

HP Server Automation

Software Version: 10.23

SA Administration Guide

Document Release Date: July 2017
Software Release Date: June 2016



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 Notice

© Copyright 2001-2016 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

Documentation Updates

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

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

To check for recent updates or to verify that you are using the most recent edition of a document, go to: **<http://h20230.www2.hp.com/selfsolve/manuals>**

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to: **<http://h20229.www2.hp.com/passport-registration.html>**

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

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: **<http://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

HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is **<http://h20230.www2.hp.com/sc/solutions/index.jsp>**

Contents

Chapter 1 User and User Group Setup and Security	1
About SA Users and User Groups	1
About Permission Types - Action, Resource and Folder Permissions	2
About Action Permissions	4
Grouping Action Permissions	4
About Resource Permissions	5
Types of Access to Resources	6
About Facility Permissions	6
About Customer Permissions	6
About Device Group Permissions	6
Examples of Resource Permissions	7
Resource and Action Permissions Combined - Example	8
Other Types of Resources	9
About Folder Permissions	9
Types of Folder Permissions	9
Folder Permissions and Action Permissions	10
Folders, Customer Constraints, and Software Policies	11
Default Folder Permissions	11
Membership in Multiple User Groups	11
Restricted Views in the SA Client Based on Permissions	13
Predefined User Groups	13
About Private User Groups	15
About Super Administrators and Super Users	16
About Super Users	16
About Customer Administrators and Customer Groups	16
Comparing Customer Administrators with Super Administrators	17
A Customer Administrator is Defined by a Customer Group	17

Example Customer Group	17
Process Overview for Security Administrators	18
About Global File System Permissions	20
Managing Users - SA Client	21
Creating a New User	22
Changing a User's Permissions	23
Changing a User's Password	23
Users Changing Their Own Password and Other Properties	23
Changing a User	26
Deleting a User	26
Finding the User Group a Particular Action Permission Comes From	26
Suspending a User	27
Activating a Suspended User	27
Assigning a User to a User Group	28
Importing Users from an LDAP Directory	28
Managing User Groups - SA Client	28
Creating a New User Group	29
Viewing User Groups	29
Copying a User Group	30
Changing a User Group	30
Deleting a User Group	31
Adding a User to a User Group	32
Setting Permissions on User Groups - SA Client	32
Setting Resource Permissions - Facilities, Customers, and Device Groups	32
Setting Action Permissions	33
Setting Folder Permissions	34
Setting OGFS Permissions	35
Setting Private User Group Permissions	37
Setting Password, Account, and Session Security Policies - SA Client	37
Resetting Initial Passwords	38
Setting Password Expiration	39

Prohibiting Reuse of Old Passwords	39
Suspending User Accounts After Login Failures	39
Suspending Inactive User Accounts	40
Locking Inactive Sessions	40
Displaying a User Login Agreement	40
Displaying a Banner on the SA Client Screen	41
Managing Super Administrators - SA Client	42
Viewing All SA Super Administrators	42
Creating a Super Administrator	42
Deleting a Super Administrator	43
Managing Customer Administrators and Customer Groups - SA Client	43
Viewing All Customer Administrators	44
Viewing All Customer Administrators for a Customer Group	44
Viewing All Customers for a Customer Group	44
Creating a Customer Group	44
Deleting a Customer Group	45
Creating a Customer Administrator from the Customer Group View	45
Creating a Customer Administrator from the User View	46
Deleting a Customer Administrator from the Customer Group View	46
Deleting a Customer Administrator from the User View	47
Specifying Password Character Requirements	47
Authenticating with an External LDAP Directory Service	48
Users Imported into SA from an LDAP Server	48
SSL and External Authentication	49
Supported External LDAP Directory Servers	49
Importing a Server Certificate from the LDAP into SA	50
Extracting the Server Certificate from Microsoft Active Directory	50
Extracting the Server Certificate from Novell eDirectory	50
Extracting the Server Certificate from SunDS	51
Importing External LDAP Users and User Groups	51
Importing LDAP Users and Groups Using LDAP Authentication Configuration	51

LDAP Authentication Configuration Prerequisites	52
The LDAP Authentication Configuration Process	53
Example LDAP Authentication Configuration Session	57
Synchronizing LDAP Users	62
Additional Steps Required on FIPS-Enabled Cores	64
SA Common Access Card (CAC) and Personal Identity Verification (PIV) Smart Card Integration	64
Smart Card/SA Integration Authentication Basics	65
SA Smart Card Integration Architecture	66
Setting Up SA/Smart Card Integration	67
Setting Up Smart Card Certificates	67
Setting Up Smart Card Certificates on All Slice Hosts	67
Creating a New Smart Card User	68
Initial Login to the SA Client as a Smart Card User	68
SA/RSA SecurID® Integration	71
RSA SecurID/SA Integration Overview	71
SA Support for SecurID Authentication Methods	72
Restrictions	72
SecurID/SA Integration Platform Requirements	73
Configuring SA/SecurID Integration	73
Phase 1: The RSA SecurID Authentication Configuration File	73
Phase 2: Enable RSA SecurID Authentication in SA	73
Phase 3: Create/Modify SA Users to Use SecurID Authentication	74
Troubleshooting	74
User and Security Reports	74
Chapter 2 SA Core and Component Security	76
Introduction to SA Core and Component Security Architecture	76
Enforcing Strict Control and Accountability	77
Stronger Controls and Accountability	77
Read-only, Digitally Signed Audit Trails	78
Signed SHA Checksums for Packages in the Software Repository	79
Role-Based Authorization	79

Audit Logging of User Activities	80
Custom Certificate Authority (CA)	80
Securing SA Internal Communications	80
Communication Between Components in an SA Core	81
Communication Between Agents and SA Core Components	82
Communication Between SA Cores	83
SA Satellite Architecture and Security	84
The SA Network: Enabling Risk Mitigation	84
SA Compatibility with Other Security Tools	85
SA Core Recertification	85
Agent versus Core Recertification	86
Adding a New Core or Slice to a Recertified Core Multimaster Mesh	87
Core Recertification Phases	87
Agent Recertification Phases	89
Agent Recertification Jobs	90
Agent Recertification Job Flow	93
SA Core Recertification Tool Usage	94
Arguments to the Core Recertification Tool	94
Security Considerations	96
Crypto Database File	96
Core Recertification Users	97
Creating the Core Recertification User	97
Removing a Core Recertification User	98
Core Recertification Prerequisites	98
Requirements for custom Certificate Authority (CA)	98
Select a New Password to Protect the Crypto Materials	99
Configuring Core Recertification	99
Ensure that All Cores are Running/Resolve Conflicts	106
Ensure That the Core Recertification Tool Correctly Recognizes the Mesh Setup ...	106
Recertifying SA Cores	107
Agent Recertification	112

Chapter 3 Multimaster Mesh Administration	114
Built-In Redundancy of the Multimaster Mesh	114
What Are Multimaster Mesh Conflicts?	114
How SA Handles Mesh Conflicts	115
Best Practices for Preventing Mesh Conflicts	115
Users	115
Administrators	116
Viewing the State of the Multimaster Mesh - SA Client	116
Resolving Mesh Conflicts - SA Client	120
Advanced Types and Causes of Mesh Conflicts	122
User Overlap Conflicts	122
Conflicts from User Duplication of Actions	123
Conflicts from Out of Order Transactions	123
Database Conflicts	124
Guidelines for Resolving Each Type of Conflict	125
Identical Data Conflict	125
Identical Data Conflict (Locked)	125
Simple Transaction Conflict	125
Unique-Key Constraint Conflict	126
Foreign-Key Constraint Conflict	126
Network Administration for a Multimaster Mesh	127
Multimaster Email Alerts	128
Facility Administration	130
Viewing Facility Information	130
Changing the Customers Associated with a Facility	132
Adding or Modifying Custom Attributes for a Facility - SA Client	132
Modifying a Facility Name - SA Client	133
Chapter 4 Satellite Administration	134
Starting/Restarting a Satellite	134
Stopping a Satellite	134
Verifying Satellite Communication with the Primary Core	135

Permissions Required for Managing Satellites	135
Viewing Satellite Information	135
Viewing Satellite Facilities and Realms	136
Viewing the Realm of a Satellite Managed Server	136
Viewing and Managing Satellite Gateway Information	136
Viewing Gateway Diagnostic and Debugging Information	137
Identifying the Source IP Address and Realm for a Connection	139
Changing the Bandwidth Usage or Link Cost Between Gateways	139
Viewing the Gateway Log or Change the Log Level	140
Restarting or Stopping a Gateway Process	140
Satellite Monitoring	140
Bandwidth Management of Remote Connections	141
The SA Bandwidth Configuration Management Tool	141
Invoking the Bandwidth Management Configuration Tool	142
Enabling/Disabling Remote Connection Bandwidth Management	144
Bandwidth Configuration Grammar	144
Satellite Software Repository Cache Management	145
Availability of Satellite Software Repository Cache Content	146
Updating Software in the Satellite Software Repository Cache	146
Setting the Software Repository Cache Update Policy	147
On-Demand Updates	148
Manual Updates	148
Emergency Software Repository Cache Updates	149
Software Repository Cache Size Management	149
Creating Software Repository Cache Manual Updates	149
Creating a Manual Update Using the DCML Exchange Tool (DET)	150
Applying a Manual Update to a Software Repository Cache	152
Staging Files to a Software Repository Cache	152
Running the Staging Utility	152
Microsoft Utility Uploads and Manual Updates	153
SA Satellite Installation and Topologies	153

Chapter 5 SA Remote Communications Administration	155
Bandwidth Management of Remote Connections	155
The SA Bandwidth Configuration Management Tool	156
Invoking the Bandwidth Management Configuration Tool	157
Enabling/Disabling Remote Connection Bandwidth Management	158
Bandwidth Configuration Grammar	158
IPv6 in SA	160
IPv4/IPv6 Dual-Stack Implementation	160
IPv6 Support in HP SA	160
SA Agent Installation	161
OS Provisioning	161
SA Managed Server Peer Content Caching	161
Requirements	161
Installing a Peer Cache	162
Configuring the Peer Cache and SA Servers	162
Remediation with Peer Caching Enabled	163
Retrieve Objects from the Peer Cache	163
Possible Errors	163
Viewing the Peer Cache Status Page	164
Concepts: SA Core Communications Infrastructure	164
Communication Between SA Cores	165
Advanced: Communication Between Agents and SA Core Components	168
SA Gateway Properties File Syntax	169
opswgw Command-Line Arguments	178
Chapter 6 SA Maintenance	179
The SA Start/Stop Script	179
Dependency Checking by the Start/Stop Script	179
Start/Stop Script Logs	179
Start/Stop Script Syntax	180
Starting the Oracle Database (Model Repository)	181
Starting a Standalone SA Core	181

Starting a Multiple-Server SA Core	181
Core Component Hosts Powered Up	181
Core Component Hosts Powered Down	182
Starting Individual SA Core Components	183
Start Order for Individual SA Core Components	183
Stopping an SA Core with Multiple Hosts	184
Multiple Data Access Engines	184
Overview of Multiple Data Access Engines	185
Reassigning the Data Access Engine to a Secondary Role	185
Designating the Multimaster Central Data Access Engine	186
Scheduling Audit Results and Snapshot Removal	186
Web Services Data Access Engine Configuration Parameters	187
Changing a System Configuration Parameter	187
Web Services Data Access Engine Configuration File	188
Increasing the Web Services Data Access Engine Maximum Heap Memory Allocation	190
Changing Software Repository Mirroring Parameters	190
Changing a System Configuration Parameter	191
Software Repository Mirroring Configuration Parameters	191
Chapter 7 Monitoring SA Core Components	192
Overview of SA Monitoring	192
Agent Monitoring	193
Agent Port	193
Monitoring Processes for Agents	193
Agent Logs	195
Agent Cache Monitoring	195
Monitoring Processes for the Agent Cache	196
Agent Cache Logs	196
Command Center Monitoring	196
Command Center Ports	196
Monitoring Processes for the Command Center	196
Command Center Logs	197

Data Access Engine Monitoring	197
Data Access Engine Port	197
Multimaster Central Data Access Engine Port Forwarding	197
Monitoring Processes for the Data Access Engine	198
Data Access Engine URLs	198
Data Access Engine Logs	199
Web Services Data Access Engine Monitoring	199
Web Services Data Access Engine Port	199
Monitoring Processes for the Web Services Data Access Engine	199
Web Services Data Access Engine URL	200
Web Services Data Access Engine Logs	200
Command Engine Monitoring	201
Command Engine Port	201
Monitoring Processes for the Command Engine	201
Command Engine Logs	201
Software Repository Monitoring	202
Software Repository Ports	202
Monitoring Processes for the Software Repository - Linux	202
Software Repository Logs	203
Software Repository Mirroring - SA Client	203
Model Repository Monitoring	205
Model Repository Port	205
Monitoring Processes for the Model Repository	205
Model Repository Logs	206
Table Space Usage	206
Multimaster Conflicts	206
Model Repository Multimaster Component Monitoring	207
Model Repository Multimaster Component Port	207
Monitoring Processes for the Model Repository Multimaster Component	207
Model Repository Multimaster Component Logs	208
Global File System Monitoring	208

Monitoring Process for the Global File System	209
Global File System Logs	210
Monitoring Processes for FUSE (Linux Only)	211
Spoke Monitoring	211
Spoke Ports	211
Monitoring Processes for the Spoke	212
Spoke Logs	212
Gateway Monitoring	212
Gateway Ports	212
Monitoring Processes for the Gateway	212
Gateway URL	214
Gateway Logs	214
OS Build Manager Monitoring	214
OS Build Manager Ports	214
Monitoring Processes for the OS Build Manager	214
OS Build Manager URL	215
OS Build Manager Logs	215
OS Boot Server Monitoring	215
OS Boot Server Ports	215
OS Boot Server Logs	215
OS Media Server Monitoring	216
OS Media Server Ports	216
OS Media Server Logs	216
Chapter 8 Troubleshooting SA - Diagnostic Tests	217
SA Core Component Internal Names	217
Core Health Check Monitor (HCM)	218
Overview of HCM Local Tests	218
Syntax of the Script for HCM Local Tests	219
Running HCM Local Tests	219
Overview of HCM Global Tests	220
Running HCM Global Tests	221

