remediate the managed servers with the patch policy. When a change needs to be made to a patch policy, a policy setter simply changes the baseline defined in the policy and the incremental differences are applied across the target servers. For more information, see [Using Solaris Patch Policies](#) on page 355.

Quick Start to Solaris Patching

Perform the following steps to set up and initialize Solaris patching in SA. Detailed information on these steps is in the rest of this chapter.

- 1 Create an SA user with the following permissions:
 - Read and write permissions on the folder /Opware/Tools/Solaris Patching
 - Read and write permission on “Manage Patch” feature permission
 - “Allow Install Patch” feature permission set to Yes
 - “Allow Uninstall Patch” feature permission set to Yes
 - “Manage Patch Compliance Rules” feature permission set to Yes.

For more information on creating users and setting permissions, see the *SA Administration Guide*.
- 2 Log in as root to an SA slice core server or a master core server.
- 3 Update the configuration file located at /etc/opt/opsware/solpatch_import/solpatch_import.conf as follows:
 - Add your SA user name and password to the lines with “hpsa_user” and “hpsa_pass”. For example:

```
hpsa_user=my_sa_username
hpsa_pass=<password>
```
 - Add your Sun online account user name and password to the lines with “download_user” and “download_pass”. For example:

```
download_user=my_sun_username
download_pass=<password>
```

This configuration file is used by the solpatch_import command.



You can create a separate, private copy of the configuration file and use the -c option or the --conf option to solpatch_import to specify your configuration file.

- 4 Optionally run the following command to encrypt your passwords in the configuration file:

```
solpatch_import --hide_passwords
```

The solpatch_import command is located in /opt/opsware/solpatch_import/bin.
- 5 If this is the first time you are using Solaris patching in SA, you must create a new Solaris patch database. The following command creates the Solaris patch database, downloads Patch information from Sun (in the patchdiag.xref file) and uploads the patch information into the database:

```
solpatch_import -a create_db
```

If you already have a patchdiag.xref file, you can use the following command to create the Solaris patch database and upload the patch information from your patchdiag.xref file into the database:

```
solpatch_import -a create_db -x <local patchdiag.xref file>
```

Note that this command can take up to a few hours to run depending on how many Solaris patches are already in your SA Library.

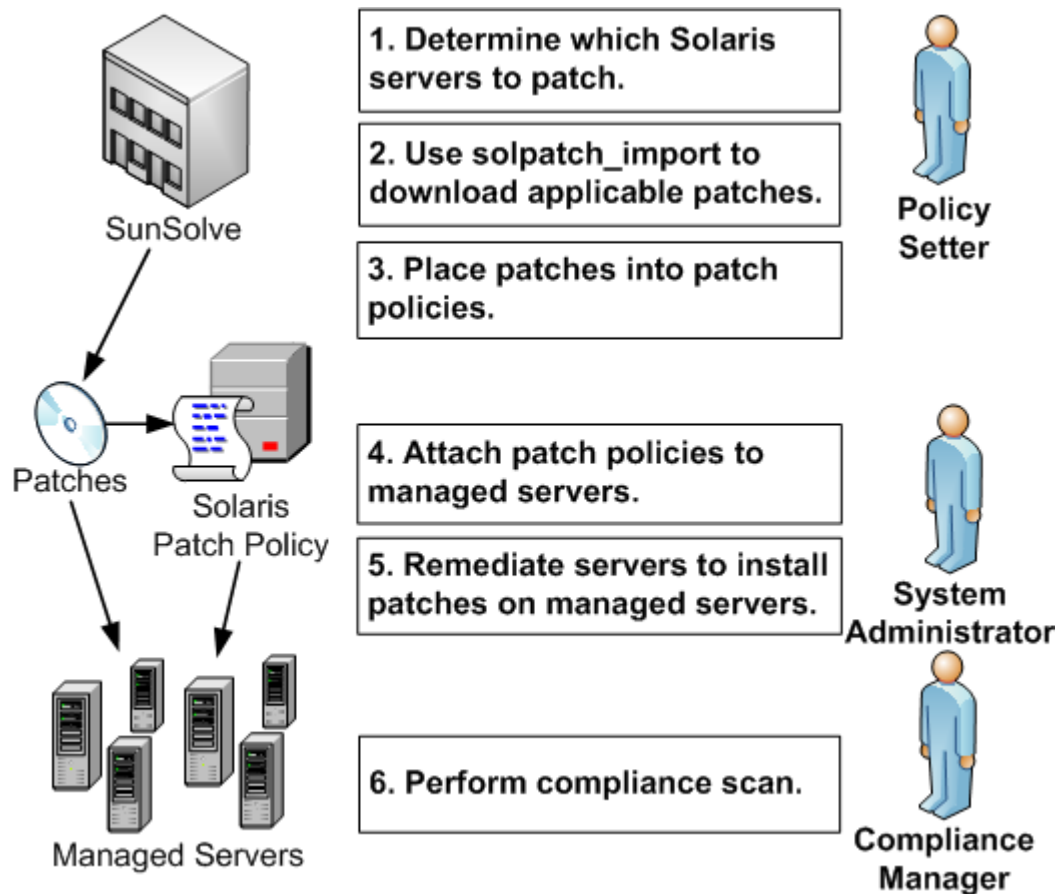
SA is now ready for you to download Solaris patches and patch clusters and install them on your servers as described in the following sections.

- 6 To ensure your Solaris patch database contains the latest patch information, see [Maintaining the Solaris Patch Database](#) on page 370.

Solaris Patch Workflow - Patching Specific Servers

Use the following steps when you know which Solaris servers you want to patch and you need to determine all the patches those servers need. The following diagram shows the steps to downloading and installing Solaris patches on a set of Solaris managed servers.

Figure 68 Solaris Patch Workflow for Patching a Set of Servers



- 1 The Policy Setter determines which Solaris servers need to be patched. For example, you may want to patch one specific Solaris server, all your servers running 5.10, all the servers used by a particular department or some other subset of your Solaris servers.
- 2 The Policy Setter uses the `solpatch_import` command to download the patches and patch clusters from Sun that are needed by the given set of Solaris servers. The `solpatch_import` command automatically determines which patches are needed by the given servers, resolves all patch dependencies and includes all applicable patches. See [Finding the Right Solaris Patches](#) on page 371 and [Automatically Importing a Solaris Patch or Patch Cluster](#) on page 374.

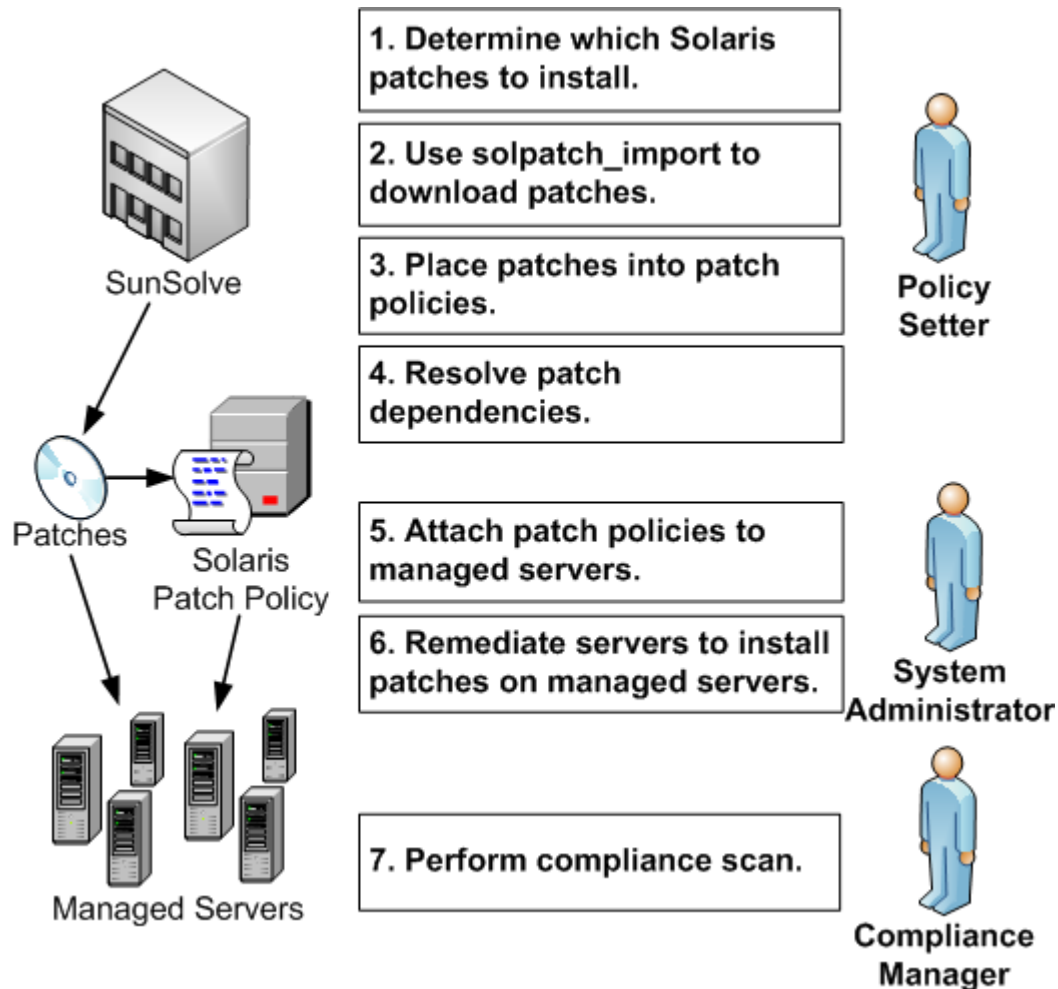
- 3 The Policy Setter places the patches into a Solaris patch policy. Note that this step can be accomplished with the `solpatch_import` command as part of [step 2](#) above, or you can manually place the Solaris patches into a patch policy with the SA Client. See [Creating a Solaris Patch Policy with the solpatch_import Command](#) on page 358 and [Adding Solaris Patches to a Patch Policy](#) on page 360.
- 4 The System Administrator attaches the patch policies to managed servers. See [Attaching a Solaris Patch Policy to a Server](#) on page 385 and [Attaching a Server to a Solaris Patch Policy](#) on page 386.

The System Administrator can test the patches by attaching the patch policy to one or more test servers to make sure they behave as expected. If problems are encountered, you can add or remove patches from the patch policy and test again. Once testing is complete, the System Administrator can attach the patch policy to all other Solaris servers.
- 5 The System Administrator remediates patch policies which installs the patches on the managed servers. See [Remediating a Server Against a Solaris Patch Policy](#) on page 386.
- 6 The Compliance Manager performs a compliance scan to determine which servers do not have the required patches installed. See [Performing a Solaris Patch Compliance Scan](#) on page 393.

Solaris Patch Workflow - Installing Specific Patches

Use the following steps when you know which Solaris patches you want to install and you need to determine all the dependent patches. The following diagram shows the steps to downloading and installing one or more specific Solaris patches.

Figure 69 Solaris Patch Workflow for Installing a Specific Patch



- 1 The Policy Setter determines the Solaris patch or patches that need to be installed. For example, you may want to install one specific Solaris security patch or one specific patch that fixes a known problem that affects your servers.
- 2 The Policy Setter uses the `solpatch_import` command to download specific patches and patch clusters from Sun. See [Automatically Importing a Solaris Patch or Patch Cluster](#) on page 374.
- 3 The Policy Setter places the patches into a Solaris patch policy. Note that this step can be accomplished with the `solpatch_import` command as part of [step 2](#) above, or you can manually place the Solaris patches into a patch policy with the SA Client. See [Creating a Solaris Patch Policy with the `solpatch_import` Command](#) on page 358 and [Adding Solaris Patches to a Patch Policy](#) on page 360.
- 4 The Policy Setter uses the **Resolve Dependencies** button in the SA Client to resolve all the dependencies on the set of patches in the patch policy, including determining dependent patches, superseding patches, obsolete patches, incompatible patches and withdrawn patches. See [Resolving Solaris Patch Dependencies](#) on page 362.
- 5 The System Administrator attaches the patch policies to managed servers. See [Attaching a Solaris Patch Policy to a Server](#) on page 385 and [Attaching a Server to a Solaris Patch Policy](#) on page 386.

The System Administrator can test the patches by attaching the patch policy to one or more test servers to make sure they behave as expected. If problems are encountered, you can add or remove patches from the patch policy and test again. Once testing is complete, the System Administrator can attach the patch policy to all other Solaris servers.

- 6 The System Administrator remediates patch policies which installs the patches on the managed servers. See [Remediating a Server Against a Solaris Patch Policy](#) on page 386.
- 7 The Compliance Manager performs a compliance scan to determine which servers do not have the required patches installed. See [Performing a Solaris Patch Compliance Scan](#) on page 393.

Solaris Patching Features

With Solaris patching you can:

- **Determine which patches your managed servers need**

SA can determine the exact set of patches your managed Solaris servers need by examining the OS version, the applications installed on your servers and the patches already installed on your servers. SA then examines all the available Solaris patches and determines the exact set of patches your servers need, the required installation order and the boot requirements, relieving you of this painstaking and error-prone chore.

- **Define Solaris Patch Policies - A model-based approach to managing your Solaris servers**

HP Server Automation enables a policy setter to create a model of their IT environment in a Solaris patch policy. Solaris patch policies specify patches, patch clusters, and scripts to be installed on the managed servers. A system administrator can then apply the patch policies to the Solaris servers in their environment.

- **Automatically download Solaris patches and related information from Sun**

SA can automatically import Solaris patches and related data, such as reboot specifications, from Sun's web site and add them to Solaris patch policies. The patch policies are stored in the SA Library and accessible from the SA Client.

- **Automatically resolve all dependent patches for a set of Solaris patches**

SA can automatically examine all the Solaris patch metadata and determine obsolete patches, superseded patches, incompatible patches, required dependent patches and withdrawn patches and update your patch policy. It also automatically places the patches in the correct install order.

- **Install Solaris patches and patch clusters on managed servers**

SA allows you to install Solaris patches and patch clusters on managed servers directly or by using Solaris patch policies. In the SA Client you can set the installation order among the patches and patch clusters in the patch policy. SA automatically includes the reboot settings from the Solaris patches in the policy.

SA also automatically ensures that each patch is applicable to each server. For example, if the package or application the patch applies to is not installed on the server, or if a newer patch is already installed on the server, then SA will not install that patch on the server.

- **Install Solaris patches in single user mode**

SA will automatically install Solaris patches in single user mode if it is required by the patch metadata published by Sun. After the patch installation is completed, SA will return to normal multiuser mode.

- **Patch Solaris zones**

With the SA Client, you can install patches on Solaris global and non-global zones by using Solaris patch policies.

- **Establish a patch installation process**

With Solaris patching, you can separate and independently schedule the various stages of patch management: analysis, download, and installation. You can be notified of job status via email upon completion of each stage and associate a ticket ID with each job.

- **Verify compliance status of servers to patch policies**

The Compliance View allows you to determine if servers are configured according to the patch policy and to remediate non-compliant servers.

- **Search for software resources and servers**

In the SA Client, the Library provides a way to search for Solaris patches, clusters and patch policies using powerful and flexible search criteria such as by availability, architecture, operating system, reboot options, version, and many other parameters. You can also search for Solaris patch policies by name, folder name, availability, operating system, and so on. See [SA Client Search](#) on page 97 for information on the search feature.

Supported Operating Systems for Solaris Patching

SA supports Solaris patching for Sun Solaris patches and Sun Solaris patch clusters. SA does not support Solaris patching for Fujitsu Solaris patches using patch policies.

The following table lists the operating system versions supported for Sun Solaris patching by SA.

Table 21 Operating System Version

OS Versions	Architecture
Solaris 6, 7, 8, 9	SUN SPARC
Solaris 10 (Update 1 through Update 7)	SUN SPARC, 64bit x86, 32 bit x86, and Niagara

Using Solaris Patch Policies

In HP Server Automation, Solaris patch policies allow you to install Solaris patches and patch clusters on servers and groups of servers. After creating a patch policy you can attach it to servers or groups of servers. When you remediate a server or group of servers, the patches specified in the attached policy are automatically installed. The remediation process compares what is actually installed on a server to the patches that should be installed on the server according to the policy. HP Server Automation then determines what operations are required to make the server conform to the policy.



You can also use software policies to manage and install patches. For more information, see [Chapter 5, Software Management](#), on page 257 of this guide.

Patch Installation Order

After you add Solaris patches and patch clusters to a patch policy, you can specify the order in which you want them to be installed. When you attach the patch policy to a server and remediate the server, SA installs the patches and patch clusters in the patch policy in the specified order.



Solaris patch policy cannot include other patch policies, but a software policy can include Solaris patch policies. You can use software policies to install Solaris patches instead of patch policies. See [Software Management](#) on page 257 and “Software Management Setup” in the *SA Policy Setter’s Guide* for more information on software policies.

Using Patch Policies with OS Provisioning

With the SA Client you can attach a Solaris patch policy to an OS sequence. When you run the OS sequence, if the remediate option is enabled (in the Remediate Policy window), then all the patches in the patch policy will be installed on the server where the OS sequence is being installed. If the remediate option is disabled, then none of the patches will be installed on the server. See [Installing \(Provisioning\) an Operating System](#) on page 494 for more information.

Solaris Patch Policy Management Tasks

Solaris Patch Policy Management consists of the following tasks:

- [Creating a Solaris Patch Policy](#) on page 357
- [Ways to Open a Solaris Patch Policy](#) on page 358
- [Editing Solaris Patch Policy Properties](#) on page 359
- [Adding Solaris Patches to a Patch Policy](#) on page 360
- [Removing Patches from a Solaris Patch Policy](#) on page 362
- [Resolving Solaris Patch Dependencies](#) on page 362
- [Adding Custom Attributes to a Solaris Patch Policy](#) on page 366
- [Deleting Custom Attributes from a Patch Policy](#) on page 366
- [Viewing the History of a Solaris Patch Policy](#) on page 367
- [Viewing all the Software Policies Associated with a Solaris Patch Policy](#) on page 367
- [Viewing OS Sequences Associated with a Solaris Patch Policy](#) on page 367
- [Viewing Servers Attached to a Solaris Patch Policy](#) on page 367
- [Locating Solaris Patch Policies in Folders](#) on page 368
- [Installing Solaris Patches Using a Solaris Patch Policy](#) on page 385.

Creating a Solaris Patch Policy

In the SA Client, you can create a patch policy in the following ways:

- [Creating a Solaris Patch Policy from the By Type View in the Library](#)
- [Creating a Solaris Patch Policy from the By Folder View in the Library](#)
- [Creating a Solaris Patch Policy with the solpatch_import Command](#)



You must have a set of permissions to create and manage a Solaris patch policy. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Creating a Solaris Patch Policy from the By Type View in the Library

Perform the following steps to create a software policy in the SA Client:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies ► Solaris**. The list of patch policies appears in the Content pane. By default, the patch policies are organized by operating system families.
- 2 Select a specific operating system.
- 3 From the **Actions** menu, select **New**. The Solaris Patch Policy window appears.
- 4 In the Name field, enter the name of the Solaris patch policy.
- 5 In the Description field, enter text that describes the purpose or contents of the policy.
- 6 Click **Browse** to specify the location for the Solaris patch policy in the folder hierarchy. The Select Folder window appears. Select a folder in the Library to specify the location of the Solaris patch policy and then click **Select**.
- 7 From the Availability drop-down list, select the SA server life cycle values for the Solaris patch policy.
- 8 From the OS drop-down list, select the operating system family or specific operating systems in that family.
- 9 To save the changes, select **Save** from the **File** menu.

Creating a Solaris Patch Policy from the By Folder View in the Library

Perform the following steps to create a Solaris patch policy in the SA Client:

- 1 From the Navigation pane, select **Library ► By Folder**. The folder hierarchy in the Library appears in the Content pane.
- 2 Select the folder that should contain the Solaris patch policy.
- 3 From the **Actions** menu, select **New Solaris Patch Policy**. The Solaris Patch Policy window appears.
- 4 In the Name field, enter the name of the Solaris patch policy.
- 5 In the Description field, enter text that describes the purpose or contents of the policy.
- 6 Click **Browse** to change the location for the Solaris patch policy in the folder hierarchy. The Select Folder window appears. Select a folder in the Library to specify the location of the Solaris patch policy and then click **Select**.

- 7 From the Availability drop-down list, select the SA server life cycle values for the Solaris patch policy.
- 8 From the OS drop-down list, select the operating system family or specific operating systems in that family.
- 9 To save the changes, select **Save** from the **File** menu.

Creating a Solaris Patch Policy with the `solpatch_import` Command

You can create Solaris patch policies with the `solpatch_import` command and place patches in the policies. For more information on this command, see [Running the `solpatch_import` Command](#) on page 369.

Example 1 - Show Sun Recommended Patches

The following command displays all the Solaris patches that are recommended for all your managed servers running Solaris 5.8:

```
solpatch_import --action=show --filter="rec,OS=5.8"
```

Example 2 - Sun Recommended Patches and Security Patches in a Policy

The following command downloads all the Sun recommended and security patches for all your managed servers running Solaris 5.8, places these patches into the SA library, and places them into the patch policy named “Sol/SolPatches” in the SA library:

```
solpatch_import --action=policy --policy_path=/Sol/Solpatches \  
--filter="rec,sec,OS=5.8"
```

Example 3 - Patch Cluster in a Policy

The following command downloads the Solaris patch cluster named “Solaris 10 SPARC Sun Alert Patch Cluster” and places all the patches in that cluster into the policy named “SolClusterPatches”. Note that the cluster is not placed into the policy, but all the patches in the cluster are placed into the policy.


```
echo "Solaris 10 SPARC Sun Alert Patch Cluster" | solpatch_import \  
-a policy --policy_path="/Sol/SolClusterPatches"
```

Ways to Open a Solaris Patch Policy

In the SA Client, there are several ways to open a Solaris patch policy.

- [Opening a Solaris Patch Policy from Search](#) on page 358
- [Opening a Solaris Patch Policy from Devices](#) on page 359
- [Opening a Solaris Patch Policy from the By Type View in the Library](#) on page 359
- [Opening a Solaris Patch Policy from the By Folder View in the Library](#) on page 359

Opening a Solaris Patch Policy from Search

- 1 From the Navigation pane, select **Search**.
- 2 Select Solaris Patch Policy from the drop-down list and then enter the name of the policy in the text field.
- 3 Select . The search results appear in the Content pane.

- 4 From the Content pane, select the Solaris patch policy and then select **Open** from the **Actions** menu. The Solaris Patch Policy window appears.

Opening a Solaris Patch Policy from Devices

- 1 From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.
Or
From the Navigation pane, select **Devices** ► **Device Groups**. The device groups list appears in the Content pane.
- 2 From the Content pane, select a server and then from the **Actions** menu, select **Open**. The Server Explorer window opens.
- 3 From the Views pane, select **Management Policies** ► **Patch Policies**. The patch policies attached to the server appear in the Content pane.
- 4 From the Content pane, select the patch policy and then select **Open** from the **Actions** menu. The Solaris Patch Policy window appears.

Opening a Solaris Patch Policy from the By Type View in the Library

- 1 From the Navigation pane, select **Library** ► **By Type** ► **Patch Policies** ► **Solaris**. The Solaris patch policies appear in the Content pane.
- 2 From the Content pane, select the Solaris patch policy and then select **Open** from the **Actions** menu. The Solaris Patch Policy window appears.

Opening a Solaris Patch Policy from the By Folder View in the Library

- 1 From the Navigation pane, select **Library** ► **By Folder**. The folder hierarchy in the Library appears in the Content pane.

From the Content pane, select the Solaris patch policy in a folder and then select **Open** from the **Actions** menu. The Solaris Patch Policy window appears.

Editing Solaris Patch Policy Properties

After you create a Solaris patch policy, you can view and modify its properties. You can view properties such as the SA user who created the Solaris patch policy, the date when it was created, and the SA ID of the Solaris patch policy. You can also modify the name, description, availability, the location of the Solaris patch policy in the Library and the operating systems of the Solaris patch policy.

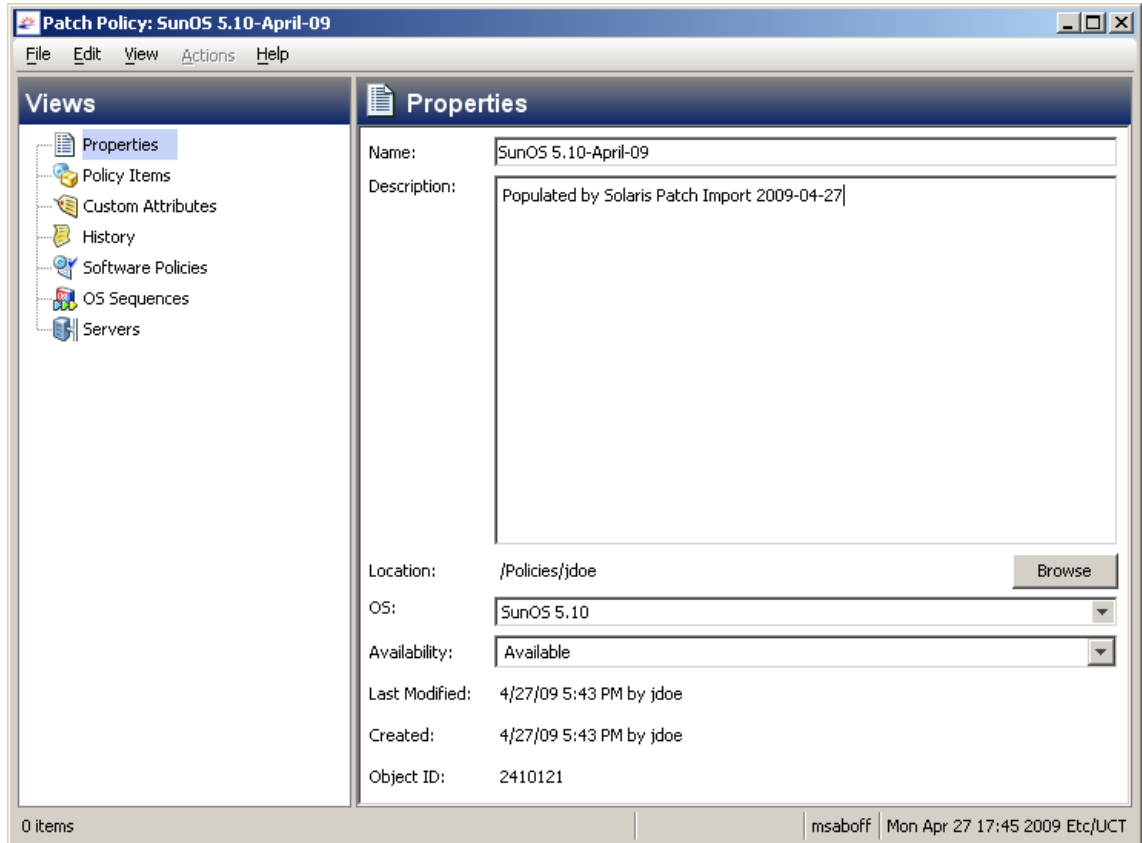


You must have a set of permissions to manage Solaris patch policies. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Perform the following steps to define the properties of a Solaris patch policy:

- 1 From the Navigation pane, select **Library** ► **By Type** ► **Patch Policies** ► **Solaris**.
- 2 From the Content pane, select the Solaris patch policy and open it. The Solaris Patch Policy window appears as shown in [Figure 70](#).

Figure 70 Solaris Patch Policy Window



- 3 From the Views pane, select Properties. You can edit the name, description, location, life cycle, and operating systems for the Solaris patch policy in the Content pane.
- 4 In the Name field, edit the name for the Solaris patch policy.
- 5 In the Description field, edit the text that describes the purpose or contents of the policy.
- 6 Click **Browse** to change the location for the Solaris patch policy in the folder hierarchy. The Select Folder window appears. Select a folder in the Library to specify the location of the Solaris patch policy and then click **Select**.
- 7 From the Availability drop-down list, select the SA server life cycle values for the Solaris patch policy.
- 8 From the OS drop-down list, select the operating system family or specific operating systems in that family.
- 9 To save the changes, select **Save** from the **File** menu.

Adding Solaris Patches to a Patch Policy


After you create a Solaris patch policy, you can add Solaris patches, Solaris patch clusters, and server scripts to it. Adding Solaris patches, Solaris patch clusters, and server scripts to a Solaris patch policy does not install them on a managed server. After you add these to a Solaris patch policy, you must attach the policy to a managed server and then remediate the server. See [Installing Solaris Patches Using a Solaris Patch Policy](#) on page 385.

You can also use the `solpatch_import` command to place patches in a patch policy. For more information, see [Creating a Solaris Patch Policy with the `solpatch_import` Command](#) on page 358.






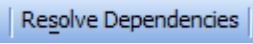
You must have a set of permissions to add Solaris patches, Solaris patch clusters, and Server Scripts to a Solaris patch policy. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Perform the following steps to add software resources to a software policy:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies ► Solaris**.
- 2 From the Content pane, select the Solaris patch policy and open it. The Solaris Patch Policy window appears.
- 3 From the Views pane, select Policy Items.
- 4 Either select , or from the **Actions** menu, select **Add...**. The Select Library Item window appears.
- 5 Select the Browse Types tab to display items that can be added to the Solaris patch policy. Select one or more items you want to add to the policy and click **Select**. The items are added to the policy.

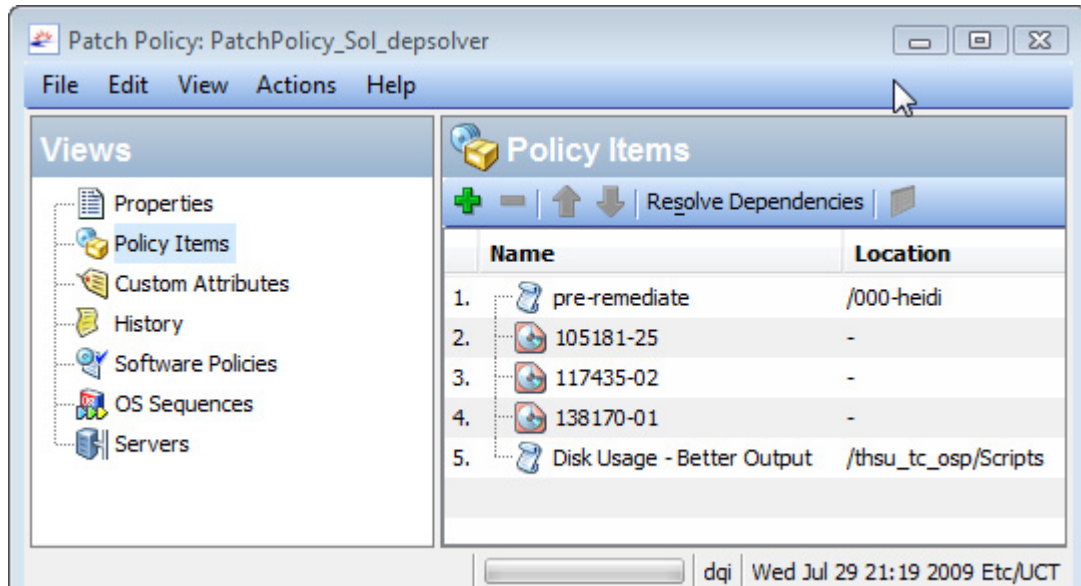
Or

Select the Browse Folders tab to display the folder hierarchy in the Library and the list of items contained in the folders. Select one or more items you want to add to the policy and click **Select**. The items are added to the policy.

- 6 To change the order in which the patches are installed, use the arrows  .
- 7 To remove a patch from the policy, select the patch and click .
- 8 To determine all dependent, obsolete, superseding, incompatible and withdrawn patches, select **Actions ► Resolve Dependencies** or select . For more information, see [Resolving Solaris Patch Dependencies](#) on page 362.
- 9 To save the changes to the policy, select **Save** from the **File** menu.

The following diagram shows a Solaris patch policy containing five patches in the SA Client.

Figure 71 Solaris Patch Policy - Policy Items View




Removing Patches from a Solaris Patch Policy

Removing Solaris patch or patch clusters from a Solaris patch policy does not uninstall it from a managed server. It only removes the patch or patch cluster from the policy. To uninstall the Solaris patch or patch cluster from a managed server, you must directly uninstall the Solaris patch or patch cluster from the managed server. See [Uninstalling Solaris Patches](#) on page 390.



You must have a set of permissions to remove Solaris patches or patch clusters from a Solaris patch policy. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Perform the following steps to remove a Solaris patch or patch cluster from a Solaris patch policy:

- 1 From the Navigation pane, select **Library > By Type > Patch Policies > Solaris** and select a version of Solaris.
- 2 From the Content pane, select the Solaris patch policy and open it. The Solaris Patch Policy window appears.
- 3 From the Views pane, select Policy Items.
- 4 Select the items that you want to remove from the list of policy items displayed in the Content pane.
- 5 Either click , or from the **Actions** menu, select **Remove**.
- 6 To save your changes, select **Save** from the **File** menu.

Resolving Solaris Patch Dependencies

When you use the `solpatch_import` command with the `filter` option, the command resolves all the patch dependencies, resulting in a complete set of installable patches.

When you add patches manually to a patch policy, SA can automatically determine the dependencies for all the patches in the patch policy. For each patch in the Solaris patch policy, SA determines the following:

- Any patches that supersede or obsolete a given patch and should be installed instead of the patch.
- Any patches that are a prerequisite to a given patch and must be installed before the patch.

For all the patches in the Solaris patch policy, SA also determines the following:

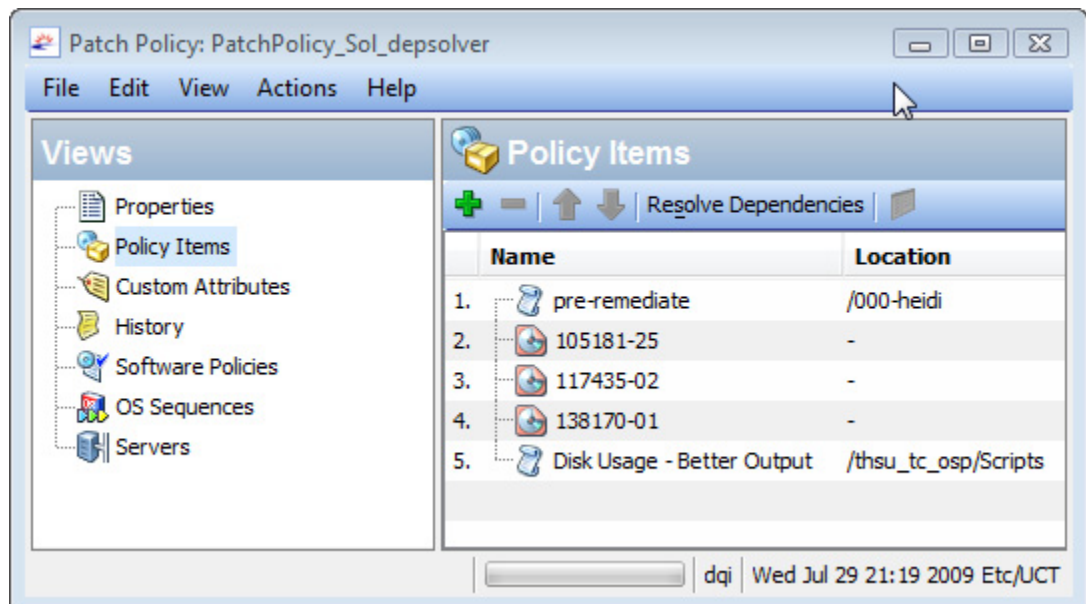
- Any patches in the set that are incompatible with each other and cannot be installed together. You must select which of the incompatible patches you want to install.
- Any patches that have been withdrawn by Sun.
- The valid installation order of all the patches, preserving the installation order of the original patches that were in the policy unless a change is required.

To determine patch dependencies, you must place the patches in a Solaris patch policy. For instructions, see [Adding Solaris Patches to a Patch Policy](#) on page 360.

To determine all the dependent patches for the patches in a patch policy, perform the following steps:

- 1 In the SA Client, locate the desired Solaris patch policy. (Select Library, the By Type tab, Patch Policies, Solaris, the version of SunOS, and the desired patch policy.)
- 2 Open the Solaris patch by double clicking on it or by selecting **Actions** ► **Open**. This displays the Patch Policy window.
- 3 In the Patch Policy window, select Policy Items in the View pane. This displays the list of Solaris patches in the patch policy, as shown below.

Figure 72 Solaris Patch Policy - Resolve Dependencies



- 4 In the Patch Policy window, select **Actions** ► **Resolve Dependencies** or select **Resolve Dependencies**. This examines the Solaris patch database of information in SA and determines all the dependencies and displays the result, showing the resulting list of patches that need to be installed.

Example - Resolving Solaris Patch Dependencies

Figure 72, "Solaris Patch Policy - Resolve Dependencies" above shows a Solaris patch policy that contains two scripts and three patches. The order shown is the order in which the scripts will be executed and patches installed.

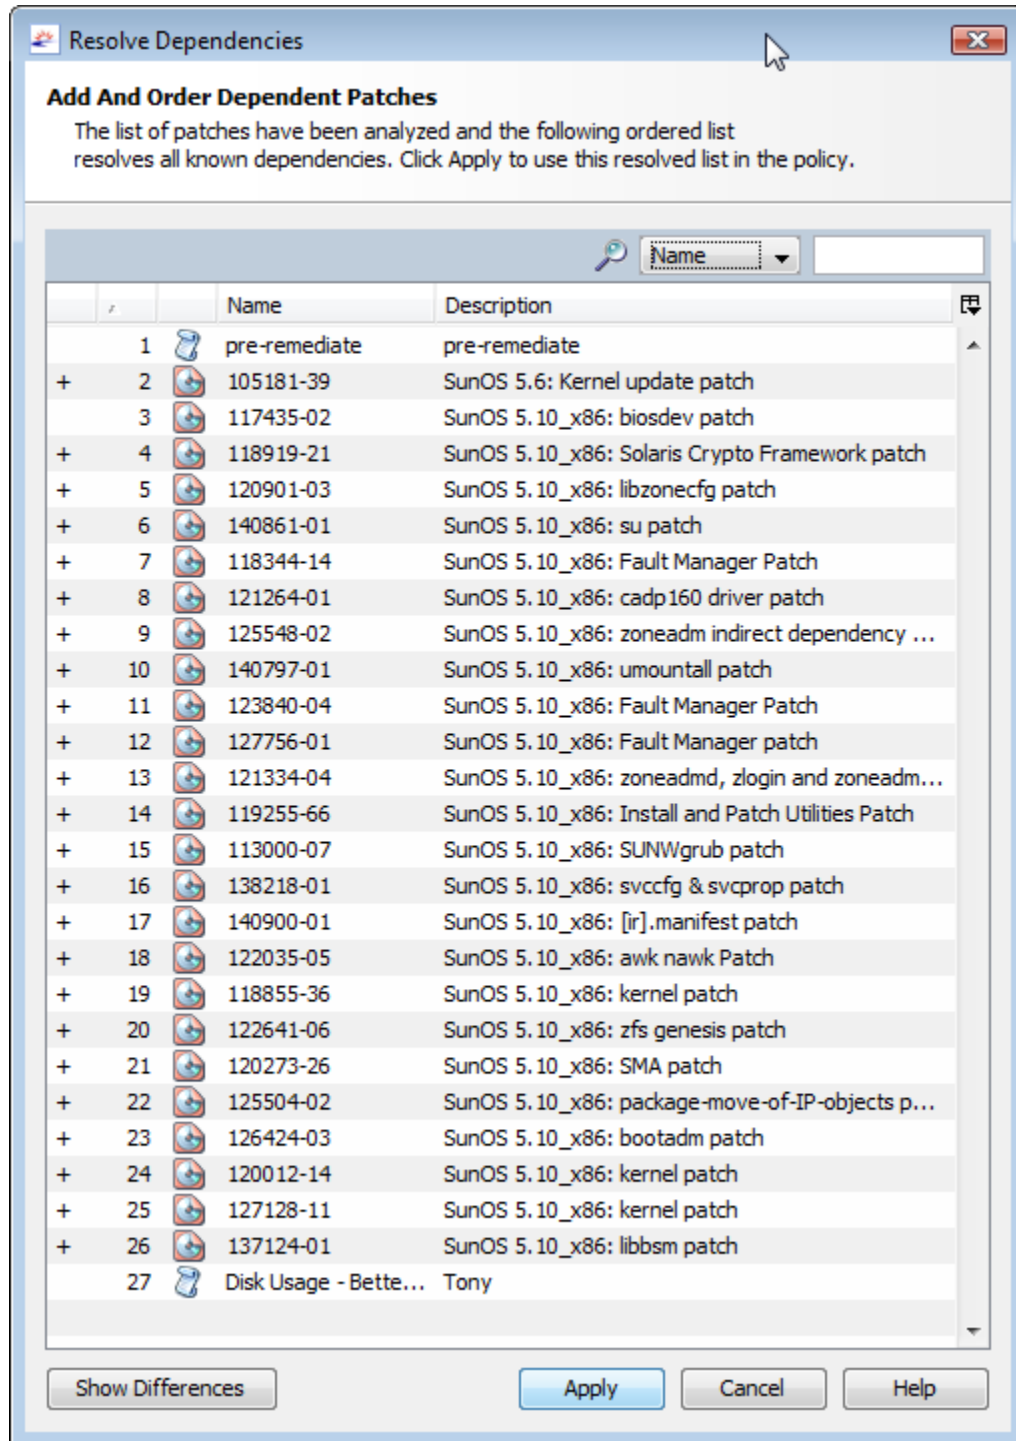
Figure 73, "Patch Dependencies for all Patches in a Patch Policy" below shows the results of selecting **Resolve Dependencies** for this patch policy. The following changes have been made to this patch policy:

- Patch 105181-25 has been replaced with a newer version, 105181-39.
- Patch 117435-02 remains in the policy.
- Patch 137124-01 replaces patch 138170-01.
- 23 additional patches have been added because they are required by 137124-01.
- The two scripts remain in the policy, in their respective positions in the policy.