Syntax of the Script for HCM Global Tests	221
Setting Up Passwordless SSH for Global Tests	222
Extending the Health Check Monitor	223
Requirements for Extensions to HCM Local Tests	223
Categories and Local Test Directories	225
local_probes/<component>/verify_pre	225
local_probes/<component>/verify_post	225
local_probes/<component>/verify_functionality	225
Directory Layout for HCM Local Tests	226
HCM Local Test Example	226
Requirements for Extensions to HCM Global Tests	227
HCM Global Test Example	228
Directory Layout for HCM Global Tests	229
HCM Global Test Directories	230
global_probes/verify_pre	230
global_probes/verify_post	230
Running a System Diagnosis	230
System Diagnostic Tests	231
Core Components Tested by the System Diagnosis Tool	232
Data Access Engine Tests	232
Standalone Tests	232
Comprehensive Tests	233
Errors Caused By Additional Database Privileges	233
Software Repository Tests	233
Standalone Tests	233
Comprehensive Tests	233
Web Services Data Access Tests	234
Standalone Tests	234
Comprehensive Tests	234
Command Engine Tests	234
Standalone Tests	234

Comprehensive Tests	234
Model Repository Multimaster Component Tests	235
Standalone Tests	235
Comprehensive Tests	235
Chapter 9 Troubleshooting SA - Log Files	236
Viewing Log Files	236
Where Log Files Are Stored	236
Product Areas and Related Component Log Files	238
About Log File Sizes	239
About Component Log Levels	239
Changing Component Log Levels	240
Boot Server Logs	240
Build Manager Logs	240
Command Engine Logs	240
Changing Log Levels	240
Data Access Engine Logs	240
Media Server Logs	240
Model Repository Logs	241
Model Repository Multimaster Component Logs	241
Changing Logging	241
Agents Logs	241
SA Client Logs	242
Changing Log Levels	242
Software Repository Logs	242
Changing Log Levels	242
Web Services Data Access Engine Logs	242
Changing Log Levels	243
Gateway Logs	243
Changing Log Levels	243
Global File System Logs	243
Changing Log Levels - OGFS Hub Component	244

Changing Log Levels - OGFS Spoke Component	244
HTTPS Server Proxy Logs	244
APX Proxy Logs	244
Changing Log Levels	245
SSHD Logs	245
Changing Log Levels	245
Global Shell Audit Logs	245
Shell Event Logs	246
Shell Stream Logs	247
Shell Script Logs	247
Example of Monitoring Global Shell Audit Logs	247
Digital Signatures in the Global Shell Audit Logs	248
Storage Management for the Global Shell Audit Logs	248
Configuring the Global Shell Audit Logs	250
Extracting Session Data	251
Listing Recent Sessions	251
Sample Output	252
dump_session Command Reference	252
Syntax	252
Options	252
Chapter 10 SA Notification Configuration	254
Configuring SA Administrator Contact Information in SA Help	254
Configuring the Mail Server for a Facility	255
Configuring the Command Engine Notification Email	255
Configuring Email Alert Addresses for an SA Core	256
Configuring Email Alert Addresses for a Multimaster Mesh	256
Chapter 11 Global Shell: Windows Subauthentication Package	258
Microsoft Windows Authentication Process	258
Microsoft Windows Subauthentication Package	259
SA Subauthentication Package	259
SA Agent Installation Changes	261

SA Agent Uninstallation Changes	265
Appendix A Permissions Reference	266
Server Objects Permission	266
Server Property and Reboot Permissions	267
Device Group Permissions	267
Server Agent Deployment Permissions	268
Virtualization Service Management Permissions	268
Virtualization Container Permissions and Server Resource Permissions	269
Virtualization Tasks and Required Permissions	270
Solaris Virtualization Permissions	276
OS Provisioning Permissions	276
Manage Boot Clients Permissions	283
Software Management Permissions	284
Chef Cookbook Management Permissions	294
Permissions for Running a Chef Recipe from a Cookbook with No Dependencies	294
Permission Management for Cookbooks with Dependencies	295
Multi-tenancy	296
Application Configuration Management Permissions	297
Patch Management for Windows Permissions	305
Patch Management for Ubuntu Permissions	309
Patch Management for Solaris Permissions	311
Solaris Patch Policy Management Permissions	313
Patch Management for Other UNIX Permissions	316
Audit and Remediation Permissions	319
Server Permissions for Audit and Remediation	319
“Allow Create Task Specific Policy Permission” for Audit and Remediation	319
OGFS Permissions for Audit and Remediation	319
Audit and Remediation User Action Permissions	320
Compliance View Permissions	334
Job Permissions	336
Script Execution Permissions	337

Flow Permissions - HP Operations Orchestration	345
Service Automation Visualizer Permissions	345
Viewing Storage in SAV and SA Permissions	347
Storage Visibility and Automation Permissions	348
Appendix B Managed Platform Support	349
Importing the New Platform Package	350
Deploying Support for the New Platform	350
Required Manage Platforms Permission	350
Using the Platform Installer	350
Running a Platform Installer	351
Deleting a Platform Installer	352

User and User Group Setup and Security

SA provides a role-based security model that allows only authorized users to perform specific operations on specific servers. Intended for security administrators, this chapter explains how to set up a role-based security structure for SA.

About SA Users and User Groups

An SA user group represents a role and defines the set of permissions needed to perform that role. You grant a set of permissions to each user group and then assign users to one or more user groups. Each user group grants a set of permissions to all the users who belong to that group.

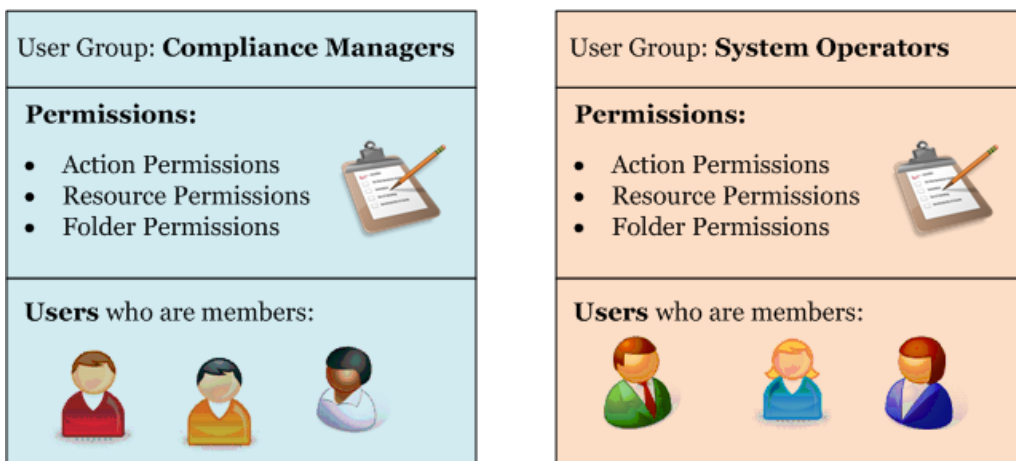
All users can belong to one or more SA user groups. The tasks that a user is authorized to perform are defined by the user groups of which the user is a member.

Each SA user group:

- **Represents a role**, which is a set of tasks and responsibilities.
- **Defines a set of permissions** that enable the set of tasks needed to perform that role.
- **Contains the set of SA users** who perform that role.

Figure 1 shows two example user groups. One is for compliance managers whose role is to run audit reports and ensure compliance of servers to corporate policies; the other example user group is for system operators whose role is to monitor servers and install software and patches. Each user group contains a set of permissions and a set of users:

Figure 1. Contents of User Groups, Based on Roles



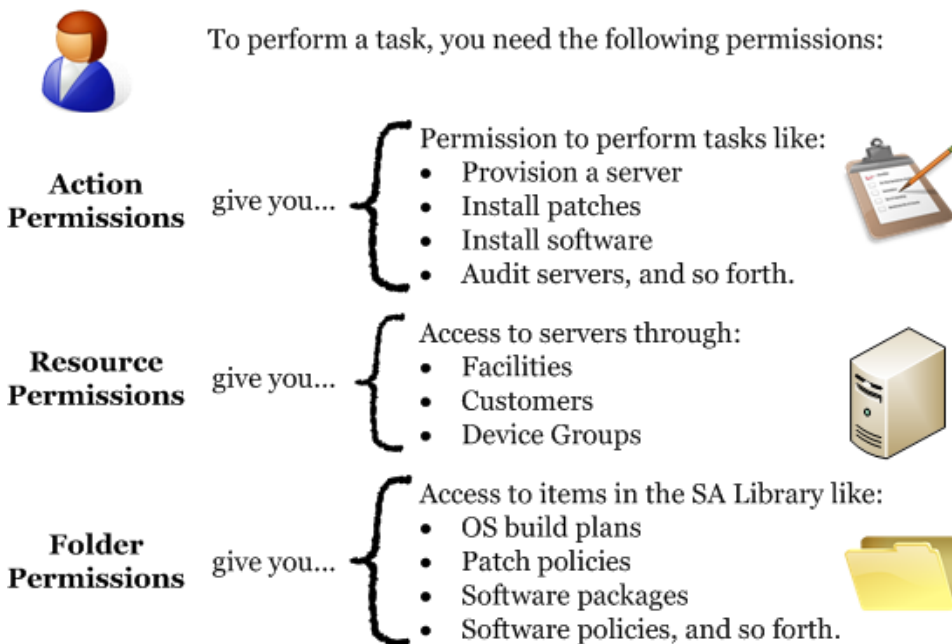
SA provides a set of predefined user groups, but you can create your own user groups to match the roles in your organization. For more information, see [Predefined User Groups](#).

About Permission Types - Action, Resource and Folder Permissions

SA provides three types of permissions needed to perform any action on servers:

- **Action permissions** specify the actions or tasks that users can perform.
- **Resource permissions** specify the servers on which users can perform these actions. All servers are grouped by facility, by customer, and by device groups. You set resource permissions by specifying access to facilities, customers, and device groups.
- **Folder permissions** specify access permissions to items in the SA Library, such as OS build plans, software packages, software policies, patch policies, audit policies, and so forth.

Figure 2. SA Permission Types Needed to Perform a Task



For example, to install software using a software policy, a user would need (at least) the permissions shown in [About Permission Types - Action, Resource and Folder Permissions](#):

Figure 3. Permissions Needed to Install Software

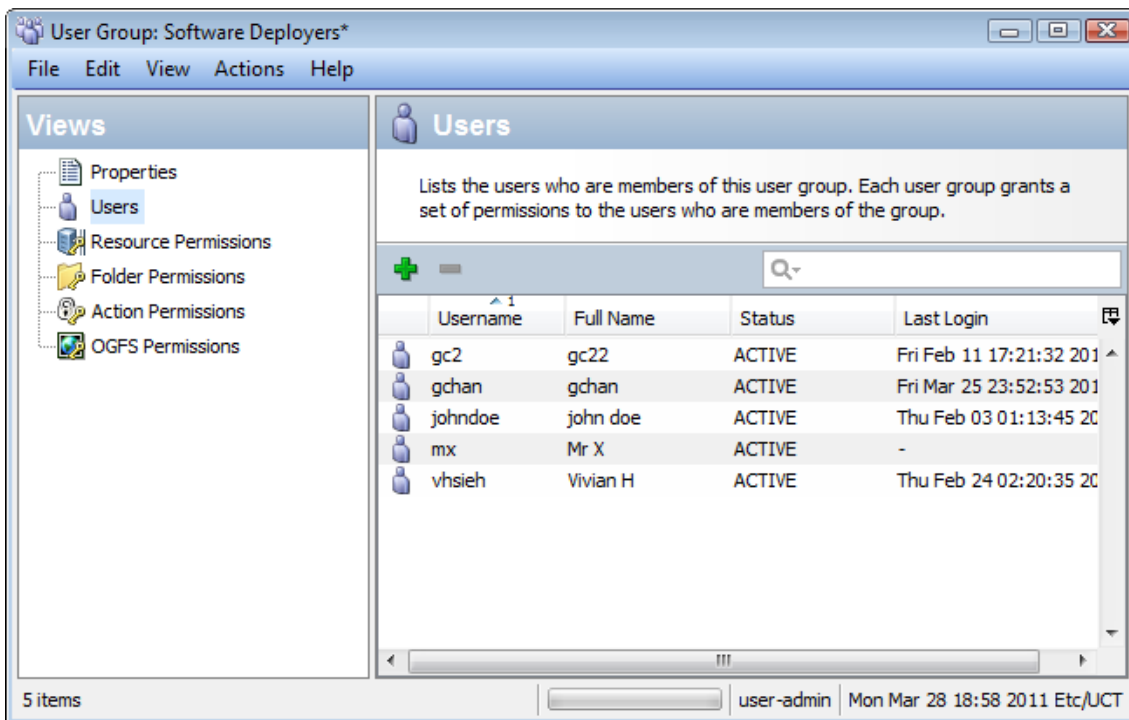
 To **install software**, you need the following permissions:

Action Permissions:	Allow Install Software: Yes Manage Software Policy: Read Allow Attach Software Policy: Yes Manage Services: Read & Write Managed Servers and Groups: Yes	
Resource Permissions:	Facility and Customer and Device Group: Read & Write	
Folder Permissions:	/software/my_app: Read	

These permissions (and others) are set in the predefined user group Software Deployers. For more information, see [Predefined User Groups](#).

Figure 4 shows the predefined user group named Software Deployers and the SA users who are members of the group. The Views navigation panel also shows the Resource Permissions, Folder Permissions, Action Permissions, and OGFS Permissions of this user group.

Figure 4. User Group Browser Showing Users Who are Members



About Action Permissions

Action permissions define the tasks that can be performed by users. Some action permissions specify the following types of access:

- **Read:** Users can perform the task but in a read-only mode.
- **Read & Write:** Users can fully perform the task.
- **None:** The task does not appear in the SA Client. Users cannot view or perform the task.