Because of the iterative nature of resolving the dependencies for a set of patches, it is not always obvious how the changes to a patch policy were made.

Figure 73 Patch Dependencies for all Patches in a Patch Policy



You can display more details about the differences between the original patch policy and the proposed new set of patches by selecting the **Show Differences** button. Use the Export button on the Show Differences dialog to save the differences between the policies to a file. You can use this information with the `solpatch import` command to import the new patches into SA. For details on importing patches into SA, see [Automatically Importing a Solaris Patch or Patch Cluster](#) on page 374 and [Manually Importing a Solaris Patch or Patch Cluster](#) on page 374.

Adding Custom Attributes to a Solaris Patch Policy



Custom attributes are simply named data values you can create and set for patch policies. They provide a way you can save additional information about patch policies. You can use them in a variety of ways including in scripts, network and server configuration, notifications, and CRON script configurations. When you set a custom attribute for a patch policy, it is available to all the servers attached to the policy. For more information on custom attributes, see “Custom Attributes for Servers” in the *SA Users Guide: Server Automation*.

When you add a custom attribute to a Solaris patch policy, the attribute values affect the servers attached to the policy. After you add a custom attribute to a Solaris patch policy, you must attach the policy to a managed server and then remediate the server against the policy.




You must have a set of permissions to add custom attributes to a Solaris patch policy. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Perform the following steps to add a custom attribute to a Solaris patch policy:

- 1 From the Navigation pane, select **Library > By Type > Patch Policies > Solaris** and select a version of Solaris.
- 2 From the Content pane, select the Solaris patch policy and open it. The Solaris Patch Policy window appears.
- 3 From the Views pane, select Custom Attributes.
- 4 Either select , or from the **Actions** menu, select **Add...** A new custom attribute is added named “New Attribute”.
- 5 Enter the name of the custom attribute and select Enter.
- 6 To give a value to the custom attribute, either double click on the row under the Value column and enter the value, or click  and enter the value in the Input Dialog.
- 7 To save the changes, select **Save** from the **File** menu.

Deleting Custom Attributes from a Patch Policy

Perform the following steps to delete a custom attribute:

- 1 From the Navigation pane, select **Library > By Type > Patch Policies > Solaris** and select a version of Solaris.
- 2 From the Content pane, select the Solaris patch policy and open it. The Solaris Patch Policy window appears.
- 3 From the Views pane, select Custom Attributes. This displays the custom attributes defined for the policy.
- 4 From the Content pane, select the custom attribute that you want to delete and then click , or from the **Actions** menu, select **Remove**.
- 5 To save the changes, select **Save** from the **File** menu.

Viewing the History of a Solaris Patch Policy

Perform the following steps to view the events associated with a Solaris Patch policy:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies ► Solaris** and select a version of Solaris.
- 2 From the Content pane, select the Solaris patch policy and open it. The Solaris Patch Policy window appears
- 3 From the Views pane, select History. The events associated with the Solaris patch policy will display in the Content pane. You can view the action performed on the policy, the user who performed the action, and the time when the action was performed.
- 4 From the Show drop-down list, select the time period you want to see the events from.

Viewing all the Software Policies Associated with a Solaris Patch Policy

A software policy can contain Solaris patch policies. In the Solaris patch policy window, you can view all the software policies that include the selected Solaris patch policy as one of the items to be installed.

Perform the following steps to view the software policies that contain the selected Solaris patch policy:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies ► Solaris** and select a version of Solaris.
- 2 From the Content pane, select the Solaris patch policy and open it. The Solaris Patch Policy window appears.
- 3 From the Views pane, select Software Policies. The list of software policies that contain the selected Solaris patch policy as one of the items to be installed appears in the Content pane.

Viewing OS Sequences Associated with a Solaris Patch Policy

In the Solaris Patch Policy window, you can view all the OS Sequences that contain the selected patch policy as one of the items to be installed. Perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies ► Solaris** and select a version of Solaris.
- 2 From the Content pane, select the Solaris patch policy and open it. The Solaris Patch Policy window appears.
- 3 From the Views pane, select OS Sequences. The list of OS Sequences that contain the selected patch policy as one of the items to be installed appears in the Content pane.

Viewing Servers Attached to a Solaris Patch Policy

In the SA Client, you can view the list of all servers and device groups that have a selected Solaris patch policy attached to them.

Perform the following steps to view the servers that have a selected Solaris patch policy attached to them:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies ► Solaris** and select a version of Solaris.
- 2 From the Content pane, select the Solaris patch policy and open it. The Solaris Patch Policy window appears.
- 3 From the Views pane, select Servers. The list of servers that have the selected Solaris patch policy attached to them appears in the Content pane.

Locating Solaris Patch Policies in Folders

Perform the following steps to locate a Solaris patch policy in the folder hierarchy:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies ► Solaris** and select a version of Solaris.
- 2 From the Content pane, select the Solaris patch policy.
- 3 From the **Actions** menu, select **Locate in Folders**. The folder hierarchy for the Solaris patch policy appears in the Content pane.

Solaris Patch Management Tasks

Solaris Patch management consists of the following tasks:

- [Initializing the Solaris Patch Database](#) on page 369
- [Finding the Right Solaris Patches](#) on page 371
- [Automatically Importing a Solaris Patch or Patch Cluster](#) on page 374
- [Manually Importing a Solaris Patch or Patch Cluster](#) on page 374
- [Exporting a Solaris Patch or Patch Cluster](#) on page 375
- [Ways to Open a Solaris Patch](#) on page 376
- [Viewing Solaris Patch or Patch Cluster Properties](#) on page 376
- [Editing Solaris Patch or Patch Cluster Properties](#) on page 380
- [Viewing Vendor Documentation for a Solaris Patch or Patch Cluster](#) on page 380
- [Importing Custom Documentation for a Solaris Patch or Patch Cluster](#) on page 381
- [Viewing the Contents of a Solaris Patch Cluster](#) on page 382
- [Viewing Patch Clusters Associated with a Solaris Patch](#) on page 382
- [Viewing all Software Policies Associated with a Solaris Patch or Patch Cluster](#) on page 382
- [Viewing all Patch Policies Associated with a Solaris Patch or Patch Cluster](#) on page 383
- [Viewing Servers Associated with a Solaris Patch or Patch Cluster](#) on page 383
- [Deleting a Solaris Patch or Patch Cluster](#) on page 383
- [Installing Solaris Patches](#) on page 384

- [Uninstalling Solaris Patches](#) on page 390

Running the `solpatch_import` Command



In a multimaster mesh environment, you must not run the `solpatch_import` command simultaneously on more than one core system as this could result in lost data. It is recommended that you only run `solpatch_import` on one of your core servers.

Some Solaris patch management tasks use the `solpatch_import` command. This section describes how to use the `solpatch_import` command.

You must have the following permissions to run the `solpatch_import` command:

Table 22 Permissions Required for Using `solpatch_import`

Type of Permission	Permission Setting
Permissions on the folders <code>/Opware</code> , <code>/Opware/Tools</code> and <code>/Opware/Tools/Solaris Patching</code> in the SA library	You must have full permissions on these folders. This is where SA stores Solaris patch information.
“Manage Patch” feature permission	You must have “Read & Write” permission.
“Allow Install Patch” feature permission	This must be set to “Yes”.
“Allow Uninstall Patch” feature permission	This must be set to “Yes”.
“Manage Patch Compliance Rules” feature permission	This must be set to “Yes”.

For more information, see “Folder Permissions” and “Patch Management for Solaris - Permissions” in the *SA Administration Guide*.

To use the `solpatch_import` command, you must log in to the SA core server as root.

To run the command, log into the core server running the Software Repository component (part of the Slice Component bundle) and as root run the `solpatch_import` command located in the following directory:

```
/opt/opware/solpatch_import/bin/
```

The complete documentation for the `solpatch_import` command is available by running the command with the following option:

```
solpatch_import --manual
```

The sections below give examples of using this command.

Initializing the Solaris Patch Database

Before downloading patches and patch data from Sun, you must perform the following steps to set up and initialize the Solaris patch database in SA.

- 1 Create a configuration file that specifies information needed by the `solpatch_import` command. The default location for this file is `/etc/opt/opware/solpatch_import/solpatch_import.conf`. If you do not use the default location, you must use the `-c` or `--conf` option. If you use the default location, you do not need the `-c` or `--conf` option.

For details on the contents of this configuration file, see the `solpatch_import` man page by running `solpatch_import --manual`. The following example shows partial contents of a configuration file.

```
[main]
hpsa_user=<SA user name>
hpsa_pass=<SA user password>
download_user=<Sun account user name>
download_pass=<Sun account password>
```

- 2 Run the following command to initialize SA for Solaris patch information:

```
solpatch_import -a create_db
```

This command downloads the `patchdiag.xref` file from Sun (or you can specify a local copy of this file if you previously downloaded it), examines the patch information and places the data in SA.



You only need to use the `-a create_db` option once to initialize the Solaris patch information in SA.

- 3 See the section below for how to make sure your Solaris patch database contains the latest patch information.

Maintaining the Solaris Patch Database

This section describes how you can make sure your Solaris patch database contains the latest patch information.

Obtaining the Latest Patch Data from Sun

Sun typically updates their patch information daily Monday through Friday. To obtain the latest Solaris patch information from Sun (in the `patchdiag.xref` file) and upload it into the SA patch database, you should run the command below periodically, based on your company policy. For example you could place the following command in a cron job.

```
solpatch_import -a update_db
```

Obtaining the Solaris Patch Supplementary Data File

SA gets its information about Solaris patches from Sun (from the `patchdiag.xref` file). However, HP SA provides valuable supplementary data about Solaris patches that you can obtain automatically from the HP BSA Network. Whenever HP updates this supplementary data, you can configure the HP BSA Network to automatically upload it into the SA Solaris patch database.

To obtain the supplementary data file whenever it is updated and automatically upload it into the SA Library, perform the following steps:

- 1 Obtain a BSA Essentials Network account from <http://www.hp.com/go/bsanetwork>.
- 2 Install and configure the HP Live Network connector (LNC) on the core server where the SA Software Repository component is installed. For complete instructions, see "Configuring the Live Network connector" in the *HP Live Network connector Installation and Configuration Guide* which is available from the BSA Essentials web site, <http://www.hp.com/go/bsanetwork>.

- 3 On the system where the LNC is installed, execute the following command to enable the Solaris patching service:

```
live-network-connector write-config --setting=content.solaris_patching=1
```

- 4 (Optional) To disable the Solaris patching service, execute the same command but set the value to 0:

```
live-network-connector write-config --setting=content.solaris_patching=0
```

Alternatively, you can manually download the supplementary Solaris patch data file from the HP BSA Network and upload it into the SA database as described below. Whichever method you use, it is recommended that you regularly check for updates and install them into the SA patch database.

Manually Downloading the Solaris Patch Supplementary Data File

This section describes how to manually download the supplementary Solaris patch data file from the HP BSA Essentials Network and upload it into the SA patch database. It is recommended that you set up the LNC to automatically upload this file whenever it changes as described in [Obtaining the Solaris Patch Supplementary Data File](#) on page 370. However, if you download the file manually, you should regularly check for updates and install them into the SA patch database as described here.

To obtain the supplementary data file, perform the following steps:

- 1 Obtain a BSA Essentials Network account from <http://www.hp.com/go/bsanetwork>.
- 2 Log in to the BSA Essentials Network web site at <http://www.hp.com/go/bsanetwork>.
- 3 Locate the package “Solaris Patching for Server Automation” under the Standard Content, under Server Automation Community”.
- 4 Download the Solaris patching package, named `solpatchdb_supplement.zip`.
- 5 Place the `solpatchdb_supplement.zip` file on a core slice server in any temporary directory such as `/tmp`.
- 6 Unzip the `solpatchdb_supplement.zip` file.
- 7 Run the file `install.sh` which was in the `solpatchdb_supplement.zip` file. This uploads the Solaris patch supplementary data into the SA patch database.
- 8 Since HP updates the Solaris patch supplementary data file, it is recommended that you periodically check this file for updates and when this file changes, follow these steps again to download the latest supplementary patch information into your SA patch database.

Finding the Right Solaris Patches

With SA you can quickly and easily determine which patches your Solaris servers need. With the `solpatch_import` command, you can:

- Display the Solaris patches required by your Solaris servers, with all dependent patches and listed in the correct install order.
- Download those patches and import them into the SA Library.
- Place those patches into a Solaris patch policy.

The following table lists the options to the `solpatch_import` command to display patch information, download patches and import them into the SA Library, and place them in a Solaris patch policy.

Table 23 Specifying Actions for the `solpatch_import` Command

Option to <code>solpatch_import</code> Command	Description
-a show or --action show	Displays information about the specified patches.
-a import or --action import	Downloads the specified patches and imports them into the SA Library.
-a policy or --action policy	Downloads the specified patches, imports them into the SA Library and places them in the specified Solaris patch policy. This action requires you to specify a Solaris patch policy with the <code>--policy_path</code> option.

The `solpatch_import` command finds all the patches that are applicable to your managed servers, excluding patches that are not applicable for example because you do not have certain software applications installed, and includes all dependent patches. The resulting set of patches are complete and in the required install order.

Use the following filters to the `solpatch_import` command to specify which Solaris patches you want:

Table 24 Specifying Desired Patches with the Filter Option to `solpatch_import`

Desired Set of Patches	Filter Options to Use	Example Filter Option	Description of Example Filter Option
All patches recommended by Sun for a particular server	rec server	-f "rec,server=sys01.hp.com"	Specifies all the patches recommended by Sun for the managed server sys01.hp.com.
All patches recommended by Sun for a set of servers	rec platform	-f "rec,OS=5.10"	Specifies all the patches recommended by Sun for all your managed servers running Solaris 5.10.
All Sun security patches for a particular server	sec server	-f "sec, server=sys01.hp.com"	Specifies all the Sun security patches for the managed server sys01.hp.com.
All Sun security patches for a set of servers	sec OS	-f "sec, OS=5.9"	Specifies all the Sun security patches for all your managed servers running Solaris 5.9.
All Sun security patches and all Sun recommended patches for a server.	rec sec server	-f "rec, sec, OS=5.8"	Specifies all the Sun security patches and all the Sun recommended patches for all your managed servers running Solaris 5.8.

The examples below show a few of the ways you can use the `solpatch_import` command to determine which patches are needed by your Solaris servers. For complete information, run `solpatch_import --manual` as described in [Running the solpatch_import Command](#) on page 369.

Finding All Patches Needed by a Particular Server

The following example command finds all the patches needed by the server named “sys01.hp.com”. The first command just displays the list of patches. The second command downloads the patches and places them into the SA Library. The third command places them into the Solaris patch policy names “SolPatches/MyPolicy”.

```
solpatch_import --action=show --filter="server=sys01.hp.com"
solpatch_import --action=import --filter="server=sys01.hp.com"
solpatch_import --action=policy --policy_path="SolPatches/MyPolicy"\
--filter="server=sys01.hp.com"
```

Finding the Sun Recommended Patches for Your Servers

The following example command finds the Sun recommended patches for all your managed servers running Solaris 10. The first command just displays the list of patches. The second command downloads the patches and places them into the SA Library. The third command places them into the Solaris patch policy named MySolPolicy.

```
solpatch_import --action=show --filter="rec,OS=5.10"
solpatch_import --action=import --filter="rec,OS=5.10"
solpatch_import --action=policy --policy_path="MySolPolicy"\
--filter="rec,OS=5.10"
```

Finding the Sun Security Patches for Your Servers

The following example command displays the Sun security patches for all your managed servers running Solaris 9:

```
solpatch_import --action=show --filter="sec,OS=5.9"
```

Finding a Specific Set of Patches

You can display information about one or more patches by providing the patch names to the `solpatch_import` command or in a text file. This example assumes the file `my_sol_patches.txt` contains the following lines:

```
120900-04 121133-02 119254-67
119317-01 121296-01 127884-01
```

The following example command displays the set of patches listed in the file `my_sol_patches.txt`:

```
solpatch_import --action=show my_sol_patches.txt
```

The following command downloads the set of patches listed in the file `my_sol_patches.txt` and places the patches into the SA Library:

```
solpatch_import --action=import my_sol_patches.txt
```

The following example command downloads the set of patches listed in the file `my_sol_patches.txt`, places the patches into the SA Library, and places the patches into a Solaris patch policy named “/SolPatches/SolPatchPolicy”:

```
solpatch_import --action=policy --policy_path=/SolPatches/SolPatchPolicy \
  my_sol_patches.txt
```

For more information on the `solpatch_import` command, see [Running the solpatch_import Command](#) on page 369.

Automatically Importing a Solaris Patch or Patch Cluster

With the `solpatch_import` command you can automatically download Solaris patches and patch clusters from Sun, import them into SA, place them into Solaris patch policies, and store the patch policies in a folder in the SA Library. The `solpatch_import` command also downloads reboot settings and patch dependencies and saves them with the patch. For example commands, see [Finding the Right Solaris Patches](#) on page 371 and [Creating a Solaris Patch Policy with the solpatch_import Command](#) on page 358.



For information on using the `solpatch_import` command, see [Running the solpatch_import Command](#) on page 369.

Manually Importing a Solaris Patch or Patch Cluster



It is recommended that you use the `solpatch_import` command to import Solaris patches and patch clusters from Sun as described in [Automatically Importing a Solaris Patch or Patch Cluster](#) on page 374. However, you can also manually import patches as described here.

Solaris patches are downloaded from Sun and stored in SA. To see if a patch has been imported, view the patch’s Availability property in the SA Client. The Availability property of an imported patch can be set to one of the values listed in the following table:

Table 25 Patch Availability Property Settings

Patch Availability Setting	Description
Available	The patch has been imported into SA, has been tested, and can be installed on managed servers.
Limited	The patch has been imported into SA but requires additional permissions (Manage Patch: Read & Write) to be installed. This is the default patch availability. For more information on permissions, see the <i>SA Administration Guide</i> .
Deprecated	The patch cannot be added to patch policies but can still be installed.
Not Imported	The patch is not stored in the SA library.



You must have a set of permissions to import Solaris patches or patch clusters. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Perform the following steps to manually import a Solaris patch or patch cluster from a file into SA:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches are organized by operating system.
- 2 From the **Actions** menu, select **Import Software....** The Import Software window appears.
- 3 Click **Browse** to locate and select the patch or patch cluster to import.

Before clicking Open in the Open window, select the character encoding to be used by the patch or patch cluster from the Encoding drop-down list.

You need to specify the character encoding so that SA can extract the metadata contained in the patch or patch cluster and correctly display the information in non-ASCII characters in the SA Client (for example, in the Patch Properties pages). Patch metadata includes comments, READMEs, scripts, descriptions, and content lists.

- 4 Click Open.
- 5 In the Import Software window, select the Type, either Solaris Patch or Solaris Patch Cluster from the Type drop-down list.

This grays out the Folder edit field because Solaris patches and patch clusters are not stored in folders.

- 6 From the Platform drop-down list, select the applicable Solaris operating system.
- 7 Click **Import** to import the Solaris patch or patch cluster into SA.
- 8 Run the following command to update the Solaris patch information in SA:

```
solpatch_import -a restore_defaults
```

For more information on the `solpatch_import` command, see [Running the solpatch_import Command](#) on page 369

Exporting a Solaris Patch or Patch Cluster

You can download a Solaris patch or patch cluster to your local computer so that you can check the installation of the patch or patch cluster on a test or staging machine.

Perform the following steps to download a patch or patch cluster:


- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches are organized by operating system in the Content pane. Navigate to the desired operating system version.
- 2 From the Content pane, select a patch or patch cluster to export.
- 3 Right click or from the **Actions** menu, select **Export....** The Export Patch window appears.
- 4 Specify the location for the package to be exported to.
- 5 Click **Export**.

Ways to Open a Solaris Patch

In the SA Client, you can open a Solaris patch or patch cluster in the following ways:

- [Opening a Patch from Search](#) on page 376.
- [Opening a Patch from the By Type View in the Library](#) on page 376.

Opening a Patch from Search

- 1 From the Navigation pane, select Search.
- 2 Select Patch from the drop-down list and then enter the name of the Solaris patch or patch cluster in the text field.
- 3 Select . The search results appear in the Content pane.
- 4 From the Content pane, select the patch or patch cluster and then select **Open** from the **Actions** menu. The Patch or Patch Cluster window appears.

Opening a Patch from the By Type View in the Library

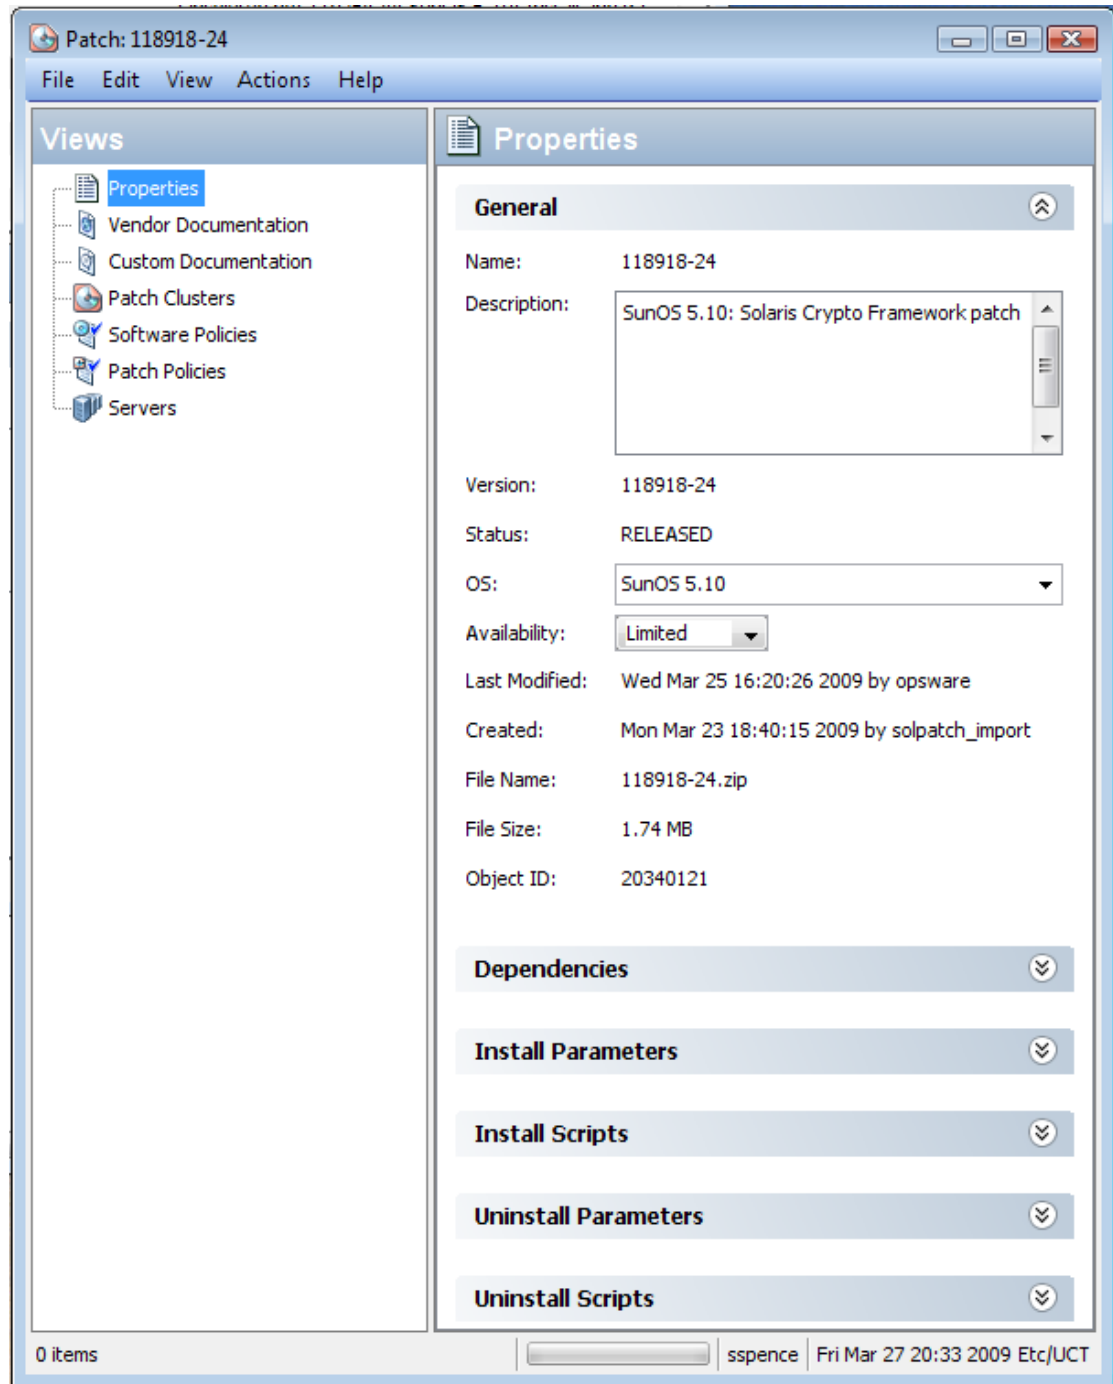
- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches appear in the Content pane.
- 2 From the Content pane, select the patch or patch cluster and then select **Open** from the **Actions** menu. The Patch or Patch Cluster window appears.

Viewing Solaris Patch or Patch Cluster Properties

Perform the following steps to view the properties of a Solaris patch or patch cluster:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches organized by operating system appear in the Content pane. Navigate to the desired OS version.
- 2 From the Content pane, select the Solaris patch or patch cluster to view.
- 3 Right click or from the **Actions** menu, select **Open**. The Patch window appears.
- 4 From the Views pane, select Properties. The following patch properties appear in the Content pane:

Figure 74 Patch Properties Window



General Properties

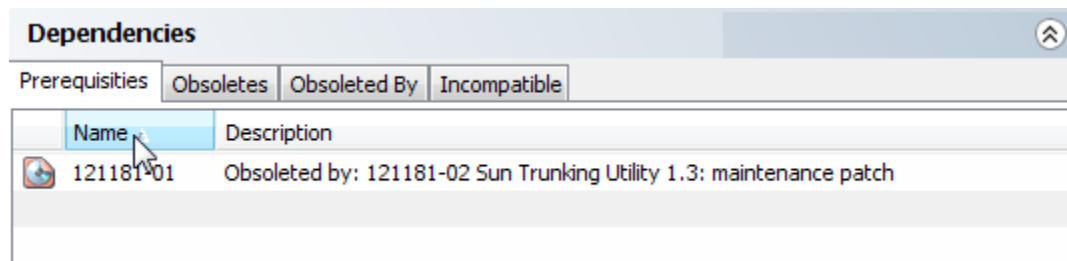
- **Name:** The name of the patch or patch cluster as defined by Sun.
- **Description:** The description of the patch or patch cluster's contents.
- **Version:** The version number of the patch or patch cluster as defined by Sun.
- **Status:** The status of the patch as defined by Sun.
- **OS:** The operating systems associated with the patch or patch cluster.

- **Availability:** The availability of the patch or patch cluster to SA users.
- **Last Modified:** The date and time when the patch was last modified and the SA user who last modified the patch.
- **Created:** The date and time when the patch or patch cluster was created by an SA user.
- **File Name:** The file name of the package.
- **File Size:** The file size of the package.
- **Object ID:** The unique SA identifier for the package.

Dependencies

The following shows the Dependencies of a patch in the patch browser in the SA Client.

Figure 75 Patch Dependencies in the Patch Properties Window



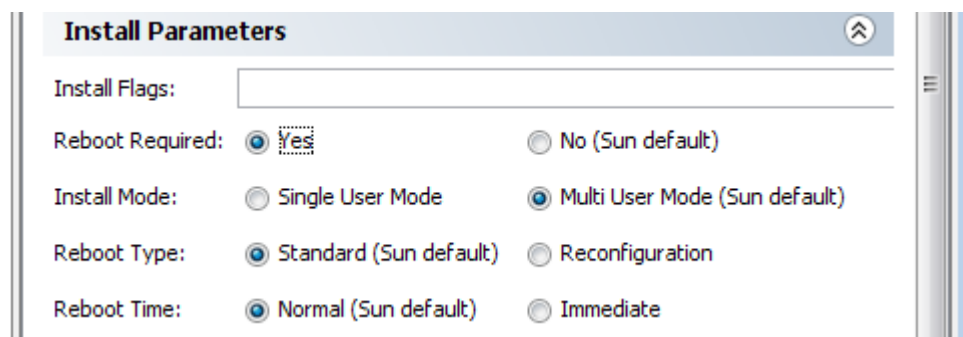
- **Prerequisites:** The patches that must be installed before this patch can be installed.
- **Obsoletes:** The older patches that are made obsolete by this patch.
- **Obsoleted by:** The newer patches that make this patch obsolete.
- **Incompatible:** The patches that cannot be installed with this patch.

Install Parameters

The Install Parameters list the actual settings for the patch and the settings that Sun specifies for the patch. The selected radio buttons are the actual settings that will be used when the patch is installed. The settings Sun recommends are labeled “Sun default”. The Sun default settings are the values that were downloaded with the patch.

The settings specified by the selected radio buttons will be used when the patch is installed. However, when you remediate a server against a patch policy or install a patch, you can override these settings. For more information, see [Overriding Reboot Settings for a Policy](#) on page 388 and [Overriding Reboot Settings for a Patch](#) on page 389.

Figure 76 Install Parameters in the Patch Properties Window



- **Install Flags:** Optional arguments to be used when the patch or patch cluster is installed on a managed server.
- **Reboot Required:** Determines if the managed server will be rebooted when the patch or patch cluster is successfully installed. Sun’s recommendation is labeled “Sun default”.
- **Install Mode:** Determines if the patch or patch cluster will be installed in single user mode or multiuser mode. Sun’s recommendation is labeled “Sun default”. Note that the Solaris system is rebooted to get in to single user mode, then the patch is installed, then the system is rebooted to get to multiuser mode.
- **Reboot Type:** Determines if a standard reboot or a reconfiguration reboot will be performed after installing the patch or patch cluster. Sun’s recommendation is labeled “Sun default”.
- **Reboot Time:** Specifies if the server will be rebooted immediately after installing the patch or at some later time after the patch or patch cluster is installed. Sun’s recommendation is labeled “Sun default”.

When installing a patch with the setting “Reboot Time: Normal”, the reboot will occur at the end of the job, unless another patch in the job requires an immediate reboot before the end of the job. However, the Job Preview and the Job Status windows will display the message “Install and Reboot” for the patch, which indicates the reboot will occur sometime after the patch is installed, not immediately after the patch is installed.

Install Scripts

- **Pre-Install Script:** A script required to run on a managed server before the patch or patch cluster is installed.
- **Post-Install Script:** A script required to run on a managed server after the patch or patch cluster is installed.
- **If script returns an error:** Specifies whether or not to stop the installation of the patch or patch cluster if the script fails.

Uninstall Parameters

- **Uninstall Flags:** The optional arguments to be used when the patch or patch cluster is uninstalled from servers.
- **Reboot Required:** Determines if the managed server will be rebooted when the patch or patch cluster is successfully uninstalled. Sun’s recommendation is labeled “Sun default”.
- **Uninstall Mode:** Determines if the patch or patch cluster will be uninstalled in single user mode or multiuser mode. Sun’s recommendation is labeled “Sun default”. Note that the Solaris system is rebooted to get in to single user mode, then the patch is uninstalled, then the system is rebooted to get to multiuser mode.
- **Reboot Type:** Determines if a standard reboot or a reconfiguration reboot will be performed after uninstalling the patch or patch cluster. Sun’s recommendation is labeled “Sun default”.
- **Reboot Time:** Specifies if the server will be rebooted immediately or at some later time after the patch or patch cluster is uninstalled. Sun’s recommendation is labeled “Sun default”.

Uninstall Scripts

- **Pre-Uninstall Script:** A script required to run on a managed server before the patch or patch cluster is uninstalled.
- **Post-Uninstall Script:** A script required to run on a managed server after the patch or patch cluster is uninstalled.
- **If script returns an error:** Specifies whether or not to stop the uninstallation of the patch or patch cluster if the script fails.

Editing Solaris Patch or Patch Cluster Properties

After you upload a new Solaris patch or patch cluster or select an existing one, you can add or edit many of the properties of the patch or patch cluster in the SA Client.



You must have a set of permissions to edit the properties of a patch or patch cluster. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Perform the following steps to edit the properties of a Solaris Patch or patch cluster:

- 1 Open the patch in the patch browser as described in [Viewing Solaris Patch or Patch Cluster Properties](#) on page 376.
- 2 Edit any of the properties that are editable in the SA Client. For a list of properties, see [Viewing Solaris Patch or Patch Cluster Properties](#) on page 376.

Viewing Vendor Documentation for a Solaris Patch or Patch Cluster

The SA Client gives you access to patch information from Sun using the URL provided with the downloaded patch.

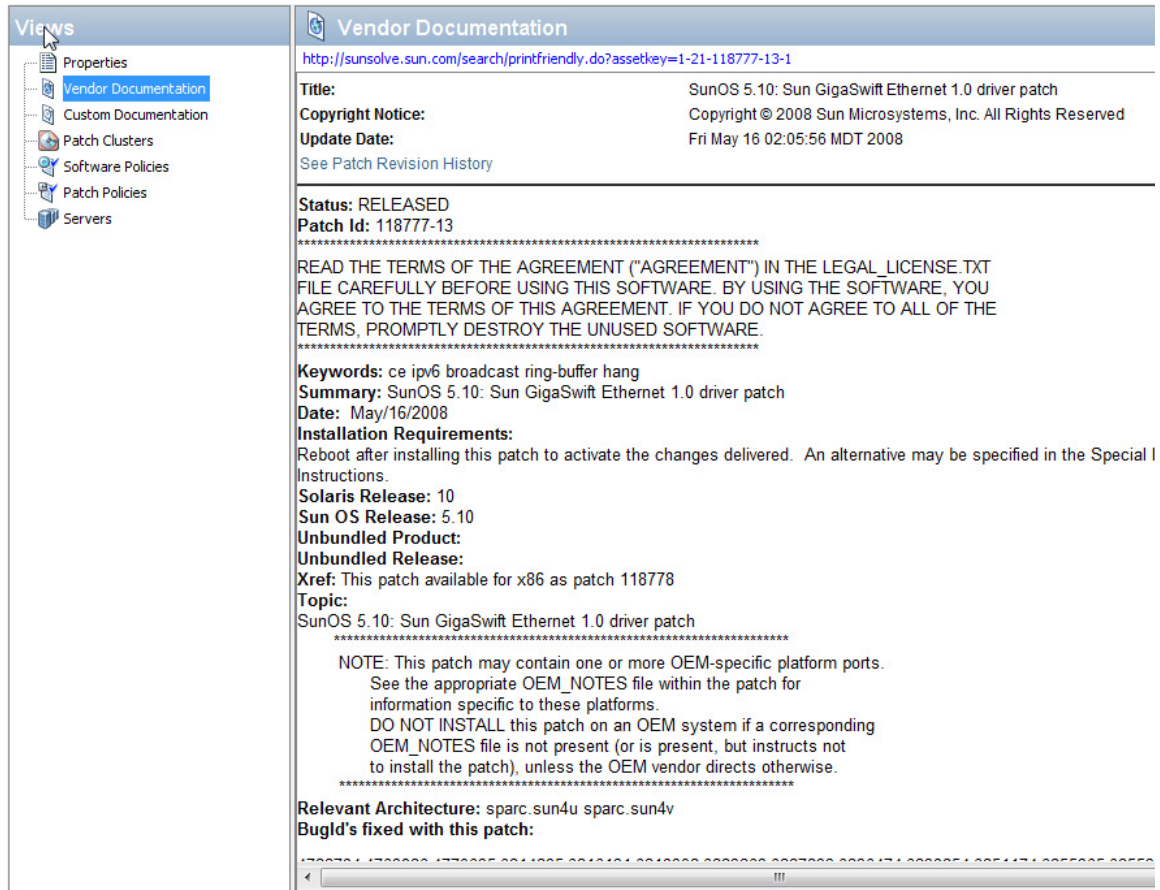


If you do not have internet access, the vendor documentation will not be displayed. However, the SA Client will display the URL you can use with an internet browser to view the documentation.

Perform the following steps to view vendor documentation for a Solaris Patch or patch cluster.

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches organized by operating systems appear in the Content pane. Navigate to the desired OS version.
- 2 From the Content pane, select the Solaris patch or patch cluster to view.
- 3 From the **Actions** menu, select **Open**. The Patch window appears.
- 4 From the Views pane, select Vendor Documentation. The documentation along with the link for the patch or patch cluster appear in the Content pane as shown.

Figure 77 Vendor Documentation Window



- 5 Click the link to view the vendor documentation.
- 6 From the **Actions** menu, select **Export** to export the contents of the vendor documentation to a desired location.

Importing Custom Documentation for a Solaris Patch or Patch Cluster

To import custom documentation for a Solaris patch or patch cluster using the SA Client, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches are organized by operating system. Navigate to the desired OS version.
- 2 From the Content pane, select the Solaris patch or patch cluster to view.
- 3 Right click or from the **Actions** menu, select **Open**. The Patch or Patch Cluster window appears.
- 4 From the Views pane, select Custom Documentation. The contents of the Custom Documentation for the patch or patch cluster appear in the Content pane.
- 5 From the **Actions** menu, select **Import**. The Import Custom Documentation window appears.
- 6 In the Import Custom Documentation window, locate the text file and specify the encoding.
- 7 Click **Import**.

Viewing the Contents of a Solaris Patch Cluster

Perform the following steps to view the contents of a Solaris patch cluster:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches organized by operating systems appear in the Content pane. Navigate to the desired OS version.
- 2 From the Content pane, select the Solaris patch cluster.
- 3 From the **Actions** menu, select **Open**. The Patch Cluster window appears.
- 4 From the Views pane, select Contents. The list of patches included in the patch cluster appears in the Content pane.
- 5 Select a patch in the Content pane, and from the **Actions** menu, select Open to view the patch properties.

Viewing Patch Clusters Associated with a Solaris Patch

Perform the following steps to view the patch clusters that contain the Solaris patch:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches organized by operating systems appear in the Content pane. Navigate to the desired OS version.
- 2 From the Content pane, select the Solaris patch.
- 3 From the **Actions** menu, select **Open**. The Patch window appears.
- 4 From the Views pane, select Patch Clusters. The list of patch clusters that contain the patch appears in the Content pane.
- 5 Select a patch cluster in the Content pane, and from the **Actions** menu, select Open to view the properties of the patch cluster.

Viewing all Software Policies Associated with a Solaris Patch or Patch Cluster

Perform the following steps to view the software policies that contain the patch or patch cluster as one of the policy items:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches organized by operating systems appear in the Content pane. Navigate to the desired OS version.
- 2 From the Content pane, select the Solaris patch or patch cluster.
- 3 From the **Actions** menu, select **Open**. The Patch or Patch Cluster window appears.
- 4 From the Views pane, select Software Policies. The list of software policies that contain the patch or patch cluster as one of the policy items appear in the Content pane.
- 5 Select a software policy in the Content pane, and from the **Actions** menu, select Open to view the properties of the software policy.

Viewing all Patch Policies Associated with a Solaris Patch or Patch Cluster

Perform the following steps to view the patch policies that contain the patch or patch cluster as one of the policy items:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches organized by operating systems appear in the Content pane. Navigate to the desired OS version.
- 2 From the Content pane, select the Solaris patch.
- 3 From the **Actions** menu, select **Open**. The Patch or Patch Cluster window appears.
- 4 From the Views pane, select Patch Policies. The list of patch policies that contain the patch or patch cluster as one of the policy items appear in the Content pane.
- 5 Select a software policy in the Content pane, and from the **Actions** menu, select Open to view the properties of the patch policy.

Viewing Servers Associated with a Solaris Patch or Patch Cluster

Perform the following steps to view the servers that have the patch or patch cluster installed using SA:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches organized by operating systems appear in the Content pane. Navigate to the desired OS version.
- 2 From the Content pane, select the Solaris patch.
- 3 From the **Actions** menu, select **Open**. The Patch or Patch Cluster window appears.
- 4 From the Views pane, select Servers. The list of servers that have the patch or patch cluster installed appear in the Content pane.
- 5 Select a server in the Content pane, and from the **Actions** menu, select Open to view the properties of the server.

Deleting a Solaris Patch or Patch Cluster

When you delete a Solaris patch or patch cluster, it is removed from SA, but it is not uninstalled from managed servers. A patch or patch cluster cannot be deleted if it is attached to a patch policy or a software policy.



You must have a set of permissions to delete a patch or patch cluster. To obtain these permissions, contact your SA administrator. See the *SA Administration Guide* for more information.

Perform the following steps to delete a patch or patch cluster:

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches organized by operating system appear in the Content pane. Navigate to the desired OS version.
- 2 From the Content pane, select a patch or patch cluster to delete.
- 3 From the **Actions** menu, select **Delete**.

Installing Solaris Patches

You can install Solaris patches directly on managed servers or on all servers in a device group or you can add Solaris patches to a Solaris patch policy (or to a software policy), attach the policy to managed servers or device groups and then remediate the servers against those policies. When you remediate the servers or device groups, the Solaris patches specified in the attached policy are automatically installed on the managed servers.

SA provides the following ways to install Solaris patches on managed servers:

- Install Solaris patches directly on managed servers using the Install Patch window. See [Installing a Unix Patch](#) on page 411.
- Install Solaris patches directly on managed servers using the Install Software window. See [Installing/Uninstalling Software without a Software Policy](#) on page 262.
- Install Solaris patches and patch clusters on managed servers using a Solaris Patch Policy. See [Installing Solaris Patches Using a Solaris Patch Policy](#), below.
- Install Solaris patches and patch clusters on managed servers using a software policy. See [Install Software Using a Software Policy](#) on page 267.