Other types of action permissions specify the following types of access:


- **Yes:** Users can perform the task.
- **No:** Users cannot perform the task.

For a complete list of action permissions, see [Permissions Reference](#) and [Setting Action Permissions](#).

Grouping Action Permissions

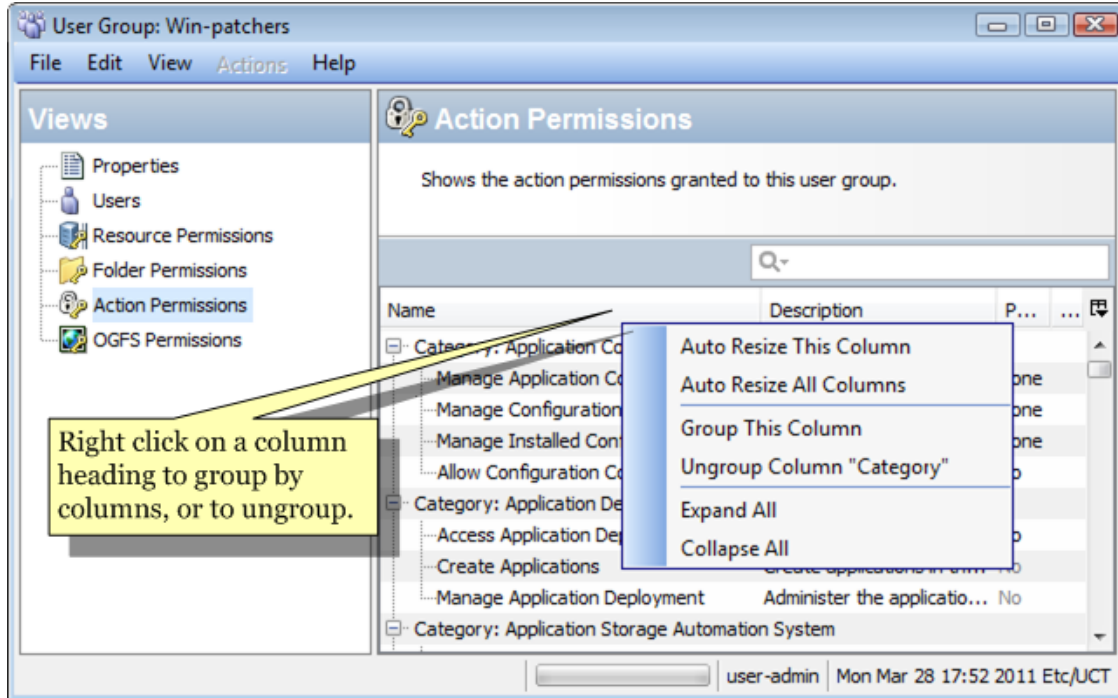
The SA Client displays the action permissions for a user group when you open the user group. The action permissions are grouped by category, as shown in **Figure 5**.

Figure 5. User Group Window - Action Permissions View, Grouped by Category



You can ungroup the action permissions or group them by other columns by right-clicking on any column, as shown in **Figure 6**.

Figure 6. User Group Window - Action Permissions View, Grouping Menu



About Resource Permissions

A *resource* is one or more managed servers. Server resources are organized into the following categories:

- **Facilities:** The servers associated with an SA Facility. Every managed server belongs to one and only one of your facilities.
- **Customers:** The servers associated with a customer. You create customers and assign each server to one customer. Every server belongs to one and only one customer, which may be the "Not Assigned" customer group.
- **Device Groups:** The servers belonging to a device group. You create device groups and assign servers to them. Every server can belong to one or more device groups.

Resource permissions for a user group determine if the users in the user group can view or modify the servers. A user group only has access to the servers in the facilities, customers, and device groups for which it has been granted resource permissions. Because every server belongs to one facility, one customer, and at least one device group, to have access to servers, a user group must have permissions to at least one facility, at least one customer, and at least one device group.

You can combine customer, facility, and device group permissions to implement security policies. For example, you can restrict access to servers that are associated with the Acme Corp. customer, reside in the Fresno facility, and belong to a device group that contains only Windows servers (see [Examples of Resource Permissions](#)).

Any one server is in a facility, is associated with a customer and is in one or more device groups. A user needs access to that facility, as well as to that customer and to at least one device group containing that server to get access to that server. See also [Setting Resource Permissions - Facilities, Customers, and Device Groups](#).

Types of Access to Resources

Resource permissions must specify one of the following types of access:

- **Read:** Users can view the resource only.
- **Read & Write:** Users can view, create, modify or delete the resource.
- **None:** The resource does not appear in the SA Client. Users cannot view or modify the resource.

About Facility Permissions

Every server is in one and only one facility. To modify a server in a particular facility, a user must belong to a user group that has Read & Write permission for the facility. For example, if you want the users of a group to be able to view (but not modify) the servers in the London facility, set the permission to Read.

The facility permissions also control access to the facility object itself. For example, to modify a property of a facility, a user must belong to a group that has Read & Write permission to the facility and the action permission to modify facilities.

About Customer Permissions

Every server is associated with one and only one SA Customer, even if it is the “Not Assigned” Customer group. An SA Customer is a logical group into which you can place servers. You can then perform IT management tasks on all servers belonging to an SA Customer as long as you have Read and/or Write privileges to that Customer, thus providing security and authorization boundaries. For example, if you want the users of a group to be able to view (but not modify) the servers associated with the Widget Inc. customer, set the permission to Read.

The customer permissions also control access to the customer object itself. For example, to add a custom attribute to a customer, a user must belong to a group that has Read & Write permission to the specific customer and the action permission to modify customers.

About Device Group Permissions

Every server can belong to one or more device groups. By setting the device group permissions, you control the access that the users in the user group have to the servers in the device group. For example, if you want the users of a group to be able to view (but not modify) the servers in the Windows Server 2008 device group, set the permission to Read.

By default, each server belongs to a public device group based on its operating system. You can view these device groups in the SA Client by selecting the Devices tab and selecting Device Groups > Public > Opsware > Operating Systems.

If a server belongs to more than one device group, the user group needs permission to only one of the device groups to get access to that server.

While a device group can contain other device groups, permissions are not inherited by the contained device groups.

You cannot control access to a private device group. Private device groups are visible only to the user who created them.

The device group permissions control access to servers that belong to device groups. However, these permissions do not control the management of the device groups. To create, modify, or delete device groups, a user must belong to a user group that has the Manage Public Device Groups and the Model Public Device Groups action permissions and the Managed Servers and Groups action permission. To add devices to a device group being used as an Access Control Group, the user must be a Super Administrator.

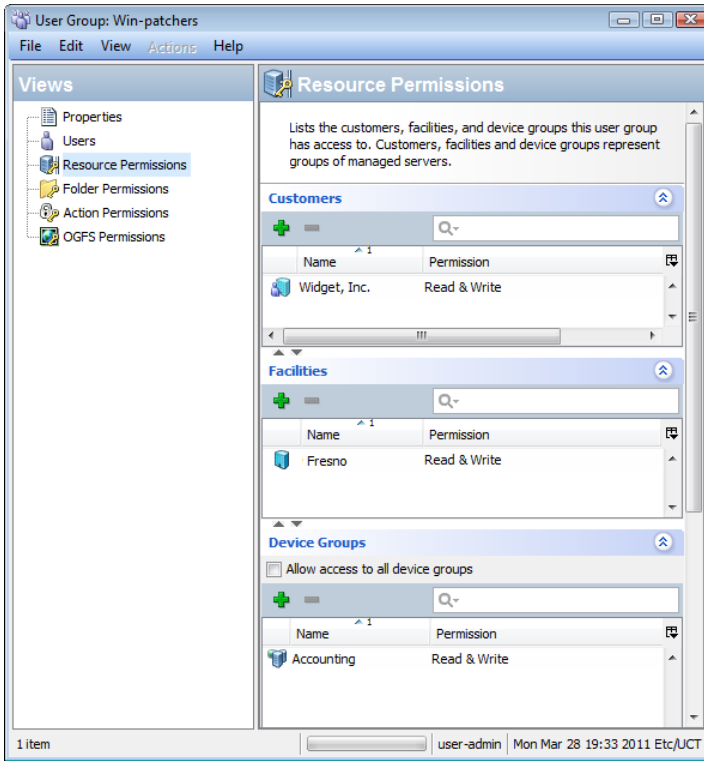
Examples of Resource Permissions

Suppose that a server resides in the Fresno facility, is associated with the Widget, Inc. customer, and belongs to the Accounting device group. To modify the server, the user group could have the permissions listed in **Table 1**. These permissions are also shown in **Figure 7** for the user group named Win-patchers.

Table 1. Example of Resource Permissions

Resource	Access Permission
Facility: Fresno	Read & Write
Customer: Widget, Inc.	Read & Write
Device Group: Accounting	Read & Write

Figure 7. Resource Permissions View in the User Group Screen



If the access permissions for the facility, customer, or device group do not match, then the **most restrictive** permissions are enforced.

For example, as **Table 2** shows, if the permission for the customer and the device group is Read & Write but the permission for the facility is Read, then the Read permission is enforced and the user will not be able to modify the servers.

If the permission for the customer is None, then the server cannot be viewed, even if the other permissions for the user group specify Read, or Read & Write.

Table 2. Example of Mismatched Resource Permissions

Resource	Permission
Facility: Fresno	Read
Customer: Widget, Inc.	Read & Write
Device Group: Accounting	Read & Write

Resource and Action Permissions Combined - Example

To perform an action on a resource, the user must belong to a group that has the necessary permissions for both the action and the resource (server). For example, suppose that a server is associated with these resources: the Widget, Inc. customer and the Fresno facility and the Red Hat AS 4 device group. To install a patch on this server, the user could belong to a group with the permissions listed in **Table 3**.

Table 3. Example of Resources Permissions and Action Permissions

Resource and Action	Permission
Customer: Widget, Inc.	Read & Write
Facility: Fresno	Read & Write
Device Group: Red Hat AS 4	Read & Write
Action: Install Patch	Yes

Other Types of Resources

Managed servers are the most common resources. Other types of resources are:

- Hardware definitions
- Realms
- OS installation profiles

Each of these resources can be associated with customers.

Folders can also be associated with customers, but access to folders is controlled in a different way (see [About Folder Permissions](#)).

About Folder Permissions

Folder permissions control access to the contents of folders in the SA Library, such as software policies, patch policies, OS build plans, server scripts, and subfolders. A folder's permissions apply only to the items directly under the folder. They do not apply to items lower down in the hierarchy, such as the subfolders of subfolders. See [Setting Folder Permissions](#).

Types of Folder Permissions

In the Folders Properties window of the SA Client, you can assign the following permissions to an individual user or a user group:

- **List Contents of Folder:** Navigate to the folder in the hierarchy, click on the folder, view the folder's properties, see the name and type of the folder's contents (but not the attributes of the contents).
- **Read Objects Within Folder:** View all attributes of the folder's contents, open object browsers on folder's contents, use folder's contents in actions.

For example, if the folder contains a software policy, users can open (view) the policy and use the policy to remediate a server. However, users cannot modify the policy. (For remediation, action and resource permissions are also required.)

Selecting this permission automatically adds the List Contents of Folder permission.

- **Write Objects Within Folder:** View, use, create, and modify the folder's contents.

This permission permits actions such as New Folder and New Software Policy. To perform most actions, action permissions are also required.

Selecting this permission automatically adds the List Contents of Folder and the Read Objects Within Folder permissions.

- **Execute Objects Within Folder:** Run the scripts contained in the folder and view the names of the folder's contents.

This permission allows users to run scripts, but not to read or write them. To view the contents of scripts, users need the Read Objects Within Folder permission and the appropriate action permission. To create scripts, they need the Write Objects Within Folder permission and the appropriate action permission.

Selecting the Execute Objects Within Folder permission automatically adds the List Contents of Folder permission.


- **Edit Folder Permissions:** Modify the permissions or add customers to the folder.

This permission enables users to delegate the permissions management of a folder (and its contents) to another user group.

Selecting this permission automatically adds the List Contents of Folder permission.

Figure 8 shows the user group named Win-patchers with the Folder Permissions view selected. This user group has list, read, write, and execute permissions to the folder named /Library/A-WinPatch.

Figure 8. Folder Permissions View in the User Group Window



Folder Permissions and Action Permissions

Action permissions determine what actions users can perform with the SA Client. Folder permissions specify which folders in the SA Library users have access to.

To perform most actions on folders and the items they contain, users need both folder and action permissions. For example, to add a software policy to a folder, users must belong to a group that has the Write Objects Within Folder permission on a particular folder and the Manage Software Policy action permission (Read & Write).

Folders, Customer Constraints, and Software Policies

If a customer is assigned to a folder, the customer constrains some of the actions on the software policies contained in the folder. These constraints are enforced through filtering: The objects that can be associated with the software policies must have a matching customer.

For example, suppose that you want to add the `quota.rpm` package to a software policy. The package and the software policy reside in different folders. The customer of the policy's folder is Widget and the customer of the package's folder is Acme. When you perform the Add Package action on the policy, the packages that you can choose will not include `quota.rpm`. The customer of the policy's folder (Widget) acts as a filter, restricting the objects that can be added to the policy. If you add the Widget customer to the folder of `quota.rpm`, then you can add `quota.rpm` to the policy.

The following list summarizes the customer constraints for software policy actions. These constraints are invoked only if the software policy's folder has one or more customers. Software policy actions not listed here, such as New Folder, do not have customer constraints.

- **Attach Software Policy:** The customer of the server being attached must be one of the customers of the software policy's folder.
- **Install Software Policy Template:** The customer of the server must be one of the customers of the folder of each software policy contained in the template.

Default Folder Permissions

When SA is first installed, the predefined user groups are assigned permissions to the top-level folders such as Package Repository. When you create a new folder, it has the same permissions and customer as its parent.

Membership in Multiple User Groups

If a user belongs to more than one user group, the user's permissions are derived from the resource and action permissions of all of the groups. The way the permissions are derived depends on whether or not the resources are folders.

If the resources are not folders, then the derived permissions are a cross-product of the resource and action permissions of all groups to which the user belongs. With a cross-product, all action permissions apply to all resource permissions. For example, Jane Doe belongs to both of the Atlanta and Portland groups, which have the permissions listed in **Table 4**. Because the derived permissions are a cross-product, Jane can perform the System Diagnosis task on the managed servers associated with the Widget Inc. customer, even though neither the Atlanta nor Portland group has this capability.

Table 4. Example of Cross-Product Permissions

Resource or Action	Atlanta User Group Permission	Portland User Group Permission
Resource: Customer: Widget, Inc.	Read & Write	None
Resource: Customer: Acme Corp.	None	Read & Write
Action: System Diagnosis	No	Yes

If the resources are virtualization containers, then the derived permissions for the user are cumulative but do not cross user groups. For example, John Miller belongs to both the San Diego and Raleigh groups shown in **Table 5**. If John has Write permissions to Server X in Virtualization Inventory Folder A, John can run power control operations on it. If John has Write permissions to Server Y in Virtualization Inventory Folder B, he can Modify the VM configuration. However, he cannot run a power control on Server Y or Modify the VM configuration of Server X.

Table 5. Example of Permissions for Virtualization Containers

Resource or Action	San Diego User Group Permission	Raleigh User Group Permission
Resource: Hypervisor Container B	None	List
Resource: Virtualization Inventory Folder A	Read	None
Resource: Virtualization Inventory Folder	None	Read & Write
Action: VM Lifecycle Management: Power Controls	Yes	None
Action: VM Lifecycle Management: Modify VM	None	Yes

If the resources are folders (or their contents), then the derived permissions for the user are cumulative but do not cross user groups. For example, Joe Smith belongs to both the Sunnyvale and Dallas groups shown in **Table 6**. Joe can create packages under the Webster folder because the Sunnyvale group has Read & Write permissions for that folder and for the Manage Package action. However, Joe cannot create packages under the Kiley folder, because neither user group can do so. Joe can create OS Sequences under the Kiley folder, but not under the Webster folder.

Table 6. Example of Cumulative Permissions

Resource or Action	Sunnyvale User Group Permission	Dallas User Group Permission
Resource: Folder Webster	Read & Write	None
Resource: Folder Kiley	None	Read & Write
Action: Manage Packages	Read & Write	None
Action: Manage OS Sequences	None	Read & Write

Restricted Views in the SA Client Based on Permissions

The SA Client displays only those resources for which the user's group has Read or Read & Write permissions.

For example, John Smith belongs to the Basic Users group, which has the permissions listed in **Table 7**. When John logs in, the SA Client displays only the servers for Widget Inc., but not those of Acme Corp.

Table 7. Example of Permissions and Restricted Views

Resource or Action	Basic Group Permission
Customer: Widget, Inc.	Read & Write
Customer: Acme Corp.	None
Wizard: Prepare OS	Yes
Wizard: Run Scripts	No

To locate or view a server, a user must belong to a user group that has Read (or Read & Write) permission to the customer and the facility and at least one device group associated with the server. Otherwise, the user cannot see the server in the SA Client.

Predefined User Groups

During an SA installation or upgrade, SA creates a set of predefined user groups based on user roles. You must grant read and/or write permissions to the Facility and Customer and other

appropriate permissions to these user groups. Use of the predefined user groups is optional. SA recommends that you copy and modify the permissions of the predefined user groups to create your own customized user groups rather than modify the default groups. Your modification or deletion of predefined user groups is not affected by SA upgrades. **Table 8** shows the predefined user groups:

Table 8. Predefined User Groups

User Group Name	Description
Opware System Administrators	Opware System Administration privileges.
Superusers	Complete access to all SA-managed objects and operations.
Viewers	Read-only access to all SA-managed objects and operations.
Reporters	Access to reporting only.
OS Policy Setters	Access to import & define OS build plans.
OS Deployers	Access to provision servers.
Patch Policy Setters	Access to set patching policy.
Patch Deployers	Access to install patches.
Software Policy Setters	Access to set software policy.
Software Deployers	Access to install software.
Compliance Policy Setters	Access to define compliance policies.
Compliance Auditors	Access to execute compliance scans.
Compliance Enforcers	Access to remediate compliance failures.
Virtualization Administrators	Access to add, edit, and remove virtualization services, manage lifecycle of VMs and VM Templates, and administer permissions for virtualization inventory.
Hypervisor Managers	(If core was upgraded from SA 9.1x) Access to create, delete, and register VMs. For more information about upgrade paths, see the SA 10.0 Upgrade Overview guide.
Virtual Machine Managers	Access to start and stop VMs.
VM Lifecycle Managers	Access to manage lifecycle of VMs, including create, modify, migrate, clone, and delete VMs, VM power controls, and deploy VM Templates.

VM Template Deployers	Access to deploy VMs from VM Templates, clone VMs, and VM power controls.
VM Template Managers	Access to manage lifecycle of VMs and VM Templates, including create, modify, migrate, clone, delete VMs, VM power controls, convert VMs to VM Templates, deploy VMs from VM Templates, and delete VM Templates.
Command Line Administrators	Shell access to servers.
Server Storage Managers	Access to manage server storage.
Storage System Managers	Access to manage storage systems.
Storage Fabric Managers	Access to manage storage fabrics.
Chef Group	A group having execute, read, write, list access to Chef objects and operations.
Command-logger Group	A Group having execute, read, write, list access to /Extensible Discovery Folder.

About Private User Groups

Note: Private user groups are intended for migrating scripts into folders in the SA Library. You should not assign permissions to users using private user groups. You should use regular user groups. For more information, see [About SA Users and User Groups](#).

When an SA administrator creates a new user, SA automatically creates a private user group for the new user and assigns the new user to the private user group. The name of the private user group is the user name.

A private user group can contain only one SA user and every SA user can belong to only one private user group. The SA administrator can then assign action and resource permissions to the private user group. The permissions that you specify for a private user group determine what the user can do with SA. Action permissions specify what actions the user can perform; resource permissions indicate the servers on which the user can perform the actions. Global File System (OGFS) permissions cannot be assigned to a private user group.

For example, when an SA Administrator creates a new user with user name john, a private user group john is also created, and a default folder called john is created in the Home directory. The SA Administrator can then assign action and resource permissions to the private user group john.

An SA user can be a member of multiple user groups and belong to the user's private group. But then the derived permissions of the private user group is not a cross-product of the resource and action permissions of all groups to which the user belongs.

When a user is deleted, SA automatically deletes the corresponding private user group and the default folder for that user is moved to the location `/Home/deleted_users` in the SA Library.

For more information, see [Setting Private User Group Permissions](#).

About Super Administrators and Super Users

A **Super Administrator** is an SA user who can create users and user groups, specify permissions for user groups, and assign users to user groups. Super administrators can also manage customers and facilities, as well as set folder permissions. To perform most of the tasks described in this chapter, you must log in to the SA Client as a super administrator.

The SA installer creates a single default user, the super administrator named `admin`. The password for `admin` is specified during the installation and should be changed immediately afterward.

Tip: As a best practice, you should not add the `admin` user to other user groups.

About Super Users

A **Super User** is different from a Super Administrator and is not automatically a Super Administrator. A Super User is any user who belongs to the predefined Superusers group. A Super User has full permissions to perform all actions, except create and modify users and user groups.

However, a super user does not automatically have access to any servers. You would need to give access to facilities, customers, and device groups as described in [Setting Resource Permissions - Facilities, Customers, and Device Groups](#).

To create a super user, add an existing user to the Superusers predefined user group. For more information, see [Predefined User Groups](#) and [Adding a User to a User Group](#).

About Customer Administrators and Customer Groups

One way to organize your servers and provide access control boundaries is to segregate your managed servers by customer. A customer represents a set of servers associated with a business organization, such as a division or a company. Typically a server is associated with a customer, because it runs applications for that customer. For more information on creating and managing customers, see the SA User Guide: Server Automation.

Comparing Customer Administrators with Super Administrators

The super administrator can delegate the management of specific user groups to a customer administrator. Like a super administrator, a customer administrator can assign users and permissions to user groups. However, a customer administrator can only modify user groups that have access to the specified customers.

A **customer administrator** is the same as a super administrator with the following constraints:

- While a super administrator can add or remove users from all user groups, a customer administrator can add or remove users only from some user groups – those that have Read and Write access to the specific customers listed in the customer group.
- While a super administrator can modify permissions on all user groups, a customer administrator can modify permissions only on some user groups – those that have Read and Write access to the specific customers listed in the customer group.
- While a super administrator can create new SA users or delete SA users, a customer administrator cannot create or delete users.

A Customer Administrator is Defined by a Customer Group

You create a customer administrator by creating a customer group. A **customer group** contains one or more SA users and one or more customers. Each user in the customer group becomes a customer administrator for the customers in the customer group. The user groups that a customer administrator can manage are the user groups with Read and Write permission to the customers listed in the customer group.

Example Customer Group

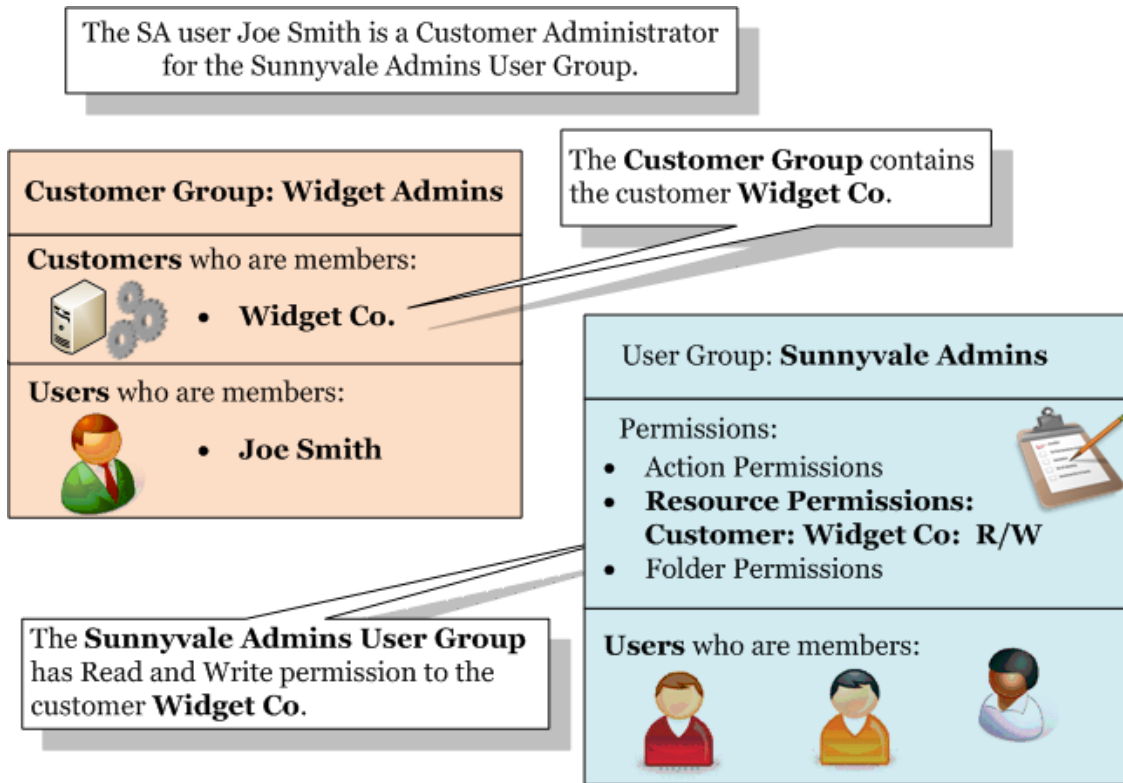
The following example shows a customer named Widget Co and a user group named Sunnyvale Admins. The Sunnyvale Admins user group has Read and Write permission to the customer Widget Co, meaning the Sunnyvale Admin users are responsible for managing the servers assigned to the Widget Co customer.

Figure 9 shows how to make the SA user Joe Smith a customer administrator for the Widget customer. The Widget Admins customer group lists Joe Smith and the customer Widget Co, which defines Joe Smith as a customer administrator for the Widget customer. Joe Smith can modify (add and remove users and change permissions in) the Sunnyvale Admins user group.

The figure shows the relationships required for Joe Smith to manage the Sunnyvale Admins user group:

- The Sunnyvale Admins user group has Read and Write permission to the Widget Co customer.
- The Widget Admins customer group contains the Widget Co customer.
- The Widget Admins customer group contains the user Joe Smith.

Figure 9. Defining a Customer Administrator



For more information, see [Managing Customer Administrators and Customer Groups - SA Client](#).

Process Overview for Security Administrators

The person responsible for the security of SA creates and maintains users and user groups, sets permissions on user groups and assigns users to user groups. This person must be able to log in to the SA Client as a user who is a super administrator. For more information, see [About Super Administrators and Super Users](#).

The following steps provide an overview of security administration for SA:

1. Identify the people in your organization who will manage SA security.
2. For each user identified in the preceding step, create a super administrator.

For instructions, see [Creating a Super Administrator](#).

3. Note the facility to which the managed servers belong.

A facility represents a data center or physical location. Depending on your organization, you may want to name the facility after the city, building, or room where the servers reside. The person who installs SA specifies the name of the facility for the core.

4. Associate managed servers with customers.

In SA, a customer represents a set of servers associated with a business organization, such as a division or a company. Typically, a server is associated with a customer, because it runs applications for that customer.

For more information on grouping servers by customer, see the SA User Guide: Server Automation.

5. (Optional) Create device groups and assign servers to the groups. Device groups are another way to organize your managed servers.

For more information on device groups, see the SA User Guide: Server Automation.

6. Plan your user groups.

Decide which SA tasks specific groups of users will perform and on which servers. Usually a user group represents a role or a job category. Examples of user groups are: UNIX System Admins, Windows Admins, DBAs, Policy Setters, Patch Admins, and so forth. See the [Predefined User Groups](#).

7. If the predefined user groups do not meet your needs, create your own user groups.

For instructions, see [Creating a New User Group](#).

8. Set the resource permissions on the user groups.

These permissions specify read and write access to servers associated with facilities, customers, and device groups. Resource permissions control which servers the members of a user group can access. For more information, see [Setting Resource Permissions - Facilities, Customers, and Device Groups](#).