If you install or remove Solaris patches outside of SA, you must perform a software registration and a compliance scan to ensure that SA has complete and up-to-date information about the server. For more information, see [Performing a Solaris Patch Compliance Scan](#) on page 393.

Installing Solaris Clusters

SA can install all Solaris patch clusters including clusters that require passcodes. Some clusters may need to reboot the server during the install process one or more times. SA will automatically perform the reboots as long as the cluster Install Parameters have “Reboot Required” set to “Yes” and the remediate job options for “Rebooting” are set to either “Reboot servers as specified by individual software items” or “Reboot servers after each installation or uninstallation”.

If any of the above options are not set, the cluster will install up to the point where a reboot is required (if one is required). The cluster status at the end of the remediation job will display “Not installed”, the job status will be “Failed”, and the output of the job will contain a message indicating that the server must be rebooted before any more patches can be installed. After rebooting the server, the rest of the cluster can be installed by running the job again. If the cluster requires a reboot, no other patches can be installed until the server is rebooted.

It is highly recommended that you read the Readme for each Solaris patch cluster before attempting to install the cluster. For clusters that require a passcode, SA does not require you to manually enter the passcode that is in the Readme file, but this does not eliminate the need to read the Readme file.

For more information see [Viewing Solaris Patch or Patch Cluster Properties](#) on page 376 and [Specifying Options for Remediation](#) on page 273.

Installing Manual Patches

SA uses the patchadd utility to install Solaris patches. However, some patches, such as firmware updates, cannot be installed with patchadd. These “manual patches” have special installation instructions in their Readme files and must be installed manually on your Solaris servers.

While you can import these patches into the SA software repository and install them manually on servers, if you attempt to remediate a manual patch, the job will result in a Warning status. The patch status will be “Will Not Install” and the output will state that the patch requires a special installation procedure and must be installed manually.

SA cannot determine if these manual patches have been installed. A compliance scan on a patch policy that contains a manual patch will report that the policy is non-compliant. In this case, you should install the patch manually and remove the patch from the policy. For more information, see [Solaris Patch Compliance](#) on page 391.

Installing Solaris Patches Using a Solaris Patch Policy

Installing Solaris patches by using a Solaris patch policy includes the following steps:

- [Attaching a Solaris Patch Policy to a Server](#) on page 385 or [Attaching a Server to a Solaris Patch Policy](#) on page 386.
- Remediate a server against a Solaris patch policy. See [Remediating Policies](#) on page 271.

Attaching a Solaris Patch Policy to a Server

When you attach a Solaris patch policy to a server or group of servers, the Solaris patch policy is associated with that server or group of servers. This action does not install the patches and patch clusters contained in the Solaris patch policy. To install the patches and patch clusters, you must remediate the server with the Solaris patch policy. See [Remediating Policies](#) on page 271 for more information.



You must have a set of permissions to attach a Solaris patch policy to a server. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide* for more information.

Perform the following steps to attach a Solaris patch policy to a server:

- 1 From the Navigation pane, select **Library ► By Type ► Patch Policies ► Solaris** and select the desired version of Solaris. The patch policies appear in the Content pane.
- 2 (Optional) From the Content pane, select the Solaris patch policy.
 - a Open the Solaris patch policy. The Solaris Patch Policy window appears.
 - b From the View pane, select Attached Servers.
 - c From the Content pane, select a server.
- 3 From the **Actions** menu, select **Attach Server**.
- 4 In the Attach Server window, select servers or device groups and then click **Attach**. You can only select servers that are not in italics. Servers in italics indicate that you do not have the permission to attach a Solaris patch policy to the server.

- 5 (Optional) Select “Remediate Servers Immediately” to remediate the servers against the Solaris patch policy. Selecting this option displays the Remediate window. This option is only available if you have the Remediate Servers permission. See [Remediating Policies](#) on page 271.

Attaching a Server to a Solaris Patch Policy

When you attach a server or group of servers to a Solaris patch policy, the policy is associated with that server or group of servers. This action does not install the patches or patch clusters contained in the Solaris patch policy. To install the patch and patch clusters, you must remediate the server with the Solaris patch policy. See [Remediating Policies](#) on page 271 for more information.



You must have a set of permissions to attach a server to a Solaris patch policy. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide* for more information.

- 1 From the Navigation pane, select **Devices ► Servers ► All Managed Servers**. The server list appears in the Content pane.
Or
From the Navigation pane, select **Devices ► Device Groups**. Navigate to a device group. The device group list displays in the Content pane.
- 2 From the Content pane, select a server or a device group.
- 3 From the **Actions** menu, select **Attach ► Patch Policy**. The Attach Solaris Patch Policy window appears.
- 4 Select Browse Solaris Patch Policies and then select one or more policies from the list.
Or
Select Browse Folders and then select one or more policies from the folder hierarchy.
- 5 Click **Attach**.
- 6 (Optional) Select “Remediate Servers Immediately” to remediate the servers against the Solaris patch policy. Selecting this option displays the Remediate window. This option is only available if you have the Remediate Servers permission. [Remediating Policies](#) on page 271.

Remediating a Server Against a Solaris Patch Policy

To install Solaris patches in a patch policy on a Solaris server, you remediate the server against the policy. To remediate Solaris servers against a Solaris patch policy, perform the steps described in [Remediating Policies](#) on page 271.

Analyzing Patch Applicability during Remediation

Before patches are downloaded and installed on each managed Solaris server, SA verifies that the patch is needed on the server. This applicability analysis checks the following:

- 1 Verifies that the server platform matches the supported platform listed for the patch.
- 2 Verifies that the patch (or a superseding patch) is not already installed on the server.

- 3 Verifies that the package the patch applies to is installed on the server.

If any of the above is not true, the patch is non-applicable and will not be downloaded to or installed on the managed server. Non-applicable patches do not impact the overall job status. That is, the job can still complete successfully.

Reboot Settings and Other Patch Install Parameters

Each Solaris patch has reboot settings specified by Sun. These reboot settings are in the Install Parameters display in the SA Client, shown in the figure below. The Sun settings are marked with “Sun default”. The actual settings that will be used are the selected radio buttons.

The figure below shows a patch for which Sun does not require a reboot after installing the patch, but the server will actually be rebooted. The policy setter has decided to override Sun’s recommendation and reboot the system anyway after this patch is installed.

Figure 78 Reboot Settings and Other Patch Install Parameters

Install Parameters

Install Flags:

Reboot Required: ☒ Yes ☐ No (Sun default)

Install Mode: ☐ Single User Mode ☒ Multi User Mode (Sun default)

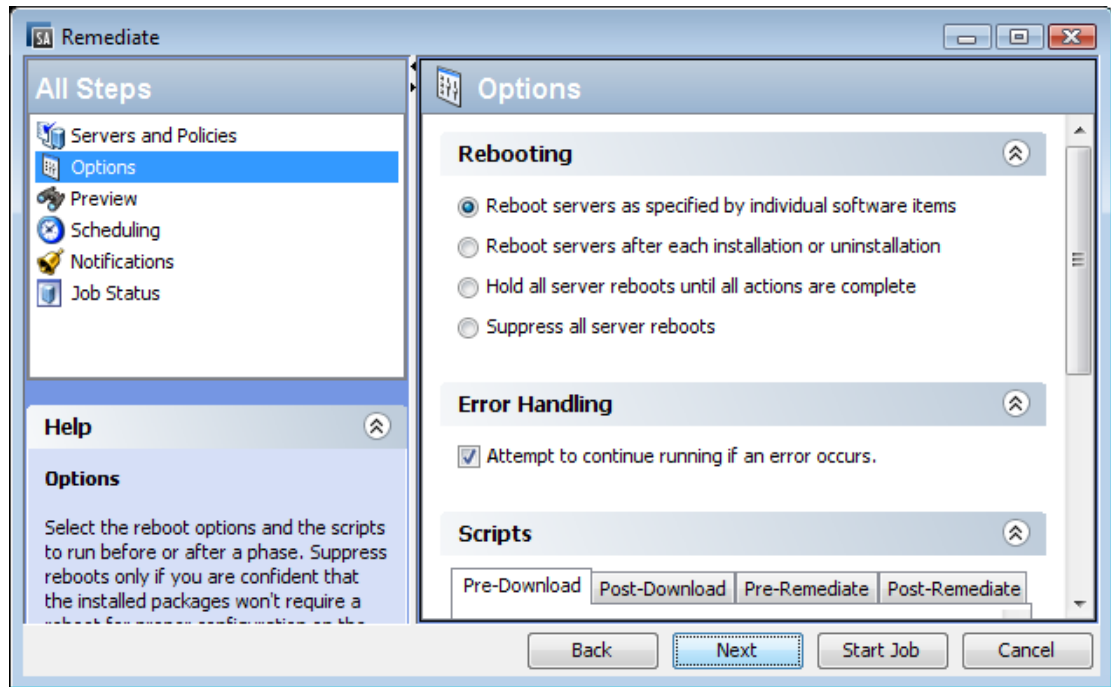
Reboot Type: ☒ Standard (Sun default) ☐ Reconfiguration

Reboot Time: ☒ Normal (Sun default) ☐ Immediate

Overriding Reboot Settings for a Policy

When you remediate a Solaris server against a Solaris patch policy, SA installs the patches and uses the reboot settings specified for each patch. However, you can override these settings when starting the remediate job. Below are the Options settings for the Remediate patch policy job.

Figure 79 Reboot Settings when Remediating a Patch Policy

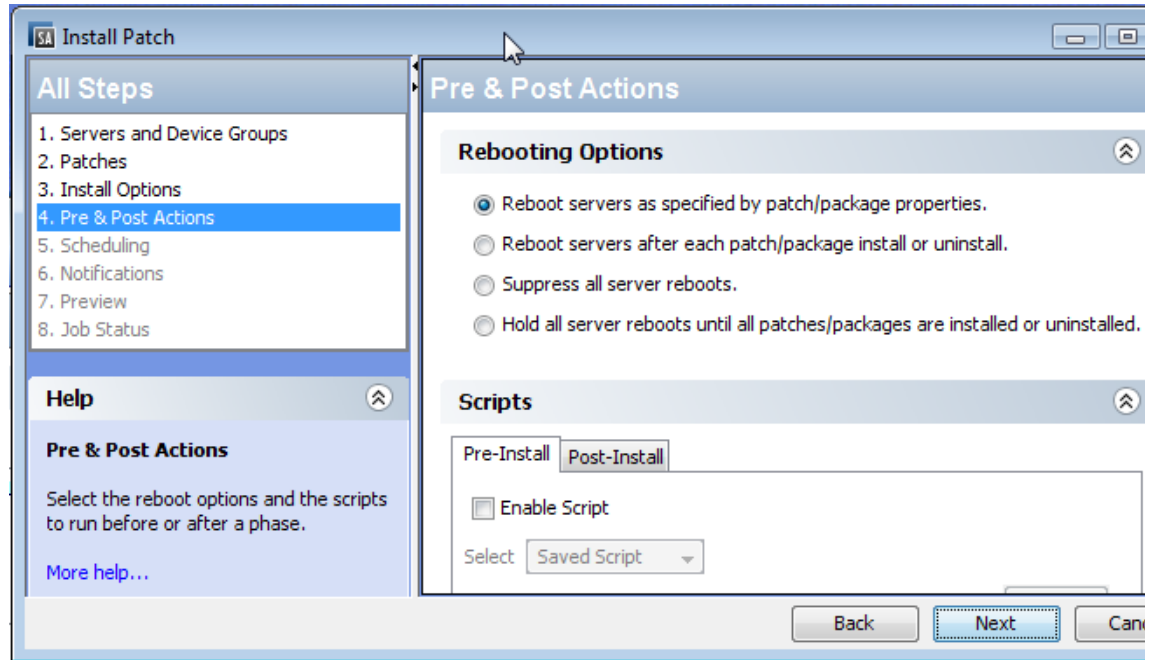


- Reboot servers as specified by individual server items: This selection uses the reboot setting specified with each patch. This selection does not override the patch settings.
- Reboot servers after each installation or uninstallation: This selection reboots the server after every individual patch is installed. This selection overrides the patches that specify no reboot.
- Hold all server reboots until all actions are complete: This selection does not reboot the server until all patches are installed on the server. This selection overrides the patches that specify an immediate reboot.
- Suppress all server reboots: This selection prevents the server from rebooting even if the patches specify a reboot. This setting overrides the patches that specify a reboot.

Overriding Reboot Settings for a Patch

When you install one or more Solaris patches, SA installs the patches and uses the reboot settings specified for each patch. However, you can override these settings when starting the install job. Below are the Pre & Post Actions settings from the Install Patch job.

Figure 80 Reboot Settings when Installing Patches



- Reboot servers as specified by patch/package properties: This selection uses the reboot setting specified for each patch. This selection does not override the patch settings.
- Reboot servers after each patch/package install or uninstall: This selection reboots the server after each patch is installed. This selection overrides the patch settings that specify no reboot.
- Suppress all server reboots: This selection prevents the server from rebooting even if some patches specify a reboot. This setting overrides the patch settings that specify a reboot.
- Hold all server reboots until all patches/packages are installed or uninstalled: This selection does not reboot the server until all the patches are installed on the server. This setting overrides the patches that specify an immediate reboot.

Troubleshooting Patch Installation

When remediating a Solaris patch that has the Install Mode (under Install Parameters in the Properties view) set to “Single User Mode” the server will be rebooted into single user mode before installing the patch. If the remediation fails for some reason (for example a network outage or a hardware failure), the system will remain in single user mode. To return the system to multi-user mode, you can perform the following steps:

- 1 Log into the Solaris server console.
- 2 Change to the directory using one of the following commands, depending on the Solaris version.

```
cd /etc/rcS.d/      # On Solaris 5.10
cd /etc/rc1.d       # On Solaris 5.6 - 5.9
```

- 3 Run the following command.

```
./S99zOpswPatching exit_single_user_mode
```

- 4 Reboot the server using the following command or another method. This will reboot the server into multi-user mode.

```
shutdown -y -g 0 -i 6
```

If you do not have access to a server console on your Solaris server, you can use the SA Global Shell (OGSH) rosh utility. Perform the following steps:

- 1 Using an SA user who has the OGFS permission “Log in to Server”, open an OGSH session. For example, you could use an ssh command like the following.

```
ssh -p 2222 <user-name>@<ogfs-host>
```

- 2 Navigate to your Solaris server using a command like the following:

```
cd /opsw/Server/@/<server name>/files/root
```

- 3 Launch the rosh utility.

- 4 Change to the directory using one of the following commands, depending on the Solaris version.

```
cd /etc/rcS.d/      # On Solaris 5.10
cd /etc/rc1.d       # On Solaris 5.6 - 5.9
```

- 5 Run the following command:

```
./S99zOpswPatching exit_single_user_mode
```

- 6 Reboot the server using the following command or another method. This will reboot the server into multi-user mode.

```
shutdown -y -g 0 -i 6
```

Note that when you reboot the server your rosh process will be terminated. Make sure the server is configured to auto-reboot.

For more information on the SA Global Shell and the rosh utility, see the chapter “SA Global Shell” and “Appendix C: Global Shell Utilities Syntax” in the *SA Users Guide: Server Automation*.

If a patch requires single user mode and fails to install for some other reason such as a dependent patch is not installed, the Solaris host will be rebooted to single user mode, the patch installation will be attempted and the host will be rebooted to multi-user mode. These two reboots occur even if the path installation fails.

Uninstalling Solaris Patches

Removing a Solaris patch or patch cluster from a Solaris patch policy does not uninstall it from a managed server. It only removes the Solaris patch or patch cluster from the Solaris patch policy. To uninstall a Solaris patch from a managed server, you must directly uninstall the Solaris patch from the managed server. To remove a patch cluster, you must remove each of the patches in the patch cluster from the managed server.

SA provides the following ways to uninstall Solaris patches from managed servers or device groups:

- Uninstall Solaris Patches directly from managed servers using the Uninstall Patch window. See [Uninstalling a Unix Patch](#) on page 418 in Chapter 8 for more information.
- Uninstall Solaris Patches directly from managed servers using the Uninstall Software window. See [Installing/Uninstalling Software without a Software Policy](#) on page 262 in Chapter 5 for more information.

Solaris Patch Compliance

A Solaris Patch compliance scan compares the Solaris patches that are installed on a managed server with the patches listed in the Solaris patch policies that are attached to the server and reports the results. If the actual server configuration does not match the Solaris patch policies attached to the server, then the server is out of compliance with the Solaris patch policies.

Patches that are not applicable to a particular Solaris server will not impact the compliance status of the server. For example:

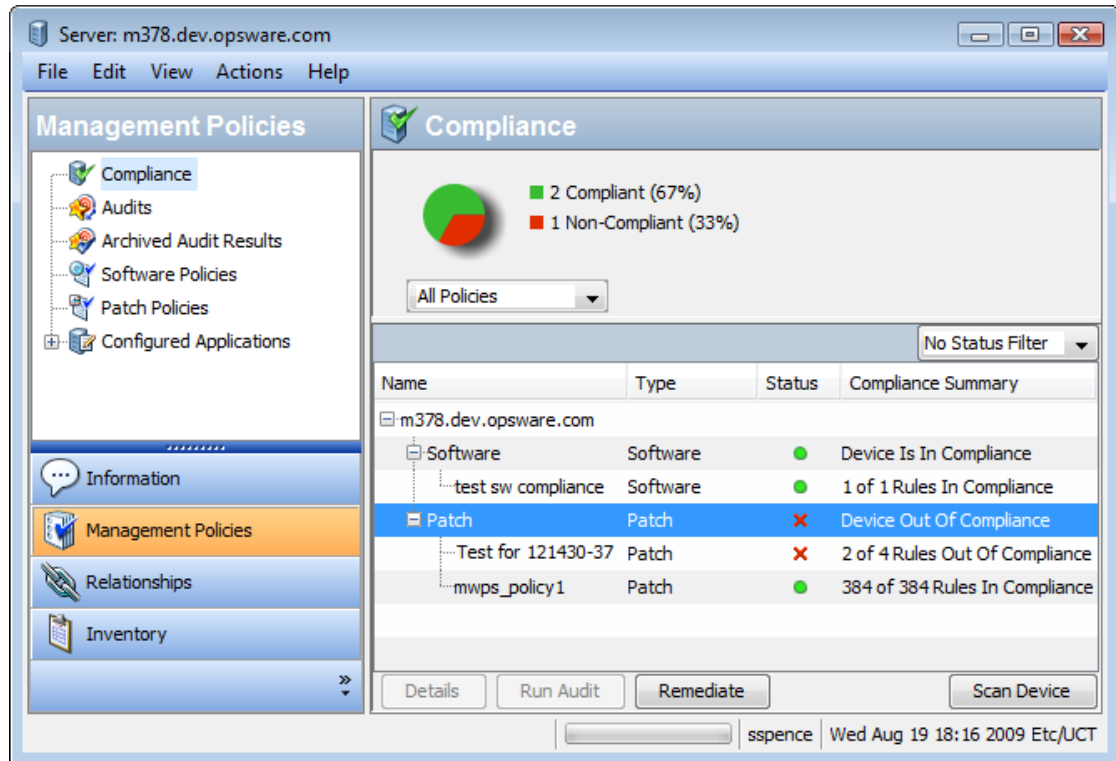
- If a policy contains a patch for the package “SUNWpkg”, but “SUNWpkg” is not installed on a particular server, the patch is not applicable to that server and that patch will not impact the results of the compliance scan for that server. The Compliance Summary does not include non-applicable patches. For example, if a policy contained 5 patches but only 3 were applicable to a given server and those 3 were installed on that server, the Compliance Summary would report “3 of 3 Rules In Compliance”, ignoring the 2 non-applicable patches.
- If a particular patch in the patch policy has been superseded by a newer patch and the newer patch is installed on a server, that server will be marked as compliant. (In essence, the patch policy is out of date. You can update the policy as described in [Resolving Solaris Patch Dependencies](#) on page 362.)
- Manual patches are always shown as out of compliance because SA cannot determine if manual patches are installed on Solaris servers. For more information, see [Installing Manual Patches](#) on page 385.

In the SA Client, when you perform a patch compliance scan, the results indicate the server’s overall compliance with all the Solaris patch policies attached to the server. Even if only one Solaris patch policy attached to the server is not compliant, the server is considered non-compliant. You can then view the non-compliant server and remediate the server against the applicable patch policy.

[Figure 81, "Compliance Results for a Solaris Server"](#) below shows the compliance view for a Solaris server. Notice that the server is out of compliance because some patches are not installed on the server:

- Patch policy “Test for 121430-37” contains 4 applicable patches, but only 2 are installed on the server.
- Patch policy “mwps_policy1” contains 384 applicable patches and all are installed on the server.

Figure 81 Compliance Results for a Solaris Server



The values for the Status column are described in the table below.

Table 26 Compliance Status for a Managed Server

Compliance Icon	Compliance Status	Description
	Compliant	All the patch policies attached to a server are compliant. That is, all the patches specified in all the patch policies are installed on the server.
	Non-compliant	At least one of the patch policies attached to the server is not compliant, which means at least one patch in the policy is not installed on the server.
	Scan Started	The patch compliance information is currently being gathered.
	Scan Failed	The patch compliance scan was unable to run.
	Scan Needed	The patch compliance information needs to be gathered or the compliance information may be inaccurate.
—	Not Applicable	The patch compliance information does not apply.

In the SA Client, you can check for patch compliance on an individual server or view overall compliance levels for all servers and groups of servers in your facility.

See [Server Compliance](#) on page 207 for information about compliance scans for all the servers in your data center.

Performing a Solaris Patch Compliance Scan



You must have a set of permissions to perform a patch compliance scan. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide* for more information.

Perform the following steps to scan a server for Solaris patch compliance:

- 1 From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.
- 2 From the Content pane, select a Solaris server.
- 3 Right click or from the **Actions** menu, select **Scan** ► **Patch Compliance**. The Patch Compliance Scan Status window appears and begins the patch compliance scan.
- 4 Click on the status icon in the Status column for more information on the current status.
- 5 When the scan finishes, view the results in the Status column of the Patch Compliance Scan Status window. See also [Table 26, "Compliance Status for a Managed Server"](#) on page 392.
- 6 (Optional) From the Content pane, select Compliance from the View drop down list to view the patch policies that are not compliant. This displays all the patch policies attached to the server and the compliance status of each policy.

Patching Solaris Zones

The Virtualization Director enables you to perform the same operations on virtual servers as you can on physical servers, including audit, remediation, application configuration, software management, patching management, and more. See *SA Users Guide: Server Automation* for information on Virtualization Director.

You can install patches on Solaris global and non-global zones by using Solaris patch policies or by installing patches directly on the virtual server. In the SA Client, you can view the Solaris zones from either the Managed Servers lists or Virtual Servers list.

Viewing Solaris Zones

Perform the following steps to view Solaris zones:

- 1 From the Navigation pane, select **Devices**.
- 2 Expand **Servers**.
- 3 Select Virtual Servers. The list of virtual servers appear in the Content pane.

Or

Select Managed Servers. To see if a server is a hypervisor or a virtual server in the **All Managed Servers** list, choose **Virtualization** from the column selector.

- 4 From the View drop-down list, select **Virtualization**. The Content pane displays the configuration properties of the virtual server.

Installing Patches Using Offline Volumes

This section describes how you can install Solaris patches using offline volumes. This section presumes you are familiar with Solaris Volume Manager.



A sample script is available that you can modify and use to install Solaris patches using offline volumes.

To install Solaris patches using offline volumes, perform the following steps:

- 1 Create a Solaris patch policy that contains the patches you want to install on the server. See [Creating a Solaris Patch Policy](#) on page 357.
- 2 Create a disk mirror on the server being patched.
- 3 Split the mirror.
- 4 Mount the offline disk.
- 5 Create a text file on the server named `/etc/opt/opsware/agent/offline_disk`.
- 6 Edit this file and enter the mount point of the offline disk, for example `/alt`.
- 7 Remediate the server against the patch policy to install the patches on the server.
SA installs the patches to the offline disk at the offline disk mount point listed in the file `/etc/opt/opsware/agent/offline_disk`.
- 8 Reboot the server to the now patched offline disk.
- 9 Verify that the are patches installed correctly on the patched disk and the server is running correctly.
- 10 If the patched disk is behaving as expected, sync the mirror.
If the patched disk is not behaving as expected, reboot the system to the original disk and syncs the mirrors.

8 Patch Management for Unix

Overview of Patching Unix Systems

With SA you can identify, install, and remove patches, and maintain a high level of security across managed servers in your organization. With the SA Client, you can identify and install patches that protect against security vulnerabilities for the AIX and HP-UX operating systems.



For more information on Solaris patching, see [Patch Management for Solaris](#) on page 349. For information on Windows patching, see [Patch Management for Windows](#) on page 287.

This section contains information about how to install and uninstall Unix patches using software policies. It also contains information about generating patch policy compliance reports.

HP Server Automation automates the key aspects of patch management, while offering a fine degree of control over how and under what conditions patches are installed.

Because patches are often released to address grave security threats, an organization needs to be able to roll out patches quickly, before systems become compromised. At the same time, however, patches can cause serious problems, from performance degradation to server failures.

SA allows you to react quickly to newly discovered threats, but it also provides support for strict testing and standardization of patch installation. And, if patches cause problems even after being tested and approved, SA allows you to uninstall the patches in a safe and standardized way.

SA stores patch information in the SA Library that includes detailed information about every server under management, the patches and software installed on the servers, and the patches and software available for installation. You can use this data to determine the severity of your exposure to a newly discovered threat, and to help assess the benefits of rolling out a patch versus the costs in downtime and testing requirements.

By automating the patching procedure, SA can reduce the amount of downtime required for patching. SA also allows you to schedule patch activity, so that patching occurs during off-peak hours.

HP Server Automation automates patch management by providing the following features:

- The SA Library where patches are stored and organized in their formats
- A database that includes information on every patch that has been applied
- Customized scripts that can be run before and after a patch is installed
- Advanced search abilities that identify servers that require patching
- Auditing abilities that enable security personnel to track the deployment of important patches

These features enable you to browse patches by a certain operating system, schedule patch downloads and installations, set up email notifications, preview a patch installation, use software policies and remediation to install and uninstall patches, and export patch information to a reusable file format.

Types of Patch Browsing

The HP Server Automation Client interface organizes Unix patches by operating systems and displays detailed vendor security information about each patch. You can browse patches by patch type, availability, platform version, and so on. You can also browse all patches that are installed on a server, and view and edit patch metadata.

Scheduling and Notifications

You can schedule when patches are uploaded into the SA Library and when they are downloaded to managed servers. As a best practice, patch installations are typically scheduled for a time that causes minimal disruption to an organization's business operation. If you are installing one patch on one server, the installation operation will start only after the download operation has completed.

You can set up email notifications that alert you whether the download and installation operations completed, succeeded, or failed. When you schedule a patch installation, you can also specify reboot preferences to adopt, override, postpone, or suppress the vendor's reboot options.

Using Software Policies to Manage Patches

Software policies enable you to customize patch distribution in your environment. They define the Unix patches that should be installed or not installed on certain managed servers. See [Software Management](#) on page 257 for more information about creating software policies to install Unix patches.

Previewing Patch Installation

While SA allows you to react quickly to newly discovered security vulnerabilities, it also provides support for strict testing and standardization of patch installation. After you have identified patches to install, SA allows you to simulate or preview the installation before you actually install a patch. This preview process tells you whether the servers that you selected for the patch installation already have that patch installed. In some cases, a server could already have a patch installed if a system administrator had manually installed it. The preview process provides an up-to-date report of the patch state of servers.

Software Policy Remediation

SA also provides a solution for remediating servers that are not operating properly due to installed patches. If installed patches cause problems, even after being tested and approved, SA allows you to uninstall the patches in a safe and standardized way. SA allows you to specify uninstall options that control server reboots and the execution of uninstall commands, and pre-uninstall and post-uninstall scripts. Similar to previewing a patch installation, you can also preview a patch uninstall. See [Software Management](#) on page 257 for more information about remediating software policies.

Exporting Patch Data

To help you track the patch state of servers or groups of servers, SA allows you to export this information. This information can be exported in a comma-separated value (.csv) file and includes details about when a patch was last detected as being installed, when a patch was installed by HP Server Automation, the patch compliance level, what patch policy exceptions exist, and so on. You can then import this information into a spreadsheet or database to perform a variety of patch analysis tasks. For more information, see [Exporting a Patch](#) on page 407.

Tracking Patches on Managed Servers

When a server is brought under management by SA, the SA Agent installed on the server registers the server's hardware and software configuration with SA. This information includes installed software and patches, is recorded in the SA Library. The SA Agent repeats this registration every 24 hours.

When a new patch is issued, you can use HP Server Automation to immediately identify the servers that require patching. The SA Library stores patches and other software. You can access the SA Library from the SA Client to install patches on the appropriate servers.

After a server is brought under management, you should install all required patches. If you install a patch manually, HP Server Automation does not have data about that patch until the next SA Agent registration. If you install a patch manually, it can take up to 24 hours until the data about that server in the SA Library is up-to-date.

Whenever you install or uninstall software or patches with HP Server Automation, however, SA immediately updates the information about the server in the SA Library.

Support for Unix Patch Testing and Installation Standardization

With SA you can minimize the risk of rolling out patches. First, when a patch is uploaded into the SA Library, its status is marked as untested and only administrators with special privileges can install it.

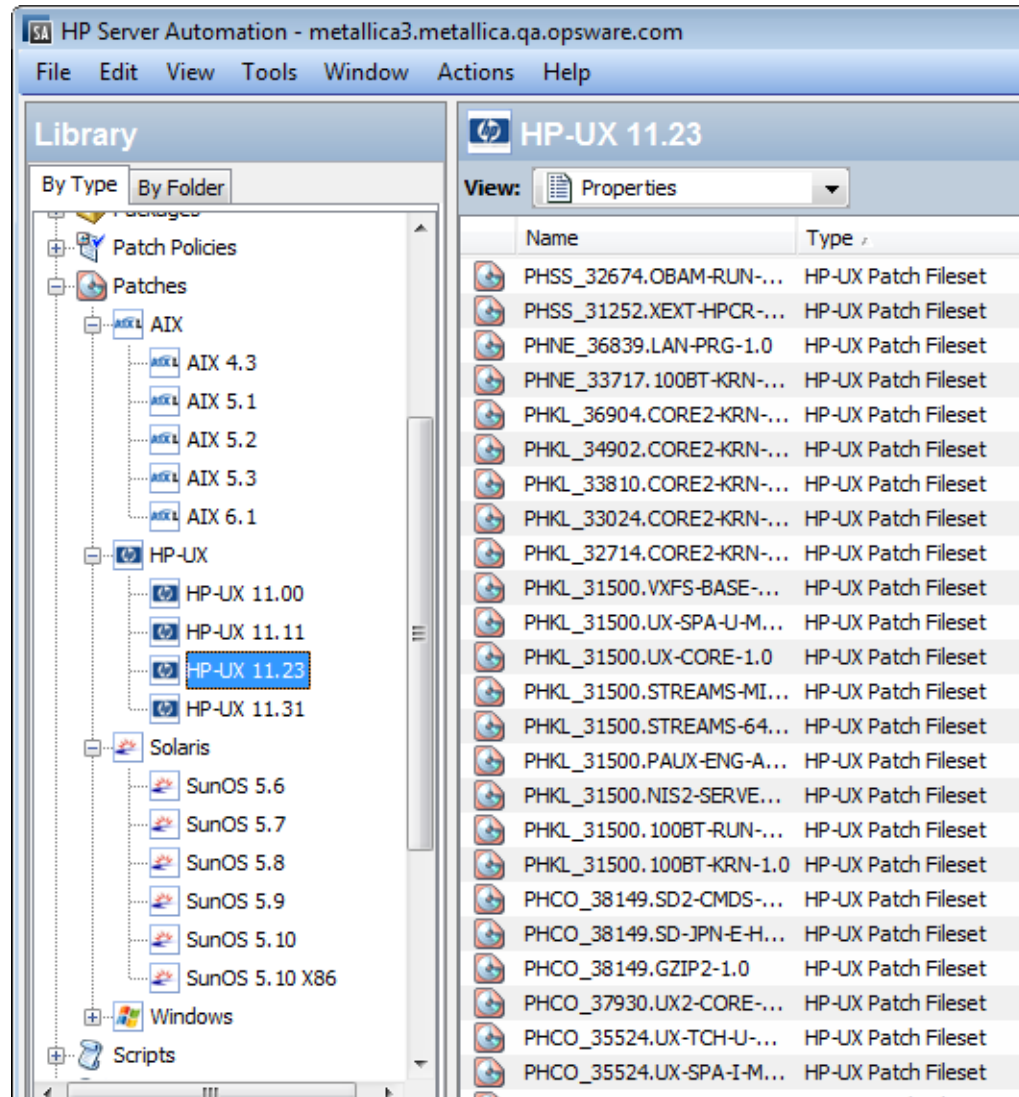
The patch administrator then defines patch installation and uninstallation options and tests the patch. Only after the patch is tested and the patch administrator marks it as available for use can other administrators install the patch.

SA allows you to standardize the way that patches are installed and uninstalled, thereby preventing ad-hoc installation procedures. Patch administrators standardize patch installation by providing pre-install and post-install scripts, install and uninstall flags, reboot instructions, and how to handle error codes from the pre-install and post-install scripts.

Viewing Patches in the SA Client

The SA Client lets you search for and display Unix patches by name, type of patch, operating system, relationship to other packages, and so on. [Figure 82](#) below shows a list of patches for HP-UX 11.23. Use the column selector to right of the column headers to control which columns of patch data to display. For more information, see [Unix Patch Information](#) on page 403 and [Viewing and Editing Unix Patch Properties](#) on page 406.

Figure 82 Unix Patches in the SA Library



Searching for Patches

In the SA Client, you can search for any information about your operational environment that is available in HP Server Automation using the SA Client. The SA Client enables you to search for patches, software policies, servers, and so on. See “SA Client Search” in the *SA Users Guide: Server Automation*.

Patch Management Roles for Unix

HP Server Automation provides support for rigorous change management by assigning the functions of patch management to the patch administrator and the system administrator:

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



Only the patch administrator should have the Patches permission, which gives access to advanced features. To obtain these permissions, contact your SA Administrator. See the Permissions Reference appendix in the *SA Administration Guide*.

Patch Administrator

In most organizations, patch administrators are responsible for reviewing the latest security threats and the patches that vendors have released to address these problems. The patch administrators are generally experts in the operating systems and applications that they manage, and are able to assess the necessity of applying patches issued by vendors. They are able to diagnose common problems that arise after patches are installed, allowing them to thoroughly test the patch application process.

In HP Server Automation, patch administrators are granted specific permissions that allow them to upload patches into HP Server Automation to test the patches and then mark them as available for use. Basic users can upload patches, but they cannot install them or mark them as available. Patch administrators are also able to edit patch options (such as installation scripts) through patch management. Other types of users are not allowed to upload or edit patches.

Typically, the patch administrator uploads patches and then tests them on non-production reference hardware. After testing the patches and determining that the patches are safe to apply to production systems, they mark the patches as available in the HP Server Automation Client, and then advise the system administrators that they must apply the approved patches.

System Administrator

System administrators are responsible for the day-to-day maintenance of the servers in a deployment. These users are not required to have the same level of expertise in low-level system details as the patch administrator.

Because the patch administrator has set up the patch installation, the system administrators can apply the patches to a large number of servers with a few mouse clicks. They are responsible for searching for the servers that require the approved patch, installing the patch, and verifying that the patches are installed successfully.

Patch Management for Specific Unix Operating Systems

The types of patches and their underlying technologies can vary according to the vendor of the operating system. This section discusses the vendor-specific details for Unix patch management in HP Server Automation.

Supported Unix Versions and Patch Types

SA supports all of the operating system versions that HP Server Automation supports, except for Linux.

Linux does not support patches in the ordinary sense. The packages are not patchable. Instead, new versions of the RPM are delivered. Linux systems that HP Server Automation manages are therefore not viewable through the patch interfaces. New Linux packages and updates should be managed and applied through the software policy. See the *SA Policy Setter Guide*, section “RPM Deployment” for information about importing and installing RPMs using a software policy.

To see the Unix versions and patch types that SA supports, perform the following steps.

- 1 In the SA Client, select the Library tab.
- 2 Select the By Type tab.
- 3 Locate and open the Patches node. This displays all the operating systems on which SA supports patches.
- 4 Select an operating system and open the node for that operating system. This displays all the versions of that operating system that SA supports. For an example, see [Figure 82 Unix Patches in the SA Library on page 398](#).

Underlying Technologies for Patch Management on Unix

Although the utilities vary, HP Server Automation enables you to perform patching tasks by using a single interface. HP Server Automation models the way it treats patches by the way the underlying utility treats patches. For example, if the Solaris patchadd utility is not able to install one patch contained in a patch cluster, the Solaris utility continues to install the remaining patches in the patch cluster. HP Server Automation respects this behavior and allows that patch installation operation to continue. Any patches that are not installed are reported at the end of the installation operation.

The following table shows the patch management and installation tools that are used for each of the supported Unix systems.

Table 27 Supporting Technologies for Patch Management on Unix

Solaris	AIX	Hp-UX
Patchadd installs Solaris patches	Installp installs and uninstalls filesets	Swlist lists patch products, files, products, and filesets
Patchrm uninstalls Solaris patches	Lslpp lists installed LPPs	Swinstall installs a depot
Showrev lists installed Solaris patches	Instfix lists installed APARs	Swremove removes a depot
Pkgadd installs Solaris packages		
Pkginfo lists installed Solaris packages		

AIX Patches

AIX periodically releases Authorized Program Analysis Reports (APARs), which specify what update filesets (contained in LPPs) are necessary to fix an identified problem. An APAR only specifies the minimum version of an update fileset required to fix a problem; an APAR can therefore be satisfied with later versions of the same filesets. To maintain compatibility, however, HP Server Automation always adopts the fileset with the lowest version number that meets the minimum version that APAR specifies. If a later version of the update fileset is uploaded, HP Server Automation still associates the earlier version of the fileset with the APAR.

When uploading an LPP, HP Server Automation recognizes which APARs the filesets contained in the LPP belong to. An entry is created for the APAR in the SA Library when the first fileset associated with an APAR is uploaded. (In some cases, a fileset is associated with more than one APAR. An entry is created for each APAR the fileset is associated with, if the entry does not already exist.)

If you want to install all LPPs that APAR specifies, you must make certain to upload all of the specified LPPs into the SA Library.

If you do not upload all of the LPPs that APAR specifies, it is still possible for the system administrator to browse for an APAR and install the partial set of LPPs that are uploaded. In such cases, the administrator receives a warning that the filesets for the APAR are not all installed.



The Patch Administrator must first upload and test an LPP before it is generally available in HP Server Automation. The new fileset is integrated into the APAR only after the LPP is tested and approved. Even though the APAR is updated automatically, you still maintain control over the exact filesets that are allowed to be installed on your managed servers.



APAR update filesets cannot be installed on a server if the server does not already have the base filesets for which the update filesets are intended.

If, however, a server has a partial set of the base filesets, the APAR can be applied and only the applicable filesets for the base filesets are installed. For example, if an APAR specifies four update filesets to update four base filesets, and you attempt to apply the APAR to a server that has only three of the base filesets, three of the four update filesets from the APAR are installed.

When installing an AIX update fileset, the SA normally applies the fileset, which allows it to be rejected (uninstalled.) If you want to commit the fileset instead (so that it cannot be removed), use the `-c` option here.



Since update filesets can be included in folders, global read permissions are required to view and edit AIX update filesets. See “Software Management Setup” in the *SA Policy Setter Guide* for information about how to use folders.

Solaris Patches

A Solaris patch cluster contains a set of selected patches for a specific Solaris release level. Ordinarily, after a patch cluster is installed, it is not possible to search for a particular patch cluster. The patches do not contain any metadata that relate them to the patch cluster in which they were originally bundled. You can only search for the individual patches.

If you install a Solaris patch cluster, however, HP Server Automation keeps track of the patch cluster in the SA Library. You can therefore search for a patch cluster to determine if a full patch cluster is installed. If you installed the patch cluster, you can uninstall individual patches in the cluster. You cannot uninstall a patch cluster.

For more information on Solaris patching, see [Patch Management for Solaris](#) on page 349

HP-UX Patches

HP-UX patches are delivered exclusively as depots, which are patch products that contain patch filesets. The depot is uploaded directly into HP Server Automation.

If a depot is already uploaded and attached to a node, it cannot be uploaded by SA. If you want to upload the depot with SA, you must detach a depot from any nodes that it is attached to, and then delete it from the SA Library.

Uploading Unix Patches into the SA Library

Before a Unix patch can be installed on a managed server, the patch must be downloaded from the server vendor and uploaded into the SA Library. For more information, see the *SA Administration Guide* and the *SA Policy Setter Guide*.

To upload Unix patches in to the SA Library, perform the following steps.

- 1 From the Navigation pane, select **Library ► By Type ► Patches**. The patches are organized by operating system.
- 2 Navigate to the desired operating system version.

- 3 From the **Actions** menu, select **Import Software...**. The Import Software window appears.
- 4 Click **Browse** to locate and select the patch to import.