9. Set the action permissions on the user groups.

To determine which action permissions are required to perform a specific task, see the tables in [Permissions Reference](#). For example, if you have a user group named Software Managers, see [Table 45. Software Management Permissions Required for User Actions](#). For more information, see [Setting Action Permissions](#).

10. Set the OGFS permissions on the user groups.

OGFS permissions are required for certain actions; for example, for actions that require access to a managed server's file system. The OGFS permissions are included in the tables in [Permissions Reference](#).

For instructions, see [Setting OGFS Permissions](#).

11. Create the folder hierarchy in the SA Library using the SA Client.

For more information on the SA Library, see the SA User Guide: Server Automation.

12. Set the folder permissions.

In general, you need read permission on a folder to use its contents in an operation, write permission to create or modify folder contents, and execute permission to run scripts that reside in a folder. For more information, see [Setting Folder Permissions](#).

13. (Optional) Delegate the management of folder permissions to certain user groups.

For instructions, see [Setting Folder Permissions](#).

14. Create new users in SA or import existing users from an external Lightweight Directory Access Protocol (LDAP) directory.

For instructions, see [Creating a New User](#) and [Authenticating with an External LDAP Directory Service](#).

15. Assign users to the appropriate groups.

For instructions, see [Adding a User to a User Group](#).

About Global File System Permissions

To use the OGFS, you need to grant OGFS permissions. OGFS permissions are separate but related to the action permissions, resource permissions, and folder permissions described in [About Permission Types - Action, Resource and Folder Permissions](#) (see also [Setting OGFS Permissions](#)).

The OGFS is a virtual file system that gives you access to all your managed servers and all their file systems. It underlies many SA Client actions, such as browsing managed server file systems and scanning servers for compliance. To perform actions that use the OGFS, you must belong to a user group that has OGFS permissions. **Table 9** lists the operations you control with OGFS permissions.

Table 9. OGFS Permissions

OGFS Permission	Tasks Allowed by this Permission
Launch Global Shell	Launch the Global Shell.
Log In To Server	Open a shell session on a UNIX server. In the SA Client, open a Remote Terminal. In the Global Shell, you can use the <code>rosh</code> command.
Read COM+ Database	Read COM Plus objects as a specific login. In the SA Client, use the Device Explorer to browse these objects on a Windows server.
Read Server File System	Read a managed server as a specific login. In the SA Client, use the Device Explorer to browse the file system of a managed server.
Read IIS Metabase	Read IIS Metabase objects as a specific login. In the SA Cli-

OGFS Permission	Tasks Allowed by this Permission
	ent, use the Device Explorer to browse these objects on a Windows server.
Read Server Registry	Read registry files as a specific login. In the SA Client, use the Device Explorer to view the Windows Registry.
Relay RDP Session To Server	Open an RDP session on a Windows server. In the SA Client, this is the Remote Terminal menu that opens an RDP client window for a Windows server.
Run Command On Server	Run a command or script on a managed server using the <code>rssh</code> utility, where that command or script already exists. In the SA Client, this is used for Windows Services accessed by the Device Explorer.
Write Server File System	Modify files on a managed server as a specific login. In the SA Client, you can use the Device Explorer to modify the file system of a managed server.

When setting an OGFS permission, in addition to specifying an operation such as Write Server File System, you also specify the managed servers to which the operation can be applied. You specify the managed servers by selecting a facility or a customer or a device group. You also specify the login name for the managed server where the operation runs. (The Launch Global Shell operation is an exception.)

For example, suppose you specify the Read Server File System permission. For the servers, you select a device group named Sunnyvale Servers. For the login name, you select the SA user name. Later, in the SA Client, the SA user `jd` opens a server belonging to the Sunnyvale Servers device group in the Device Explorer. In the Views pane, the string `jd` appears in parentheses next to the File System label. When the user drills down into the file system, the Device Explorer displays the files and directories to which the UNIX user `jd` has access.

If you specify a super user such as `root` for the login name, make sure that the resource you select only allows access to the correct set of servers. For `root`, you should limit access to servers by customer or device group, not by facility.


For the Launch Global Shell permission, you do not specify the managed servers, because a Global Shell session is not associated with a particular server. Also, you do not specify the login user for this permission. If you open a Global Shell session with the SA Client, you do so as your current SA login. If you open it with the `ssh` command, you are prompted for an SA login (user name).

Managing Users - SA Client

This section describes how to manage users with the SA Client. To manage users, you must log in to the SA Client as a super administrator (`admin`) and select the Administration tab, as shown in

Figure 10.

Figure 10. Users Listed Under the Administration Tab



Creating a New User

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

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your SA users.
4. Select the **Actions > New** menu or select the New User icon. This displays the New User window.
5. Enter the first name, last name, and full name of the user.
6. To allow the new user to administer users and user groups, select the check box labeled Super Administrator. For more information, see [About Super Administrators and Super Users](#).
7. Enter the contact information for the new user. An email address is required.
8. Enter the log-in information for the new user.
 - The user credentials can be stored in HP SA or on an RSA SecurID server connected to SA. You can change the user password in the SA Client only if the credential store is HP SA.
 - The SA user name must be made up of letters, numbers, periods, hyphens, and underscores. SA user names are not case sensitive.

- The password must be at least six ASCII characters long and may not include the “\” or “^” characters.
- 9. Enter the locale, time zone, and date format preferences.
- 10. Select the User Groups view to assign the user to one or more user groups. Assigning the user to user groups grants the corresponding permissions to the user. Use the “+” button to add the user to a user group. Use the “-” button to remove the user from the selected user group.
- 11. Select **File > Revert** to discard your changes.
- 12. Select **File > Save** to save the new user.

Changing a User’s Permissions

All permissions are contained in user groups. Each user’s permissions are determined by the user groups to which they belong. To modify user permissions you must modify the permissions defined in the user groups to which the user belongs or change the user groups to which the user belongs. For more information, see [Assigning a User to a User Group](#) and [Setting Permissions on User Groups - SA Client](#).

Changing a User’s Password

Only a super administrator (`admin`) can change the passwords of other SA users. If the user name has been imported from an external LDAP directory, then the password cannot be changed with the SA Client. For more information, see [Authenticating with an External LDAP Directory Service](#).

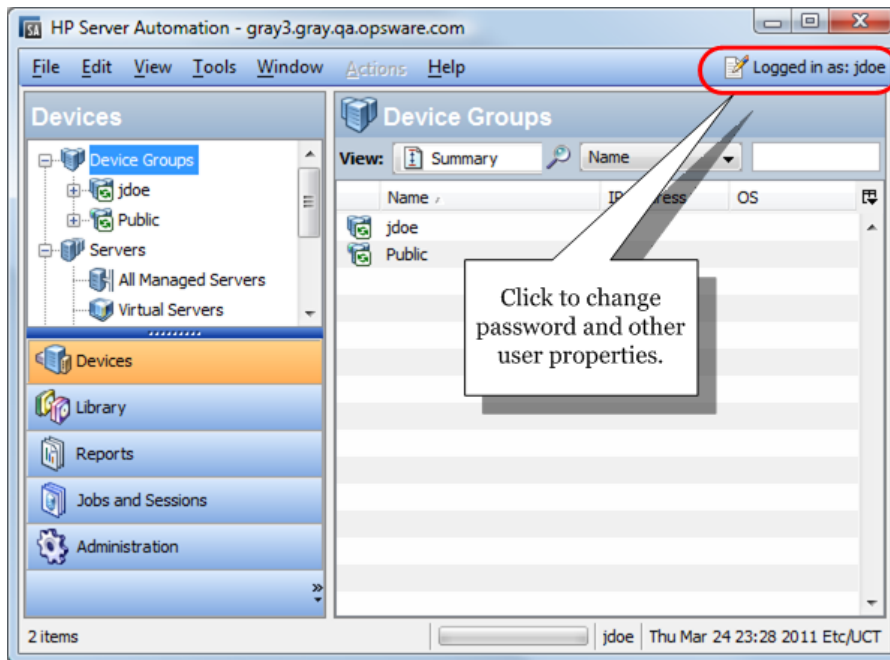
To change a user’s password, you need to open the user in a user window and select the Properties view. Perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your SA users.
4. Select the user you want to modify.
5. Select the **Actions** menu, or right-click and select **Open**. This displays the user information in a new window.
6. Select the Properties view. This displays the user’s login information, including a Change Password link.
7. Select the Change Password link. This displays the Change Password dialog.
8. Enter the new password. Note that when you modify the user’s password, the change takes effect immediately.
9. Select OK. This modifies the user’s password.

Users Changing Their Own Password and Other Properties

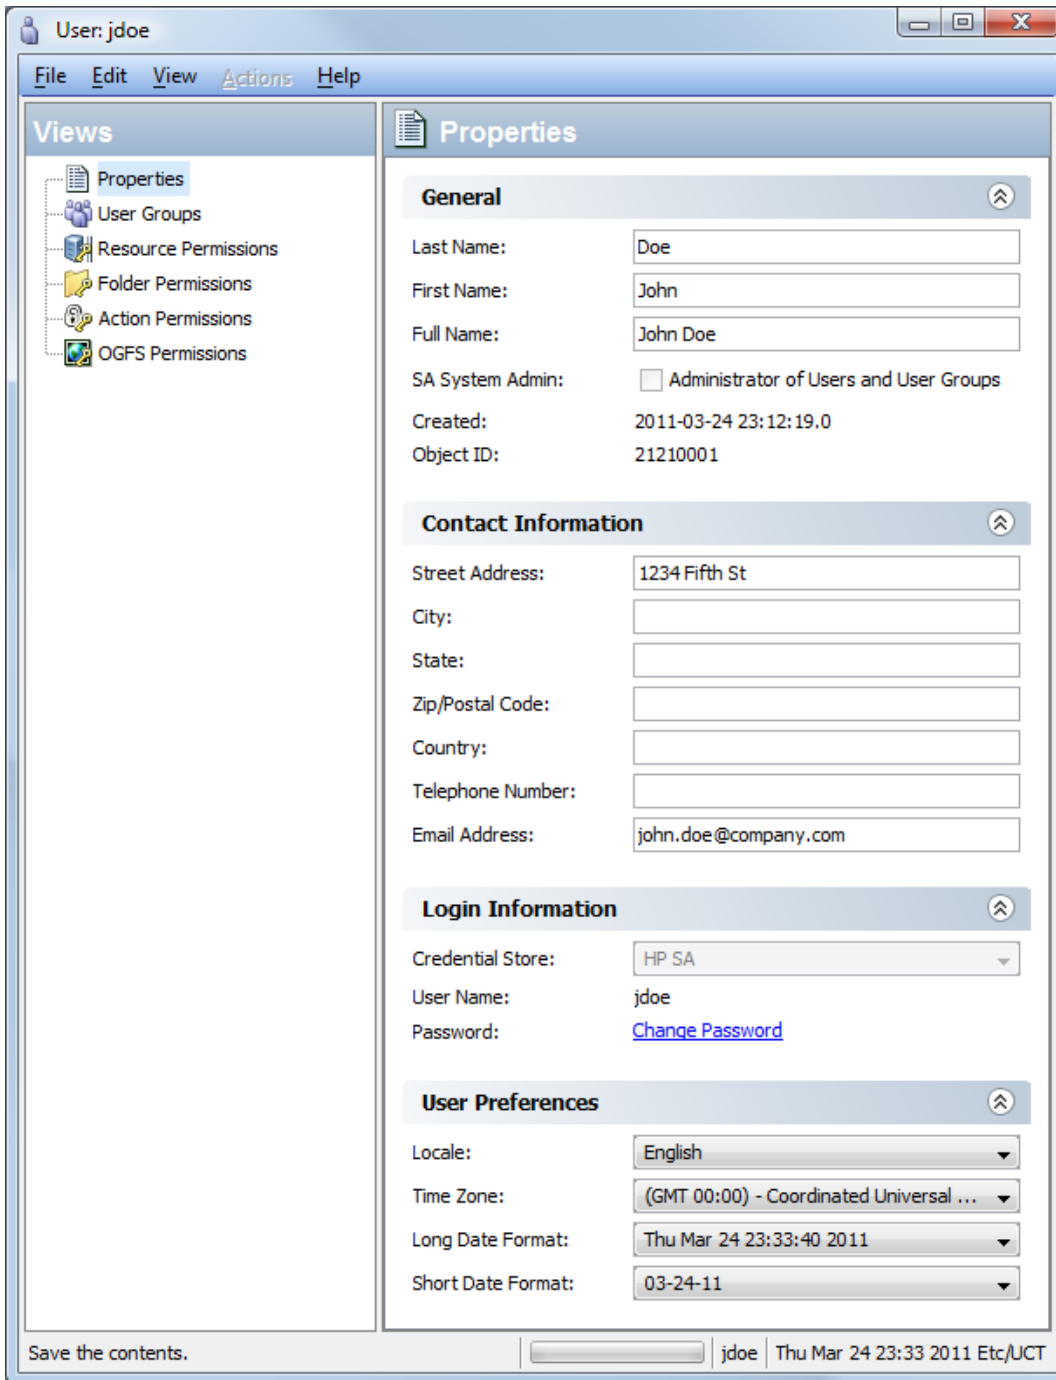
Any user can change their own password and their profile information.

Figure 11. Users Changing Their Own Password



1. From the SA Client screen, select the “Logged in as” link in the upper right corner, as shown in the previous figure. This displays the user properties window, as shown in **Figure 12**.

Figure 12. User Properties Window and Change Password Link



2. To change password, select the Change Password link. Note that when modifying a password, the change takes effect immediately.
3. Change other properties as needed.
4. If any properties were changed, select **File > Save**.
5. Select **File > Close**.

Changing a User

To modify an SA user from the SA Client, perform the following steps.

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your SA users.
4. Select the user you want to modify.
5. Select the **Actions** menu, or right-click and select **Open**. This displays the user information in a new window.
6. Optionally modify any of the user's properties. The **Properties** view lists the user's name, contact information, login information, where their credentials are stored, their user name, a link to change their password, and their date and time settings. Note that when you modify the user's password, the change takes effect immediately.
7. Optionally add or remove the user from a user group. The **User Groups** view lists the user groups to which the user belongs. Each user group grants a set of permissions to all the users who belong to the group.
8. The permissions are viewable but not modifiable from the user window. To modify permissions, you need to modify user groups as described in [Setting Permissions on User Groups - SA Client](#).
9. Select **File > Revert** to discard your changes.
10. Select **File > Save** to save the changes.

Deleting a User

To delete an SA user from the SA Client, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your SA users.
4. Select one or more users you want to delete.
5. Select the **Actions > Delete** menu, or select the delete icon.

Finding the User Group a Particular Action Permission Comes From

If a user belongs to more than one user group, you can determine which user group grants a particular action permission as follows.


1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your SA users.

4. Select the user you want to view.
5. Select the **Actions** menu, or right-click and select **Open**. This displays the user information in a new window.
6. Select the Action Permissions view. This displays all the action permissions organized by the user groups to which the user belongs.
7. You can also right-click on any column header and ungroup the User Group column, then use the column selector at the far right of the column headers to display the User Group column. This will show each permission followed by the user group that grants that permission.

Suspending a User

A suspended user cannot log in to SA, but the user name has not been deleted. A suspended user is indicated by a status of Suspended in the SA Client. A user can be suspended in the following ways:


- **Login Failure:** If you select the check box labeled Login Failure on the Security Settings tab, and someone tries to log in with the wrong password a specified number of times, the user account is suspended. For instructions on accessing the Security Settings tab, see the first two steps of [Resetting Initial Passwords](#).
- **Account Inactivity:** If you select the check box labeled Account Inactivity on the Security Settings tab, and the user has not logged on for the specified number of days, the user account is suspended.
- **Expired Password:** A user can be suspended if the password has expired and the expiration count is full.
- **Suspend:** You can suspend a user's account as described below. If the user is logged in, a message will be displayed and they will be logged out.

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your users.
4. Select the user you want to suspend.
5. Select the  **Suspend** button or select **Actions > Suspend**.

Activating a Suspended User

To activate a suspended user, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your users.

4. Select the suspended user you want to activate.
5. Select the  **Activate** button or select **Actions > Activate**.

Assigning a User to a User Group

Assign each SA user to a group reflecting the user's role in your organization. To assign an SA user to a user group, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your SA users.
4. Select the user you want to assign.
5. Select the **Actions** menu or right-click and select **Open**. This displays the user information in a new screen.
6. Select the User Groups view. This displays the user groups that the user is a member of.
7. Select the "+" button or select the **Actions > Add** menu. This displays all the user groups.
8. Select one or more user groups.
9. Select the Select button. This adds the user to the user groups.
10. Select **File > Revert** to discard your changes.
11. Select **File > Save**.

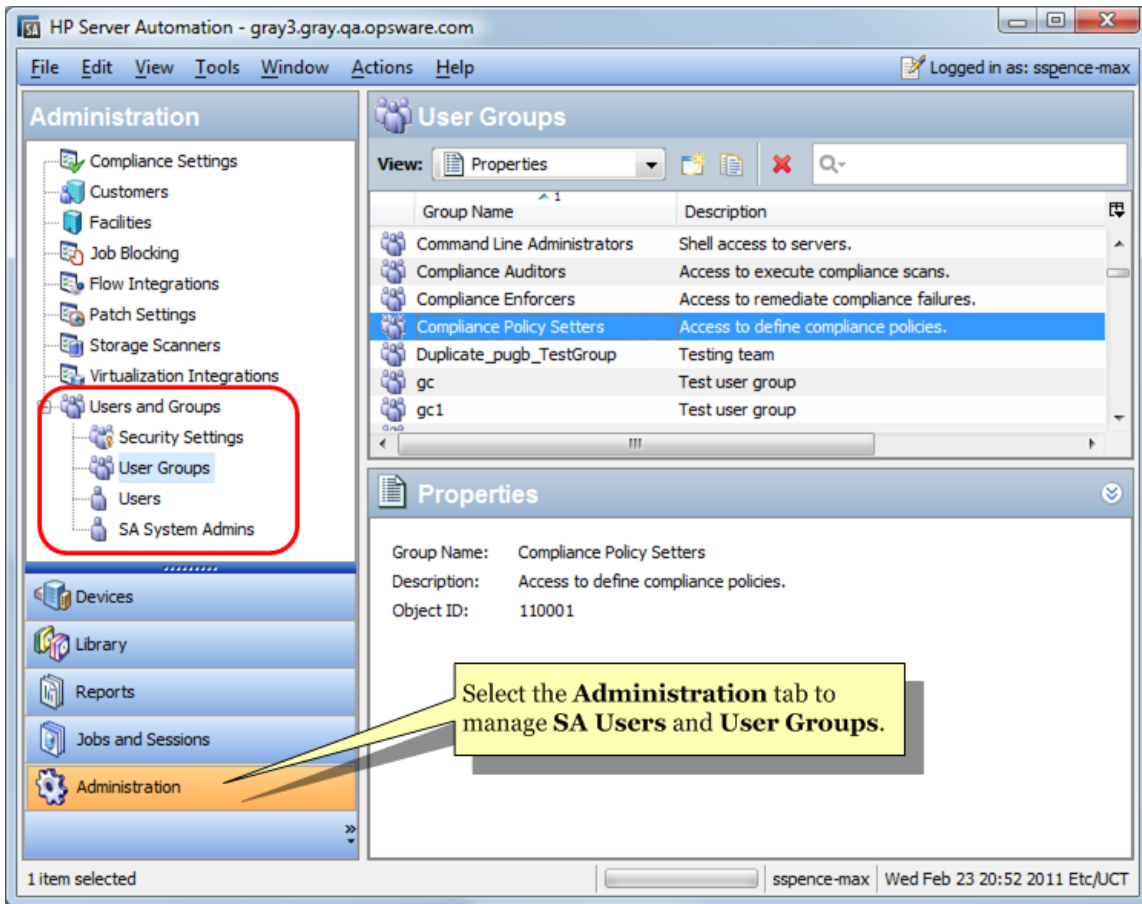
Importing Users from an LDAP Directory

You can import user information from an LDAP directory and use the LDAP directory for authentication when logging into SA. For more information, see [Authenticating with an External LDAP Directory Service](#).

Managing User Groups - SA Client

This section describes how perform tasks with user groups. To manage user groups, you must log in to the SA Client as a super administrator (`admin`) and select the Administration tab, as shown in Figure 13.

Figure 13. User Groups Listed Under the Administration Tab



Creating a New User Group

To create a new user group from the SA Client, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.
3. Select the User Groups node. This displays all your user groups.
4. Select the Actions menu or right-click and select the **New** menu. This displays the new user group window.
5. Select the Properties view. Enter the name and a description for the user group.
6. Select **File > Save** to save the new user group.
7. Set the permissions for the user group and add users to the user group as described in [Setting Permissions on User Groups - SA Client](#).
8. Select **File > Revert** to discard your changes.
9. Select **File > Save** to save your changes.

Viewing User Groups

To view your user groups from the SA Client, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.
3. Select the User Groups node. This displays all your user groups.
4. Select a user group to display information about that user group.
5. In the View drop-down list, select any of the following:
 - **Properties** displays the name, description, and SA object ID of the selected user group.
 - **Users** displays all the SA users who are members of the selected user group.
 - **Resource Permissions** displays the customers, facilities, and device groups members of the user group have access to. It also lists the type of access to each customer, facility, and device group: Read access or read and write access.
 - **Folder Permissions** shows the access permissions to folders in the SA Library granted to members of the group.
 - **Action Permissions** show the actions that members of the user group can perform with the SA Client.
 - **OGFS Permissions** show the Global Shell and Global File System actions that members of the user group can perform, the resources they have access to, Global File System, and what user name they will use to log in to managed servers to perform those actions.

Copying a User Group

You can duplicate an existing user group as follows.

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.
3. Select the User Groups node. This displays all your user groups.
4. Select the user group that you want to copy.
5. Select the duplicate icon or select the **Actions > Duplicate** menu, or right-click on the user group and select the Duplicate menu. This displays the Duplicate User Group screen.
6. Enter the name and a description of the new user group. The name must be unique.
7. Select the Duplicate button. This creates a new user group that is a copy of the existing user group.

Changing a User Group

User groups define resource, folder, action, and OGFS permissions. Every user who is a member of the user group has those permissions. To modify a user group from the SA Client, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.
3. Select the User Groups node. This displays all your user groups.
4. Select a user group. This displays information about that user group in the lower part of the screen.
5. Select the **Actions** menu or right-click and select the **Open** menu. This displays the user group in a new window.
6. In the navigation pane, select any of the following views:
 - **Properties** displays the name, description, and SA object ID of the selected user group. You can change the name and description of the user group.
 - **Users** displays all the SA users who are members of the selected user group. Use the “+” and “-” buttons to add and remove users from the user group. For more information, see [Adding a User to a User Group](#).
 - **Resource Permissions** displays the facilities, customers, and device groups to which members of the user group have access. It also lists the type of access granted to each facility, customer, and device group: read access or read and write access. Use the “+” and “-” buttons to add and remove facilities, customers, and device groups from the user group and to set the type of access. For more information, see [Setting Resource Permissions - Facilities, Customers, and Device Groups](#).
 - **Folder Permissions** displays the folders in the SA Library and the access permission granted to each folder for the user group. Select a folder, select the **Actions** menu or right-click and select the **Folder Properties** menu to display the folder properties window. Select the Permissions tab to view and modify the permissions. For more information, see [Setting Folder Permissions](#).
 - **Action Permissions** displays the tasks that can be performed by members of the user group. Select the Permission column next to the permission you want to change and select the new permission. For more information, see [Setting Action Permissions](#).
 - **OGFS Permissions** displays the OGFS and Global Shell (OGSH) permissions. Select the “+” and “-” icons to add and remove permissions. For more information, see [Setting OGFS Permissions](#).
7. Select **File > Revert** to discard your changes.
8. Select **File > Save**.

Deleting a User Group

You can delete one or more existing user groups as follows.

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.

3. Select the User Groups node. This displays all your user groups.
4. Select one or more user groups that you want to delete.
5. Select the delete icon, select the **Actions > Delete** menu, right-click on the user group and select the **Delete** menu, or press the Delete key on your keyboard.

Adding a User to a User Group

You can add one or more users to any user group as follows.

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.
3. Select the User Groups node. This displays all your user groups.
4. Select a user group. This displays information about that user group in the lower part of the screen.
5. Select the **Actions** menu or right-click and select the **Open** menu. This displays the user group in a new screen.
6. In the navigation pane, select the Users view. This displays all the users who are members of the group.
7. Select the "+" icon or the **Actions > Add** menu. This displays all the SA users.
8. Select one or more users.
9. Select the Select button. This adds the users to the user group.
10. Select **File > Revert** to discard your changes.
11. Select **File > Save**.

Setting Permissions on User Groups - SA Client

This section describes how to set **action permissions**, **resource permissions**, **folder permissions** and **OGFS permissions** for a user group. All those permissions are granted to the users who are members of the user group.

Setting Resource Permissions - Facilities, Customers, and Device Groups

All managed servers are grouped by customers, facilities, and device groups. The **Resource Permissions** view lists the **customers**, **facilities**, and **device groups** the user group has access to. For more information, see [About Resource Permissions](#).

To modify resource permissions for a user group, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.

3. Select the User Groups node. This displays all your user groups.
4. Select a user group. This displays information about that user group in the lower part of the screen.
5. Select the **Actions** menu, or right-click and select the **Open** menu. This displays the user group in a new screen.
6. In the navigation pane, select the Resource Permissions view. This displays all the facilities, customers, and device groups to which the user group has access.
7. To add access to a customer, perform the following steps:
 1. Select the “+” icon under the Customers heading. This displays a list of all customers in a separate window.
 2. Select one or more customers.
 3. Select the access, either Read or Read & Write.
 4. Select the Add button.
8. To remove access to a customer, select the customer and select the “-” button.
9. To add access to a facility, perform the following steps:
 1. Select the “+” icon under the Facilities heading. This displays a list of all facilities in a separate window.
 2. Select one or more facilities.
 3. Select the access, either Read or Read & Write.
 4. Select the Add button.
10. To remove access to a facility, select the facility and select the “-” button.
11. To add access to all device groups, select the check box labeled Allow access to all device groups.
12. To add access to a subset of device groups, perform the following steps:
 1. Clear the check box labeled Allow access to all device groups. This displays the “+” icon.
 2. Select the “+” icon under the Device Groups heading. This displays a list of all public device groups in a separate window.
 3. Select one or more device groups.
 4. Select the access, either Read or Read & Write.
 5. Select the Add button.
13. To remove access to a device group, select the device group and select the “-” button.
14. Select **File > Revert** to discard your changes.
15. Select **File > Save**.

Setting Action Permissions

This section describes how to set action permissions for a user group. For more information, see [About Action Permissions](#).

To modify action permissions for a user group, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.
3. Select the User Groups node. This displays all your user groups.
4. Select a user group. This displays information about that user group in the lower part of the screen.
5. Select the **Actions** menu or right-click and select the **Open** menu. This displays the user group in a new screen.
6. In the navigation pane, select the Action Permissions view.
7. Locate the permission you want to modify using the Name and Description columns. You can right-click on any column to group or ungroup by that column for easier browsing.
8. Select the current value for the permission in the Permission column. This displays a drop-down list of the available values. Select the desired value.

Tip: You can select and set multiple permissions simultaneously. Select multiple permissions by dragging the mouse, or by using the Shift and Control keys on your keyboard and the mouse. Right-click to display the available permission values, then select the desired values. If a permission value is grayed out, that permission is controlled by another, related permission that needs to be changed first. For example, the permissions “Create Applications” and “Manage Application Deployment” both require that the permission “Access Application Deployment” be set to Yes before they can be set.

9. Select **File > Revert** to discard your changes.
10. Select **File > Save**.

Setting Folder Permissions

This section describes how to set folder permissions for a user group. For more information, see [About Folder Permissions](#).