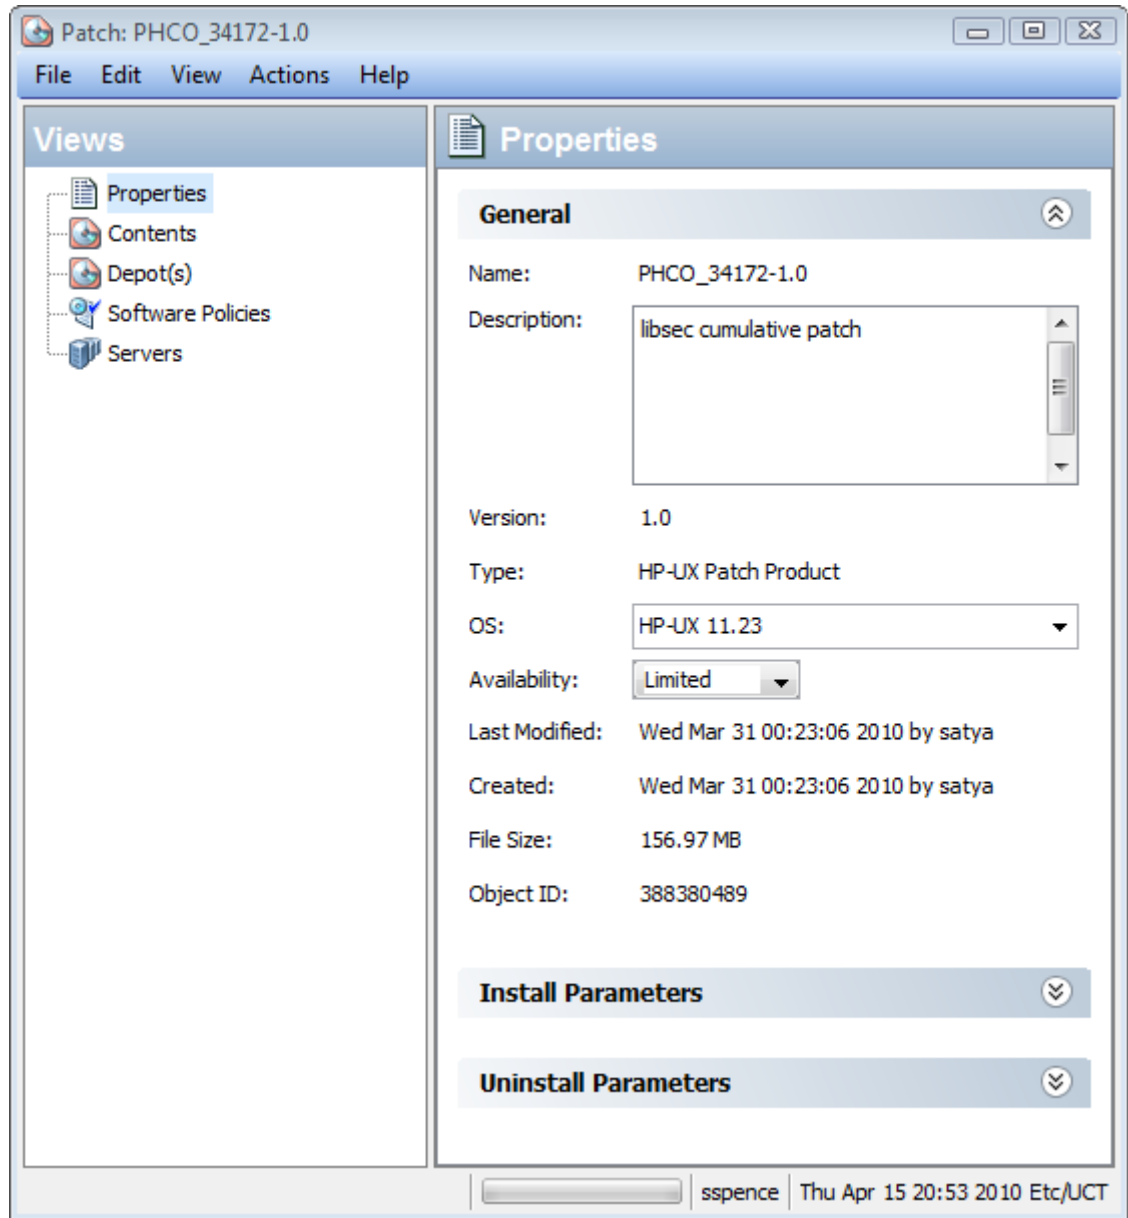
Before clicking Open in the Open window, select the character encoding to be used by the patch from the Encoding drop-down list.

You need to specify the character encoding so that SA can extract the metadata contained in the patch and correctly display the information in non-ASCII characters in the SA Client (for example, in the Patch Properties pages). Patch metadata includes comments, READMEs, scripts, descriptions, and content lists.
- 5 Click Open.
- 6 In the Import Software window, select the appropriate type from the Type drop-down list.
- 7 In the Folder field, enter the path in the SA Library where you want to store the patch. Or select the Browse button.
- 8 From the Platform drop-down list, select all the operating system versions that the patch applies to. You can only install the patch on servers that are running those versions of the operating system.
- 9 Click **Import** to import the patch into the SA Library.

Unix Patch Information

The SA Client displays detailed information about a patch in several different views. For example, [Figure 83](#) below shows the Properties view of an HP-UX patch. Note that the details about each patch vary depending on the type and OS of patch. To view or edit patch properties, see [Viewing and Editing Unix Patch Properties](#) on page 406.

Figure 83 Unix Patch Properties in the SA Client



The Patch Properties View

Patch properties include the following information. Note that some information is only displayed for certain operating systems and not others.

- **Version:** The version number of the patch.
- **Status:** The vendor's status for the patch.
- **Type:** The type of Unix patch. Some examples are HP-UX Patch Product, HP-UX Patch Fileset, Solaris Patch, Solaris Cluster, AIX APAR and AIX Update Fileset.
- **OS:** The Unix operating systems that are known to be affected by this patch.
- **Availability:** The status of a patch within HP Server Automation, which can be one of the following:

- **Limited:** The patch has been imported into SA but requires additional permissions (Manage Patch: Read & Write) to be installed. This is the default patch availability. For more information on permissions, see the *SA Administration Guide*.
- **Available:** The patch has been imported into HP Server Automation, tested, and has been marked available to be installed on managed servers.
- **Deprecated:** The patch cannot be added to patch policies or set as a patch policy exception but can still be installed.
- **Object ID:** The HP Server Automation unique ID for the patch.
- **Dependencies:** When present, lists the dependencies on the selected patch. This is only provided for some patch types and some platforms. For more information, see [Viewing Solaris Patch or Patch Cluster Properties](#) on page 376.
- **Install Parameters:** When present, lists the actual install settings for the patch and the settings that the patch vendor specifies for the patch. This is only provided for some patch types and some platforms.
- **Install Scripts:** When present, lists scripts that will run on a managed server before or after the patch is installed. This is only provided for some patch types and some platforms.
- **Uninstall Parameters:** When present, lists the actual uninstall settings for the patch and the settings that the patch vendor specifies for the patch. This is only provided for some patch types and some platforms.
- **Uninstall Scripts:** When present, lists scripts that will run on a managed server before or after the patch is uninstalled. This is only provided for some patch types and some platforms.

The Contents View

Patch Contents are displayed only for certain types of patch containers such as HP-UX Patch Products, AIX APARs and Solaris Clusters. The Contents view lists all the patches included in the selected patch container.

The Patch Depots View - HP-UX Only

Patch Depots are only displayed for HP-UX Patch Products. The Depots view displays the HP-UX depots that contain the selected patch product. SA displays HP-UX depots as SA packages.

The Patch Products View - HP-UX Only

Patch Products are only displayed for HP-UX Patch Filesets. The Patch Products view displays the HP-UX patch products that contain the selected HP-UX patch fileset.

The Patch Clusters View - Solaris Only

Patch Clusters are only displayed for Solaris patches. The Patch Clusters view displays the Solaris patch clusters that contain the selected Solaris patch. For more information on Solaris patches, see [Patch Management for Solaris](#) on page 349.

The LPPs/APARs View - AIX Only

The LPPs/APARs view is only displayed for AIX patches. This view displays the LPPs and APARs that contain the selected patch.

The Software Policies View

The Software Policies view displays all the software policies that include the selected patch.

The Patch Policies View

The Patch Policies view displays all the patch policies that include the selected patch. The Patch Policies view is only displayed for some platforms.

The Servers View

The Servers view displays all the servers where the selected patch is installed.

Viewing and Editing Unix Patch Properties

The SA Client displays information about Unix patches that have been imported into HP Server Automation as described in [Unix Patch Information](#) on page 403. You can edit some of a patch's properties in the properties view. Some properties are not editable.

You can set the install and uninstall parameters on either the patch properties page or when you are install or uninstall the patch. The parameters on the Properties view are saved in the SA Library, but the parameters specified during a patch install or uninstall are used only for that action. The parameters specified during an install or uninstall override those on the patch Properties view.

To view or edit information about a patch, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Unix operating system.
- 3 (Optional) Use the column selector to sort the patches according to Name, Type, Availability, and Description.
- 4 In the content pane, select a patch.
- 5 Right click the patch or select the Actions menu and select the Open menu. This displays the patch in a separate screen.
- 6 If you have modified any properties, select **File ► Save** to save your changes.

Finding Servers That Have a Unix Patch Installed

To find out which servers have a particular patch installed, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Unix operating system. The content pane will display all patches associated with that operating system.

- 3 From the content pane, select a patch.
- 4 From the View drop-down list in the content pane, select Servers. This shows all the servers where the selected patch is installed.

Exporting a Patch

You can export patches to the local file system. However, not all patch types can be exported. If you attempt to export a patch and find that the Export menu is grayed out, that patch cannot be exported.

To export a patch from the SA Library to the local file system, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Unix operating system. The content pane will display all patches associated with that operating system.
- 3 From the content pane, select a patch.
- 4 From the **Actions** menu, select **Export**. If the Export menu is grayed out, that patch cannot be exported.
- 5 In the Export Patch window, enter the folder name that will contain the patch file in the File Name field.
- 6 Click **Export**.

Deleting a Patch

This action removes a patch from HP Server Automation, but does not uninstall the patch from managed servers. A patch cannot be deleted if it is attached to a policy.



Do not delete all of the patches from HP Server Automation. If you do so accidentally, contact your support representative for assistance in uploading all of the patches back into SA.

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Unix operating system. The content pane will display all patches associated with that operating system.
- 3 From the content pane, select a patch.
- 4 From the **Actions** menu, select **Delete Patch**.
- 5 In the Delete Patches windows, click **Delete**.

Using Software Policies to Manage Patches

Patch Policies for Windows and Solaris are the best way to manage patches for the Windows and Solaris platforms. For more information see [Patch Management for Windows](#) on page 287 and [Patch Management for Solaris](#) on page 349.

For other platforms, software policies enable you to customize patch distribution in your environment. Software policies define which Unix patches should be installed or not installed on certain managed servers.

If you use software policies and you also perform ad hoc patch installs, you must run the remediate process to install all applicable patches on servers. See [Software Management](#) on page 257 for more information about creating and remediating software policies to install Unix patches.

Patch Compliance Reports

To troubleshoot and resolve patch compliance problems, you can run and examine several patch compliance reports in the SA Client. The following patch compliance reports identify whether all patches in a software policy were installed successfully on managed servers in your environment.

Patch Policy Compliance (All Servers)

This report groups all managed servers by their patch policy compliance level to show compliant and non-compliant servers.

Patch Policy Compliance by Customer

This report lists all servers by the customer they belong to and then by the patch policy compliance level.

Patch Policy Compliance by Facility

This report groups all managed servers by the facility they belong to and then by the patch software policy compliance level.



See the *SA Users Guide: Server Automation* for information about how to run, export, and print these reports.

Patch Administration for Unix

You can customize patch administration for Unix to best support your environment by setting the availability flag.

Setting the Default Patch Availability

You can set the default patch availability with the SA Client. The default used by the script overrides the default set by the SA Client. See the *SA Administration Guide* for information about the script.

To set the default value for the Availability of a newly imported patch, perform the following steps:

- 1 From the Navigation pane, select Administration.
- 2 Select Patch Configuration.
- 3 For the Default Availability for Imported Patches, select either Available or Limited. The default is Limited.

If the patch is Available, it can be installed on managed servers. If the patch is Limited, it has been imported into HP Server Automation and can be installed only by a patch administrator who has the required permissions (Manage Patch: Read & Write). To obtain these permissions, contact your SA Administrator. See also the *SA Administration Guide*.

Patch Installation

The patch installation process consists of the following two phases:

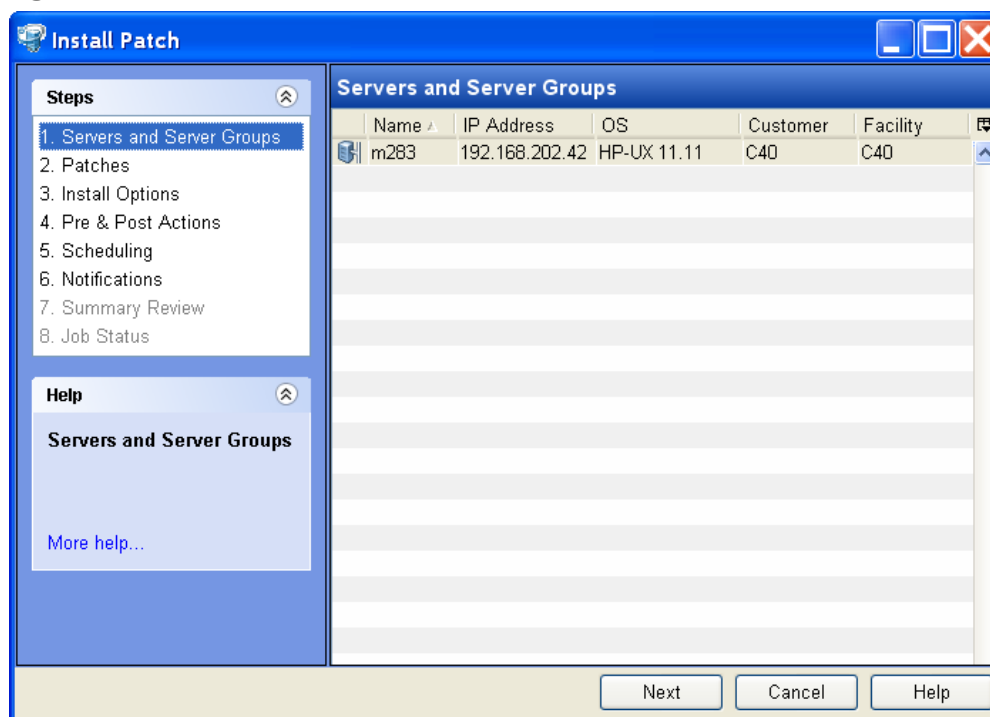
- **Download Phase:** This is when the patch is downloaded from HP Server Automation to the managed server. This phase is commonly referred to as the staging phase.
- **Installation Phase:** This is when the patch is installed on the managed server. This phase is commonly referred to as the deployment phase.

You can specify whether you want the installation to occur immediately after the patch is downloaded (staged) or you can schedule the installation to occur at a later date and time. SA also supports best-effort installations of multiple patches by allowing you to specify that the patch installation process will continue even when an error occurs with one of the patches.

SA displays the name of the command that installs the patch. The SA Agent runs this command on the managed server. You can override the default command-line arguments that you want to perform the installation.

To optimally manage Unix patch installations, SA allows you to manage server reboot options, and pre and post installation scripts, simulate (preview) a patch installation, and set up email notifications to alert you about the status of the installation process. The Install Patch window guides you through setting up these conditions.

Figure 84 Install Patch Window



Installation Flags

You can specify installation flags that are applied whenever a Unix patch is installed. However, HP Server Automation also uses default installation flags and requires that patches are installed with these flags. You must therefore be certain that you do not specify any

installation flags that override or contradict the default flags passed in by HP Server Automation. See [Setting Unix Install Options](#) on page 412 for information about how to specify commands.

The following table lists the default installation flags that HP Server Automation uses.

Table 28 Default Installation Flags

Unix Patch Type	Flags
AIX	-a -Q -g -X -w
HP-UX	None

Application Patches

SA does not allow you to apply a patch to an operating system for which the patch is not intended. When you are installing an application patch, SA does not automatically filter out servers that do not have the corresponding application installed. Although SA does not prevent you from doing so, you should not attempt to apply application patches to servers that do not have the necessary applications installed. If a patch is for an application that is not installed on the server, the patch will not be applied and an error message will display, such as “There was an error with package <name of the package>”.

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

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

Installing a Unix Patch

Before a patch can be installed on a managed server, it must be imported into HP Server Automation and its status must be Available. Administrators who have the required permissions can install patches that are marked Limited.



You must have a set of permissions to manage patches. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide*.

You can perform the installation by explicitly selecting patches and servers.

To install a patch on a managed server, perform the following steps:

- 1 From the Navigation pane, select Library and then select Patches.
- 2 Expand the Patches and select a specific Unix operating system.
- 3 From the content pane, select a patch.
- 4 From the View drop-down list, select Servers (or Server Groups).

- 5 From the Show drop-down list, select Servers without Patch Installed (or Server Groups without Patch Installed).
- 6 From the Preview pane, select one or more servers.
- 7 From the **Actions** menu, select **Install Patch**.

The first step of the Install Patch window appears: Servers and Server Groups. For instructions on each step, see the following sections:

- [Setting Unix Install Options](#)
- [Setting Reboot Options for a Unix Patch Installation](#)
- [Specifying Install Scripts for a Unix Patch Installation](#)
- [Scheduling a Unix Patch Installation](#)
- [Setting Up Email Notifications for a Unix Patch Installation](#)
- [Previewing a Unix Patch Installation](#)
- [Viewing Job Progress of a Unix Patch Installation](#)

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8 When you are ready to launch the installation job, click **Start Job**.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If the Install Patch window remains open until the job completes, SA updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select **Refresh** from the **View** menu to update information in the Patch Preview pane.

Setting Unix Install Options

You can specify the following types of patch installation options:

- Perform the patch installation immediately after the patch is downloaded or at a later date and time.
- Do not interrupt the patch installation process even when an error occurs with one of the patches.
- Use different command-line options to perform the installation.

To set these options, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Install Options step.
- 2 Select one of the following Staged Install Options:

Continuous: This allows you to run all phases as an uninterrupted operation.

Staged: This allows you to schedule the download and installation to run separately.

- 3 Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.
- 4 In the Install Command text box, enter command-line arguments for the command that is displayed.

- 5 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Setting Reboot Options for a Unix Patch Installation

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is installed on it, completely suppress all server reboots, or postpone reboots until all patches have been installed.



When you are selecting reboot options in the Install Patch window, Hewlett Packard recommends that you use the Unix reboot recommendations, which is the “Reboot servers as specified by patch properties” option. If you cannot use the Unix reboot setting, select the single reboot option, which is the “Do not reboot servers until all patches are installed” option.

The following options determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Install Patch window; they do not change the Reboot Required option, which is on the Install Parameters tab of the patch properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install:** Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Do not reboot servers until all patches are installed:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select one of the Rebooting Options.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Specifying Install Scripts for a Unix Patch Installation

For each patch, you can specify a command or script to run before installation or after installation. A pre-install script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-install script fails, the patch would not be installed. A pre-install script could also be used to shut down a service or application before it is patched. A post-install script could be used to perform a certain cleanup process on the managed server.

You can also specify the following types of scripts to run on the managed server before or after an installation or download phase:

- **Pre-Download:** A script that runs before patches are downloaded from SA to the managed server. This is available only if you select Staged in the Install Options step.

- **Post-Download:** A script that runs after patches are downloaded from SA to the managed server and before the patch is installed. This is available only if you select Staged in the Install Options step.
- **Pre-Install:** A script that runs before patches are installed on the managed server.
- **Post-Install:** A script that runs after patches are installed on the managed server.

To specify a pre-install script, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select the Pre-Install tab. You may specify different scripts and options on each of the tabs.
- 3 Select Enable Script. This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
- 4 Select either Saved Script or Ad-Hoc Script.
A Saved Script has been previously stored in HP Server Automation with the SAS Web Client. To specify the script, click **Select**.
- 5 If the script requires command-line flags, enter the flags in the Command text box.
- 6 Specify the information in the User section. If you choose a system other than Local, enter the User Name, Password, and Domain. The script will be run by this user on the managed server.
- 7 To stop the installation if the script returns an error, select the Error check box.
- 8 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

Scheduling a Unix Patch Installation

Since the two phases of patching can be decoupled, you can schedule when you want patches installed (deployed) to occur independently of when patches are downloaded (staged).

To schedule a patch installation, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Scheduling step.
By default, the Scheduling step displays only the scheduling options for the install phase. If you selected Staged in the Install Options step, the scheduling options for the download phase will also be displayed.
- 2 Select one of the following Install Phase options:
 - **Run Task Immediately:** This enables the system to perform a preview analysis in the Summary Review step. The scheduling option for the download phase is **Run Immediately Following Download**.
 - **Run Task At:** This enables you to specify a later date and time that you want the installation or download performed.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.





A scheduled patch installation can be cancelled (prior to its execution), even if the patch download has already completed.

Setting Up Email Notifications for a Unix Patch Installation

You can set up email notifications to alert users when the download and installation operations complete successfully or with errors.

To set up email notifications, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Notifications step.
- 2 To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase.
- 3 Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 4 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.



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

Previewing a Unix Patch Installation

The installation preview process provides an up-to-date report about the patch state of servers. The installation preview is an optional step that lets you see what patches will be installed on managed servers and what type of server reboots are required. This preview process verifies whether the servers you selected for the patch installation already have that patch installed. In some cases, a server could already have the patch installed if a system administrator had manually installed it, which means that SA does not know about it.

The preview process also reports on dependency information, such as patches that require certain Unix products, and patches that obsolete other patches or are obsoleted by other patches. If a dependency is not met, SA will display an error message indicating this condition.



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

To preview a patch installation, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Summary Review step.
- 2 Verify the information displayed for the Servers, Server Groups, and Patches at the top of the window.
- 3 (Optional) Click **Preview** to see the separate actions that will be performed when the patch is installed. To view the details of a previewed action, select a row in the table.
- 4 Click **Start Job** to launch the installation job or click **Cancel** to close the Install Patch window without launching the installation.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Unix Patch Installation

You can review progress information about a patch installation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

- 1 From the Install Patch window, click **Next** to advance to the Job Progress step. This will start the installation job.

The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:

- **Analyze:** HP Server Automation examines the patches needed for the installation, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.
 - **Download:** The patch is downloaded from HP Server Automation to the managed server.
 - **Install:** After it is downloaded, the patch is installed.
 - **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
 - **Pre/Post Install/Download Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstallation.
 - **Install & Reboot:** When a patch will be installed is also when the server will be rebooted.
 - **Verify:** Installed patches will be included in the software registration.
- 2 To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select *Jobs and Sessions* to review detailed information about the job. See the *SA Users Guide: Server Automation* for more information about browsing job logs.
 - 3 Click **Stop Job** to prevent the job from running or click **Close** to close the Install Patch window.

Patch Uninstallation

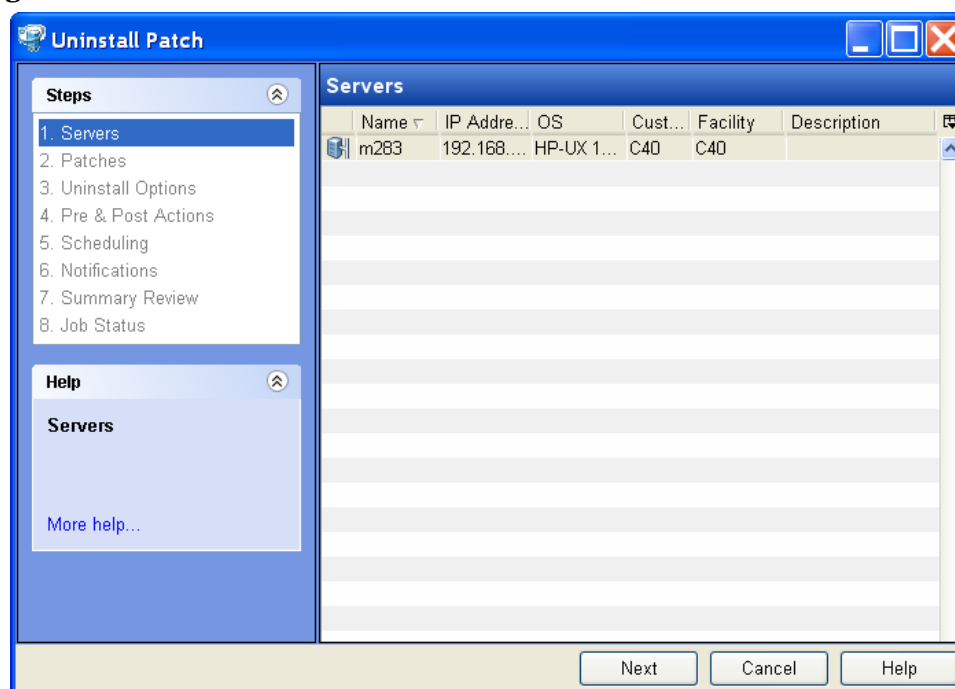
SA provides granular control over how and under what conditions Unix patches are uninstalled (removed) from managed servers. To minimize problems, you can only uninstall one patch at a time. You cannot use SA to uninstall a patch that was not installed using SA.

To help you optimally manage these conditions, SA allows you to do the following:

- Manage server reboot options, and pre and post installation scripts.
- Simulate (preview) a patch uninstallation.
- Set up email notifications to alert you about the status of the uninstallation process.

The Uninstall Patch window guides you through setting up these conditions.

Figure 85 Uninstall Patch Window



Uninstallation Flags

You can specify uninstallation flags that are applied whenever a Unix patch is uninstalled. However, HP Server Automation also uses default uninstallation flags and requires that patches are uninstalled with these flags. You must therefore be certain that you do not specify any uninstallation flags that override or contradict the default flags passed by HP Server Automation.

The following table lists the default uninstallation flags that HP Server Automation uses.

Table 29 Default Uninstallation Flags

Operating System/Patch Types	Flags
AIX	-u -g -X
AIX Reject Options	-r -g -X
HP-UX	None

Uninstalling a Unix Patch

To remove a patch from a managed server, perform the following steps:

- 1 From the Navigation pane, select Library and then select Patches.
- 2 Expand the Patches and select a specific Unix operating system.
- 3 From the content pane, select a patch.
- 4 From the View drop-down list, select Servers.
- 5 From the Show drop-down list, select Servers with Patch Installed.
- 6 From the Preview pane, select one or more servers.
- 7 From the **Actions** menu, select **Uninstall Patch**.

The first step of the Uninstall Patch window appears: Servers.

For instructions on each step, see the following sections:

- [Setting Reboot Options for a Unix Patch Uninstallation](#)
- [Specifying Pre and Post Install Scripts for a Unix Patch Uninstallation](#)
- [Scheduling a Unix Patch Uninstallation](#)
- [Setting Up Email Notifications for a Unix Patch Uninstallation](#)
- [Viewing Job Progress of a Patch Uninstallation](#)

After you have completed a step, select **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

- 8 When you are ready to launch the uninstallation job, select **Start Job**.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

- [If the Uninstall Patch window remains open until the job completes, SA updates the Patch Compliance column in the All Managed Servers window with the revised compliance count \(in parenthesis\) for affected servers. Press F5 or select Refresh from the View menu to update information in the Patch Preview pane.](#)

Setting Uninstall Options

You can specify the following types of patch uninstallation options:

- Do not interrupt the patch uninstallation process even when an error occurs with one of the patches.
- Use different command-line options to perform the uninstallation.

To set these options, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Uninstall Options step.
- 2 Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.
- 3 In the Uninstall Command text box, enter command-line arguments for the command that is displayed.
- 4 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Setting Reboot Options for a Unix Patch Uninstallation

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is removed from it, completely suppress all server reboots, or postpone reboots until all patches have been uninstalled.



When you are selecting reboot options in the Uninstall Patch window, Hewlett Packard recommends that you use the Unix reboot recommendations, which is the “Reboot servers as specified by patch properties” option in the window. If it is not possible to use the Unix reboot setting, select the single reboot option, which is the “Do not reboot servers until all patches are installed” option in the window.

The following options determine whether the servers are rebooted after the patch is uninstalled. These options apply only to the job launched by the Uninstall Patch window; they do not change the Reboot Required option, which is on the Uninstall Parameters tab of the patch properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties:** By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install:** Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots:** Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Do not reboot servers until all patches are installed:** If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select one of the Rebooting Options.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Specifying Pre and Post Install Scripts for a Unix Patch Uninstallation

For each patch, you can specify a command or script to run before uninstallation or after uninstallation. A pre-uninstall script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-uninstall script fails, the patch would not be removed from a server. A pre-uninstall script could also be used to shut down a service or application before it is removed from a server. A post-uninstall script could be used to perform a certain cleanup process on the managed server.

You can specify the following types of scripts to run on the managed server before or after a patch uninstallation:

- **Pre-Uninstall:** A script that runs before the patch is removed from a managed server.
- **Post-Uninstall:** A script that runs after the patch is removed from a managed server.

To specify a script, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Pre & Post Actions step.
- 2 Select the Pre-Uninstall or Post-Uninstall tab.
You may specify different scripts and options on each of the tabs.
- 3 Select Enable Script.
This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
- 4 Select either Saved Script or Ad-Hoc Script.
A Saved Script has been previously stored in HP Server Automation with the SAS Web Client. To specify the script, click **Select**.
- 5 If the script requires command-line flags, enter the flags in Commands.
- 6 Specify the information in the User section. The script will be run by this user on the managed server.
- 7 To stop the uninstallation if the script returns an error, select Error.

Scheduling a Unix Patch Uninstallation

You can schedule that a patch will be removed from a server immediately, or at a later date and time.



To schedule a patch uninstallation, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Scheduling step.
- 2 Select one of the following Install Phase options:
 - **Run Task Immediately:** This enables you to perform the uninstallation in the Summary Review step.
 - **Run Task At:** This enables you to specify a later date and time that you want the uninstallation performed.
- 3 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Setting Up Email Notifications for a Unix Patch Uninstallation

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

To set up email notifications, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Notifications step.
- 2 To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the uninstallation phase.
- 3 Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 4 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

Previewing a Unix Patch Uninstallation

The uninstallation preview process provides an up-to-date report about the patch state of servers. The uninstallation preview is an optional step that lets you see what patches will be removed from managed servers. This preview process verifies whether the servers you selected for the patch uninstallation have that patch installed.



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

To preview a patch uninstallation, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Summary Review step.
- 2 Verify the information displayed for the Servers, Server Groups, and Patches at the top of the window.
- 3 (Optional) Click **Preview** to see the separate actions that will be performed when the patch is uninstalled. To view the details of a previewed action, select a row in the table.
- 4 Click **Start Job** to launch the job or click **Cancel** to close the Uninstall Patch window without launching the uninstallation.

If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Patch Uninstallation

You can review progress information about a patch uninstallation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

- 1 From the Uninstall Patch window, click **Next** to advance to the Job Progress step. The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:
 - **Analyze:** HP Server Automation examines the patches needed for the uninstallation, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.

- **Uninstall:** The patch is uninstalled.
 - **Final Reboot:** If this action is specified in the Pre & Post Actions step, the server is rebooted.
 - **Pre/Post Uninstall Script:** If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstallation.
 - **Uninstall & Reboot:** When a patch will be installed is also when the server will be rebooted.
 - **Verify:** Installed patches will be included in the software registration.
- 2 To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select **Jobs and Sessions** to review detailed information about the job. See the *SA Users Guide: Server Automation* for more information on browsing job logs.
 - 3 Click **Stop Job** to prevent the job from running or click **Close** to close the Uninstall Patch window.

9 Script Execution

Overview of Script Execution

The Script Execution feature allows you to automate the management and execution of scripts in the SA Client. It also allows you to organize your scripts in folders and define security permissions around them. From the SA Client, you can create or upload a script, set it up to run simultaneously across multiple Unix or Windows servers, and monitor it as it executes on each server. After a script is executed, you can view the results for every server and then export the script results. You can also modify, delete, and rename a script. You can also execute scripts in the Global Shell using the SA Client.

Script Execution Features

The Script Execution feature in the SA Client enables you to perform the following functions:

- Organize your scripts into folders and define security permissions to control access of their contents across different users and user groups.
- Create or upload scripts in the SA Client.
- Run scripts across multiple Unix or Windows servers or server groups.
- Execute scripts in the Global Shell.
- Schedule one time or recurring script execution jobs.
- Notify the status of the script execution job via email.
- Approve script execution jobs.
- View the script output against multiple servers in a tabular format.
- Export the script execution results.
- Search for scripts and script execution jobs.

Script Execution Process

The script execution process involves defining permissions, managing scripts, and executing scripts.

- **Defining Permissions**

In this phase, an SA Administrator assigns Folder permissions, Client feature permissions, and Customer constraints to define the security boundaries across various user groups. The permissions determines the actions the users in a user group can perform with the SA Client.

See the *SA Administration Guide* for more information about defining security permissions.

- **Managing Scripts**

In this phase, a policy setter or an advanced system administrator performs script management tasks such as creating or importing scripts, editing script properties, exporting scripts, and deleting scripts. See [Managing Scripts](#) on page 426 for more information.

- **Executing Scripts**

In this phase, a system administrator executes server scripts directly on servers or server groups and OGFS scripts in the Global Shell. A system administrator can also execute scripts by adding the scripts to a software policy and then remediating the servers against the software policy. See [Executing Scripts](#) on page 433 and [Install Software Using a Software Policy](#) on page 267 for more information.

Types of Scripts

In the SA Client, the Script Execution feature supports two main types of scripts : Server scripts and OGFS scripts.

The Server script allows you to execute scripts on Unix and Windows servers managed by SA. The SA Client supports the following types of Server scripts for Unix and Windows operating systems: Unix/Linux shell, Windows batch (.BAT), Windows Visual Basic (VBScript), and Windows PowerShell.

The OGFS scripts allows you to execute scripts in the Global Shell from the SA Client. You can specify the directory path in the OGFS to execute the scripts. See the *SA Users Guide: Server Automation* for information about the Global Shell.

The server scripts are further classified to Saved Scripts, and Ad-hoc Scripts.

- Saved scripts are accessible to all the users, if they have the appropriate permissions. You are required to have the appropriate permissions to create, view, edit, and execute shared scripts. Private scripts are only accessible to the user who created them. They can only be created, edited, deleted, or executed by the user who created the script.
- Ad-Hoc scripts are created or uploaded for one-time use and is not stored in HP Server Automation. Ad- Hoc script is created or uploaded and then immediately executed by a user and during this process , only one user has access to the script.

After you create a script and save it as a specific type of script in HP Server Automation, you cannot convert the script to the other type of script.

In the SA Client, you can specify to run a Sever script as a Super User or as a specified user. A Super User script allows you to execute the script as root on UNIX or Local System on Windows servers without entering a password. If the script is not designated as a Super User Script, then you need to enter a username and password to run the script. You also require the appropriate permissions to manage and run Super User Server Scripts. See the *SA Administration Guide* for information on the permissions required to run the Super User Server Scripts. All the OGFS scripts can only be executed as an SA User.

Managing Scripts

The script management tasks include:

- [Creating a Script](#)
- [Opening a Script in the SA Client](#)
- [Editing Script Properties](#)
- [Locating Scripts in Folders](#)
- [Exporting a Script](#)
- [Renaming a Script](#)
- [Deleting a Script](#)



You must have a set of permissions to create and manage a script. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide* for more information.

Creating a Script

In the SA Client, you can create a script from either the By Type or the By Folder view in the Library.

Script Creation Guidelines

HP Server Automation supports the following types of Server scripts for Unix and Windows operating systems: Unix/Linux shell, Windows batch (.BAT), Windows Visual Basic (VBScript), and Windows PowerShell.

When creating scripts you must adhere to the following guidelines:

- 4 MB is the maximum size allowable for a script.
- When you create a Unix shell script with a language other than the Bourne (sh) shell, use the sh-bang (!) format at the top of the script to specify the correct command interpreter. The command interpreter needs to be present on the managed server.

For example, if you are using Perl, the beginning of the script would contain the following line:

```
#!/usr/bin/perl
```

The following example shows a short Perl script (it displays “hello world”):

```
#!/usr/bin/perl
print "hello world\n"
```

- VBScripts are executed by the VBScript interpreter on the Windows server.
- To access command line parameters with Unix shell commands, use the following convention: \$1 \$2...
- To access command line parameters with Windows .BAT, use: %1 %2...

- Script lines do not need to be terminated in a specific way. But with Windows scripts, HP Server Automation converts all `\n` to `\r\n`. With Unix scripts, all `\r\n` are converted to `\n`.
- Scripts should be written to send error output to standard error.
- Scripts should use the standard convention of returning a zero code to indicate success. For other return codes, there is no standard code system to follow. Create unique non-zero return codes to handle each type of error.

Creating a Script from the By Type View in the Library

To create a script perform the following steps:

- 1 From the Navigation pane, select **Library** ► **By Type** ► **Scripts**. The three main types of scripts appear in the content pane.
- 2 Select the script type and then from the **Actions** menu, select **New script**. The Script window appears as shown in [Figure 86](#).

Figure 86 Script Window

Properties

Name: Simple BAT

Type: Windows .BAT

Location: / Select

Changes Server: ☒ Yes ☐ No

Run as super user: ☒ Yes ☐ No

Script Contents: Enter the script contents or import a script file Import Script File

dir c:\temp

Description: dir

Last Modified: Thu Aug 02 18:16:24 2007

Last Modified By: paul

Created: Thu Aug 02 18:16:18 2007

Created By: paul

Opsware ID: 195380040

- 3 In the Name field, enter the name of the script.
- 4 (Windows only) Select the script type from the Type drop-down list.

- 5 Click **Select** to specify the location for the script in the folder hierarchy. The Select Folder window appears. Select a folder in the Library to specify the location of the script and then click **Select**.
- 6 In the Changes Server field, select Yes, if the script causes a change in the server configuration when executed.
- 7 In the Run as Super User field, select Yes if the script can be run as a Super User when executed. Selecting yes, allows you to run the script as a Super User without providing a password for the script.

This option is enabled only if you have to appropriate permission. See the *SA Administration Guide* for more information about script execution permissions.
- 8 In the Script Contents field, enter the contents of the script or click **Import Script File** to import a script.



If you import a script that uses Unicode (UTF8) encoding and your computer's regional language settings are set to English, and then you export the script and attempt to execute it, you may encounter errors because Unicode (UTF8) encoding may add a "." or other special character at the beginning of the script. If this occurs, simply edit the script to remove the extraneous characters.

- 9 In the Description field, enter text that describes the purpose or contents of the script.
- 10 To save the changes, select **Save** from the **File** menu.

Creating a Script from the By Folder View in the Library

To create a script perform the following steps:

- 1 From the Navigation pane, select **Library ► By Folder**. The folder hierarchy in the Library appears in the Content pane.
- 2 Select the folder that should contain the script.
- 3 From the **Actions** menu, select **New ► Script**. The Script window appears.
- 4 In the Name field, enter the name of the script.
- 5 Select the script type from the Type drop-down list.
- 6 Click **Select** to change the location for the script in the folder hierarchy. The Select Folder window appears. Select a folder in the Library to specify the location of the script and then click **Select**.
- 7 In the Changes Server field, select Yes, if the script causes a change in the server configuration when executed.
- 8 In the Run as Super User field, select yes if the script can be run as a Super user when executed. OGFS Scripts can only be executed as an SA User.

This option is enabled only if you have to appropriate permission. See the *SA Administration Guide* for more information about script execution permissions.
- 9 In the Script Contents field, enter the contents of the script or click **Import Script File** to import a script. In the Open window, select the script to import and then click **Open**.



If you import a script that uses Unicode (UTF8) encoding and your computer's regional language settings are set to English, and then you export the script and attempt to execute it, you may encounter errors because Unicode (UTF8) encoding may add a "." or other special character at the beginning of the script. If this occurs, simply edit the script to remove the extraneous characters.

- 10 In the Description field, enter text that describes the purpose or contents of the script.
 - 11 To save the changes, select **Save** from the **File** menu.
-




The Library in the SA Client contains a Home directory and each user has a folder in the Home directory. You can save private scripts in this folder and later execute the script on managed servers.

Opening a Script in the SA Client

In the SA Client, there are several ways to open a script. You can open a script from:

- The Search option in the Navigation pane
- The By Type view in the Library
- The By Folder view in the Library
- The Device list in the Navigation pane

Opening a Script from Search

- 1 From the Navigation pane, select **Search**.
- 2 Select Server Script or OGFS Script from the drop-down list and then enter the name of the script in the text field.
- 3 Select . The search results appear in the Content pane.
- 4 From the Content pane, select the script and then select **Open** from the **Actions** menu. The Script window appears.

Opening a Script from the By Type view in the Library

- 1 From the Navigation pane, select **Library** ► **By Type** ► **Scripts**. The scripts appear in the Content pane.
- 2 From the Content pane, select the script and then select **Open** from the **Actions** menu. The Script window appears.

Opening a Script from the By Folder view in the Library

- 1 From the Navigation pane, select **Library** ► **By Folder**. The folder hierarchy in the Library appears in the Content pane.
- 2 From the Content pane, select the script in a folder and then select **Open** from the **Actions** menu. The Script window appears.

Opening a Script from Devices

- 1 From the Navigation pane, select **Devices ► Servers ► All Managed Servers**. The server list appears in the Content pane.
Or
From the Navigation pane, select **Devices ► Device Groups**. The device groups list appears in the Content pane.
- 2 From the Content pane, select a server and then from the **Actions** menu, select **Open**. The Server Explorer window opens.
- 3 From the Views pane, select **Management Policies ► Software Policies**. The software policies attached to the server appear in the Content pane.
- 4 From the Content pane, select the software policy and then select **Open** from the **Actions** menu. The Software Policy window appears.
- 5 From the Views pane, select Policy Items. The policy items appear in the Content pane.
- 6 From the Content pane, select the script and then select **Open** from the **Actions** menu. The Script window appears.

Editing Script Properties

After you create a script, you can view and modify its properties. You can view properties such as the SA user who created the script, the date when it was created, and the Object ID of the script. You can also modify the name, description, contents, the Library folder location of the script and the script options.