To modify folder permissions for a user group, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.
3. Select the User Groups node. This displays all your user groups.
4. Select a user group. This displays information about that user group in the lower part of the screen.
5. Select the **Actions** menu or right-click and select the **Open** menu. This displays the user group in a new screen.
6. In the navigation pane, select the Folder Permissions view. This displays all the folders in the SA Library and their current permissions.
7. Locate and select the folder you want to modify.

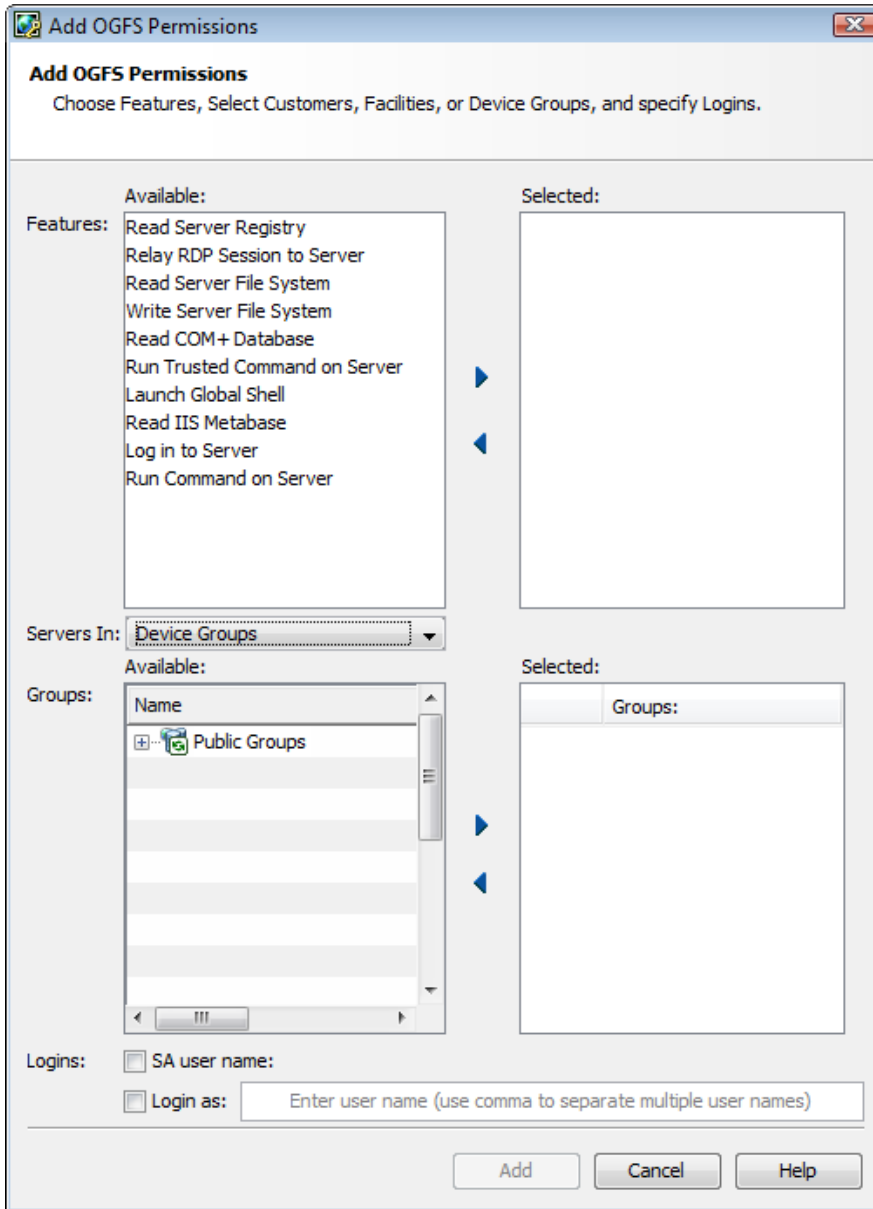
8. Select the **Actions** menu or right-click and select the **Folder Properties** menu. This displays the folder properties in a new window.
9. Select the Permissions tab. This displays all the users and user groups that have access to the folder.
10. Select a user or a user group. This displays the current access permissions at the bottom of the window.
11. Set the access permissions at the bottom of the screen.
12. To optionally give access to other users or user groups, select the Add button, select one or more users or user groups and select the Add button.
13. To optionally remove access for a user or user group, select the user or user group and select the Remove button.
14. Select the OK button.
15. Select **File > Revert** to discard your changes.
16. Select **File > Save**.

Setting OGFS Permissions

This section describes how to set OGFS permissions for a user group. For more information, see [About Global File System Permissions](#).

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the User Groups node.
3. Select the User Groups node. This displays all your user groups.
4. Select a user group. This displays information about that user group in the lower part of the screen.
5. Select the **Actions** menu or right-click and select the **Open** menu. This displays the user group in a new window.
6. In the navigation pane, select OGFS permissions. This displays the current OGFS permissions.
7. To add permissions, select the “+” icon. This displays the Add OGFS Permissions window, as shown in **Figure 14**. This screen has three main parts:
 - **Features** lists the action permissions for performing tasks with the OGFS and OGS.
 - **Groups** lists the servers that the actions can be performed on. Servers are grouped by facilities, customers or device groups.
 - **Logins** specifies the login name to be used when connecting to servers using the OGFS and OGS.

Figure 14. Add OGFS Permissions Window



8. In the Features section, select the OGFS actions you want to grant under the Available list. Select the arrow to move those actions to the Selected list.
9. In the Groups section, first select the type of server group you want to select from in the Servers In drop-down list. Select either Customers, Facilities or Device Groups.
10. Select one or more customers, facilities or device groups. Select the arrow to move them to the Selected list.
11. In the Logins section, select the check box labeled SA user name if you want OGFS users to log in with their SA user name. Otherwise select the check box labeled Login as and enter one or more user names for logging into servers with the OGFS.

12. Select the Add button.
13. To remove permissions, select one or more permissions and select the “-” button.
14. Select **File > Revert** to discard your changes.
15. Select **File > Save** to save your changes.

For more information on OGFS permissions, see [About Global File System Permissions](#).

Setting Private User Group Permissions

Note: Private user groups are intended for migrating scripts into folders in the SA Library. You should not assign permissions to users using private user groups. You should use regular user groups. For more information, see [About SA Users and User Groups](#).

For information about private user groups, see [About Private User Groups](#). To modify a private user group, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your SA users.
4. Select the user you want to set private user group permissions for.
5. Select the **Actions** menu or right-click and select **Open**. This displays the user information in a new window.
6. Select the User Groups view. This displays all the user groups the user is a member of, including the private user group. The private user group has the same name as the user.
7. Select the private user group.
8. Select the **Actions** menu or right-click and select **Open**. This displays the private user group in a new window.
9. To modify resource permissions, select the Resource Permissions view. For more information, see [Setting Resource Permissions - Facilities, Customers, and Device Groups](#).
10. To modify action permissions, select the Action Permissions view. For more information, see [Setting Action Permissions](#).
11. Select **File > Revert** to discard your changes.
12. Select **File > Save** to save the changes.

Setting Password, Account, and Session Security Policies - SA Client

You can set several policies to keep your SA user passwords secure, automatically disable inactive user accounts, and automatically lock inactive user sessions. Perform the following steps:

1. In the SA Client, select the Administration tab.
2. In the navigation panel, open the Users and Groups node. This displays the Security Settings node.
3. Select the Security Settings node. This displays the password policy settings.
4. Set any of the following policies:
 - **Reset** forces each user to reset their password the first time they log in to SA.
 - **Expiration** forces each user to change their password after the specified number of days. You can also specify how many times the user can postpone the change before it is required by specifying a number for “Allow graceful logins.”
 - **Retention** specifies how many previous passwords to save. This setting prohibits users from reusing passwords. For example, if you specify 10, users cannot reuse their previous ten passwords.
 - **Login Failure** specifies how many times someone can attempt to log in with the wrong password before the user account is suspended. When a user account is suspended you can reactivate it by selecting **Administration > Users and Groups**, selecting the user and selecting the Activate button. For more information, see [Suspending a User](#).
 - **Account Inactivity** specifies how long a user account can be unused before it is suspended. When a user account is not used for the specified number of days, the user account is suspended. When a user account is suspended you can reactivate it by selecting **Administration > Users and Groups**, selecting the user and selecting the Activate button. For more information, see [Suspending a User](#).
 - **SA Client Session Inactivity** specifies how long a user session can be idle before the SA Client is locked. Specify a value in minutes.
5. To revert to the previously saved settings, select the **View > Refresh** menu or press the F5 key on your keyboard.
6. After setting the values you want, select the Save button.

Resetting Initial Passwords

To require users to reset their passwords the first time they log in to SA, perform the following steps:

1. In the SA Client, select the Administration tab.
2. In the navigation panel, open the Users and Groups node. This displays the Security Settings node.
3. Select the Security Settings node. This displays the password policy settings.
4. Set the check box labeled “Reset password on first login.”
5. Select the Save button.

Setting Password Expiration

To require SA users to change passwords after a certain number of days, perform the following steps:

1. In the SA Client, select the Administration tab.
2. In the navigation panel, open the Users and Groups node. This displays the Security Settings node.
3. Select the Security Settings node. This displays the password policy settings.
4. Select the check box labeled Expiration.
5. Enter the number of days before password expiration.
6. Enter the number of graceful logins with the old password that will be allowed before the user is suspended.
7. Select the Save button.

To activate a suspended user, see [Activating a Suspended User](#).

Prohibiting Reuse of Old Passwords

To save a copy of users' old passwords and prevent them from reusing their old passwords, perform the following steps.

1. In the SA Client, select the Administration tab.
2. In the navigation panel, open the Users and Groups node. This displays the Security Settings node.
3. Select the Security Settings node. This displays the password policy settings.
4. Set the check box labeled Retention.
5. Enter the number of old password to save and prohibit.
6. Select the Save button.

Suspending User Accounts After Login Failures

You can suspend a user account if someone attempts to log in with the wrong password after a certain number of tries as follows.

1. In the SA Client, select the Administration tab.
2. In the navigation panel, open the Users and Groups node. This displays the Security Settings node.
3. Select the Security Settings node. This displays the password policy settings.
4. Set the check box labeled Login Failure.
5. Enter the number of failed login attempts. If someone tries to log in to any account and fails after the specified number of tries, the user account will be suspended.
6. Select the Save button.

To activate a suspended user, see [Activating a Suspended User](#).

Suspending Inactive User Accounts

You can automatically suspend user account if they do not log in for a certain period of time.

1. In the SA Client, select the Administration tab.
2. In the navigation panel, open the Users and Groups node. This displays the Security Settings node.
3. Select the Security Settings node. This displays the password policy settings.
4. Set the check box labeled Account Inactivity.
5. Enter the number of days. If any user does not log in for the specified number of days, the user account will be suspended.
6. Select the Save button.

To activate a suspended user, see [Activating a Suspended User](#).

Locking Inactive Sessions

You can automatically lock any SA Client session if the user has been inactive for a certain period of time. The user must enter their password to unlock the session.

1. In the SA Client, select the Administration tab.
2. In the navigation panel, open the Users and Groups node. This displays the Security Settings node.
3. Select the Security Settings node. This displays the password policy settings.
4. Set the check box labeled SA Client Session Inactivity.
5. Enter the number of minutes. If any logged in user does use the SA Client for the specified number of minutes, the SA Client will be locked and the user will have to enter their password.
6. Select the Save button.


Displaying a User Login Agreement

You can display a message whenever a user logs in and require that they acknowledge the message. Perform the following steps:

1. In the SA Client, select the Administration tab.
2. In the navigation panel, open the Users and Groups node. This displays the Security Settings node.
3. Select the Security Settings node. This displays the user agreement settings and the banner settings.
4. Under User Agreement Settings, select “Enable display.”
5. Enter the text you want displayed in the user agreement.
6. Select the Save button.

Whenever any user logs in to the SA Client, the specified message is displayed and the user must acknowledge the message, as shown in **Figure 15**.

Figure 15. User Login Confirmation Dialog



Displaying a Banner on the SA Client Screen

You can display a message at the top of each SA Client screen in any background color. Perform the following steps:

1. In the SA Client, select the Administration tab.
2. In the navigation panel, open the Users and Groups node. This displays the Security Settings node.
3. Select the Security Settings node. This displays the user agreement settings and the banner settings.
4. Under Banner Settings, select “Enable banner display.”
5. Select either a color from the drop-down list or specify a hexadecimal color code between 000000 and FFFFFFFF. The first 2 digits are the red component, the second 2 digits are the green component and the last 2 digits are the blue component.
6. Enter the text you want displayed in the banner.
7. Select the Save button. This displays the banner at the top of all SA Client screens as shown in **Figure 16**.

Figure 16. SA Client Banner Settings



Managing Super Administrators - SA Client

Super administrators can assign permissions to user groups and assign users to user groups. To manage super administrators, you must log in to the SA Client as a super administrator. When SA is first installed, the default super administrator is the `admin` user. See also [About Super Administrators and Super Users](#).

Viewing All SA Super Administrators

To view all SA super administrators, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Super Administrators node.
3. Select the Super Administrators node. This displays all your super administrators.

Creating a Super Administrator

An SA super administrator is an SA user who can create and modify SA users and user groups. To create an SA super administrator, follow the steps described in [Creating a New User](#) and check the

box labeled “Super Administrator.”

To make an existing user into a Super Administrator, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Super Administrators node.
3. Select the Super Administrators node. This displays all your super administrators.
4. Select the **Actions > Add** menu, or select the New User icon. This displays a list of all SA users.
5. Select one or more users that you want to make super administrators.
6. Click the Select button. This changes the selected users into super administrators.

Deleting a Super Administrator

To remove super administrator privileges from an SA user and leave that user’s other permissions, follow the steps described in [Changing a User](#) and clear the check box labeled Super Administrator. Alternatively, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Super Administrators node.
3. Select the Super Administrators node. This displays all your super administrators.
4. Select one or more users.
5. Select the **Actions > Remove** menu, right-click and select Remove, or select the remove button.