To view and edit script properties, perform the following steps:

- 1 Open a script in the SA Client. See [Opening a Script in the SA Client](#) on page 429 for ways to open a script. The Script window appears.
- 2 In the Name field, edit the name of the script.
- 3 Click **Select** to change the location for the script in the folder hierarchy. The Select Folder window appears. Select a folder in the Library to specify the location of the script and then click **Select**.
- 4 In the Changes Server field, select Yes, if the script causes a change in the server configuration when executed.
- 5 In the Run as Super User field, select yes if the script can be run as a Super User when executed. Selecting yes, allows you to run the script as a Super User without providing a password for the script.

This option is enabled only if you have the appropriate permission. See the *SA Administration Guide* for more information about script execution permissions.

- 6 In the Script Contents field, edit the contents of the script or click **Import Script File** to import another script. In the Open window, select the script to import and then click **Open**.



If you import a script that uses Unicode (UTF8) encoding and your computer's regional language settings are set to English, and then you export the script and attempt to execute it, you may encounter errors because Unicode (UTF8) encoding may add a "." or other special character at the beginning of the script. If this occurs, simply edit the script to remove the extraneous characters.

- 7 In the Description field, edit the text that describes the purpose or contents of the script.
- 8 To save the changes, select **Save** from the **File** menu.

Viewing All the Software Policies Associated with a Script

In the SA Client, Server scripts can be added to a software policy. In the Scripts window, you can view all the software policies that contain the selected Server script. You cannot add OGFS script to a software policy.

To view the policy usage for a script, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Scripts**.
- 2 From the Content pane, select the script and open it. The Scripts window appears.
- 3 From the Views pane, select Policy Usage. The list of software policies associated with the scripts appears in the Content pane.

Viewing Script Version History

To view the version history of a script perform the following steps:



- 1 From the Navigation pane, select **Library ► By Type ► Scripts**.
- 2 From the Content pane, select the script and open it. The Scripts window appears.
- 3 From the Views pane, select Version History. The events associated with the script will display in the Content pane. You can view the script content from different versions of a script. The  indicates the current version of the script. You can view the script content from different versions of a script. See the *SA Users Guide: Server Automation* for more information on server history.
- 4 To make any of the previous version of script current, select the script version and from the **Actions menu**, select **Set as Current Version** as shown in figure [Figure 87](#).

Figure 87 Script Version History



Locating Scripts in Folders

To locate a script in the folder hierarchy, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Scripts**.
- 2 From the Content pane, select the script and then select **Locate in Folders** from the **Actions** menu. The folder hierarchy for the script appears in the Content pane.

Exporting a Script

To download a script, perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Scripts**. The scripts appear in the Content pane.
Or
From the Navigation pane, select **Library ► By Folder** and then select the folder which contains the script.
- 2 From the Content pane, select a script to export.
- 3 From the **Actions** menu, select **Export Script**. The Export Software window appears.
- 4 In the Browse window, specify the location for the script to be exported to.
- 5 Click **Export**.

Renaming a Script

To rename a script perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Scripts**.
- 2 From the Content pane select the script, and then from the **Actions** menu select **Rename**.
- 3 Enter the new name for the script in the Content pane.

Deleting a Script

To delete a script perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Scripts**.
- 2 From the Content pane select the script, and then from the **Actions** menu select **Delete**. The Confirmation window appears.
- 3 Click **Delete** to delete the script.

Executing Scripts

In the SA Client, you can execute scripts in the following ways:

- Execute a server script directly on servers or server groups and execute scripts in the Global Shell. See [Running a Server Script \(Saved Script or Ad-Hoc Script\)](#) on page 434 and [Running an OGFS Script](#) on page 438 for more information.
- Add a script to a software policy and execute the script by attaching the software policy to the server and then remediating the server against the software policy. See [Install Software Using a Software Policy](#) on page 267 for more information.

A software policy allows you to execute multiple scripts on a servers or server groups simultaneously, and execute a sequence of scripts on a server by specifying an install order in the software policy. See the *SA Policy Setter's Guide* for information about software policy.



You must have a set of permissions to execute a script. To obtain these permissions, contact your SA Administrator. See the *SA Administration Guide* for more information. For security purposes, several permission-based scenarios can be experienced to run or copy scripts in folders, run super user scripts, run non-super user scripts, etc.

Ways to Open the Run Script Window

The Run Script window allows to you execute a script on managed servers. In the SA Client you can launch the Run Script window in the following ways:

- [From the Device List](#)
- [From the Device Explorer](#)
- [From the Library](#)

From the Device List

- 1 From the Navigation pane, select **Devices ► Servers ► All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices ► Device Groups**. The device group list appears in the Content pane.

- 2 From the Content pane, select a server or device group.
- 3 From the **Actions** menu, select **Run Script**. The Run Script window appears.

From the Device Explorer

- 1 From the Navigation pane, select **Devices ► Servers ► All Managed Servers**. The server list appears in the Content pane.
- 2 From the Content pane, select a server.
- 3 From the Action menu, select **Open**. The Device Explorer appears.
- 4 From the **Actions** menu, select **Run Script**. The Run Script window appears.

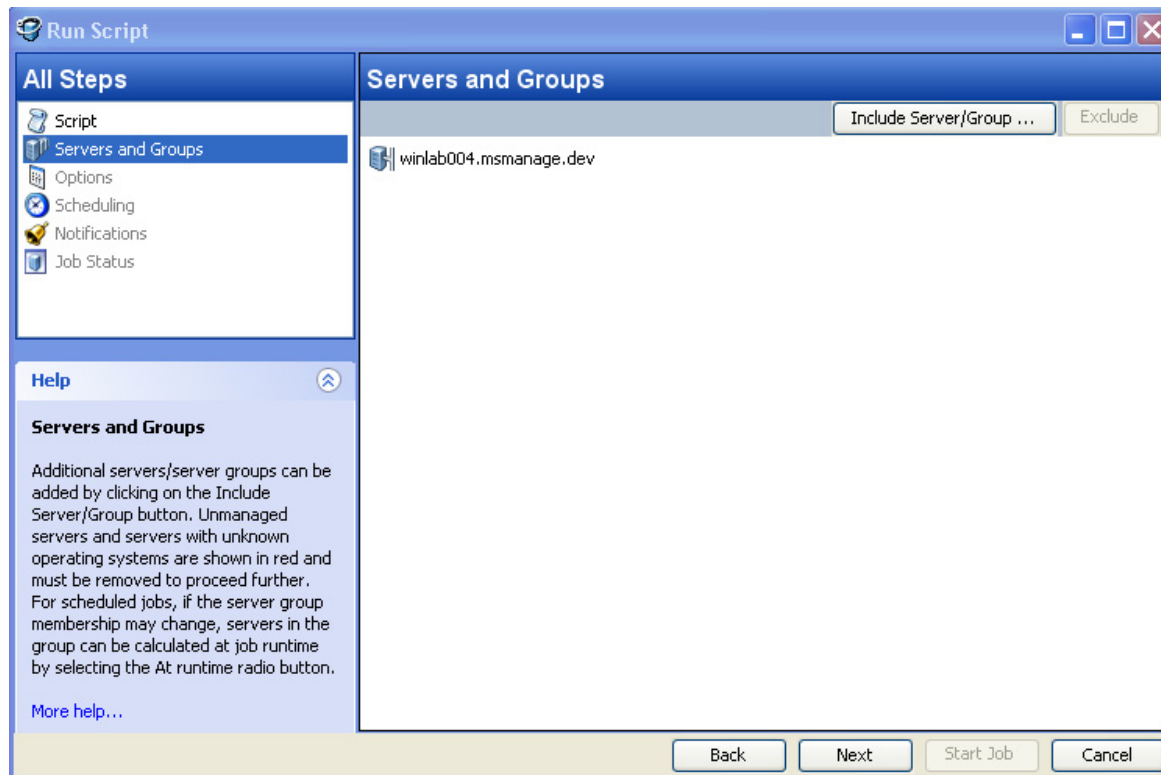
From the Library

- 1 From the Navigation pane, select **Library** ► **By Type** ► **Scripts**. The scripts list appears in the Content pane.
- 2 From the Content pane, select a script.
- 3 From the **Actions** menu, select **Run**. The Run Script window appears.

Running a Server Script (Saved Script or Ad-Hoc Script)

The Run Script, as shown in [Figure 88](#), allows you to run a script on managed servers and consists of the following steps:

Figure 88 Run Server Script Window



- Servers and Groups
- Script
- Options
- Scheduling
- Notification
- Job Status

See [Ways to Open the Run Script Window](#) on page 433 on how to access the Run Script window. If you access the Run Script window from the Device list or Device Explorer, the first step in the window is Script. If you access the Run Script window from the Library, the first step in the Run Script window is Servers and Groups.

Servers and Groups

This step allows you to specify the servers or server groups for executing the script. In this step, you can add and remove servers or server groups from the list.

If you choose the option **Now**, then the membership is determined based on the time when you made the selection. As a result the script is executed on the servers that were in the group when you selected the option. Changes to the group membership does not affect the list of servers on which the script will be executed.

If you choose the option **Runtime**, then the membership is determined when the script execution job is run. The script is executed on the servers present in the server group when the job is run. Changes to group membership is reflected in the list of servers when is script is executed.



To be able to select the **Runtime** option, the “Allow Run Refresh Jobs” permission is required. See the *SA Administration Guide* for more information on permissions.

To select servers and groups perform the following steps :

- 1 Open the Run Script window from one of the methods described in [Ways to Open the Run Script Window](#) on page 433.
- 2 In the Run Script window, select the step Servers and Groups.
- 3 (Optional) Click **Include Server/Group** to add servers or server groups to the list or select a server or server group and click **Exclude** to remove servers from the list.
- 4 For a server group, in the Server Group Calculation field, select the option **Now** to execute the script on the servers that were in the group when you made the selection. Select the option **Runtime** to execute the script on the servers when the job is run.
- 5 Click **Next** to proceed to the Script step.

Script

This step allows to select a saved script or define an ad-hoc script to be executed on managed servers. See [Types of Scripts](#) on page 424 for information on the script types.

Saved Script

To select a saved script perform the following steps:

- 1 To select a saved script, select the option **Select Saved Script**.
- 2 From the Name drop-down list select the script or click **Select Script** to open the Select Script window. Select the script from the Select Script window.
- 3 The script properties such as version, type, location are displayed in the content pane. To view the contents on the script, click **View Script**. The contents of the script are displayed in the Run Script window.
- 4 Click **Next** to proceed to the Options step.

Ad-Hoc Script

To define an ad-hoc script perform the following steps:

- 1 To select an ad-hoc script, select the option **Define Ad-hoc Script**.

- 2 (Windows only) From the Type drop-down list, select the script type.
- 3 Enter the contents of the script in the Script Contents field or click **Import Script File** to import a script.



If you import a script that uses Unicode (UTF8) encoding and your computer's regional language settings are set to English, and then you export the script and attempt to execute it, you may encounter errors because Unicode (UTF8) encoding may add a "." or other special character at the beginning of the script. If this occurs, simply edit the script to remove the extraneous characters.

- 4 Click **Next** to proceed to the Options step.

Options

This step allows you to specify the runtime options and output options for executing a script. In this step you can specify whether to execute the script as root or Local System or as a specified user. You can also specify the script time-out value, any additional parameters for executing the script, and the output options for the script.

To specify the runtime and output options for a script perform the following steps:

- 1 In the Runtime User field select root (for Unix) or Local System (for Windows) to execute the script as root or local system. To execute the script as root or local system, you require the appropriate permissions. See the *SA Administration Guide* for information about the permissions required for executing scripts.
Or
 - a Select Name and enter user name and password to execute the script as a specified user. To execute the script simultaneously across multiple servers or server groups, you must use the same user name and password across all the servers.
 - b (Windows only) Enter the domain name in the domain field.
- 2 In the Script timeout field enter the script timeout value in minutes. The time out value is the amount of time required for a script to complete execution activities on a server. If the script is not executed when the timeout value is reached, then the script is stopped by SA and a script error occurs. Select a timeout value greater than the time required for execution to complete.
- 3 In the Specify any needed parameters for this script execution field, enter any parameters if required.
- 4 In the Output Options, select Discard all script output to discard script output or else select Retain script output.
- 5 Select the output size of the script from the Size of the output to retain drop-down list.
- 6 Click **Next** to proceed to the Scheduling step.

Scheduling

This step allows you to schedule the script execution job. You can choose to run the script execution job immediately, or on a specified date and time, or on a recurring basis.

To schedule a script execution job, perform the following steps:

- 1 In the Schedule Frequency section, choose to run the script once, daily, weekly, monthly, or on a custom schedule. Select any one of the following options:

- **Once:** Choose this option to run the job immediately or only once at a specified date and time.
 - **Daily:** Choose this option to run the job on a daily basis at a specified time.
 - **Weekly:** Choose this option to specify the day or days of the week to run the job.
 - **Monthly:** Choose this option to specify the months to run the job, and the days of the month.
 - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values.
- 2 In the Time and Duration section, for each type of schedule, specify the start time for the job. You must also specify the start date and end date for the job. The Time Zone is set according to the time zone set in your user profile.
 - 3 Click **Next** to proceed to the Notifications step.

Notifications

This step allows you to set email notifications to alert users on the success or failure of a job. You can also associate a Ticket ID with the job. This setting is optional.

To set email notifications, perform the following steps:

- 1 Click **Add Notifier**.
- 2 Enter the addresses in the Email Address of Recipient field.
- 3 To send email to the address if the job succeeds, select the checkbox On Success.
- 4 To send email if the job fails, select the checkbox On Failure.
- 5 Enter an ID to be associated with this job in the Ticket ID field.
- 6 Click **Next** to proceed to the Job Status step.

Job Status

This step allows you to start the job, view the job progress, the job results, the script output for a managed server, and export the script output from all the servers.

SA supports the following file formats for exporting script output results:

- A Zip file with folders for each managed server
- A Zip file containing no folders
- Consolidated raw text file
- Consolidated formatted text file
- Consolidated CSV file

You can also view jobs in the Jobs Log window of the SA Client. See the *SA Users Guide: Server Automation* for information about Job Logs.

To start a job, perform the following steps:

- 1 To start the job, click **Start Job**.

If you selected Immediately in the Scheduling step, the job will begin now. If you scheduled the job for a later time, the job will run later. You can then view the job in the Jobs Log window of the SA Client.

- 2 The job's progress information appears in the Job Status window. You can view the server on which the script was executed, the job status, and the exit code. If the exit code is zero, then it indicates that the script is executed successfully. If the exit code is non-zero, then it indicates an error during script execution.

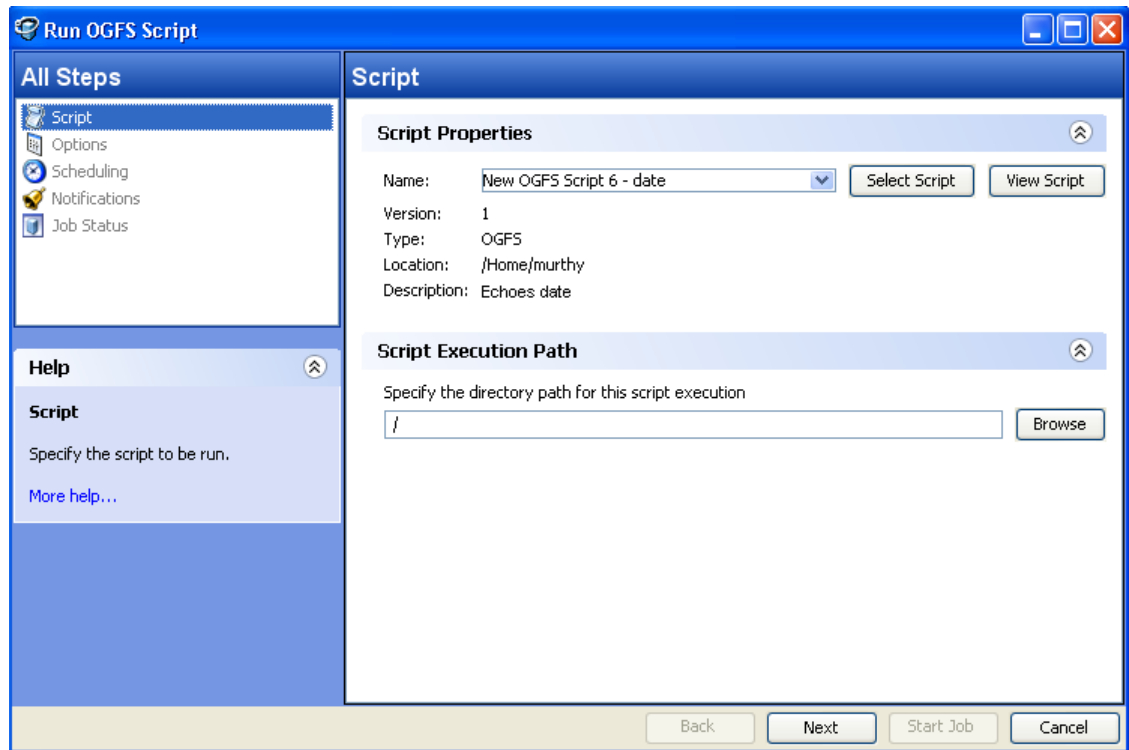
If the job status is displayed as Pending Approval, then the job is blocked until it is approved by a process that is external to SA. See the *SA Users Guide: Server Automation* for information about job status.

- 3 (Optional) To view the script output from a managed server, select the managed server and script output appears below the table.
- 4 (Optional) To view the script output from all the managed servers, select the option Show output in table. The output for each server appears in the Output column in the table.
- 5 (Optional) To view the output for all the servers in separate columns, select the option Show output in table and enter the delimiter character in the Delimiter checkbox. The output for each server appears in separate columns in the table.
- 6 (Optional) To export the script output results, click **Export All Results**. In the Browse window specify the location and the file type and click **Export**.
- 7 Click **Close** to exit the Run Script window.

Running an OGFS Script

The Run OGFS Script, as shown in [Figure 89](#), allows you to run an OGFS script and consists of the following steps:

Figure 89 Run OGFS Script Window



- Script
- Options

- Scheduling
- Notification
- Job Status

Script

This step allows to specify an OGFS script for execution.

To select an OGFS script perform the following steps:

- 1 From the Navigation pane, select **Library ► By Type ► Scripts**. The scripts list appears in the Content pane.
Or
From the Navigation pane, select **Library ► By Folder**. The folder hierarchy in the Library appears in the Content pane.
- 2 From the Content pane, select an OGFS script.
- 3 From the **Actions** menu, select **Run**. The Run OGFS Script window appears.
- 4 In the Script Properties section, select script from the Name drop-down list or click **Select Script** to open the Select Script window. Select the script from the Select Script window.
- 5 The script properties such as version, type, location, description are displayed in the content pane. To view the contents on the script, click **View Script**. The contents of the script are displayed in the Run OGFS Script window.
- 6 In the Script Execution Path section, enter the OGFS directory path for executing the script or click **Browse** to specify the directory path in the OGFS.
- 7 Click **Next** to proceed to the Options step.

Options

This step allows you to specify the runtime options and output options for executing a script. In this step you can specify the script time-out value, any additional parameters for executing the script, and the output options for the script.

To specify the runtime and output options for a script perform the following steps:

- 1 In the Script timeout field enter the script timeout value in minutes. The time out value is the amount of time required for a script to complete execution activities. If the script is not executed when the timeout value is reached, then the script is stopped by SA and a script error occurs. Select a timeout value greater than the time required for execution to complete.
- 2 In the Specify any needed parameters for this script execution field, enter any parameters if required.
- 3 In the Output Options, select Discard all script output to discard script output or else select Retain script output.
- 4 Select the output size of the script from the Size of the output to retain drop-down list.
- 5 Click **Next** to proceed to the Scheduling step.

Scheduling

This step allows you to schedule the script execution job. You can choose to run the script execution job immediately, or on a specified date and time, or on a recurring basis.

To schedule a script execution job, perform the following steps:

- 1 In the Schedule Frequency section, choose to run the script once, daily, weekly, monthly, or on a custom schedule. Select any one of the following options:
 - **Once:** Choose this option to run the job immediately or only once at a specified date and time.
 - **Daily:** Choose this option to run the job on a daily basis at a specified time.
 - **Weekly:** Choose this option to specify the day or days of the week to run the job.
 - **Monthly:** Choose this option to specify the months to run the job, and the days of the month.
 - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values.
- 2 In the Time and Duration section, for each type of schedule, specify the start time for the job. You must also specify the start date and end date for the job. The Time Zone is set according to the time zone set in your user profile.
- 3 Click **Next** to proceed to the Notifications step.

Notifications

This step allows you to set email notifications to alert users on the success or failure of a job. You can also associate a Ticket ID with the job. This setting is optional.

To set email notifications, perform the following steps:

- 1 Click **Add Notifier**.
- 2 Enter the addresses in the Email Address of Recipient field.
- 3 To send email to the address if the job succeeds, select the checkbox On Success.
- 4 To send email if the job fails, select the checkbox On Failure.
- 5 Enter an ID to be associated with this job in the Ticket ID field.
- 6 Click **Next** to proceed to the Job Status step.

Job Status

This step allows you to start the job, view the job progress, view the job results, view the script output for a managed server, and export the script output from all the servers.

SA supports the following file formats for exporting script output results:

- A Zip file with folders for each managed server
- A Zip file containing no folders
- Consolidated raw text file
- Consolidated formatted text file
- Consolidated CSV file

You can also view jobs in the Jobs Log window of the SA Client. See the *SA Users Guide: Server Automation* for information about Job Logs.

To start a job, perform the following steps:

- 1 To start the job, click **Start Job**.
If you selected Immediately in the Scheduling step, the job will begin now. If you scheduled the job for a later time, the job will run later. You can then view the job in the Jobs Log window of the SA Client.
- 2 The job's progress information appears in the Job Status window. You can view the server on which the script was executed, the job status, and the exit code. If the job status is displayed as Pending Approval, then the job is blocked until it is approved by a process that is external to SA. See the *SA Users Guide: Server Automation* for information about Job Logs.
- 3 (Optional) To view the script output from all the managed servers, select the option Show output in table. The output for each server appears in the Output column in the table.
- 4 (Optional) To view the output for all the servers in separate columns, select the option Show output in table and enter the delimiter character in the Delimiter checkbox. The output for each server appears in separate columns in the table.
- 5 (Optional) To export the script output results, click **Export All Results**. In the Browse window specify the location and the file type and click **Export**.
- 6 Click **Close** to exit the Run OGFS Script window.

10 Running SA Extensions

Server Automation (SA) gives you the capability to extend its functionality by creating **Automation Program Extensions (APXs)**. This section describes the APX extensions feature and how to run extensions.



For information on how to create APX extensions, see “Extending SA with Automation Platform Extensions (APXs)” in the *HP SA Platform Developer’s Guide*.

APX extensions provide a framework that allows anyone familiar with script-based programming tools such as shell scripts, Python, Perl, and PHP, to extend the functionality of SA and create applications that are tightly integrated into SA. SA provides two types of APX extensions:

- **Program APX Extensions** run in the Global File System (OGFS) and can use all of the OGFS functionality. You can use typical programming practices to leverage the SA API and access a core’s Managed Servers to implement new custom functionality. For example, you could write an APX extension that gathers BIOS information from managed servers and populates custom fields using shell commands.
- **Web APX Extensions** allow you to create a web-based application, where either an Apache 2.x process or a CGI/PHP script is called using GET or POST URL. Web APX extensions can contain static web resources such as images, and can employ CGI or PHP for dynamic content generation.

APX extensions allow you to access data about your managed environment and share and process that data with web applications, scripts, programs and other custom applications.

Methods of Running Extensions

You can run extensions in any of the following ways.

- **From Managed Servers, by selecting a server first:** You can select one or more servers in the SA Client, then right click or select the **Actions** menu and select the **Run Extension** menu item. This applies only to certain program extensions. For details, see [Run Extensions on Managed Servers](#) on page 444.
- **From the SA Library, By Type:** In the SA Client you can run an extension by selecting Library, By Type, Extensions, Web or Program, selecting an extension, and doing one of the following.
 - Double click the extension.
 - Right click the extension and select **Run...**
 - Under the **Actions** menu, select **Run...**

- **From the SA Library, By Folder:** In the SA Client you can run an extension by selecting Library, By Folder, navigating to and opening an extension in the Library. Under the **Actions** menu, select **Run....**
- **From the Global Shell (program extensions only):** From the Global Shell, run `/opsw/apx/bin/<extension name>` where `<extension name>` is the unique name of the extension. Provide any parameters the extension requires. For more information on the Global Shell, see the *SA Users Guide: Server Automation*.
- **From a web browser (web extensions only):** From a web browser, enter the URL `https://<SA core>/webapp/<extension name>` where `<SA core>` is the IP address or host name of your SA core and `<extension name>` is the unique name of the web extension.
- **From the SA API (program extensions only):** From the SA API, use the method `ProgramAPXService.startProgramAPX()`. For more information on the SA API, see the *HP SA Platform Developer's Guide*.



Most extensions cannot run on VMware ESXi hypervisor servers because SA does not install an Agent on ESXi servers. Instead, SA manages ESXi servers remotely using a web services interface. For more information, see [Chapter 7, Virtual Server Management](#), on page 9 of this guide.

Run Extensions on Managed Servers

Some extensions take one or more managed servers as input and perform some operation or gather some information from those managed servers. For example, an extension could gather certain information about devices on the servers using a script. You can run this extension by selecting one or more servers and selecting the extension to run. This extension describes how to run extensions that take one or more managed servers as input to the extension.



Only program extensions that implement the “com.hp.client.server.RightClickToRun” interface can be run using this method. This interface indicates that the extension takes one or more servers as input parameters.

Only extensions you have permission to execute will be shown in the SA Client. For more information on permissions, see the *SA Administration Guide*.

For details on creating extensions, the RightClickToRun interface and permissions, see “Extending SA with Automation Platform Extensions (APXs)” in the *HP SA Platform Developer's Guide*.

To run an SA program extension by first selecting one or more servers, perform the following steps.

- 1 From the SA Client Navigation pane, select **Devices ► All Managed Servers**.
- 2 Select one or more servers in the Contents pane.
- 3 Right click the server or select the **Actions** menu, then select **Run Extension ► Select Extension....**

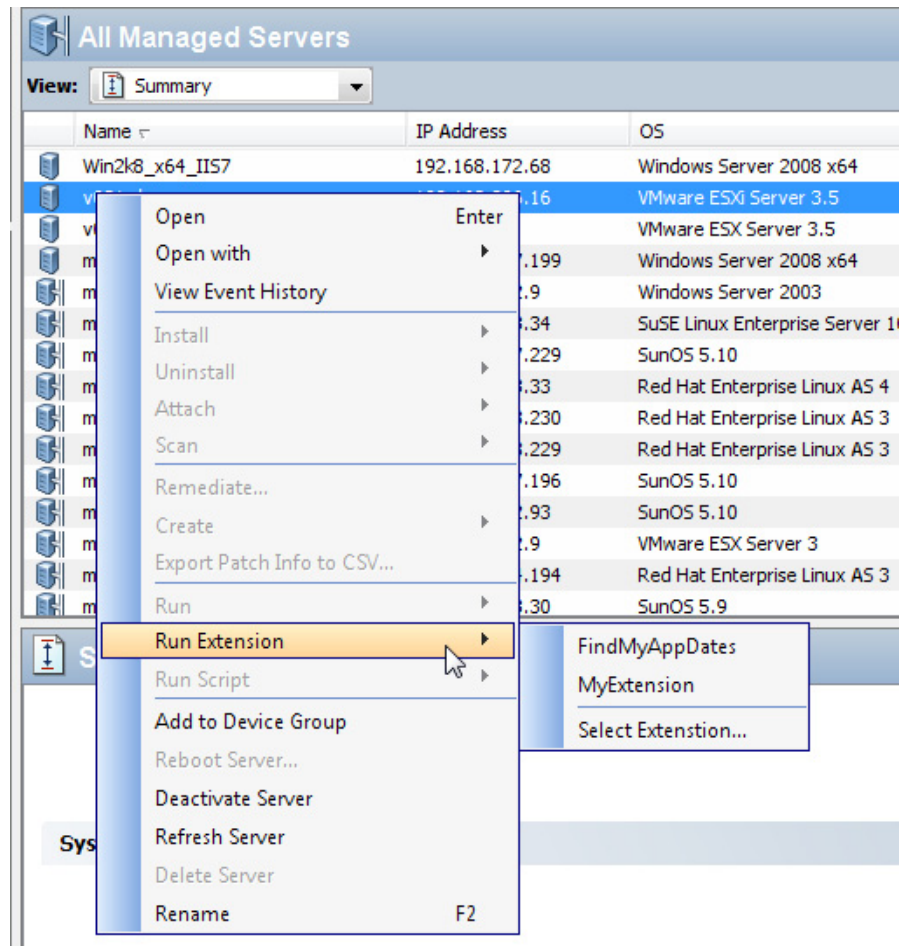
Or

Right click the server or select the **Actions** menu, then select **Run Extension ►** and select a named extension, if any are shown in the menu.

The **Run Extension** menu lists all of the program extensions that have been run at least once before. You can select one of these extensions without having to use the **Select Extension...** window.

For example, the following shows the **Run Extension** menu item selected, two sample program extensions named **MyExtension** and **FindMyAppDates**, and the **Select Extension...** menu item.

Figure 90 The Run Extension Menu




➤ Extensions appear in the **Run Extension** menu only after they have been run at least once using the **Select Extension...** menu. That is, to make your extension appear in the **Run Extension** menu, you must first run your extension by choosing **Select Extension...** and then choosing your extension as described in step 4 below.

- 4 If you selected a particular extension from the **Run Extension** menu in [step 3](#) on page 444, skip ahead to [step 5](#) below.

If you chose the **Select Extension....** menu item, the SA Client displays the available extensions. The following screen shows three extensions MyExtension, FindMyAppDates and Extensible Discovery. Select the extension you want to run and click OK.

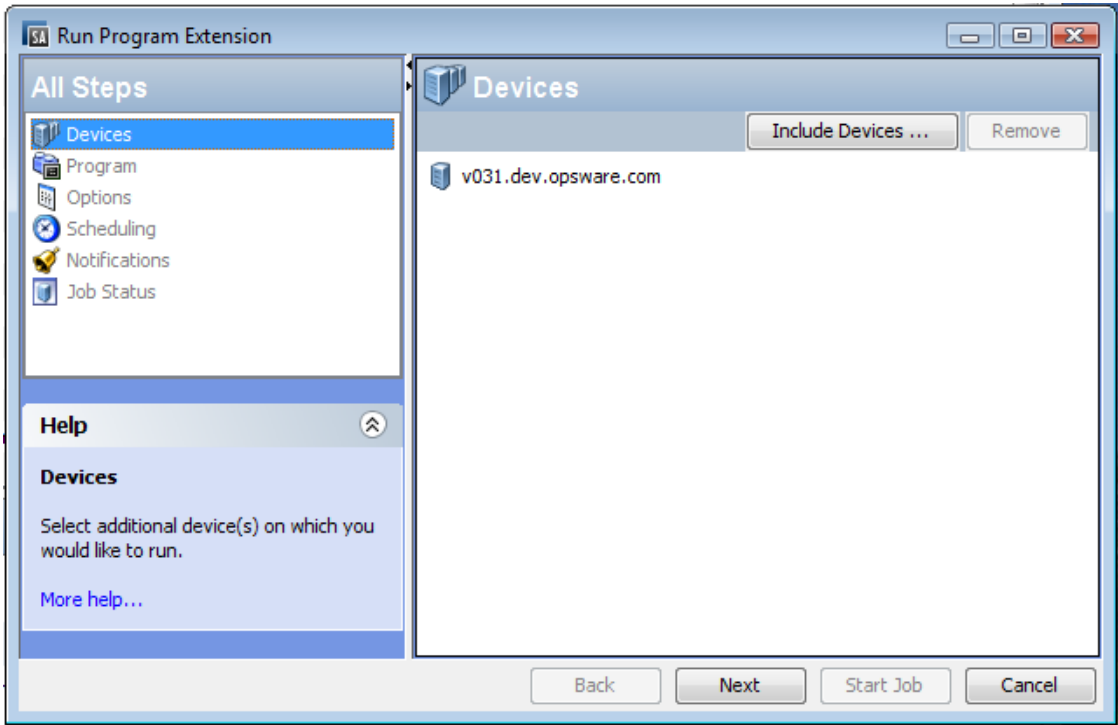
Figure 91 The Select Extension Window



- 5 Once you have chosen a particular extension to run, the SA Client displays the Run Program Extension screen as shown below. This screen shows the devices the extension will run against.

If you want to run the extension against more devices, select the Include Devices... button and select additional devices or device groups. To remove devices, select the devices and click Remove.

Figure 92 Select Devices on the Run Program Extension Window



- 6 Click the Next button. This displays the Program Properties and Program Execution Path. Optionally edit the program execution path and click Next to display the Runtime Options and Output Options.
- 7 Optionally change the Runtime Options and Output Options and click Next to display the Schedule Frequency and Time and Duration.

Or you can click Start Job to skip the remaining options and run the extension with default options.
- 8 Optionally specify how frequently to run the extension and click Next to display Email Notifications and Ticket Tracking.
- 9 Optionally change the email notifications and add a ticket ID and click Next to display the Job Status screen.
- 10 Click Start Job to run the extension. The Run Program Extension screen displays the progress of the extension. When the extension finishes, click the Close button.

Running Extensible Discovery on Managed Servers

SA can gather a large amount of information from your servers by default. Extensible Discovery lets you gather additional information about your servers quickly and easily. Extensible Discovery is a feature that you can customize and use to discover and obtain custom information about servers in your managed environment.

This section describes how to run Extensible Discovery. To customize Extensible Discovery and add your own scripts, see [Adding Scripts to Extensible Discovery](#) on page 448.



To run Extensible Discovery, you must have write access to the managed servers where you want to run Extensible Discovery and you must have execute access to the Extensible Discovery folder in the SA Library. For more information on permissions, see the *SA Administration Guide*.



You cannot run Extensible Discovery on VMware ESXi hypervisor servers because SA does not install an Agent on ESXi servers. Instead, SA manages ESXi servers remotely using a web services interface. For more information, see [Chapter 7, Virtual Server Management](#), on page 9 of this guide.

To run Extensible Discovery, perform the following steps.

- 1 Run the Extensible Discovery extension as described in [Run Extensions on Managed Servers](#) on page 444. This does the following:
 - a Remediates the software policies “Customer Provided Scripts” and “HP Provided Scripts” on the selected servers. This copies all the scripts in these policies to all the managed servers. It places the scripts on the managed servers.
 - b Runs all the scripts in the Extensible Discovery directory on the selected servers.
 - c Places the output from the scripts either in custom attributes or custom fields.



The first time you run Extensible Discovery on a server will typically take longer than subsequent runs because it needs to install scripts on the server the first time you run it.

- 2 Examine the results by viewing the Custom Attributes or Custom Fields for each server.
If an error occurs, check the following custom attributes for more information.
 - **HPSW_ED_error**: If an error occurs on a server, this custom attribute will be created for the server and it will contain any relevant error messages. To determine which servers had an error, you can search for servers with this custom attribute.
If Extensible Discovery runs without errors on the server later, this custom attribute will be removed.
 - **HPSW_ED_warning**: This custom attribute will be created on servers where a non-fatal warning occurs.
If Extensible Discovery runs without warnings on the server later, this custom attribute will be removed.For more information, see [Comparing Custom Fields and Custom Attributes](#) on page 455.

Running Extensible Discovery from the OGS

You can also run Extensible Discovery from the Global Shell (OGS) using the command interface as follows:

```
/opsw/apx/bin/com/opsware/extensible_discovery -d <server ids> -g <group ids>
```

The following table describes the options to this command. For more information on the Global Shell, see the *SA Users Guide: Server Automation*.

Table 30 Options to the extensible_discovery Command

Option	Usage
-h --help	Displays help for this command.
-d <server ids> --deviceids=<server ids>	Specifies one or more server IDs or server names separated by commas. Runs Extensible Discovery on the specified servers.
-g <group ids> --groupids=<group ids>	Specified one or more device group IDs or device group names separated by commas. Runs Extensible Discovery on all the servers in the specified groups.

To find a server ID or a device group ID in the SA Client, display the server or device group and locate the ID in the Object ID column. The Object ID is an integer value.

Adding Scripts to Extensible Discovery

SA can gather a large amount of information from your servers by default. Extensible Discovery lets you gather additional information about your servers quickly and easily. Extensible Discovery is a feature that you can customize and use to discover and obtain custom information about servers in your managed environment. With Extensible Discovery you can:

- Write scripts that gather custom information from your servers and easily incorporate your scripts into Extensible Discovery.

- Schedule the execution of your scripts.
- Search and generate reports on the resulting custom information.
- Export the resulting information to other tools for further analysis and decision support.
- Automatically update all servers when your scripts change.

Scripts Provided with Extensible Discovery

The following scripts are included with Extensible Discovery. These scripts are automatically executed on the selected managed servers when you run Extensible Discovery.

Table 31 Scripts Included with Extensible Discovery

Script Name	Operating System	Description
get_oslevel.sh	AIX	Returns the operating system level. (Uses the AIX command <code>oslevel -s</code> .)
get_install_date.sh	HP-UX	Returns the date the operating system was installed. The date value is formatted to be placed in a custom attribute or a string type custom field, but not a date type custom field.
get_firmware_version.sh	Linux	Returns the BIOS version.
get_firmware_version.sh	Solaris	Returns the EEPROM (or OBP) version when run on Sparc servers. Otherwise returns the BIOS version.
get_firmware_version.vb	Windows	Returns the BIOS version.

These scripts will create custom attributes with the name `HPSW_ED_firmware_version` or `HPWS_ED_oslevel`. If you want the return values to be placed in custom fields, you must create custom fields named `HPSW_ED_firmware_version` and `HPWS_ED_oslevel` before running the extension, and the return values will be stored in the custom fields you created. If the custom fields don't exist, Extensible Discovery will create custom attributes on each server and place the data in the custom attributes.

These scripts are in the SA Library under `/Opware/Tools/Extensible Discovery/HP Provided Components`. To view these scripts, in the SA Client select **Library, By Folder** and navigate to the **HP Provided Scripts** folder. Select a .zip file and select **Actions ► Export Software**.

If you do not want these scripts to execute on managed servers, remove them from the “HP Provided Scripts” software policy as described in [Software Policies Provided with Extensible Discovery](#) below.

Software Policies Provided with Extensible Discovery

Extensible Discovery uses two software policies that contain the scripts it runs to gather information from managed servers. Extensible Discovery automatically remediates these policies on managed servers and runs the scripts in them.

- **HP Provided Scripts:** This software policy contains the HP-provided scripts listed in [Table 31, "Scripts Included with Extensible Discovery"](#). It is located in the SA Library under `/Opware/Tools/Extensible Discovery/HP Provided Components`.
- **Customer Provided Scripts:** This software policy is initially empty. You can place your custom scripts in this software policy as described below and Extensible Discovery will use your scripts. It is located in the SA Library under `/Opware/Tools/Extensible Discovery/Customer Provided Components`.

When you run Extensible Discovery, it remediates these policies on the selected managed servers and runs all the scripts in these policies. If you don't want any of the scripts in these policies to run, remove them from the policy.

Extensible Discovery copies the scripts in these policies to the following directories on the managed servers:

- `/var/opt/opware/extensible_discovery/scripts/` on UNIX systems.
- `%SYSTEMDRIVE%\Program Files\Common Files\Opware\extensible_discovery\scripts\` on Windows systems.

Extensible Discovery runs all the scripts in these directories regardless of how the scripts got there.

You can use these software policies or create your own software policies, but if you create your own, you must remediate them on managed servers where you want to run Extensible Discovery. For more information on software policies, see [Chapter 5, Software Management](#), on page 257 of this guide.

Writing Your Own Scripts for Extensible Discovery

This section describes requirements and guidelines for writing scripts for use with Extensible Discovery. For instructions on how to add your scripts to Extensible Discovery, see [Adding Your Own Scripts to Extensible Discovery](#) on page 451.

- HP provides several sample scripts you can use in the software policy "HP Provided Scripts". You can view these scripts in `/Library/Opware/Tools/Extensible Discovery/HP Provided Components`. For more information, see [Scripts Provided with Extensible Discovery](#) on page 449.
- A best practice is to store your script in a version control system.
- Your script needs to handle error situations within your script and return an accurate exit status. The following explains various error situations and how Extensible Discovery handles them.
 - The exit status of a shell script is the exit status of the last command in the script. A script that returns zero status is considered successful.
 - A script could return zero status but also send some output to `stderr`. This case is treated as a success. The output from `stderr` will be treated the same as `stdout`.
 - A script that returns zero status and outputs nothing to `stdout` or `stderr` will be considered a success and the blank value will be written to the appropriate custom attribute or custom field. You could use this method to set the custom attribute or custom field to "".
 - A script that returns a non-zero value is treated as an error. A message including the script's `stdout` and `stderr` will be stored in the `HPSW_ED_error` custom attribute.

- A script that is not executable or has syntax errors is treated as an error. A message including the script's stdout and stderr will be stored in the HPSW_ED_error custom attribute.
- If a script determines that it has nothing to collect on a server (for example, if the `get_eeprom_version.sh` is run on a Solaris x86 server, where there is no EEPROM), it should return an exit status of 3, which will be interpreted by Extensible Discovery as not applicable, and will return nothing for that particular data item.
- For UNIX shell scripts, if you want the script to fail when any individual command in the script fails, start the script with the line `#!/bin/sh -e`. For more information, see the documentation for the UNIX shell.

Adding Your Own Scripts to Extensible Discovery

To customize extensible discovery, you need to write one or more scripts that gather the custom data you want, then import your scripts into SA. For requirements and guidelines on writing scripts, see [Writing Your Own Scripts for Extensible Discovery](#) on page 450.



This section presumes you are familiar with software policies. For details on software policies, see “Software Management Setup” in the *SA Policy Setter’s Guide*.



To add scripts to Extensible Discovery, you must have write access to the Extensible Discovery folder and the ability to create software policies. For more information on permissions, see the *SA Administration Guide*.

To create and add your script to Extensible Discovery, perform the following steps.

- 1 Decide where you want the output of your script to go. It can go either in a custom attribute or in a custom field.
 - For more information on custom attributes, see “Custom Attributes for Servers” in the *SA Users Guide: Server Automation*.
 - For more information on custom fields, see “Custom Fields for Servers” in the *SA Users Guide: Server Automation*.
- 2 Decide on the name of the custom attribute or custom field. Extensible Discovery will place the output of your script in the named custom attribute or custom field.
 - For a custom attribute, the name will be in the following format: `HPSW_ED_<name>` where `<name>` is any string you choose.
 - For a custom field, use any value for the `<name>` of the custom field.

You must use the `<name>` string in the name of your script as described in the steps below.

- 3 Write your script such that the output of the script is the data you want to capture. You can write any of the following types of scripts:
 - UNIX shell scripts in a file ending with `.sh`.
 - Visual Basic scripts in a file ending with `.vbs`.
 - Windows batch scripts in a file ending with `.bat`.

For more information how to write these scripts for Extensible Discovery, see [Writing Your Own Scripts for Extensible Discovery](#) on page 450.

- 4 Test your script to make sure it is functioning properly.
- 5 Name your script `get_<name>.sh` or `get_<name>.vbs` or `get_<name>.bat`, where `<name>` is from [step 2](#) on page 451 above. Extensible Discovery uses `<name>` to locate the custom attribute or custom field for the output of the script.
- 6 Make sure your script file has execute permissions. For example, for UNIX scripts use the `chmod` command.
- 7 If you want the output to go into a custom attribute, skip this step and go to [step 8](#) below. Extensible Discovery creates and places the output in the custom attribute named `HPSW_ED_<name>` by default.

If you want the output to go to a custom field, you must create a custom field named “`<name>`” where `<name>` is the string you used in [step 2](#) and [step 5](#) above. Extensible Discovery first checks for an existing custom field of the specified name. If that custom field exists, Extensible Discovery places the script output there. If that custom field does not exist, Extensible Discovery creates and stores the output in a custom attribute named “`HPSW_ED_<name>`”.

For example, if you create a script named `get_mysysdata.sh`, the output from the script will be placed into the custom field named `mysysdata`, if it exists. Otherwise the output will be placed in the custom attribute named `HPSW_ED_mysysdata`.

For instruction on creating custom fields, see [Creating and Managing Custom Fields](#) on page 456.

- 8 Wrap your script into a .zip file. Include any other files your script needs in the .zip file.
As a best practice, include a version string in the name of your .zip file and increment the version string with each subsequent version of your script. For more information, see [Upgrading Your Scripts in Extensible Discovery](#) on page 453.
- 9 Import your .zip file into a package in SA.

For convenience, you can place your packages in `/Opware/Tools/Extensible Discovery/Custom Provided Components`. Make sure your imported package specifies the proper target operating system.

For detailed instructions, see “Importing a Package” in the *SA Policy Setter’s Guide*.

- 10 Open your package in the SA Client and set the Default Install Path property to one of the following and save your changes.
 - UNIX: `/var/opt/opsware/extensible_discovery/scripts`
 - Windows: `%ProgramFiles%\Common Files\Opware\extensible_discovery\scripts\`.
SA replaces `%ProgramFiles%` with the appropriate system Program Files directory.
- 11 Add your package to a software policy. Add it either to the software policy named “Customer Provided Scripts” or add it to your own software policy.

Extensible Discovery remediates the “Customer Provided Scripts” policy by default whenever it runs.

Note that any user who has write access to the “Customer Provided Components” folder can run arbitrary code on any servers that Extensible Discovery is run on. For greater security, use your own software policy and set security on your software policy to meet your security requirements.

- 12 If you added your script to your own software policy, you must remediate your policy on all servers where you want Extensible Discovery to run.

If you added your script to the policy named “Customer Provided Scripts”, you can skip this step.

- 13 Run your script as described in [Running Extensible Discovery on Managed Servers](#) below.



If you intend to use a custom field but inadvertently run your script without having created the custom field and it creates custom attributes on many servers, you can use the following OGSF command to remove the custom attributes from all servers.

```
rm /opsw/Server/@/*/*CustAttr/<custom attribute name>
```

For more information on the SA Global Shell (OGSF), see the *SA Users Guide: Server Automation*.

Upgrading Your Scripts in Extensible Discovery

This section describes how to upgrade your scripts that are used with Extensible Discovery. The following steps presume you have already created and installed a script called `get_mysysdata.sh` and wrapped it in the file `get_windows_data_v1.0.zip` and imported it into Extensible Discovery as described in [Adding Your Own Scripts to Extensible Discovery](#) on page 451. To upgrade this script, perform the following steps.

- 1 Create the new version of your script and give it the same name as the original script file. For example, use `get_mysysdata.sh`. Follow the instructions in [Writing Your Own Scripts for Extensible Discovery](#) on page 450.

- 2 Wrap your script in a .zip file and increment the version string of the .zip file. For example, you could use `get_windows_data_v1.1.zip`.

- 3 Import your .zip file into a package in SA.

For convenience, you can place your packages in `/Opware/Tools/Extensible Discovery/Custom Provided Components`. Make sure your imported package specifies the proper target operating system.

For detailed instructions, see “Importing a Package” in the *SA Policy Setter’s Guide*.

- 4 Open your package in the SA Client and set the Default Install Path property to one of the following and save your changes.

- UNIX: `/var/opt/opsware/extensible_discovery/scripts`

- Windows: `%ProgramFiles%\Common Files\Opware\extensible_discovery\scripts\`. SA replaces `%ProgramFiles%` with the appropriate system Program Files directory.

- 5 Open the software policy “Customer Provided Scripts”. If you have used another software policy, open that policy.

- 6 Remove the old .zip file from the policy, `get_windows_data_v1.0.zip` in this example.

- 7 Add your new .zip file to the policy, `get_windows_data_v1.1.zip` in this example.

- 8 If you are using the “Customer Provided Scripts” policy, run Extensible Discovery as described in [Running Extensible Discovery on Managed Servers](#) on page 447. This remediates your new script on the managed servers.

If you are using another policy, remediate the servers with that policy. This remediates your new script on the managed servers so Extensible Discovery can be run.

Removing Your Scripts from Managed Servers

When you run Extensible Discovery, it copies your scripts in “Customer Provided Scripts” to the specified managed servers. To remove the scripts from the managed servers, perform the following steps. For this scenario, assume you have a script named `get_mysysdata.sh` wrapped in the file `get_mysysdata_v2.5.zip`.

- 1 Create a new version of your .zip file with a new version number, for example `get_mysysdata_v2.6.zip`.
- 2 Copy everything from the old .zip file into the new .zip file, except for the script you want to remove from managed servers, `get_mysysdata.sh` in this example.
- 3 Import your .zip file into a package in SA.

For convenience, you can place your packages in `/Opware/Tools/Extensible Discovery/Customer Provided Components`. Make sure your imported package specifies the proper target operating system.

For detailed instructions, see “Importing a Package” in the *SA Policy Setter’s Guide*.

- 4 Open your package in the SA Client and set the Default Install Path property to one of the following and save your changes.
 - UNIX: `/var/opt/opware/extensible_discovery/scripts`
 - Windows: `%ProgramFiles%\Common Files\Opware\extensible_discovery\scripts\`. SA replaces `%ProgramFiles%` with the appropriate system Program Files directory.
- 5 Open the software policy “Customer Provided Scripts”. If you are using your own software policy, open your policy.
- 6 Remove the old .zip file, `get_mysysdata_v2.5.zip` in this example.
- 7 Add your new .zip file, `get_mysysdata_v2.6.zip` in this example.
- 8 If you are using the software policy “Customer Provided Scripts”, Run Extensible Discovery as described below. This removes your script from the managed servers.

If you are using your own software policy, remediate the managed servers. This removes the scripts from the managed servers.

Output from Extensible Discovery Scripts

Each script being used by Extensible Discovery provides output that is placed either in a custom attribute or a custom field. The maximum size of this output is 1000 bytes. To save more than 1000 bytes, perform the following steps.

- 1 Follow step 1 through [step 6](#) on page 452 under [Adding Your Own Scripts to Extensible Discovery](#) on page 451.
- 2 Create a configuration file the same name as your script except ending with “.cfg”. For example, if your script is named `get_mysysdata.sh`, create the file `get_mysysdata.cfg`.
- 3 Enter the following line in your configuration file.

```
MAXBYTESTOCAPTURE=<number of bytes>
```

where <number of bytes> is the maximum number of bytes your script will produce.

- 4 Wrap your script file and the configuration file into a .zip file.
- 5 Follow the remaining steps under [Adding Your Own Scripts to Extensible Discovery](#) on page 451.

Comparing Custom Fields and Custom Attributes

SA can store a large amount of information about your managed servers. Custom Attributes and Custom Fields provide a way for you to store additional information about your servers quickly and easily. Custom Attributes and Custom Fields are data elements you can create for servers and other objects in SA.

Custom Attributes and Custom Fields are similar but they have several differences as described in the following table. In general, you should use custom fields when all servers require the data to be stored and you should use custom attributes when only a subset of servers require the data to be stored. However, see the following table for other differences before you decide which to use.

Table 32 Comparison of Custom Attributes and Custom Fields

	Custom Attributes	Custom Fields
Data Type:	String only.	Typed. Must be one of the types listed in Table 33, "Custom Field Data Types" on page 456.
Objects Allowed for:	Allowed for any object: servers, device groups, customers, facilities, OS installation profiles, and software policies.	Allowed only for servers and device groups.
Number:	Each custom attribute is for one object only.	Each custom field creates an instance for all servers or device groups. All managed servers have the same named custom field, but the value can vary with each server. Similarly for device groups.
Searches Allowed:	Search is only allowed on the custom attribute name, not on its value. That is, you can search for all servers that define a particular custom attribute.	Search is allowed based on custom field values, including different matching criteria for different data types. For example, if you have a custom field of type date, you can search for all servers where the date value is one month old or older.
Inheritance:	Inherited from more general objects. For example, servers inherit custom attributes defined for device groups they belong to.	No inheritance.

Table 32 Comparison of Custom Attributes and Custom Fields

	Custom Attributes	Custom Fields
Permissions Required to View:	Read permission on the server, device group or other object.	Read permission on the server or device group.
Permissions Required to Modify the Value:	Write permission on the server, device group or other object.	Write permission on the server or device group.
Permissions Required to Create or Delete:	Write permission on the server or device group where you are creating the custom attribute.	Manage Virtual Columns permission. Write permission on the server or device group.

For more information on Custom Attributes and Custom Fields, see “Custom Attributes for Servers” and “Custom Fields for Servers” in the *SA Users Guide: Server Automation*.

Creating and Managing Custom Fields

SA can store a large amount of information about your managed servers. Custom Fields provide a way for you to store additional information about your servers quickly and easily. Custom Fields are data elements you can create for servers and device groups.

When you create a custom field for servers, every server in your managed environment gets an instance of the custom field. When you create a custom field for device groups, every device group gets an instance of the custom field. The value of the custom field can be different for each server or device group.

For example, if your managed environment contains 500 servers and you create a custom field for servers, you would have 500 separate custom fields, one for each server. If you had 75 device groups and you created a custom field for device groups, you would have 75 separate custom fields, one for each device group.

Data Types in Custom Fields

Custom fields are typed. Each custom field you create must be of one of the following types.

Table 33 Custom Field Data Types

Custom Field Type	Description
String	Any characters, up to a maximum of 3999 characters.
Long String	Any characters. Use this type for strings longer than 3999 characters.
URI	A string representing a Uniform Resource Identifier.
Date	A date.
Number	A positive or negative integer.
File	An attached file.

Creating a Custom Field with the Custom Field Management Web Extension

The Custom Field Management web extension lets you create and delete custom fields.



To create or delete custom fields, you must have the following permissions: Manage Virtual Columns, Execute permission on the Web Extensions folder in the Library, and Read access to at least one managed server. For more information on permissions, see the *SA Administration Guide*.

To create a custom field, perform the following steps.

- 1 In the SA Client navigation pane, select Library and the By Type tab.
- 2 Select Extensions, then select Web.
- 3 Select the Custom Field Management extension and either right click or select the **Actions** menu and select **Run...**. This displays the Custom Field Management window as shown below.

Figure 93 Custom Field Management Web Extension - Create a Custom Field

The screenshot shows a web browser window titled "Custom Field Management". The main content area has a blue header bar that says "Create a New Custom Field Definition". Below this, there are three form fields: "Create on Object Type:" with a dropdown menu showing "Server", "Custom Field Definition Name:" with a text input field, and "Custom Field Definition Type:" with a dropdown menu showing "String". Below these fields is a button labeled "Process Input". At the bottom of the form area, there is a note: "Note: The string type can hold up to 3999 characters. The long string type should be used for any string." Below the note, there are two links: "Create a new Custom Field" and "Delete a Custom Field".

- 4 In the first drop-down list, select the object you want the new custom field to be associated with. If you select Server, every server will get an instance of the custom field. If you select Device Group, every device group will get an instance of the custom field.
- 5 Enter the name of the new custom field in the text input field.
- 6 In the second drop-down list, select the data type of the custom field. See [Table 33, "Custom Field Data Types"](#) on page 456.
- 7 Select Process Input to create the custom field.

Deleting a Custom Field with the Custom Field Management Web Extension

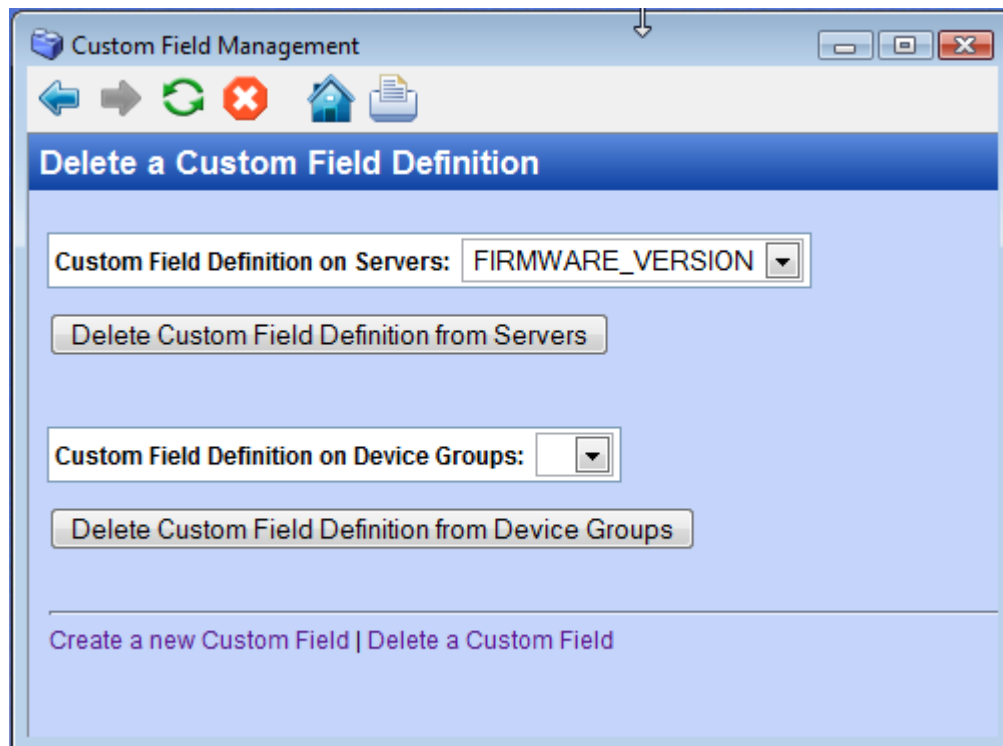
The Custom Field Management web extension lets you create and delete custom fields. To delete a custom field, perform the following steps.



When you delete a custom field, you delete all the values stored by all the servers or device groups associated with the custom field.

- 1 In the SA Client navigation pane, select Library and the By Type tab.
- 2 Select Extensions, then select Web.
- 3 Select the Custom Field Management extension and either right click or select the **Actions** menu and select **Run...**. This displays the Custom Field Management window as shown in Table 93, "Custom Field Management Web Extension - Create a Custom Field" on page 457.
- 4 Select the "Delete a Custom Field" link. This displays the Delete a Custom Field Definition window as shown below.

Figure 94 Custom Field Management Web Extension - Delete a Custom Field



- 5 To delete a custom field defined for servers, select a custom field name in the first drop-down list. The example above shows the custom field FIRMWARE_VERSION selected.

To delete a custom field defined for device groups, select a custom field name in the second drop-down list.

- 6 To delete a custom field defined for servers, select Delete Custom Field Definition from Servers.

To delete a custom field defined for device groups, select Delete Custom Field Definition from Device Groups.

11 Operating System Provisioning

HP Server Automation (SA) OS Provisioning provides you with the ability to install (or *provision*) pre-configured operating systems on servers in your facility, ensuring that each server in your facility has a standardized, default operating system configuration that you control.

SA OS Provisioning supports:

- Windows, Solaris, and Linux.
- Network or CD/DVD-based installations.
- Separation of duties between data center staff and systems administrators.
- A model-based approach — in which you create a *standard build* in SA which can then be installed on many systems.

OS Provisioning integrates with your operating system vendors' native installation technology, specifically:

- WinPE and Windows setup answer files: `unattend.txt`, `unattend.xml`, `sysprep.inf`
- Red Hat Kickstart
- SuSE YaST (Yet another Setup Tool)
- Solaris Jumpstart

You can provision an operating system on:

- A server in SA's unmanaged server pool that does not have an operating system installed (*bare metal sever*)
- A server in SA's *unmanaged server pool* with an installed operating system
- A server in SA's *managed server pool* with an installed operating system (*reprovisioning*)

SA OS Provisioning Prerequisites

Supported Operating Systems and Media for OS Provisioning

In order to provision operating systems to servers, the operating systems must be supported by OS Provisioning and the operating system media must first be made available to SA by uploading it to the Media Server. Additionally, OS Installation Profiles must be created in the SAS Web Client and/or in the SA Client. For more information about OS Provisioning set up, see the *SA Policy Setter's Guide*.

For a complete listing of all platforms supported for OS Provisioning, see *SA Supported Platforms* that is provided in the documentation directory of the distribution media.

Supported Boot Media

SA OS Provisioning works with:

- A CD-ROM for Windows via the WinPE preinstallation environment
- A CD-ROM for Linux
- Network booting for all supported operating systems.

Non-network booting is not supported for Sun Solaris (SPARC and x86).

Itanium-Based Systems

As of SA 9.0 and later, OS Provisioning supports provisioning only on Red Hat Enterprise Linux Itanium systems. Suse Linux Enterprise Server and Windows Itanium-based systems are not supported.

Solaris Servers

In order to perform PXE booting of a VMWare ESX Solaris 64-bit VM, the minimum required RAM is 1 GB.

OS Provisioning includes a DHCP-based JumpStart configuration that makes the complexity of JumpStart transparent to the end user.

For example, unlike typical JumpStart systems, OS Provisioning does not require configuration updates to the JumpStart server for each installation that you provision. Instead, OS Provisioning provides OS Installation Profile for each version of the Solaris operating system that you may install on servers.

The process for Solaris OS Provisioning generally follows the typical provisioning process.

See the *SA Policy Setter's Guide* for more information on the Solaris build process.

SPARC SUN4U Servers

In order to provision a bare metal SPARC SUN4U server with any of the supported SPARC Solaris versions, the server must support *Solaris 10 U8*.

HP-UX or AIX Operating Systems

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

Windows Servers

Windows system administrators can perform unattended, scripted installations as well as WinPE-based image installations of Windows Server 2008, Windows Server 2003, Windows XP Professional, and Windows 2000.

This installation-based approach allows system administrators to adapt to variations in hardware. OS Provisioning can use the information about correct hardware-specific software and drivers contained in a server's hardware signature file.

WinPE Memory Requirements

In order to perform PXE booting of a VMWare ESX Windows 2003 x86 or x86_64 VM using WinPE, the minimum required RAM is 512 MB (higher than the VMWare recommended RAM minimum).

See the *SA Policy Setter's Guide* for more information on the Windows build process.

Permissions

You must have been granted a specific set of feature permissions by your SA or system administrator to perform OS Provisioning. You must also have the permissions to access the servers associated with SA customers, facilities, or server groups.

For more information, see the Permissions Reference appendix in the *SA Administration Guide*.

Network Setup for OS Provisioning

- It is essential that you correctly configure any network switch ports used for OS Provisioning. These switch ports must have PortFast mode enabled and must be set for speed/duplex auto-negotiation. While provisioning using manually configured interface speed and duplex settings is possible for Solaris-based and Red Hat Linux-based boot images, It is recommended that you use auto-negotiation as it has been found to work the most consistently.
- You should configure the OS Build Agent to connect to the OS Provisioning Build Manager using the IP address as opposed to the DNS name.

If you must use DNS names, you must specify the DNS name during the SA Installer interview and save it in the Response File (see the `boot_server.buildmgr_host` parameter). You must also configure DNS so that the servers being provisioned can resolve the Build Manager host. The hostnames of all OS Provisioning Media Servers must also be resolvable.

Hardware Preparation

Before you use OS Provisioning to install an operating system, the target server must meet certain requirements which can vary according to the operating system being provisioned.

Windows Hardware Preparation Requirements

Before you provision the Windows operating system, you must prepare the hardware by performing the following tasks:

- OS Provisioning supports RAID configuration during provisioning, however you must complete certain configuration steps. For more information, see the *SA Policy Setter's Guide*.
- If there is a RAID controller installed, you may have to extend the Windows operating system media distribution (provide third party RAID drivers) based on hardware vendor-specific requirements. The Microsoft Windows operating system media might not

(depending on the version of Windows) include the necessary drivers for many RAID controllers. Also, certain newer types of SATA controllers may also require additional drivers.

- If you use a WinPE-based PXE or CD-ROM boot image to install the Windows operating system, disk partitioning is performed as part of the operating system installation. You can control the disk partitioning by editing the OS Installation Profile in SA. (For more information about creating installation profiles, see the *SA Policy Setter's Guide*.) Partitioning can also be controlled as an OS Build Plan task. See “OS Build Plans” on page 480.
- If you use a WinPE-based PXE or CD-ROM boot image and you are using a RAID or SATA controller, you might need to supply an operating system-specific Build Customization Script for OS Sequence based provisioning. These scripts enable you to load necessary hardware drivers before the operating system installation commences. For more information about Build Customization Scripts, see the *SA Policy Setter's Guide*. For OS Build Plan-based provisioning, you may be able to load necessary hardware drivers using OS Build Plan tasks.

Sun Solaris Hardware Preparation Requirements

To provision Solaris on a server, the hardware must meet the following requirements:

- The server must have a DHCP-capable PROM (older servers can be upgraded to DHCP-capable PROM).
- The server must be part of the SUN4U system architecture (platform group).

Linux and VMware ESX Hardware Preparation Requirements

- There are no special hardware requirements for Linux, however, if you have RAID drives installed, you must prepare the hardware by configuring valid, logical drives for RAID.
- VMware ESX hardware requirements are the same as Linux.

Red Hat Linux Hardware Preparation Requirements

- You must change the configuration of the managed switch for Red Hat Linux to enable PortFast. If this isn't done, when the Red Hat Linux installer attempts to use NFS to mount the media, the DHCP request could time out. (This problem is fixed in the packages listed in the advisory RHEA-2004:518-06.)
- If you will be PXE booting a Solaris 10 VM Boot image with an SA Boot Server that is hosted on a server running Linux, ensure that NFSv2 is enabled and that NFSv3 or NFSv4 is disabled. For more information, see the *SA Policy Setter's Guide*

OS Build Plan Requirements

- OS Build Plans are available only for Windows hosts. OS Build Plans require Automation Platform Extensions (APX) to complete certain tasks. In order for these APXs to operate, you must have the Adobe Flash Player installed on all machines from which you run the SA Client and OS Build Plans.

- SA provides a set of baseline OS Build Plans that you copy and use as a template for your own Build Plans. These Build Plans are not installed by default during SA installation or upgrade, rather you will be required to download and install the Build Plans using the DCML (DET) tool. These default OS Build Plans have been tested and are known to work. They simply need to be copied and adapted to your environment (location of the Media Server, required scripts, required reboots, etc.). For more information, see [OS Build Plans](#) on page 480.

The OS Provisioning Process

This section provides information about the SA OS Provisioning process and contains the following topics:

- [Affect of the OGFS Agent on Server Lifecycle](#)
- [Overview of the OS Provisioning Process](#)
- [Network Setup for OS Provisioning](#)

Overview of the OS Provisioning Process

The process for provisioning new servers typically includes tasks similar to the following:

1 Preparation

- a Physically prepare the server for operation and connect it to a network that can communicate with SA.
- b In some cases, you must prepare the server hardware for OS Provisioning.
See [Hardware Preparation](#) on page 461 in this chapter for more information.
- c OS Installation Profile(s) defined and available.
- d OS Build Plans (Windows only) and/or OS Sequences defined and available.

2 Boot the Server

Power on and boot the server using one of the following boot methods:

- a Use a bootable CD, or DVD provided by SA.



The bootable CD or DVD is not required for Intel-based servers that support PXE/WinPE/WinPE-OGFS or Unix servers as these servers can be remotely booted over a network.

- b For servers that can be booted over the network, powering on the server causes the server to initiate its network boot process.

For more information about booting servers remotely, see [Network Booting a Linux or VMware ESX Server using PXE](#) on page 465, [Network Booting a Windows Server Using PXE, WinPE, and WinPE/OGFS](#) on page 470, [Network Booting a Solaris Server](#) on page 474, and, for HP ILO servers, [The Manage Boot Clients \(MBC\) Option](#) on page 475.

- 3 After the unmanaged server boots successfully, it appears in the SAS Web Client in the list of servers ready for operating system installation.

See [Verifying That a Server is Ready for Operating System Installation](#) on page 505 in this chapter for more information.

4 Install the Operating System (Provision)

Select an operating system to install. If made available by your SA Administrator, you can also select a complete server operating system baseline (which can include a base operating system, a set of operating system patches, system utilities, and middleware software) to provision.

You can install the operating system immediately or schedule the installation later.

SA OS Provisioning-supplied CD Boot Images

SA OS Provisioning provides several service operating system boot CD images (ISOs) that you can use to record to CD. These ISO images can also be configured in virtual machine CD-ROM drives or mounted using iLO Virtual Media or similar technology. Use the SA Client Export utility to download the required image(s) and burn boot CDs.

These files are located in the SA Software Library Folder:

```
/Opware/Tools/OS Provisioning/WinPE
```

and are named using the format:

```
OPSWwinpe<arch>-<version>.iso
```

Booting Servers Remotely

On *Intel-based servers*, you can remotely boot a new server over a network using PXE. For other servers that do not support network boot technology, SA supports bootable CDs.

For *Windows and Linux or VMware ESX servers*, the SA Boot CD contains a small operating system, network drivers, the software required to mount a network drive, and the required SA communications infrastructure.

For *Solaris servers*, you can provision an operating system over the network if DHCP is available.



To boot servers over the network, the installation client must either be able to communicate with the SA DHCP server on the SA Core network or, for operating systems on which it is supported, you must supply static network configuration information at boot time. If the installation client is running on a different network than the SA core network, your environment must have a DHCP proxy (IP helper).



Before attempting to boot any server using DHCP to the SA server pool, ensure that the DHCP server's configuration file has the following line uncommented:

```
authoritative;
```

Booting from a CD

You can boot a remote server in three ways:

- 1 Mount a datastore ISO file into an ESX VMs virtual CD-ROM drive
- 2 Use iLO Virtual Media or a similar out-of-band management technology provided by your server's hardware vendor.
- 3 Record (burn) the operating system's ISO image to a physical CD and load the CD into a CD-ROM drive on the target server)

Network Booting a Linux or VMware ESX Server using PXE

The following section explains how to network boot a Linux or VMware ESX server with PXE. For information on how to boot a Windows server with WinPE, see [Network Booting a Windows Server Using PXE, WinPE, and WinPE/OGFS](#) on page 470. For more information about hardware support, see the *SA Policy Setter's Guide*.

To boot a Linux or VMware ESX server using PXE, perform the following tasks:

- 1 Prepare the server and connect it to the SA network, configure the server to boot using PXE.

See the hardware vendor's documentation for information about configuring a server to boot using PXE.

- 2 Power on the server and select the option to boot the server using PXE.
- 3 The following menu is displayed. Choose an SA boot image by entering the appropriate text (windows, winpe, linux4, etc.) at the boot prompt.

```
winpe32-ogfs - Windows Build Agent (WINPE 32-bit - OGFS based)
winpe64-ogfs - Windows Build Agent (WINPE 64-bit - OGFS based)
winpe32      - Windows Build Agent (WINPE 32-bit)
winpe64      - Windows Build Agent (WINPE 64-bit)
linux        - Linux Build Agent (RHEL 3.0-based)
linux5       - Linux Build Agent (RHEL 5.4-based)
solaris      - Solaris x86 Build Agent
localdisk    - Normal boot from localdisk (default after 10 second
```

- 4 Select linux or linux5 and press enter to begin the boot process.



If the operating system you are provisioning is Red Hat Enterprise Linux 3 IA64, you must add the custom attribute `kernel_arguments` with the value `console=ttyS1` to the OS Installation Profile.



If you are booting a VMware ESX server, select one of the `linux` options.

- 5 After the booting process finishes successfully, a message appears on the console indicating that the server is ready for OS Provisioning and the server now appears in the SA Client Unprovisioned Servers list as available for operating system installation.
- 6 (Optional) Record the MAC address and/or the serial number of the server so that you can locate the server in the SA Client Unprovisioned Servers list.

- 7 Verify that the server appears in the SA Client Unprovisioned Servers list and that it is ready to hand off for operating system installation.

See [Verifying That a Server is Ready for Operating System Installation](#) on page 505 in this chapter for more information.

Booting a Red Hat Enterprise Linux Server in a Non-DHCP Environment

If you plan to use OS Provisioning in an environment without a DHCP server, you must assign static IP information for the managed server and manually configure that server to resolve the SA Core.

There are several reasons you might need to manually specify the network information for a server being provisioned:

- You don't use DHCP and must manually specify the static IP address and the Agent's IP and Port
- You must provision a server but DHCP is inactive.
- You must provision a server but DHCP is blocked by firewall rules.

CD boot images for Linux OS provisioning in non-DHCP environments can be exported by selecting **Library > By Folder > Opsware > Tools > OS Provisioning**.

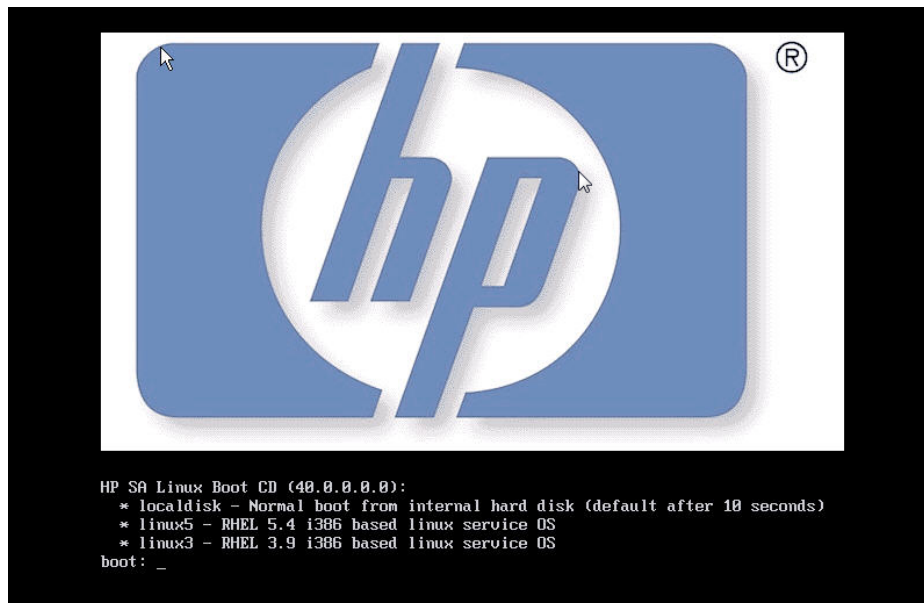
The images are named using the following format:

```
HPSA_linux_boot_cd-<version>.iso
```

This section provides details for provisioning in a non-DHCP environment.


When you boot an unmanaged server in a non-DHCP environment, you will see a boot screen similar to that shown in [Figure 95](#):

Figure 95 Red Hat Linux Boot Screen



After you select the boot method, you will see a Network Configuration dialogue that allows you to enter a static IP address for the server, the subnet mask, The host gateway IP address, and the IP address and default port for the SA Agent Gateway, [Figure 96](#):

Figure 96 Red Hat Linux Network Configuration Dialog



You can manually configure the following fields:


- Interface: the NIC to be used
- IP Address: static IP address for the server being provisioned
- Netmask: netmask for the server being provisioned
- Gateway: Gateway IP address the server being provisioned should use (network level IP router)
- DNS Server: the IP address the server being provisioned should use
- DNS Search Path: the fully qualified DNS suffix the server being provisioned should use
- Agent Gateway IP: the default SA Agent Gateway hostname or IP address
- Agent Gateway Port: the port used for the SA Agent Gateway

After the information in these fields is entered and applied, the server is able to register with the SA Core. You can now start the normal OS Provisioning process.

DHCP Custom Attribute

Servers that have been registered with the SA Core using a static IP specification will display the `hpsa_netconfig` custom attribute in the server record, as shown in [Figure 97](#):

Figure 97 hpsa_netconfig Custom Attribute in Server Record



Booting a Red Hat Enterprise Linux Itanium 64-bit Server in a Non-DHCP Environment

If you plan to use OS Provisioning in an environment without a DHCP server, you must assign static IP information for the managed server and manually configure that server to resolve the SA Core.

There are several reasons you might need to manually specify the network information for a sever being provisioned:

- You don't use DHCP and must manually specify the static IP address the Agent's IP and Port
- You must provision a server but DHCP is inactive.
- You must provision a server but DHCP is blocked by firewall rules.

You can export the Linux Itanium image by logging in to the SA Client and selecting **Library ► By Folder ► Opware ► Tools ► OS Provisioning**.

The images are named using the following format:

`HPSA_linux_boot_cd_IA64-<version>.iso`

The following section provides details for provisioning in a non-DHCP environment.

When you boot an unmanaged server in a non-DHCP environment, you will see a boot screen similar to that shown in [Figure 98](#):

Figure 98 Red Hat Linux Itanium 64-bit Boot Screen

```
HP SA Linux Boot CD (40.0.0.3):  
Enter the appropriate linux service OS  
at the 'Elilo boot:' prompt.  
  
linux5      - RHEL 5.4 ia64 based linux service OS  
  
linux5-txt  - RHEL 5.4 ia64 based linux service OS for serial consoles  
  
ELILO boot:
```

After you select the boot method, you will see a Network Configuration dialogue that allows you to enter a static IP address for the server, the subnet mask, The host gateway IP address, and the IP address and default port for the SA Agent Gateway, [Figure 99](#):

Figure 99 Red Hat Linux Itanium 64-bit Network Configuration Dialog

SANetGui Network Settings

Interface: eth0
IP Address: 192.168.208.54
Netmask: 255.255.255.240
Gateway: 192.168.208.49
DNS Server: 192.168.194.4
DNS Search Path: c126.dev.hp.com
Agent Gateway IP: 192.168.208.52
Agent Gateway Port: 8017

OK

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen



If the operating system you are provisioning is Red Hat Enterprise Linux 3 IA64, you must add the custom attribute `kernel_arguments` with the value `console=ttyS1` to the OS Installation Profile.

You can manually configure the following fields:

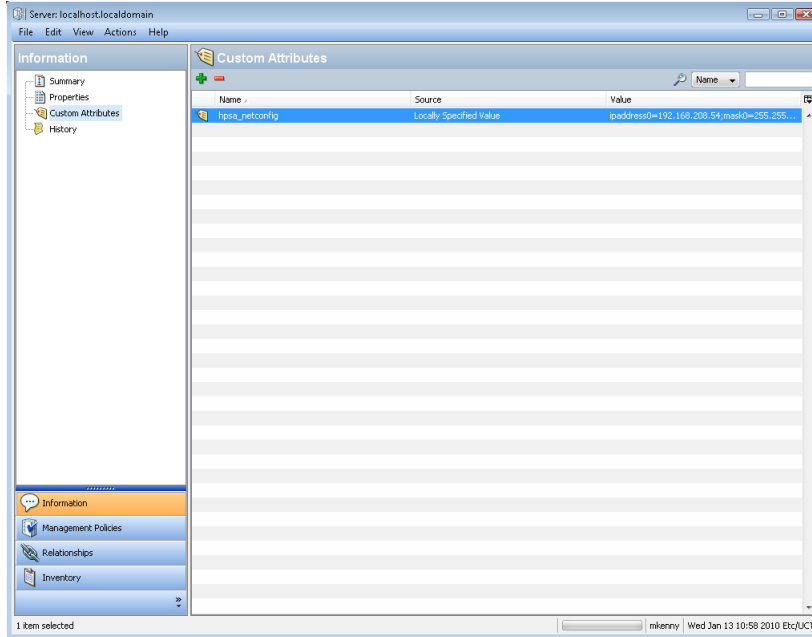
- Interface: the NIC to be used
- IP Address: static IP address for the server being provisioned
- Netmask: netmask for the server being provisioned
- Gateway: Gateway IP address the server being provisioned should use (network level IP)
- DNS Suffix: the fully qualified DNS suffix the server being provisioned should use
- Agent Gateway IP: the default SA Agent Gateway hostname or IP address
- Agent Gateway Port: the port used for the SA Agent Gateway

After the information in these fields is entered and applied, the server is able to register with the SA Core. You can now start the normal OS Provisioning process.

DHCP Custom Attribute

Servers that have been registered with the SA Core using a static IP specification will display the `hpsa_netconfig` custom attribute in the server record, as shown in [Figure 100](#):

Figure 100 `hpsa_netconfig` Custom Attribute in Server Record



Network Booting a Windows Server Using PXE, WinPE, and WinPE/OGFS

OS Provisioning supports booting a server with no operating system by PXE booting into a WinPE preinstallation environment. You can choose between either a WinPE x86 32-bit environment, a WinPE x86 64-bit environment, or WinPE x86 32-bit and 64-bit using an OGFS-specialized Server Agent that boots from the boot image.

In the WinPE x86 32-bit and WinPE x86 64-bit environments, an OS Provisioning Build Agent is loaded onto the server to be provisioned. In the WinPE-OGFS environment, the agent is a full SA Server Agent that is booted when an OGFS-enabled image is selected and is specialized for OGFS functionality including running OS Build Plans, disk repair, file restoration, and other non-destructive maintenance.

For more information about how the various OS Provisioning components like Build Agents and OGFS Server Agents function, see [Advanced SA OS Provisioning Architecture](#) on page 501.

WinPE also allows WIM-based image installation, as an alternative to unattended Windows installations.



You can PXE boot a Windows server using DHCP if that server resides on the same VLAN as an SA Boot Server (from an SA Core or Satellite). If your Windows server is not on a VLAN with PXE available, you can use a DHCP Helper IP relay configuration in your network switch. If neither configuration for network booting is available, you can boot using a CD-ROM

To boot a bare metal server with PXE into a WinPE preinstallation environment, perform the following steps:

- 1 Mount the new server in a rack and connect it to the SA build network.
- 2 Configure the server to boot using PXE.

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

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

The SA menu appears and prompts you to select the type of OS Build Agent to load on the server.

All of the SA boot image options are displayed:

```
winpe32-ogfs - Windows Build Agent (WINPE 32-bit - OGFS based)
winpe64-ogfs - Windows Build Agent (WINPE 64-bit - OGFS based)
winpe32      - Windows Build Agent (WINPE 32-bit)
winpe64      - Windows Build Agent (WINPE 64-bit)
linux        - Linux Build Agent (RHEL 3.0-based)
linux5       - Linux Build Agent (RHEL 5.4-based)
solaris      - Solaris x86 Build Agent
localdisk    - Normal boot from localdisk (default after 10 second)
```

- 4 At the boot prompt enter:

```
winpe32, winpe64, winpe32-ogfs, or winpe64-ogfs
```



If you do not select an option within 10 seconds, the server defaults to booting from the local disk. If you need more than 10 seconds to make your decision you can type anything but do not press ENTER at the command line.

- 5 A new menu displays the option to boot a WinPE x86 32 bit environment or a Windows x64 64 bit environment. Make a selection by using the arrow keys to highlight your choice, and then press ENTER.

The server will now be booted with the WinPE preinstallation environment. This may take a few minutes to complete, depending upon the speed of the network and the machine.

Once booting has finished, a new window will appear indicating that the server has had an SA Build Agent loaded and registered with the SA core.

- 6 (Optional) Record the MAC address and/or serial number of the server so that you can locate the server in the Server Pool list in the SAS Web Client or in the Unprovisioned Servers list in the SA Client.
- 7 Verify that the newly racked server shows up in the SA Client Unprovisioned Servers, or SAS Web Client Server Pool, and is ready for OS installation. See [Verifying That a Server is Ready for Operating System Installation](#) on page 505 in this chapter for more information.

Booting a Windows Server in a Non-DHCP Environment

If you plan to use OS Provisioning in an environment without a DHCP server, you must assign static IP information for the managed server and manually configure that server to resolve the SA Core.

There are several reasons you might need to manually specify the network information for a sever being provisioned:

- You don't use DHCP and must manually specify the static IP address and the Build Manager's IP and Port
- You must provision a server but DHCP is inactive.
- You must provision a server but DHCP is blocked by firewall rules.

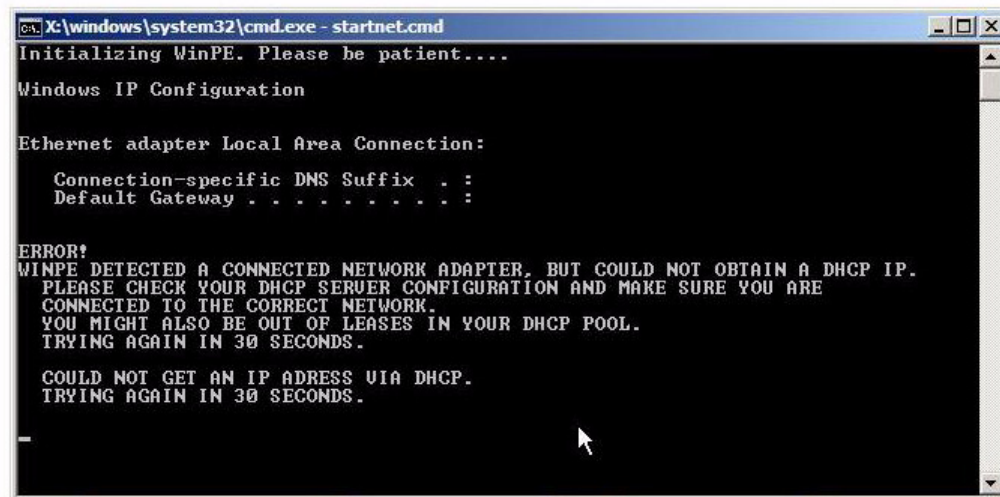
When provisioning a server using WinPE, by default, WinPE looks for a DHCP server. If a DHCP server is not found, you are prompted to enter the IP address, Subnet mask, Gateway and Name server of the host, and the Port and Hostname/IP of the SA Core.

This section provides details for provisioning in a non-DHCP environment.

Booting an Unmanaged Windows Server in a Non-DHCP Environment

When you boot an unmanaged server into a non-DHCP environment, by default WinPE looks for an available DHCP server. If WinPE does not find a DHCP server, you see a display similar to [Figure 101](#).

Figure 101 WinPE Console Display when DHCP Server Not Found



At this point, you will see a Network Configuration dialogue that allows you to enter the SA Agent Gateway IP or enter a static IP address for the server, the subnet mask, The host gateway IP address, and the IP address and default port for the Build Manager. See [Figure 102](#):

Figure 102WinPE Network Configuration

Interface*: Intel(R) PRO/1000 MT Network Connection

MAC Address: 00:0C:29:C1:65:43

☐ Static IP ☒ DHCP

IP Address: 192.168.208.61

Netmask: 255.255.255.240

Gateway: 192.168.208.49

DNS Server: 192.168.194.4

DNS Suffix:

Server Automation Agent Gateway

Hostname / IP:

Port*: 8017

OK Cancel

You can manually configure the following fields:


- IP Address: static IP address for the server being provisioned
- Subnet: Subnet mask for the server being provisioned
- Gateway: Gateway IP address the server being provisioned should use (network level IP router)
- DNS Server: the IP address the server being provisioned should use
- DNS Suffix: the fully qualified DNS suffix the server being provisioned should use
- Agent Gateway: SA Agent Gateway hostname or IP address
- Port: The port used for the Build Manager

After the information in these fields is entered and applied, the server being provisioned will be able to register with the SA Core.

DHCP Custom Attribute

Servers that have been registered with the SA Core using a static IP specification will display the `hpsa_netconfig` custom attribute in the server record, as shown in [Figure 103](#):

Figure 103hpsa_netconfig Custom Attribute in Server Record



Network Booting a Solaris Server

When SA is installed, OS Provisioning is configured so that the Boot Server listens for broadcast requests from new servers and responds using DHCP.

Perform the following steps to boot a Solaris server over the network:

- 1 Mount the new Solaris server in a rack and connect it to the network.

The installation client on this network must be able to communicate with the SA DHCP server on the SA core network. If the installation client is running on a different network than the SA core network, your environment must have a DHCP proxy (IP helper).

- 2 Enter one of the following commands at the prompt:

```
ok boot net:dhcp - install
```

or

```
ok boot net:dhcp - install <interface_setting>  
<buildmgr=hostname|IP_address>
```

where <interface_setting> is one of the following options:

```
autoneg, 100fdx, 100hdx, 10fdx, 10hdx
```

You can include an interface setting with the boot command to set the network interface to a specific speed and duplex during OS Provisioning. Specifying this boot argument allows you to override the default interface setting that was specified when SA was installed in the local facility.

You can use a variety of methods including Solaris build customization scripts or specifying the values in a Solaris Package or RPM in the operating system media to set the network interface with a specific speed and duplex.

See the *SA Policy Setter's Guide* for more information.

The Manage Boot Clients (MBC) Option

The Manage Boot Clients (MBC) option provides several services. You can:

- Remotely boot a server. You do not need console access to the server.
- Pre-create server records.
- Create custom attributes that set server configuration during OS Provisioning.
- Reconfigure services like DHCP when new servers are provisioned.
- Initiate OS Provisioning from a portal or an automated script where, typically, the user will not be available for interactive responses.

For example, you can change the default PXE image that a server uses to boot, change whether a server is assigned a DHCP lease, or specify the DHCP IP that is assigned to the server. You can also change a server's behavior when it enters the server pool, such as automatically invoking an OS Sequence when it enters the pool.

If the server is an HP ProLiant server with iLO2 enabled, and you know its iLO information, MBC can also remotely power on the server.

Any user, such as a system administrator who performs OS Provisioning and who is responsible for the base operating system, system utilities, patching, and the hand off of servers to internal business units, will find MBC quite useful.

You can access MBC functionality:

- From the SA Client
- From the Global File System command line
- From a script
- From a browser/portal form

Requirements

- The OS Provisioning infrastructure relies on SA Boot Server services for the MBC extensions.
- The OS Provisioning boot images must be served by the TFTP server that is shipped with SA.
- In order to take advantage of the DHCP reconfiguration feature, you must use the SA DHCP server.
- On a newly installed SA Core, a new user prior to running the MBC Web APX must first be granted Launch Global Shell permissions and must log in to the OGSF at least once in order to initialize the user environment (so that MBC can write temporary files to the user's home directories during use).

Required Permissions

In order to execute MBC, a user must have the *Allow Execute OS Sequence, Managed Server and Groups, Manage Customers, Server Pool, Read & Write permission to customer Not Assigned* and *Allow Configuration of Network Booting* permissions, write access to all pre-existing servers they will act on, and permissions to run the MBC APXs (thus, they need execute access on the /Opware/Tools/OS Provisioning/Manage Boot Clients folder).

Installation

The HP BSA Installer creates the MBC APXs during the SA Core installation. The installer creates a folder containing the MBC APXs in the SAS Web Client Library, and adds an MBC Configuration Software Policy as part of the data baseline.

The following four APXs are installed for MBC:

- Program APX
- Web APX
- Integration Hook APX
- DHCP Cleanup Web APX

Using the Manage Boot Clients (MBC) Option

When MBC runs, it creates new server record(s) in the SA database in the Planned lifecycle. These records are displayed with a *blueprint* icon and can optionally have custom attributes assigned to them. Some of these custom attributes change how SA handles a server or configuration of an operating system installation (for example, you can set the **ComputerName** for a Windows unattended installation).

Executing MBC will typically change the default PXE menu choice when the server PXE boots, so that the user does not need to choose a PXE image on console of the server that's booting up. MBC also allows users to associate an OS Sequence with the server record so that, when the server registers as an unprovisioned server with HP SA, a provisioning job is kicked off automatically. Running an MBC APX from the SA Client

The MBC Web Interface

You can launch MBC Web APXs in three ways:

From the SA Client

- Select **Library ► Extensions ► Web ► Manage Boot Clients Web APX**.
- or, from the Unprovisioned Servers list, right click in the server list pane (not directly on a server) and select **Manage Boot Clients**.

From a Browser

You can also use a browser and navigate to:

`https://occ.example.com/webapp/osprov.manage_boot_clients_web/`

where `occ.example.com` is the local hostname or IP address for your SA Core.

The browser interface allows you to choose whether to use a form to input data for a singular host, or whether to input a CSV to set up multiple server records. After clicking the **Submit** button, it is grayed out to prevent double-submissions and a combined Progress/Results page is displayed.

The MBC Form-Based Method (Web-based)

The Web form-based interface provides a set of four pages that guide you through setting up an MBC job. You provide the information necessary to boot and provision a server on the first three pages/forms. The final page displays the progress/results of the job. You can act only on a single server when using the form-based method. For multiple server setup, you must use the CSV method.

Using the CSV Method from the Web Interface

The CSV input method can be accessed by clicking the **Multiple Client Form...** button on the first page of the MBC Web UI. The CSV input form allows acting on multiple server records at once, where each line in the CSV represents a server record.

The MBC APX Command-Line Interface

MBC also provides a Program APX, which is available to users as an executable in the Global Shell (OGSH). This can be useful for programmatic access to MBC while integrating with other systems.

Usage:

Users who have the appropriate permissions can run MBC from OGSH with this command:

```
/opsw/apx/bin/osprov/manage_boot_clients_script
```

Running MBC from the command line with no arguments will provide a usage statement.

This is an example command line entry that executes MBC and uses an existing CSV file:

```
/opsw/apx/bin/osprov/manage_boot_clients_script -m import  
<full path to CSV file with boot clients>
```

Special Attributes for the CLI and CSV Input Form

There are several special attributes which are not stored as custom attributes (except `sequence_id`) when entered, but instead are dealt with in distinct ways. [Table 34](#) lists these special attributes and how they are dealt with.

Table 34 MBC Special Attributes for the CLI and CSV Input Form

Parameter	Description
<code>pxe_image</code>	Specifies a PXE configuration files for the server. The value should be set to one of the options seen in the default PXE menu (such as <code>winpe32</code> , <code>winpe64</code> , or <code>linux5</code>). This creates a symlink in <code>/opt/opsware/boot/tftpboot/pxelinux.cfg</code> from the MAC address to the PXE configuration file.
<code>sequence_id</code>	If specified, will invoke an OS Sequence installation (as <code>detuser</code>) as soon as the server is added to the Server Pool. Note: <code>sequence_id</code> actually is stored as a custom attribute on the server. This custom attribute is removed from the server record before the first reboot of the server.
<code>customer</code>	Sets the customer association for the server.
<code>use</code>	Sets the use field for the server. The value specified should be all caps (for example, <code>PRODUCTION</code>)
<code>stage</code>	Sets the stage field for the server. The value specified should be all caps (for example, <code>IN DEPLOYMENT</code>)
<code>facility</code>	Sets the facility association for the server. This is necessary when you run an MBC APX from a facility other than the one that the target server is associated with (necessary when you have a satellite that defines its own facility).

Table 34 MBC Special Attributes for the CLI and CSV Input Form (cont'd)

Parameter	Description
ilo.*	See “iLO Integration”.

Additional non-MBC-specific custom attributes are available for the installation of Windows, Solaris, and Linux operating systems. See the OS Provisioning chapter in the *SA Policy Setter's Guide*.

CSV Input Files

MBC's ability to accept CSV input files allows you to move servers into the Managed Server Pool and provision them with an operating system without the use of a console and an interactive session.

For example:

```
00:0c:29:e1:28:2e,hostname=testvm1,pxe_image=linux5,
sequence_id=2110061
00:0c:29:f9:12:f3,hostname=testvm2,pxe_image=winpe32
00:0c:29:0d:ab:b4,pxe_image=solaris,sequence_id=2110061
```

These CSV entries would cause MBC to create three Planned Server records and set them up to boot to the linux5, winpe32, and solaris PXE images, respectively. The servers processed by the first and third CSV entries will also have an OS Sequence applied when they register with SA. The first two entries would have specific display names shown in SA (hostname=), while the third would have an auto-generated hostname that be similar to dhcp-client-00:0c:29:0d:ab:b4. For more information on these attributes and their function, see the Special Attributes in [Table 34](#).

Example CSV Entries

```
00:13:E8:9A:93:BA,pxe_image=winpe32,dhcp.ip=10.2.3.11,
dhcp.hostname=m0011,customer=WealthManagement,
sequence_id=2030001,dns_server=10.6.4.2,
kernel_arguments=noacpi,root_password=wealth

00:13:E8:9A:93:BC,pxe_image=winpe32,dhcp.ip=10.2.3.12,dhcp.hostname=m0012,
customer=WealthManagement,sequence_id=2030001,
dns_server=10.6.4.2,kernel_arguments=noacpi,
root_password=wealth

00-13-E8-9A-93-99,pxe_image=linux

00:13:E8:9A:93:AA,pxe_image=windows,custattr1=val1,
custattr2=val2

00:13:E8:9A:93:BB,pxe_image=windows,customer=Opware

00:0c:29:23:a1:7f,pxe_image=linux,sequence_id=310005,
testca=testval

00:0c:29:af:46:6b,pxe_image=linux,sequence_id=310005,
testca=testval

00:0c:29:be:96:6e,pxe_image=winpe32,sequence_id=320005
```

```
00-13-21-DD-DD-24,pxe_image=linux,sequence_id=310001,
dhcp.hostname=danube,ilo.hostname=10.128.32.102,
ilo.username=Administrator,ilo.password=adminpass,
ilo.reboot_if_on=1
...
```

The first item on each line of CSV must be a MAC address followed by a list of arbitrary, comma-separated name/value pairs, where the names and values are separated by equal signs. Each of these name/value pairs is stored as a custom attribute on the server record which allows the user to set up many custom attributes simultaneously.

Special Attributes for DHCP Reconfiguration

MBC has the ability to add host definitions to SA DHCP configuration files. This is useful in environments where SA DHCP is used, but configured to deny unknown clients (that is, it will only provide DHCP leases to *approved* MAC addresses). When you specify a DHCP hostname's MAC address on the **General** Form, MBC adds this MAC address to DHCP configuration. You can also specify DHCP IP address if required.

Table 35 lists the DHCP reconfiguration special attributes you can use in the CSV:

Table 35 DHCP Reconfiguration Special Attributes

Attribute	Description
dhcp.hostname	Specifies the MAC address for hostname(s) that are authorized for DHCP leases.
dhcp.ip	Specifies the IP address(es) of hosts that are authorized for DHCP leases.

iLO Integration

MBC includes integration with the HP Integrated Lights-Out 2 (iLO2) Standard. This increases the level of control that SA has over servers, down to the level where the users no longer have to even power on the servers. When the user provides an iLO IP and credentials, MBC will connect to the iLO API and automatically power on the server. ILO also provides more thorough hardware discovery.

Table 36 show the special attributes used for ILO Integration:

Table 36 ILO Special Attributes

Special Attribute	Description
ilo.hostname	Hostname or IP address for the iLO. This must be accessible from the hub/OGFS server. This value is stored as a custom attribute by MBC.

Table 36 iLO Special Attributes (cont'd)

Special Attribute	Description
ilo.username	Username to use to authenticate to the iLO. This value is stored as a custom attribute by MBC.
ilo.password	Password used to authenticate to the iLO. This value is not stored as a custom attribute by MBC.
ilo.reboot_if_on	Default: power the server on only if it is currently off. If you specify this argument with a non-null value, MBC reboots the server, even if it's already on. This value is not stored as a custom attribute by MBC.

The first page of the Web APX has form inputs for the iLO parameters.

The following is an example CSV that will cause MBC to boot/reboot the server:

```
00-13-21-DD-DD-24,pxe_image=linux,sequence_id=310001,
dhcp.hostname=danube,ilo.hostname=10.128.32.102,
ilo.username=Administrator,ilo.password=adminpass,
ilo.reboot_if_on=1
```

Specifying OS Provisioning Tasks

Before you can begin provisioning operating systems, your SA Administrator must have defined *OS Installation Profiles* as described in the *SA Policy Setters Guide*.

You must also define either *OS Build Plans* (Windows) or *OS Sequences* that are used to specify and organize provisioning tasks as described in the following sections.

OS Build Plans

As of this SA release, OS Provisioning provides a new, more flexible method for Windows hosts to specify how an operating system is installed called *OS Build Plans*. You use an OS Build Plan to specify server provisioning details, such as operating system configuration information, software, customization scripts and patch policies. While similar to the OS Sequence capability, OS Build Plans provide these functional improvements over OS Sequences:

- OS Build Plans make it easier to customize the operating system installation to meet your specific needs, for example:
 - Integration with other internal systems at specific points during the operating system build phase.
 - Running a RAID configuration utility or a firmware update
 - Modifying the unattend.xml file from a script before beginning an installation process
- Simpler architecture. OS Build Plans use the same network ports and protocols as a full SA Agent. Fewer SA Core Components are involved.

- OS Build Plans use the more robust and powerful execution environment of the Global Shell (OGFS).
- A more transparent build process means easier progress monitoring and troubleshooting.
- The use of an OGFS Agent provides an easy way to configure and troubleshoot servers before or during an operating system build.
- OS Build Plans allow simpler set up:
 - Running the `import_media` utility is no longer required.
 - Defining OS Installation Profiles in the SAS Web Client is now optional, not required.
- No separate client installation is required to deploy operating systems.
 - The new Run OS Build Plan wizard is a web application.
 - The SA Client can be used to define OS Build Plans.
 - Build Plan APXs can be run from the command line or from scripts.
- Perform other tasks beyond OS Installation. For example, OS Build Plans can be created for image capture, file restore, or secure data erasure.



You can still use OS Sequences to configure your operating system installation. The functionality is still fully available. See [OS Sequences](#) on page 489. HP recommends, however, that you explore the advanced features of OS Build Plans and consider migrating from OS Sequences to OS Build Plans. OS Build Plans are currently available only for Windows operating system installations.

Affect of the OGFS Agent on Server Lifecycle

- Servers running the OGFS Agent have their state field set to `MAINTENANCE` meaning they are ready for OS Build Plans or ad-hoc OGSF access.
- When an OS Build Plan is running, the server's lifecycle is set to `PROVISIONING`. This prevents running more than one job against the same server simultaneously.
- When an OS Build Plan completes, the final state of the target server is affected by several conditions as well as the server's initial lifecycle when the OS Build Plan was first launched.

What are OS Build Plans?

When you install (provision) an operating system on a server, you may have certain configurations settings you want to specify for all similar servers, you may also want to set environment variables, install and configure applications, configure RAID settings, and so on.

SA OS Build Plans provide a framework you can use to design server installation templates that configure a server exactly as you want it configured during OS Provisioning.

OS Build Plans take advantage of the SA Global File System (OGFS) and Automation Platform Extensions (APXs) to install operating systems on unprovisioned servers.

Using SA OS Build Plans, you can specify the following tasks to be performed during OS Provisioning:

- **Run scripts**

You can create and specify OGFS scripts and server scripts (with arguments) for answer file processing, custom attribute handling, customer assignments, and so on and server scripts for reboots and general functionality to be run during OS Provisioning.

- **Install zip packages**

You can specify zip packages to be installed which can be used to deliver any additional software utilities you may need during the OS installation. For example, you might need to install some additional driver software or make use of a third party image installation program.



OS Build Plans do not process pre/post scripts nor reboot settings on ZIP packages. Such settings apply only to Software Policy usage of ZIP packages. Pre/post scripts can be defined as separate Build Plan steps. Rebooting should be handled by the boot script provided for this purpose.

- **Attach Patch Policies**

You can specify the Patch Policies that must be applied.

- **Attach Software Policies**

You can specify the Software Policies that must be applied

- **Server remediation**

You specify the remediation's reboot options and Error Handling option.

- **Add server to Static Device Groups**

You can specify a static Device Group to which the server should be added.

Baseline OS Build Plans

SA provides a set of baseline OS Build Plans that you copy and use to base your Build Plans on. These Build Plans are not installed by default during SA installation or upgrade, rather you will be required to download and install the Build Plans using the DCML (DET) tool. Instructions for installing the baseline OS Build Plans are included with the download. Installation commands will be similar to the following:

```
mkdir /tmp/osbp_import
unzip -d /tmp/osbp_import /var/tmp/OPSWosbp_content-40.X.Y.Z.zip
/opt/opsware/cbt/bin/cbt -cf /var/tmp/core.cfg -i /tmp/osbp_import
```

These default OS Build Plans have been tested and are known to work. They simply need to be copied and adapted to your environment (location of the Media Server, required scripts, required reboots, etc.). These Build Plans appear in the SA Client Library under

/Opsware/Tools/OS Provisioning/OS Build Plans/Windows.

The baseline OS Build Plans currently available include:

- *Install Windows OS Installation Profile*: allows you to use an existing OS Installation Profile by specifying its name or ID when you run this script. This allows you to perform the tasks already specified in the Installation Profile. However, it *does not handle* Build Customization Scripts, packages from the OS Installation Profile, or remediation.
- *Windows 2003 Default Install*: a build plan designed specifically to install Windows Server 2003.

- *Windows 2003 WIM Install*: a build plan designed specifically to install Windows Server 2003 using a WIM image.
- *Windows 2003 Default x64 Install*: a build plan designed specifically to install Windows Server 2003 x86_64.
- *Windows 2003 x64 WIM Install*: a build plan designed specifically to install Windows Server 2003 x86_64 using a WIM image.
- *Windows 2008 Install*: a build plan designed specifically to install Windows Server 2008.
- *Windows 2008 WIM install*: a build plan designed specifically to install Windows Server 2008 using a WIM image.
- *Windows 2008 x64 Install*: a build plan designed specifically to install Windows Server 2008 x86_64.
- *Windows 2008 x64 WIM Install*: a build plan designed specifically to install Windows Server 2008 x86_64 using a WIM image.
- *SAMPLE: Windows 2003 Install to Alternate Disk*: a build plan designed to allow you to change the default drive letter for installation.
- *SAMPLE: Windows 2008 with Static IP*: a build plan designed to allow you to specify a static IP for the server being provisioned.

OS Build Plan Requirements

- During installation, SA by default installs several Automation Platform Extensions (APX) that perform the actions specified in Build Plans. These APXs appear in the SA Client's APX Library and can be run manually or in scripts for certain tasks, but they should not be removed or modified.
- The OS Build Plan wizard requires that you install the Adobe Flash Player plug-in on all clients from which you will run the wizard.
- In order for Windows OS Build Plan OS provisioning to copy/install folders and files from the OS media's \$OEM\$ folder to the mapped destination folder on the target server, you should use one of the following options:
 - Create a new OS Build Plan by copying a baseline Windows OS Build Plan and ensure that the unattend file section in the `Configure Windows <version> Default Unattend.txt` script has the following entry:


```
OemPreInstall=Yes
```

 Add the `Inject Required Unattend.txt` setting OGFS script as a step in the OS Build Plan before the `Mount Windows Store` script.
 - Use one of the baseline Windows Default Install OS Build Plans and provide it with the name of an existing OS Installation Profile. In this case, the legacy behavior is emulated, so `OemPreInstall=Yes` is automatically added if missing.

Copying a Baseline OS Build Plan



You must always make a copy of the default baseline OS Build Plan content files. Never modify the original.

After you have downloaded and installed the baseline OS Build Plans, HP strongly recommends that you copy an appropriate plan and use that as the basis for your build plan.

To copy a plan:

- 1 Log in to the SA Client.
- 2 Navigate to **Library ► OS Build Plans**.
- 3 In the OS Build Plan pane, right click and select **New** from the context Menu.
- 4 In the OS Build Plan Wizard, give your new OS Build Plan a name and optional description on the Properties page.
- 5 Select Build Plan Items and click the Copy Plan button and navigate to the folder in which you stored the baseline OS Build Plans.
- 6 Highlight the plan you want to copy and click Select. The task scripts contained in the baseline OS Build Plan are copied into the task list for your new plan.
- 7 Modify the scripts for your environment. Note that some script have required parameters like @MediaServer@ for which you must supply the fully qualified path and filename for your media server.
- 8 Save your new OS Build Plan.

Alternatively:

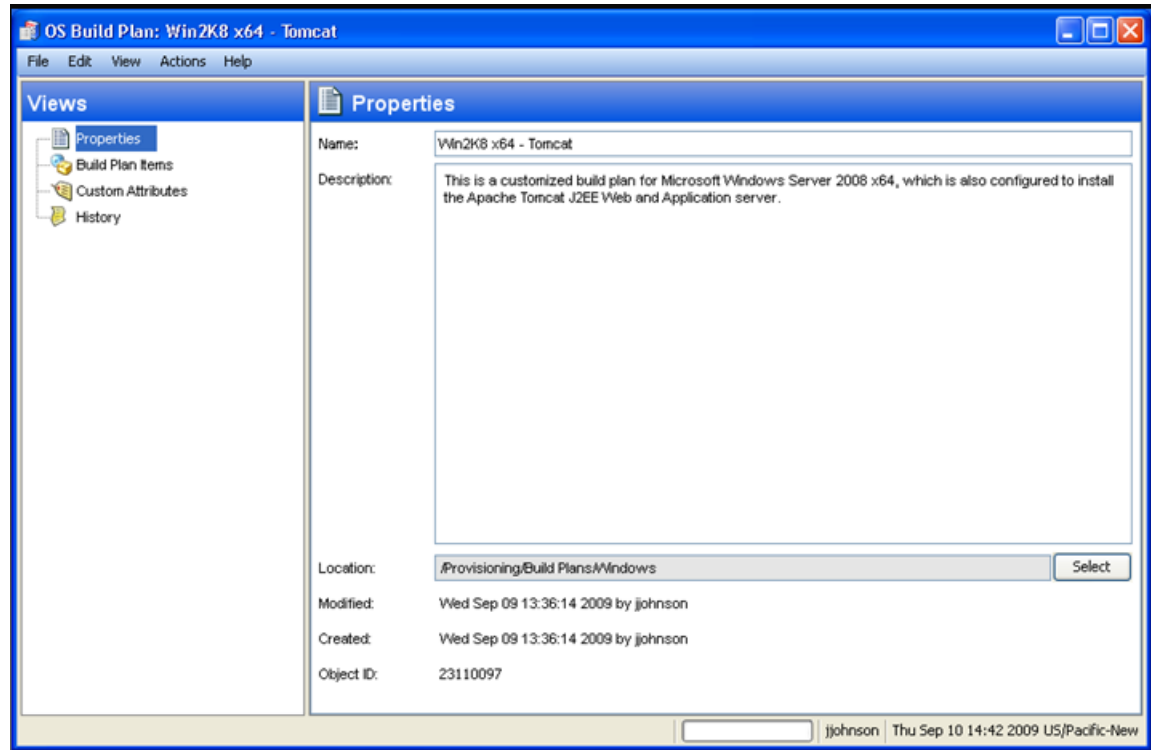
- 1 Log in to the SA Client.
- 2 In Navigation Plane open **Library ► Tools ► OS Provisioning ► OS Build Plans ► Windows**.
- 3 Highlight and copy the required build plan.
- 4 Paste and save the plan to a different folder and rename as required.
- 5 Open and modify the Build Plan for your environment.

Viewing/Modifying an OS Build Plans

You view and modify OS Build Plans using the SA Client.

- 1 Log in to the SA Client.
- 2 From the Navigation pane, select the By Type tab, then select **Library ► Tools ► OS Provisioning ► OS Build Plans ► Windows** or the folder in which you saved a modified Build Plan. Right click in the OS Build Plans list pane and select **Open**.
- 3 SA displays the Build Plan Properties page.

Figure 104Build Plan Properties Page




On this page, you see the OS Build Plan name and optional description. You can view other build plans using the Select button.

On this page you can also rename the current build plan or select a different plan. SA also displays information about the build plan itself including location, last modified date and the userid of the modifier, and the OS Build Plan's Object ID.

- 4 Select Build Plan Items in the View pane to display the Build Plan Items page in which you can see the tasks that have been assigned to the plan.

Figure 105Build Plan Items Page



You use the Build Plan Items page to add and organize the tasks performed by your OS Build Plan.

The Action pane displays a list of actions you can add to your build plan. Double click on an action to add it to the plan. Use the green arrow keys to move actions up and down in the build plan order. To remove an action you have added, highlight the action and click Remove.

Actions available are:

- **Add to Device Group:** attaches the server to a Device Group.
- **Attach Software Policy:** attach a policy that specifies software to be installed. See “Software Management” in this guide for more information.
- **Attach Patch Policy:** attach a policy that specifies patches to be applied to the server. See the Patch Management chapters for Windows, Solaris and Linux in this guide.
- **Install Zip Package:** specify a zip package to install and the installation path.




You can specify the installation path for the zip package in the OS Build Plan. If you add a zip package to an OS Build Plan but *do not* specify the installation path in the plan, SA uses the installation path specified in the zip package.

- **Run Script:** run OGFS scripts or Windows Visual Basic and batch scripts.
- **Remediate Server:** specify the remediation's reboot options and Error Handling option.

For each action, you must use Select pane at the bottom to select the specific script, policy, device group, and on to assign for the action. For some actions, for example Remediate Server, you may be required to specify additional configuration information. An item that requires additional configuration information is indicated by a red exclamation point superimposed on its icon.

- 5 Select Custom Attributes from the View pane to display the Custom Attributes page.

Figure 106Build Plan Custom Attributes Page



This page displays any custom attributes that have been specified for the OS Build Plan.



Currently, only the `timeout` custom attribute is supported.


Click the green plus sign to add a new custom attribute. Highlight an existing custom attribute and click the minus sign to remove it. You can also control the number of custom attributes displayed and search for specific attributes using the search box (wildcards supported).

When you select Add or select an existing build plan, you can specify or change the Name field and in the Value field, specify the values for the custom attribute. Clicking the ellipsis button causes the Custom Attribute editor to open. This simple editor allows you to more easily enter multiple custom attribute values.

The process of adding/modifying custom attributes is similar to that described in “Default Custom Attribute Values” in the *SA Policy Setter’s Guide*. See that section for more information about custom attributes.

- 6 Select History from the View pane to display the Build Plan History page, [Figure 107](#).

Figure 107Build Plan History Page



This page displays a chronological history of changes made to the OS Build Plan.

Minimum Baseline OS Build Plan Modification

HP strongly recommends that you copy a baseline OS Build Plan and use it as the basis for your own build plan by modifying it for your environment.

The following are the minimum modifications you must make to a baseline OS Build Plan in order to customize it for your environment.

- **Media Server Host:** the fully qualified path to the OS Provisioning Media Server host. The default is `MediaServer`. Specify from the OS Build Plan Items page, Install OS Media.
- **Product Key Custom Attribute:** using the Custom Attributes page shown in [Figure 106](#), you must specify the correct product key for your Windows product.
- **Path to Windows Setup:** provide the full path to the Windows installer executable program.
- **[WIM only] Path to the WIM Image:** provide the full path to the WIM image.
- Any requirements specific to your OS Provisioning environment such as scripts to be run, reboots, remediation, etc.

You can now run your OS Build Plan against unprovisioned servers. For more information, see [Using an OS Build Plan for OS Provisioning](#) on page 495.

OS Sequences

An OS Sequence defines what to install on a server, such as operating system configuration information taken from an OS Installation Profile that you specify, software and patch policies, and the target servers on which to install the operating system.



When you create an OS Sequence, it is saved into the Folder list in the Library. You must have permissions to the folder where you want to save the OS Sequence. For more information on how folder permissions work, see User and Group Setup in the *SA Administration Guide*.

OS Sequence Contents

You can specify the following in an OS Sequence:

- **Properties:** Allows you to name the OS Sequence and choose a location to save it in a library folder. You must have permissions to write to the folder where you save the OS Sequence, otherwise you will be unable to save it in the selected location in the library.

Install OS: Allows you to choose an OS Installation Profile. If the OS Installation Profile already has a customer associated with it, you will be unable to select a customer for the OS Sequence. If it does not have a customer associated with it, then you can select one here. Once you choose a customer, then all servers on which you install the operating system using this OS Sequence will be associated with that customer.

Attach Patch Policies is available for Windows and Solaris OS Sequences.

For more information Chapter 6, “Patch Management for Windows” or Chapter 7, “Patch Management for Solaris”.

- **Attach Device Group:** Allows you to select a device group (group of servers) for a the server once the OS Sequence has been run. You can select any public static group to attach to the OS Sequence.

A group of servers can also have software and patch policies associated with it. If you enable remediation in the OS Sequence (in Remediate Policies), then all software and patches associated with the group of servers will also be installed on the server when you run the OS Sequence. If you disable remediation, then none of the software or patches in the policies attached to the group of servers will be installed on the server.

For information on groups of servers, see Server Management in the *SA Users Guide: Server Automation*.

- **Remediate Policies:** Allows you to choose to enable or disable remediation when the server is provisioned with the OS Sequence. The Default is **Disabled**.

When remediation is disabled, running an OS Sequence installs the operating system however no policies in the OS Sequence are remediated—that is, no software or patches in any of the policies attached to the OS Sequence are installed when the sequence is run.

If you enable remediation, then all software and patches in all policies attached to the server will be installed when the OS Sequence is run. This is also true for any policies attached to the group of servers selected for the OS Sequence. You can also set reboot and pre and post installation script options.



In order to perform OS Provisioning with remediation, you must have at minimum read access to all server module policies.

Defining an OS Sequence

To create an OS Sequence, perform the following steps:

- 1 In the SA Client, from the Navigation pane, select Library and then select OS Sequences.
- 2 Choose an OS folder.
- 3 From the **Actions** menu, select **New...**
- 4 In the Views pane of the OS Sequence window, select Properties and enter a name for the OS Sequence.
- 5 Click **Change** in the Content pane to choose a location in the folder library to save the OS Sequence. You must have permissions to write to the folder where you save the OS Sequence.
- 6 From the Views pane, click **Tasks** then **Install OS** to choose an OS Installation Profile.
- 7 If the OS Installation Profile does not have a customer associated with it, then select a customer from the Assign Customer drop-down list. If the OS Installation Profile already has a customer associated with it, you will be unable to select a customer for the OS Sequence. All servers provisioned with this OS Installation Profile will be associated with the specified customer (if a customer has been assigned).
- 8 From the Views pane, select **Attach Software Policy**.
- 9 At the bottom of the Content pane, click **Add** and select a software policy to add to the OS Sequence.
- 10 From the Views pane, select **Attach Patch Policies**.
- 11 At the bottom of the Content pane, click **Add** and select a patch policy to add to the OS Sequence.
- 12 From the Views pane, select **Attach Device Group**.
- 13 At the bottom of the Content pane, click **Add**. Select a device group to place the server into, after the OS Sequence has been run. You can only select a public static group for this option.
- 14 From the Views pane, select **Remediate Policies**.
- 15 In the Content pane, choose to enable or disable remediation when the server is provisioned with the OS Sequence. If you select Disable Remediation, then when you run the OS Sequence, the operating system will be installed but no policies in the OS Sequence will be remediated — this means that no software in any of the policies attached to the OS Sequence will be installed when the sequence is run.
- 16 If you select Enable Remediation, then you will need to configure the Rebooting and Scripts parameters. For the rebooting options, you can select one of the following:
 - **Reboot servers as dictated by properties on each installed item:** Selecting this option will allow any reboot settings to run that might be set in any software or patch policies attached to the OS Sequence.
 - **Hold all server reboots until after all items are installed:** This option will override any pre-install reboot options that might be set in any software or patch policies attached to the OS Sequence. If any post-install reboots have been set, then they will execute after the operating system has been installed.
 - **Suppress all server reboots:** This option will override reboot options set in any software or patch policies attached to the OS Sequence.

- 17 Next, in the Scripts section, select either a Pre-Install/Post-Install Script. These tabs allow you to set a pre- or post-install script to be executed before the OS Sequence has been run and after the operating system has been installed. Click **Enable Script** to enable a the script parameters.
- 18 From the Select drop-down list, select either Saved Script or Ad Hoc Script. Each script type has its own settings:

Saved Script

- **Command:** Add any commands or arguments to be executed here.
- **Script Timeout:** Enter a numerical value for the number of minutes to pass until the script will timeout.
- **User:** Enter a user name and password, or choose to run the script as Local System. (If using Unix, choose root as the user.)
- **Error:** Select if you want the OS Sequence job to stop if the script returns an error.

Ad Hoc Script

- **Type:** Choose UNIX shell for Unix systems, or for Windows, select BAT or VBSCRIPT.
 - **Script:** Enter the text of the script. An Ad-Hoc script runs only for this operation and is not saved in SA. In the Script box, enter the contents of the script.
 - **Command:** If the script requires command-line flags, enter the flags here.
 - **Script Timeout:** Enter a numerical value for the number of minutes to pass until the script will timeout.
 - **User:** Enter a user name and password, or choose to run the script as Local System account. (If using Unix, choose root as the user.)
 - **Error:** Select if you want the OS Sequence job to stop if the script returns an error.
- 19 When you have finished making your selections, from the **File** menu, select **Save** to save the OS Sequence.

OS Sequence Prerequisites

Firewall Considerations

The following operating systems come with default firewall settings that must be modified during the operating system installation process in order to allow the SA Server Agent to be properly installed and configured on the target server.

- VMware ESX Server 3.0
- Windows Server 2008, Windows Server 2003 x64 and Window Server 2003 R2
- Windows XP SP2

OS Provisioning makes minor modifications to the firewall configurations on the managed server such that communication between the SA core and the Server Agent is not blocked.

VMware ESX 3.0 Firewall Settings

VMware ESX 3.0 ships by default with an IPTABLES firewall that will block communication between the core and the mini-agent or agent. In order for communications to and from the SA core to succeed, rules are added to the VMware ESX firewall by the build scripts and the SA Server Agent.

Windows Server 2008, Windows Server 2003 SP1 and Windows XP SP2 Firewall Settings

For Windows Server 2008, Windows 2003 SP1 and Windows XP SP2, in order for OS Provisioning and ongoing management to succeed, SA must ensure that the Windows firewall settings are configured to bypass the default “Security Out Of the Box” experience and allow communication over the SA ports. Thus the OS Provisioning process updates the Windows Firewall settings in the `unattend.txt`, `unattend.xml`, or `sysprep.inf` answer file as necessary for provisioning and management to work.

OS Provisioning looks for the following Windows Firewall configurations in `unattend.txt`, `unattend.xml`, or `sysprep.inf`:

- There is no Windows firewall configuration.
- There is a Windows firewall configuration, but it does not allow the ports needed by SA.
- There is a Windows firewall configuration that does allow the ports needed by SA (no changes will be made).

In any of the cases, after running an OS Sequence and installing the operating system (and agent), any predefined firewall settings remain in tact, with the exception that the SA Server Agent will have been installed and all of its required ports will have been opened.

Red Hat EL 4 and Red Hat EL 5 Firewall Settings

For Red Hat EL 4 and Red Hat EL 5, the following line in your `ks.cfg` profile will enable the firewall and allow the Server Agent to function correctly:

```
firewall --enabled --port 1002:tcp,1002:udp,1001:tcp,1023:tcp
```

Suse Linux Enterprise Server Firewall Settings

For Suse Linux Enterprise Server 9 and 10, the following lines in your `autoyast.xml` profile will enable the firewall and allow the SA Agents to function correctly.

Suse Linux Enterprise Server 9

```
<firewall>
  <fw_allow_fw_broadcast_dmz>no</fw_allow_fw_broadcast_dmz>
  <fw_allow_fw_broadcast_ext>no</fw_allow_fw_broadcast_ext>
  <fw_allow_fw_broadcast_int>no</fw_allow_fw_broadcast_int>
  <fw_dev_dmz></fw_dev_dmz>
  <fw_dev_ext>auto</fw_dev_ext>
  <fw_dev_int></fw_dev_int>
  <fw_ipsec_trust>no</fw_ipsec_trust>
  <fw_log_accept_all>no</fw_log_accept_all>
  <fw_log_accept_crit>yes</fw_log_accept_crit>
  <fw_log_drop_all>no</fw_log_drop_all>
  <fw_log_drop_crit>yes</fw_log_drop_crit>
  <fw_masq_nets></fw_masq_nets>
  <fw_masquerade>no</fw_masquerade>
```



```

<fw_protect_from_internal>yes</fw_protect_from_internal>
<fw_route>no</fw_route>
<fw_services_dmz_ip></fw_services_dmz_ip>
<fw_services_dmz_tcp></fw_services_dmz_tcp>
<fw_services_dmz_udp></fw_services_dmz_udp>
<fw_services_ext_ip></fw_services_ext_ip>
<fw_services_ext_tcp>1001 1002 1023</fw_services_ext_tcp>
<fw_services_ext_udp>1002</fw_services_ext_udp>
<fw_services_int_ip></fw_services_int_ip>
<fw_services_int_tcp></fw_services_int_tcp>
<fw_services_int_udp></fw_services_int_udp>
<enable_firewall config:type="boolean">true</enable_firewall>
<start_firewall config:type="boolean">true</start_firewall>
</firewall>

```

Suse Linux Enterprise Server 10

```

<firewall>
  <FW_ALLOW_FW_BROADCAST_DMZ>no</FW_ALLOW_FW_BROADCAST_DMZ>
  <FW_ALLOW_FW_BROADCAST_EXT>no</FW_ALLOW_FW_BROADCAST_EXT>
  <FW_ALLOW_FW_BROADCAST_INT>no</FW_ALLOW_FW_BROADCAST_INT>
  <FW_DEV_DMZ></FW_DEV_DMZ>
  <FW_DEV_INT></FW_DEV_INT>
  <FW_FORWARD_ALWAYS_INOUT_DEV></FW_FORWARD_ALWAYS_INOUT_DEV>
  <FW_FORWARD_MASQ></FW_FORWARD_MASQ>
  <FW_IGNORE_FW_BROADCAST_DMZ>no</FW_IGNORE_FW_BROADCAST_DMZ>
  <FW_IGNORE_FW_BROADCAST_EXT>yes</FW_IGNORE_FW_BROADCAST_EXT>
  <FW_IGNORE_FW_BROADCAST_INT>no</FW_IGNORE_FW_BROADCAST_INT>
  <FW_IPSEC_TRUST>no</FW_IPSEC_TRUST>
  <FW_LOG_ACCEPT_ALL>no</FW_LOG_ACCEPT_ALL>
  <FW_LOG_ACCEPT_CRIT>yes</FW_LOG_ACCEPT_CRIT>
  <FW_LOG_DROP_ALL>no</FW_LOG_DROP_ALL>
  <FW_LOG_DROP_CRIT>yes</FW_LOG_DROP_CRIT>
  <FW_MASQUERADE>no</FW_MASQUERADE>
  <FW_PROTECT_FROM_INT>no</FW_PROTECT_FROM_INT>
  <FW_ROUTE>no</FW_ROUTE>
  <FW_SERVICES_DMZ_IP></FW_SERVICES_DMZ_IP>
  <FW_SERVICES_DMZ_RPC></FW_SERVICES_DMZ_RPC>
  <FW_SERVICES_DMZ_TCP></FW_SERVICES_DMZ_TCP>
  <FW_SERVICES_DMZ_UDP></FW_SERVICES_DMZ_UDP>
  <FW_SERVICES_EXT_IP></FW_SERVICES_EXT_IP>
  <FW_SERVICES_EXT_RPC></FW_SERVICES_EXT_RPC>
  <FW_SERVICES_EXT_TCP>1001 1002 1023</FW_SERVICES_EXT_TCP>
  <FW_SERVICES_EXT_UDP>1002</FW_SERVICES_EXT_UDP>
  <FW_SERVICES_INT_IP></FW_SERVICES_INT_IP>
  <FW_SERVICES_INT_RPC></FW_SERVICES_INT_RPC>
  <FW_SERVICES_INT_TCP></FW_SERVICES_INT_TCP>
  <FW_SERVICES_INT_UDP></FW_SERVICES_INT_UDP>
  <enable_firewall config:type="boolean">true</enable_firewall>
  <start_firewall config:type="boolean">true</start_firewall>
</firewall>

```

Installing (Provisioning) an Operating System

This section describes how to install an operating system on an unprovisioned server using the SA Client, also known as *provisioning* a server.



Before you can perform the tasks described in this section, you or your SA Administrator must have set up your system for OS Provisioning as described in the *SA Policy Setters Guide*.


In order to install an operating system using the SA Client, you must create, define, and run an *OS Sequence* or an *OS Build Plan*. An OS Sequence defines what to install on an unprovisioned server, including operating system build information from an OS Installation Profile, selected software and patch policies, and remediation settings. An OS Build Plan provides the same functionality as well as extended capabilities. For information about creating OS Build Plans, see [OS Build Plans](#) on page 480. For information about creating an OS Sequence, see [OS Sequences](#) on page 489.

Your SA Administrator must have created the required OS Installation Profile(s) and granted you the required privileges for accessing OS Installation Profile(s) and for OS Provisioning.

The Unprovisioned Servers List

To provision a server and install an operating system, select an unprovisioned server from the Unprovisioned Servers list in the SA Client. Servers in the Unprovisioned Servers list have registered their presence, but do not have an operating system installed. From this location, you can install an operating system by selecting an unprovisioned server. See [Figure 108](#).


Figure 108SA Client Unprovisioned Servers List



Select an unprovisioned server in the list and the Content pane will display detailed information about the unprovisioned server that was gathered by the OS Build Agent after a network boot.

The View drop-down list enables you to view the server in the following ways:

- **Summary:** Provides information about the host name set by booting the server the first time over the network or by using an SA Boot CD. It also displays the operating system of the loaded OS Build Agent (Windows, Red Hat Linux, or Solaris), processor type, manufacturer and model of the server, and SA registration information.
- **Properties:** Displays placeholders for various management and reported information which will be filled in later once the server is provisioned.
- **Hardware:** Displays details about the hardware on the server, such as a processor type, physical and virtual memory, storage and network interfaces.
- **Custom Attributes:** Allows you to read and manage custom attributes.
- **History:** Indicates the first event associated with the server.


You can also search for an unprovisioned server using the search tool  in the upper right corner of the Content pane. You can choose a filter, then enter text to search for the server.



You also have the option of running an OS Sequence from the Library and then selecting a server or servers as you configure the Run OS Sequence window.



Some servers in the Unprovisioned Servers list are in a server lifecycle state called *Planned*, which means the server has been partially prepared for OS Provisioning. (It has a device record created for it, but no SA OS Build Agent installed yet.) An OS Sequence can't be run on Servers in the Planned state.

To display the server lifecycle stage value in the Unprovisioned Servers list, in the upper right corner of the Content pane, select the column selector  and from the list select Lifecycle. For more information, see your SA administrator.

The OS Installation Profile

An OS Installation Profile defines all necessary parameters of an operating system, including the operating system type and version, the OS Media Resource Locator (MRL), configuration or response file(s), build customization scripts, and packages required for operating system installation.

OS Installation Profiles are typically defined by an SA Administrator and made available to authorized users for application against servers in SA server pools.

For more information about defining OS Installation Profiles, see the *SA Policy Setter's Guide*.

Using an OS Build Plan for OS Provisioning

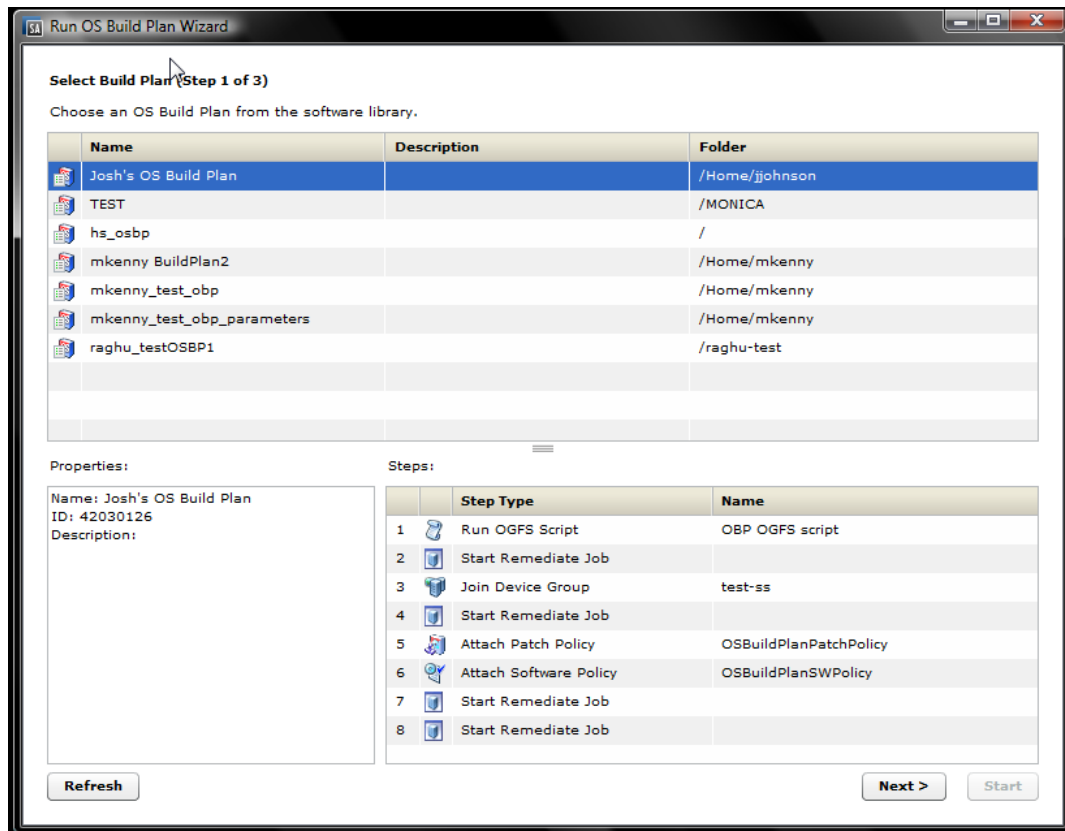
For information about creating OS Build Plans, see [OS Build Plans](#) on page 480.

To install an operating system on an unprovisioned server using an OS Build Plan, perform the following tasks:

- 1 Log on to the SA Client specifying the SA Core that manages the server you will install the operating system on.
- 2 In the Navigation pane, select Devices and select Unprovisioned Servers.

- 3 Select a server from the list of available unprovisioned servers list
- 4 Right-click on a server and select **Run OS Build Plan**. Alternatively, you can select **Action Menu ► Run ► OS Build Plan** and choose a target server in the Run OS Sequence window or use the search pane to search for a list of Build Plans.
- 5 SA displays the first page of three of the Run OS Build Plan Wizard.

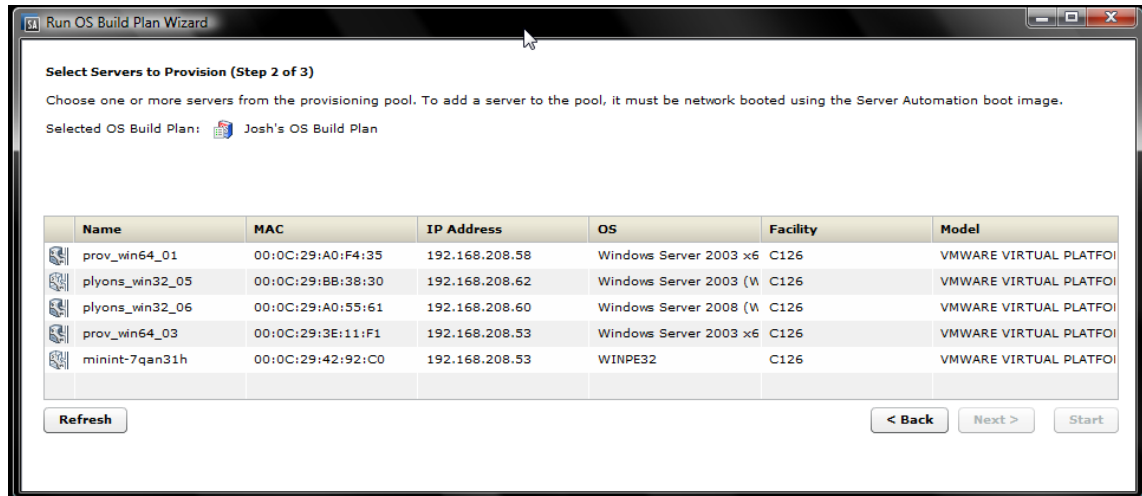
Figure 109The Run OS Build Plan Wizard - Page 1



On this page you select the OS Build Plan to run. The Properties pane provides a description of the Build Plan, the Steps pane displays the tasks assigned to the Build Plan that will be performed during OS Provisioning.

- 6 Select the OS Build Plan to Run and click Next.
- 7 The second page of the Run OS Build Plan Wizard displays.

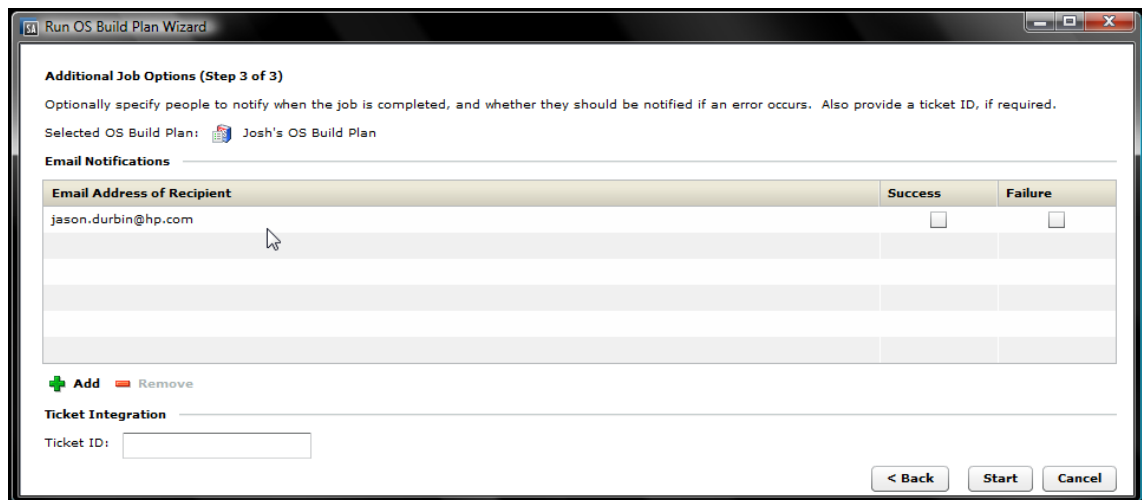
Figure 110The Run OS Build Plan Wizard - Page 2



On this page you select the server(s) you want to run the OS Build Plan on. If you select a server that is already under SA management, SA displays a warning message that all data on that server will be erased if the Build Plan is run against it.

- 8 Select the server(s) to run the OS Build Plan on and click Next.
- 9 The third page of the Run OS Build Plan Wizard displays.

Figure 111The Run OS Build Plan Wizard - Page 3



On this page you select email notification options and, if required, specify a Job Ticket ID. Click the green plus icon to add a notification and specify the email address to be notified in case of job failure, success, or both. Highlight an email address and click the red minus icon to remove it. To specify a Job Ticket ID, specify the ID in the Ticket ID field.

- 10 Click Start to begin running the OS Build Plan Job. When the job begins to run, click on the Job in the Job Status window or click **Close** to exit the Job Status window. You can also check the status of the Job by clicking on Job Logs under Jobs and Sessions in Navigation Pane.
- 11 When the OS Build Plan job completes successfully, you can check the **Devices ► All Managed Servers** list to see the newly provisioned server.

Using an OS Sequence for OS Provisioning

For information about creating an OS Sequence, see [OS Sequences](#) on page 489.

To install an operating system using an OS Sequence, perform the following tasks:

- 1 Log on to the SA Client specifying the SA Core that manages the server on which you are installing an operating system.
- 2 In the SA Client Navigation pane, select **Devices > Unprovisioned Servers**.
- 3 Select a server from the list of available unprovisioned servers list
- 4 Right-click on a server and select **Run OS Sequence**.

or


Select **Action Menu > Run > OS Sequence** and choose a target server in the Run OS Sequence window.



or

From the Navigation pane, select **Library > OS Sequences**. Select a folder that contains existing OS Sequences for the correct operating system then right-click an available OS Sequence and select **Run...**



If the **Run OS Sequence** menu item is grayed out, one or more of the unprovisioned servers is in a server lifecycle stage of **Planned**. Servers in this stage cannot be provisioned. You can display the server lifecycle stage value in the Unprovisioned Servers list.

In the upper right corner of the Content pane, select the column selector . From the list, select **Lifecycle**. For more information, see your SA administrator.

- 5 In the Select OS Sequence pane, click **Add** to add an OS Sequence or click **Next** if OS Sequence is already listed.
- 6 In the Run OS Sequence window, step one requires that you add an unprovisioned server or servers to provision. To add a server, click **Add**.
- 7 Click **Next**, and in the Scheduling pane choose if you want to run the OS Sequence, immediately, or at a later date and time.
- 8 Click **Next** and in the Notifications pane, select an email notifier. Click **Add Notifier** and enter an email address.
- 9 You can specify if you want the email to be sent upon success of the OS Sequence job () or failure of the OS Sequence job ()
- 10 The ticket ID field is only used when Professional Services has integrated SA with your change control systems. It should be left blank otherwise.
- 11 Click **Next**, and review the OS Sequence information before you run the job.
- 12 Click **Start Job** to run the OS Sequence. When the OS Sequence job begins to run, click on the Job in the Job Status window or click **Close** to exit the Job Status window. You can also check the status of the Job by clicking on Job Logs under Jobs and Sessions in Navigation Pane.
- 13 When the OS Sequence job has completed successfully, you can check the **Devices > All Managed Servers** list to see the newly provisioned server.



If you scheduled the OS Sequence job to run at a later date and would like to cancel it, from the Navigation pane, select **Jobs and Sessions ► Recurring Schedules**. Then, select the job, right-click and select **Stop**.



If an OS Sequence does not have remediation enabled, a newly provisioned servers will not immediately perform a full software registration. Full software registration occurs after a small variable delay, usually less than one hour. Thus when provisioning without remediation, the server's installed software packages and patches might not be listed immediately.

Model Base Packages Functionality

OS Provisioning provides the ability to create software policies that model the base set of packages installed during OS Provisioning.

During OS Provisioning — after the base operating system installation, agent installation, and reachability test, but before reconcile/remediate — a new script triggers software registration on the newly provisioned server, then models the installed packages as a software policy.

To activate this functionality, the server being provisioned must have a custom attribute defined (or inherited) named `model_base_packages`. The value for this attribute must either be empty or an absolute folder path to the name of the software policy to be created (or updated) with the package list.

If the `model_base_packages` value is empty, a software policy is created (or updated if it already exists) in the same folder as the OS Sequence. The software policy name will be the OS Sequence name plus `Base Packages`.

Each installed package that is successfully found in SA is added to the list of software policy items. A list of package names and versions that were not found in SA will be available as a custom attribute named `missing_packages` in the software policy. This policy is attached to the OS Sequence which has remediation enabled. Because the above occurs before remediation, this policy is included in the remediation, thus adopting the modeled packages since they are by definition already installed.

You should only specify the `model_base_packages` custom attribute value as empty when running OS Sequences from the SA Client. When running OS Provisioning from the SAS Web Client, the `model_base_packages` custom attribute value must be the path to the Software Policy.

The only valid value for the `model_base_packages` custom attribute is the path to a Software Policy. For example:

```
/Customer/OS Baselines/Solaris 10 baseline Q4 2007
```

In this case, the Software Policy will be created at the specified path and with the specified name. Any folders that are missing will automatically be created. If the Software Policy already exists, it will be updated.



When run from the SAS Web Client Install OS wizard, the Software Policy will be attached to the server being provisioned. However, since the Install OS wizard triggers a legacy reconcile, remediate is bypassed so the policy will not be remediated.

Note that it is not necessary to use the Model Base Packages feature for every OS Provisioning job. It needs only to be used once after an OS Installation Profile changes. From that point on, the Software Policy will be attached to the OS Sequence unless you remove it, and will be available for other servers as they are provisioned.

Model Base Packages Script Usage

The `model_base_packages.py` Command Engine script will function when called from another Command Engine script such as `provisionOS.py`. You can also run it as a standalone `python2` `pytwist` script. The following are valid arguments when invoking the script:

```
model_base_packages.py --opsware-username you [--opsware-password yourpass]
--server <serverID> --ossequence <ossequenceID> [--policy_path "/Some/Folder
Path/Some Policy"]:
```

Table 37 Options

Argument	Description
--version	Show the program version number and exit
-h, --help	Show this help message and exit
-u OPSWAREUSERNAME, --opsware-username=OPSWAREUSERNAME	Login username for SA
-p OPSWAREPASSWORD, --opsware-password=OPSWAREPASSWORD	Login password for SA
-s SERVER, --server=SERVER	Numeric Server ID of server to model
m POLICYPATH, --policy_path=POLICYPATH	Absolute path to the software policy that will model the packages
-e OSSEQUENCE, --ossequence=OSSEQUENCE	Numeric OS Sequence ID to link to the model software policy. If you specify an OS Sequence but not a policy path, the software policy will be created in the folder that contains the OS Sequence with the OS Sequence's name plus "Base Packages".

Reprovisioning a Managed Server

You can reprovision a managed server, but this process completely removes all data on the server as well as any network configuration settings.

Certain attributes, defined in the build script for each operating system, are preserved after you re-provision the server. For more information on OS Provisioning build scripts, see OS Provisioning Setup in the *SA Administration Guide*.



You can only re-provision a server that runs the Solaris or Linux operating system (but not Solaris x86).



For Linux re-provision, you can use the custom attribute `boot_kernel` to determine which kernel that re-provision will boot to. For more information, see the *SA Policy Setter's Guide*.

To re-provision a managed server, perform the following steps:

- 1 From the Navigation pane, select **Devices** ► **All Managed Servers**.
- 2 Select a managed server to re-provision and from the **Actions** menu, select **Run OS Sequence**.
- 3 You will be shown a warning message that you are about to re-provision a managed server. By doing so, you will lose all data on the server. Click **Yes** to proceed.
- 4 In the Run OS Sequence window, please select the appropriate option before you begin the re-provisioning:

Yes, I understand the OS installation process will erase all data on the selected servers. (Mandatory. You must select this option in order to proceed.)
- 5 Click **Next**. In the Run OS Sequence window, select an unprovisioned server or servers to provision. To add a server, click **Add**.
- 6 Click **Next**. In the Select OS Sequence pane, click **Add** to add an OS Sequence.
- 7 Click **Next**, and in the Scheduling pane, choose if you want to run the OS Sequence, immediately, or at a later date and time.
- 8 Click **Next** and in the Notifications pane, select an email notifier. Click **Add Notifier** and enter an email address.
- 9 (Optional) Specify if you want the email to be sent upon the success of the OS Sequence job or failure of the OS Sequence job.
- 10 You can also specify a Ticket Tracking ID in the Ticket ID field.
- 11 Click **Next**, and review the OS Sequence information before you run the job.
- 12 Click **Start Job** to run the OS Sequence. When the OS Sequence has run, click **View Results** to view the results of the OS Sequence job.
- 13 When the OS Sequence job has been run, you can check the Devices ► All Managed Servers list to see the newly re-provisioned server.

Advanced SA OS Provisioning Architecture

OS Provisioning Components

OS Provisioning comprises several components with distinct functions:

- The OS Build Agent
- OGFS Agent (a specialized SA Server Agent)
- The Build Manager
- Build Scripts
- The Media Server
- The Boot Server

The OS Build Agent

Similar to the SA Server Agent, the OS Build Agent is a simplified agent whose function is to run commands as instructed by the Build Manager. Newly registered OS Build Agents appear in the Server Pool.

Booting a new server for the first time loads an OS Build Agent on the server; however, the server does not have the target operating system installed and might not have access to disk resources. SA can still communicate with the server and perform commands on it remotely because the OS Build Agent is running an operating system that is loaded into memory.

The OS Build Agent performs the following functions:

- Registers the server with SA when the OS Build Agent starts.
- Listens for command requests from SA and performs them.
- Performs commands even though a target operating system is not installed.

OGFS Agent

The OGFS Agent, a specialized SA Server Agent, is part of the SA-supplied operating system boot image. When an unprovisioned server is PXE booted, the server is registered with the SA Core, an agent certificate is obtained from the core, and the Server Agent is started. Unlike a non-OGFS boot image, the Build Agent is not required and uses the Core's OGFS functionality to complete the Agent tasks.

The Build Manager

The build manager performs several functions:

- Manages newly registered OS Build Agents.
- Coordinates scripts that gather hardware inventory from OS Build Agents.
- Coordinates the scripts that perform the operating system installation with the OS Build Agent.
- Communicates with the OS Build Agents using a simple protocol.

The Media Server

The Media Server is installed as part of a typical SA Core installation when you specify that you want to install the OS Provisioning components. In order to provision operating systems, you must first upload a valid copy of the operating system's installation media to the Media Server. During OS Provisioning, SA will use the copy of the operating system installation media on the Media Server to do the provisioning.

SA provides file servers that can share operating system media using NFS and Samba if you do not have existing NFS/Samba servers that you want to use or are not familiar with configuring these servers.

The Boot Server

The Boot Server listens for broadcast requests from new servers in the server pool and responds using DHCP. Network booting requires DHCP/BOOTP, TFTP, and PXE (x86).

Build Customization Scripts

OS Provisioning build customization scripts provide hooks into the build process that allow you to modify operating system installations at specific points. These hooks call a single build customization script at the appropriate time in the operating system installation process.

Because each build customization script is specific to the operating system it installs, build customization and installation vary by operating system. Before you can use a build customization script as part of an operating system installation profile, you need to create the build customization script and import it into the SA Client.

Bootimg and Provisioning Using WinPE-OGFS Boot Images

In the WinPE x86 32-bit and WinPE x86 64-bit environments, an OS Provisioning Build Agent is loaded onto the server to be provisioned. In the WinPE-OGFS environment, the agent is a full SA Server Agent that is booted when an OGFS-enabled image is selected and is specialized for OGFS functionality including running OS Build Plans, disk repair, file restoration, and other non-destructive maintenance.

See the *SA Users Guide: Server Automation* for details about using OGFS commands.

How the OS Build Agent Locates the Build Manager

How the OS Build Agent locates the Build Manager depends on the boot method.

WinPE

- The Build Manager is located by loading the configuration file
`/opt/opsware/boot/tftpboot/DHCPOptions.ini`
which contains the OS Provisioning settings specified during SA installation.
- If the process above fails, SA retrieves DHCP options containing the agent gateway IP address and Build Manager port.
- If the processes above fail, SA defaults to the hostname `buildmgr` on port 8017.

Linux x86:

Linux x86 locates the Build Manager using kernel arguments supplied at PXE boot time. These are configured during the SA installation and stored in the file

`/opt/opsware/boot/tftpboot/pxelinux.cfg/default`

Linux IA64:

Linux IA64 locates the Build Manager using kernel arguments supplied at PXE boot time. These are configured during the SA installer and stored in the file

```
/opt/opsware/boot/tftpboot/elilo.conf
```

Solaris

For Solaris OS Provisioning, the JumpStart build script runs the OS Build Agent, which contacts the Build Manager (via the Agent Gateway in the core). The Solaris `begin` script attempts to locate the Build Manager in the following ways:

- By using information that the SA DHCP server provided
- By looking for the host name `buildmgr` in DNS as configured by the DHCP server

You can override the way that the OS Build Agent contacts the Build Manager by specifying a boot argument at the prompt when you boot a new Solaris server, for example:

```
ok boot net:dhcp - install buildmgr=buildmgr.example.com:8017
```

```
ok boot net:dhcp - install buildmgr=192.168.1.15:8017
```

Non-DHCP Environments

In both Windows and Red Hat non-DHCP environments, SA locates the Build Manager using the network configuration specifications you provide. See [Booting a Red Hat Enterprise Linux Server in a Non-DHCP Environment](#) on page 466 and [Booting a Windows Server in a Non-DHCP Environment](#) on page 472.

Loading OS Build Agents

You can load an OS Build Agent on a server by booting the server with PXE for Intel-Based machine, Elilo for Itanium-Based machine or by using the network (Solaris). After a successful installation, the server appears in the Server Pool list.

You should verify that the newly racked server shows up in the SA Client **Unprovisioned Servers** list, or SAS Web Client **Server Pool**, and is ready to hand off for operating system installation.

The SA Client's **Unprovisioned Servers** list and the SAS Web Client **Server Pool** list display the servers that have registered their existence with SA but do not yet have an operating system installed.

You can start the operating system installation process in either one of two ways:

- From the SA Client's **Unprovisioned Servers** list, right click on the server in the content pane, and choose Run OS Sequence. Please See “Installing (Provisioning) an Operating System” on page 494 for details.


From the SAS Web Client **Server Pool**, select the server and click **Install OS**. This option is only available for SA version 6.1 cores and later.

Verifying That a Server is Ready for Operating System Installation

Perform the following steps to verify that a server is ready for operating system installation:

- 1 Log into the SA Client.
- 2 Select Devices in the Navigation pane, expand Servers and select All Managed Servers. The Managed Server page appears, as shown in [Figure 112](#).

Figure 112All Managed Server List in the SA Client



- 3 (Optional) From the drop-down lists, select the manufacturer, model, or facility of the server and click **Update**.
- 4 For Intel x86 and Sun SPARC processor-based servers, locate the MAC address and Host ID of the server that you just booted.

The Lifecycle column indicates whether the server is available for OS Provisioning or if the boot into the server pool failed.

See [Affect of the OGFS Agent on Server Lifecycle](#) on page 481 for more information.

To obtain more information about the server, double click the server name.

Recovering when an OS Build Agent Fails to Load

When an OS Build Agent fails to load on a server, the server does not appear in the Managed Server list.

You can check the server console for error messages and try to boot the server again with PXE or by using the SA Boot CD.

If all errors were successfully resolved, the initial boot occurs, the OS Build Agent is loaded on the server, the server appears in the Managed Server list, and the Lifecycle column indicates that the server is available.

If you are unable to resolve the error condition, contact your SA administrator for troubleshooting assistance.

Index

A

adding

- rule exceptions to an audit, 164

ad hoc audit, running, 173

AIX

APARs

- about, 401, 402

- uploading, 401, 402

- LPPs, about, 401

APARs. *See* AIX APARs.

application and storage signatures (SAV)

- about, 80

application configuration

- Compliance Dashboard, 244

Application Configuration Management

- configuring audit rules, 127

application signature (SAV)

- defined, 27

- evaluation order, 81

attaching

- software policy to server, 268

audit

- auditing process, 111

- audit results

 - viewing and remediating, 186

- audits and the Compliance Dashboard, 107

- configuring, overview, 117

- elements of, 112

- performing

 - audit results, from, 173

- re-running from audit results, 173

- results, value based remediation, 184

- running from the library, 172

- running on a server, 173

- saving as audit policy, 170

- scheduling, 174

- scheduling recurring, 174

- searching for, 189, 203

- selection criteria

 - inclusions/exclusions, 157

- snapshot used in, 190

- sources, server or snapshot, 118

- viewing completed audit job, 176

- ways to create, 113

 - creating from a server, 114

 - from a group of servers, 114

 - from an audit policy, 115

 - from a snapshot, 114

 - from the library, 114

Audit and Remediation

- audit policies, 166

- audit process overview, 111

- audit results, 176

- capturing golden server configuration, 106

- creating an audit policy, 167

- deleting

 - snapshot specification, 203

- examples (use cases), 105

- exceptions, 164

 - adding to an audit, 164

 - editing, 165

 - rules that cannot have exceptions, 164

- linking and importing audit policies, 168

- overview, 105

- performing audit

 - audit results, from, 173

- rules

 - configuring, application configuration, 126

 - configuring, COM+, 131

 - configuring, compliance checks, 151

 - configuring, custom script, 132

 - configuring, file system, 135

 - configuring, hardware, 138

 - configuring, IIS Metabase, 139

 - configuring, operating system, 145

 - configuring, users and groups, 147

 - configuring, Windows Registry, 148

 - configuring, Windows services, 149

 - server objects, 121

- scheduling audits, 174

- selection criteria

 - inclusions/exclusions, 157

- terms and concepts, 108

- viewing

 - and remediating audit results, 186

- ways to create an audit, 113

audit policies

- creating, 167

- linking and importing, 168

- locating in the SA Client library, 171
 - overview, 166
- audit policy
 - exporting to HTML or CSV, 171
 - saving, 170
- B**
- booting
 - Solaris servers, over network, 474
- Build Manager
 - OS Build Agents, locating, 503
- C**
- COM+ object
 - configuring Audit and Remediation rule, 131
- comparing scan results (SAV)
 - heuristics used, 94
- comparing snapshots (SAV)
 - "significant" object attribute differences, 92
 - how to, 93
 - object attribute differences, 91
 - source and comparison, 90
- comparing snapshots (VAM)
 - object existence comparison, 91
- compliance
 - performing software compliance scan, 282, 393
 - software, 281
 - viewing server compliance in SAV, 62
- compliance checks
 - configuring Audit and Remediation rule, 151
 - creating custom categories, 155
 - restoring to defaults, 156
- compliance checks, editing properties, 154
- compliance checks, managing, 154
- Compliance Dashboard
 - application configuration, 244
 - audit, 234
 - compliance statuses, 211
 - general categories, 231
 - patch, 241
 - refreshing, 227
 - remediation overview, 231
 - software compliance, 237
 - terms and concepts, 216
 - viewing, 216
- compliance status for Compliance Dashboard, 211
- Compliance View
 - overview, 207
 - usage, proactive or reactive, 208
- configuring

- snapshot specification, 195
- copying
 - objects to a server from a snapshot, 205
- creating
 - application or storage signature (SAV), 83
 - application tier (SAV), 79
 - audit policy, 167
 - business application contacts in SAV, 78
 - custom compliance checks categories, 155
 - OS sequence, 489
 - scripts, 426
 - snapshot in SAV, 89
 - snapshot specification, 194
 - snapshot specification from library, 194
- custom script
 - configuring Audit and Remediation rule, 132
- D**
- deleting
 - scripts, 432
 - snapshot, 189, 204
 - snapshot job schedule, 200
 - snapshot specification, 203
- depots
 - patch management, 401
- detaching
 - software policy to server, 276
- device tree (SAV), 28
- DHCP
 - servers, booting with, 463
 - Solaris servers, usage of, 462, 464
- E**
- editing
 - audit rule exceptions, 165
 - audit schedule, 175
 - compliance checks properties, 154
 - snapshot job schedule, 199
- e-mail notifications, 317, 318, 339, 346, 415, 421
- error messages (SAV), 98
- evaluation order, SAV application signatures, 81
- Exceptions, Audit and Remediation
 - about, 164
 - adding to an audit, 164
 - considerations, 164
 - rules that cannot have exceptions, 164
- executing
 - OGFS script, 438
 - scripts, 433
 - server script, 434

executing scripts, 433

exporting

audit policy, 171

SAV properties to .csv, 57

scripts, 432

exporting a SAV map, 55

F

File System

configuration Audit and Remediation rule, 135

font, 333

G

global shell

opening from SAV, 75

H

hardware,configuring Audit & Remediation rule, 138

hardware preparation, overview, 461

hotfix chaining, 296

HP-UX

depots

patch management, 401

I

icons, toolbar icons (SAV), 36

IIS Metabase

configuring Audit and Remediation rule, 139

importing

audit policy rules, 170

installation

flags, overview, 334, 342, 410

installing

Install Patch Wizard, 411, 417

OS Build Agents

verification, 505

software using software policy, 267, 276

Install Patch Wizard, 411, 417

Install Patch wizard, 333, 410

install scripts,specifying, 315, 420

ISM Control

running, 278

J

Japanese, 331

K

Korean, 331

L

layer 2 connections displayed in SAV, 47

linking an audit policy to audit or snapshot
specification, 168

locales, 331

locating audit policies in the SA Client library, 171

M

Managing

compliance checks, 154

Microsoft patch management prerequisites, 290

Model Repository, 294, 310, 311

msiexec.exe, 295

N

network

Solaris servers, booting over, 474

O

opening

Device Explorer from SAV, 75

global shell from SAV, 75

remote terminal from SAV, 75

scripts, 429

opening a Business Application, 85

operating systems

configuring Audit and Remediation rule, 145

patch management, supported for, 400

provisioning, 459

OS Build Agents

Build Manager, locating, 503

failure to install, recovering from, 505

verifying installation, 505

OS provisioning

hardware preparation, 461

SA Client

creating an OS sequence, 489

overview, 494

reprovisioning a managed server, 500

select unprovisioned servers, 494

Solaris servers, 460

Windows servers, 460

OS sequence

attach device group, 489

creating, 489

- set remediate policy, 489
- overview
 - script execution, 423

P

- package types
 - AIX APAR, 401, 402
 - HP-UX depots, 401, 402
 - LPP, 401
 - RPM, 400
 - Windows Hotfix, 334
- patch
 - Compliance Dashboard, 241
- patch compliance, 310, 311, 312
- patch compliance scan, 293
- patches
 - installation flags, 334, 410
 - types supported, 400
 - uninstallation flags, 342
- patch installation, previewing, 339, 415
- patch installation, scheduling, 316, 338, 414
- patch management
 - Microsoft patch releases, 293
 - operating systems, supported, 400
 - patch information from Agent, 397
 - patch testing, support for, 397
 - roles, 399
 - supported Unix versions, 400
 - uploading automatically, 306
- patch policy, 310
- patch policy exception, 311
- patch reboot options, 315, 337, 413
- patch uninstallation, previewing, 346, 421
- patch uninstallation, scheduling, 345, 420
- performing
 - audit
 - audit results, from, 173
- policy setter, 297
- populate-opsware-update-library, 306
- printing a SAV map, 55
- process families (SAV)
 - defined, 27
 - properties of, 58
- properties
 - application signature (SAV), 60
 - storage signature (SAV), 60
- properties (SAV)
 - links, 65

- network device, 65
- network interface, 68
- port group, 67
- process family, 58
- server compliance, 62
- tiers, 59
- virtual server, 64
- virtual switch, 67
- property page (SAV)
 - port group, 67

Q

- qchain.exe, 295
- QNumber, 317

R

- refreshing
 - Compliance Dashboard, 227
- remediate
 - overview, 270
 - software policy, 271
- remote terminal
 - opening from SAV, 75
- renaming
 - scripts, 432
- reports
 - software policy, 282
- reprovisioning a managed server, SA Client, 500
- restoring compliance checks to defaults, 156
- RPM
 - patching, 400
- running
 - ad hoc audit, 173
 - audit from server, 173
 - audit from the library, 172
 - ISM Controls, 278
 - snapshot specification, 197

S

- SA Client
 - OS installation with, 494
- SAS Web Client
 - patch administration in, 409
- saving
 - audit or snapshot specification as audit policy, 170
 - snapshot specification as policy, 196
- scan
 - patch compliance, 323

- software compliance, 282, 393
- scan timeout (SAV), 74
- scheduling
 - audit, 174
 - audit, recurring, 174
 - snapshot in SAV, 90
 - snapshot job, 198
- script execution, 433
 - creating, 426
 - deleting, 432
 - editing scripts
 - editing
 - scripts, 430
 - executing OGFS script, 438
 - executing server script, 434
 - exporting, 432
 - opening a script, 429
 - overview, 423
 - process, 424
 - renaming, 432
 - types, 424
 - viewing script history, 431
- scripts
 - Distributed Scripts
 - overview, 423, 443
 - running on devices in SAV, 76
- search
 - audit, 189, 203
- sending email to Business Application contacts, 79
- server
 - attaching software policy, 268
 - detaching software policy, 276
- Server Agent, 294
 - registration, 397
- Server Map (SAV), 43
- server objects
 - Audit and Remediation, 121
- servers
 - booting
 - over network, 474
 - reprovisioning, SA Client, 500
 - Solaris servers, booting, 474
- Service Automation Visualizer (SAV)
 - accessing servers
 - opening a remote terminal, 75
 - opening Device Explorer, 75
 - opening global shell, 75
 - Business Application
 - opening, 85
 - saving, 85
 - saving as template, 86

- business application
 - application signatures evaluation order, 81
 - copying and cutting a tier, 80
 - creating a tier, 79
 - deleting a tier, 80
 - pasting a tier, 80
 - templates, 77
 - tiers, 79
- comparing snapshots
 - "significant" object attribute differences, 92
 - comparison types, 91
 - how to, 93
 - object attribute difference, 91
 - object existence comparison, 91
 - source and comparison, 90
- creating a snapshot, 89
- data collection and display, 24
- error messages, 98
- filtering data, 95
 - criteria to use, 97
 - using regular expressions, 98
- icons in toolbar, 36
- menus, 39
- opening a snapshot, 90
- overview, 19
- prerequisites to run, 20
- process and link connection symbols, 54
- properties
 - application signature, 60
 - connection links, 65
 - network devices, 65
 - network interfaces, 68
 - port group, 67
 - process families, 58
 - server compliance, 62
 - storage signature, 60
 - tiers, 59
 - virtual servers, 64
 - virtual switches, 67
- property pages
 - port groups, 67
 - server, 57
- SAV application, 24
 - application and storage signatures
 - explained, 80
 - application signature, 27
 - creating definition, 76
 - devices tree, 28
 - process families, 27
 - storage signatures, 27
 - tiers tree, 25

- SAV maps
 - exporting, 55
 - network map, 46
 - printing, 55
 - SAN map, 51
 - server map, 43
 - showing IPC service names, 56
 - storage map, 48
- scheduling a snapshot, 90
- supported operating systems, 20
- symbols used in maps, 53
- virtualization settings, 73
- showing IPC service names in SAV maps, 56
- signatures (SAV)
 - cutting and copying, 84
 - deleting, 84
 - editing, 84
 - pasting, 85
- snapshot
 - copying objects to server from, 204
 - deleting, 189, 204
 - template, 203
 - deleting job schedule, 200
 - difference between snapshot specification, 190
 - editing job schedule, 199
 - locating, 200
 - locating in SA Client, 200
 - process, 192
 - scheduling, 198
 - used in an audit, 190
 - used with audit policies, 191
 - viewing contents of, 201
- snapshot specification, 193
 - and audit policies, 191
 - configuring, 195
 - configuring rules for, 196
 - creating from library, 194
 - creating from server, 194
 - deleting, 203
 - elements of, 191
 - relationship to snapshots, 190
 - running, 197
 - selection criteria
 - inclusions/exclusions, 157
- software compliance
 - Compliance Dashboard, 237
 - compliance remediate options, 238
- software installation
 - attaching software policy, 268
 - detaching software policy, 276
 - installing using software policy, 267, 276
 - process, 261

- remediate
 - overview, 270
 - remediate software policy, 271
 - software policy template overview, 277
- software policy
 - attaching to server, 268
 - compliance overview, 281
 - detaching to server, 276
 - installing software, 267, 276
 - performing software compliance scan, 282, 393
 - remediate, 271
 - remediate overview, 270
 - reports, 282
 - running ISM controls, 278
 - software policy template overview, 277
- software policy template
 - overview, 277
- Software Repository, 294
- Solaris
 - booting servers over network, 474
 - OS provisioning, 460
- storage signature (SAV), 27
- Summary Review, 318, 340, 346, 347, 415, 421

T

- Tiers properties (SAV), 59
- troubleshooting
 - OS Build Agents
 - installation failure, 505
 - verifying installation, 505
- types of scripts, 424

U

- uninstallation
 - flags, overview, 334, 342, 410
- Uninstall Patch wizard, 341, 417
- unzip.exe, 296
- users and groups, configuring Audit and Remediation rules, 147

V

- verifying
 - installation of OS Build Agents, 505
- viewing
 - audit results, 186
 - audit server usage, 115
 - completed audit job, 176
 - Compliance Dashboard, 216
 - script history, 431

- snapshot contents, 201

Virtualization Director

- scan time out preference, 74
- virtualization settings, 73

W

Windows Hotfix

- installation flags, 334
- uploading, 334

Windows Registry

- configuring Audit and Remediation rule, 148

Windows servers

- OS provisioning, 460

Windows services

- configuring Audit and Remediation rule, 149

Windows Update Agent, 296

wizards

- Distributed Script Execution, 423
- Install Patch, 411, 417