Managing Customer Administrators and Customer Groups - SA Client

One way to organize your servers and provide access control boundaries is to organize your managed servers by customer. A customer represents a set of servers associated with a business organization, such as a division or a company. Typically a server is associated with a customer because it runs applications for that customer. For more information on creating and managing customers, see the SA User Guide: Server Automation.

You can delegate super administrator tasks to a customer administrator. A **customer administrator** manages the users who manage the servers assigned to a customer. A customer administrator is a super administrator with access only to certain user groups.

You create customer administrators by creating customer groups and assigning customers and users to the customer group. For more information, see [About Customer Administrators and Customer Groups](#).

Viewing All Customer Administrators

A customer administrator is a user listed in a customer group. To view all SA customer administrators, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Super Administrators node.
3. Select the Super Administrators node. This displays all your super administrators and customer administrators. You can distinguish the two types of administrators by the icon as shown below:



Customer Administrator icon



Super Administrator icon

Viewing All Customer Administrators for a Customer Group

A customer administrator is a user listed in a customer group. To view all SA customer administrators for a customer group, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Under the Users and Groups node in the navigation pane, select the Customer Groups node. This displays all your customer groups.
3. Select a customer group.
4. Select the Users view. This displays all the users who are members of the customer group. These users are customer administrators for the customers listed in the customer group.

Viewing All Customers for a Customer Group

A customer administrator is a user listed in a customer group. To view all customers in a customer group, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Under the Users and Groups node in the navigation pane, select the Customer Groups node. This displays all your customer groups.
3. Select a customer group.
4. Select the Customers view. This displays all the customers who are members of the customer group.

Creating a Customer Group

A customer group associates one or more users with one or more customers and makes those users customer administrators. An SA customer administrator is an SA user who can modify all the user groups that have access to that customer. To create an SA customer administrator, you must create a customer group. Perform the following steps:

1. Log in to the SA Client as a super administrator, such as admin.
2. Select the Administration tab in the navigation pane.
3. Under the Users and Groups node in the navigation pane, select the Customer Groups node. This displays all your existing customer groups.
4. Select the **Actions > Add** menu or select the New Item icon.
5. Enter the name and a description of the customer group.
6. Select the Customers view.
7. Select the "+" icon or the **Actions > Add** menu. This displays all your customers.
8. Select one or more customers and press Select.
9. Select the Users view.
10. Select the "+" icon or the **Actions > Add** menu. This displays all your SA users.
11. Select one or more users that you want to add to the customer group and press Select.
12. Select **File > Save**.
13. Select **File > Close**.

Deleting a Customer Group

A customer group associates one or more users with one or more customers and makes those users customer administrators. An SA customer administrator is an SA user who can modify certain user groups. To delete a customer group, perform the following steps:

1. Log in to the SA Client as a super administrator, such as admin.
2. Select the Administration tab in the navigation pane.
3. Under the Users and Groups node in the navigation pane, select the Customer Groups node. This displays all your existing customer groups.
4. Select the customer group you want to delete.
5. Select the "X" icon or the **Actions > Delete** menu or right-click and select **Delete** or press the Delete key on your keyboard. This removes the selected customer groups.

Creating a Customer Administrator from the Customer Group View

An SA customer administrator is an SA user who can modify certain user groups. To create an SA customer administrator, add an SA user to a customer group. Perform the following steps:

1. Log in to the SA Client as a super administrator, such as admin.
2. Select the Administration tab in the navigation pane.
3. Under the Users and Groups node in the navigation pane, select the Customer Groups node. This displays all your existing customer groups.
4. Select a customer group. See also [Creating a Customer Group](#).
5. Select the **Actions > Open** menu or right-click and select **Open**. This opens the customer group in a separate window.
6. Select the Users view. This displays all the SA users who are members of that customer group.

7. Select the “+” icon or the **Actions > Add** menu. This displays all your SA users. See also [Creating a New User](#).
8. Select one or more users that you want to make customer administrators and press Select.
9. Select **File > Save**.
10. Select **File > Close**.

This allows the new customer administrator to modify the user groups with resource permissions to the customer.

Creating a Customer Administrator from the User View

An SA customer administrator is an SA user who can modify certain user groups. To create an SA customer administrator, add an SA user to a customer group. Perform the following steps:

1. Log in to the SA Client as a super administrator, such as admin.
2. Select the Administration tab in the navigation pane.
3. Under the Users and Groups node in the navigation pane, select the Users node. This displays all your existing SA users.
4. Select a user (see also [Creating a New User](#)).
5. Select the **Actions > Open** menu, or right-click and select **Open**. This opens the user in a separate window.
6. Select the Customer Groups view. This displays all the customer groups the user belongs to.
7. Select the “+” icon or the **Actions > Add** menu. This displays all your customer groups (see also [Creating a Customer Group](#)).
8. Select one or more customer groups, and press Select.
9. Select **File > Save**.
10. Select **File > Close**.

This allows the new customer administrator to modify the user groups with resource permissions to the customer.

Deleting a Customer Administrator from the Customer Group View

An SA customer administrator is an SA user who can modify certain user groups. To delete an SA customer administrator, remove that SA user from the customer groups to which the user belongs. Perform the following steps:

1. Log in to the SA Client as a super administrator, such as admin.
2. Select the Administration tab in the navigation pane.
3. Under the Users and Groups node in the navigation pane, select the Customer Groups node. This displays all your existing customer groups.
4. Select a customer group.
5. Select the **Actions > Open** menu, or right-click and select **Open**. This opens the customer group in a separate window.

6. Select the Users view. This displays all the SA users who are members of that customer group.
7. Select one or more users that you want to delete from the customer group, then select the “–” icon or the **Actions > Remove** menu, right-click and select **Remove**, or press the Delete key on your keyboard. This removes the selected SA users from the customer group so they are no longer customer administrators. The users are still valid SA users, however.
8. Select **File > Save**.
9. Select **File > Close**.

Deleting a Customer Administrator from the User View

An SA customer administrator is an SA user who can modify certain user groups. To delete an SA customer administrator, remove that SA user from the customer groups to which the user belongs. Perform the following steps:

1. Log in to the SA Client as a super administrator, such as admin.
2. Select the Administration tab in the navigation pane.
3. Under the Users and Groups node in the navigation pane, select the Users node. This displays all your existing SA users.
4. Select a user.
5. Select the **Actions > Open** menu, or right-click and select **Open**. This opens the user in a separate window.
6. Select the Customer Groups view. This displays all the customer groups to which the user belongs.
7. Select one or more customer groups from which you want to remove the user, then select the “–” icon; the **Actions > Remove** menu, right-click and select **Remove**, or press the Delete key on your keyboard. This removes the user from the customer groups.
8. Select **File > Save**.
9. Select **File > Close**.

Specifying Password Character Requirements

To specify character requirements for SA user passwords, perform the following steps:

1. Select the **Administration** tab in the SA Client.
2. In the navigation pane, select **System Configuration > Configuration Parameters**. This displays the SA components, facilities, and realms that have system configuration parameters.
3. In the list of SA components, select Server Automation System Web Client (occ). This displays the system configuration parameters for this component.

4. Locate the parameter `owm.features.Min>PasswordPolicy.allow`, and set it to `true`.

This parameter must be `true` for the other password parameters on this page to take effect. To disable the other password parameters, set `owm.features.Min>PasswordPolicy.allow` to `false`.

5. Set the values for the password parameters listed in **Table 10**.
6. Select the Revert button to discard your changes, or the Save button to save your changes.
7. To apply these parameter changes to other cores in a multimaster mesh, you must restart the other cores. For instructions, see [SA Maintenance](#).

Table 10. Password Requirements on the Modify Configuration Parameters Page

Password Requirement	Parameter	Allowed Values	Default Value
Maximum number of repeating, consecutive characters	<code>owm.pwpolicy.maxRepeats</code>	Must be greater than 0	2
Minimum number of characters	<code>owm.pwpolicy.minChars</code>	Positive integer	6
Minimum number of non-alphabetic characters	<code>owm.pwpolicy.minNonAlphaChars</code>	Must be less than the value of <code>owm.pwpolicy.minChars</code>	0

Authenticating with an External LDAP Directory Service

You can configure SA to use an external LDAP directory service for user authentication. With external authentication, you do not have to maintain separate user names and passwords for SA. When users log in to the SA Client, they enter their LDAP user names and passwords.

The LDAP directory is read-only to SA. After LDAP users are imported, any changes to the user attributes in the directory will require you to reimport the users from the LDAP directory.

Note: An SA Agent must be installed on all domain controllers in order for `rosh/ttlg` using Active Directory credentials to work.

Users Imported into SA from an LDAP Server

All SA user names must be unique, regardless of the authentication mechanism.

LDAP users must be successfully imported into SA before they can log onto SA.

Importing users from an LDAP directory must be done by the SA user administrator.

Imported users are managed in the same way as users created by the SA Client. For example, use the SA Client to assign imported users to user groups and delete imported users from SA.

If you delete an imported user with the SA Client, the user is not deleted from the external LDAP directory.

With the SA Client, search for users in the external LDAP, and then import selected users into SA. You can limit the search results by specifying a filter.

The LDAP import process fetches the following user attributes from the LDAP directory:

```
firstName
lastName
fullName
emailAddress
phoneNumber
street
city
state
country
```

SA also fetches LDAP user distinguished names (DN) during the import. The user DN is mapped to the SA user name.

After the import process, you may edit the imported user information within the SA Client. However, you cannot change the user login name or password. Importing a user is a one-time, one-way process. Changes to the user attributes you make using the SA Client are not propagated back to the external LDAP directory server.

If you use external authentication, you can still create separate users with the SA Client. However, this practice is not recommended, because of the likelihood of inadvertently creating duplicate users in the LDAP directory and in the SA Client. If there are duplicate users, the user defined in the SA Client will be used, and the user in the LDAP directory will be ignored.

To see which users have been imported in the SA Client, select the Administration tab, then select Users under the Users and Groups view. Make sure the Credential Store column is displayed. Users with Directory Server in the Credential Store column have been imported from the LDAP server.

SSL and External Authentication

Although SSL is not required for external authentication, it is strongly recommended. The certificate files needed for LDAP over SSL must be in Privacy Enhanced Mail (PEM) format. Depending on the LDAP server, you may need to convert the server's Certification Authority (CA) certificate to PEM format.

Supported External LDAP Directory Servers

You can use the following directory server products with SA:

- Microsoft Active Directory (Windows Server 2000, 2003, 2008, or 2012)
- Novell eDirectory 8.7
- SunDS 5.2

Importing a Server Certificate from the LDAP into SA

For SSL, the necessary certificates must be extracted from the LDAP directory and copied to SA. To import a server certificate from the LDAP directory into SA, perform the following steps:

1. Extract the server certificate from the external LDAP directory. For instructions, see the following sections.
2. Convert the extracted certificate to PEM format.

Certificates created on Windows systems are in Distinguished Encoding Rules (DER) format. The following example converts a certificate from DER to PEM format with the `openssl` utility:

```
OpenSSL> x509 -inform DER -outform PEM -in mycert.der -out mycert.pem
```

3. Copy the server certificate to the location specified by the LDAP configuration file (`twist_custom.conf`). For example, the `twist_custom.conf` file could have the following line:

```
aaa.ldap.server-  
cert.ca.fname=/var/opt/opsware/crypto/twist/ldapcert.pem
```

Extracting the Server Certificate from Microsoft Active Directory

To extract the server certificate, perform the following steps:

1. Run either the Certificates MMC snap-in console or the Certificate Services web interface.
2. Export the Root CA certificate from the Windows CA into DER format.

Extracting the Server Certificate from Novell eDirectory

To extract the server certificate, perform the following steps:

1. Find out the name of the local CA entry. (Example: CN=CORP-TREE CA.CN=N=Security)
2. Open the eDirectory Administration utility, and click **Modify Object**.
3. Enter the entry name (CN=CORP-TREE CA.CN=Security).
4. Select the Certificates tab.
5. Click **Self Signed Certificate**.
6. Click **Export**.
7. In the dialog, click No for exporting the private key, and then click **Next**.

8. Select the appropriate format (usually DER).
9. Click **Save the exported certificate to a file**.

Extracting the Server Certificate from SunDS

Typically, instead of exporting a server CA certificate from SunDS, you obtain the certificate that was imported into SunDS.

Importing External LDAP Users and User Groups

After you complete the tasks in this section, your users will be able to log in to the SA Client with their LDAP user names and passwords.

Note: This method does not import LDAP user groups. If you want to import users and user groups, see [Importing LDAP Users and Groups Using LDAP Authentication Configuration](#).

To import external users with the SA Client, perform the following steps:

1. In the SA Client navigation pane, select the Administration tab. This displays the Users and Groups node in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your SA users.
4. Select the **Actions > Import Users** menu. This displays information from your LDAP directory.
5. Select the Import Users tab. This displays all the users in your LDAP directory.
6. Select one or more users.
7. You can optionally assign the users to one or more users groups. Select the Assign Groups tab, and select one or more user groups.
8. Select the Import Users button. This imports the users into SA.

Importing LDAP Users and Groups Using LDAP Authentication Configuration

LDAP Authentication Configuration LDAP Authentication Configuration is a command line tool used to configure LDAP and import users and user groups into SA. This can be a complex process that requires some preparation.

Once LDAP has been configured, the LDAP Users & User Groups Synchronization APX can also be used to import LDAP users and user groups into SA.

Note: You should not edit user groups being maintained by LDAP synchronization. These user groups are indicated by the description, `__DO_NOT_EDIT__MAINTAINED_BY_LDAP_SYNC_`.

LDAP Authentication Configuration Prerequisites

The LDAP Authentication Configuration tool is a script that must be run on an SA Core's Slice Component bundle host. Before running the script, you must have the following information available:

Table 11. LDAP Authentication Configuration Prerequisites

Prerequisite	Description
Hostname	Semicolon separated list of fully-qualified host name (FQHN) or IP address of the LDAP directory server that SA is to use. Only the first listed host is used for communication, the other hosts are used to handle failover scenarios.
LDAP server port	The LDAP directory server port. The default SSL port is 636 and the default non-SSL port is 389. SA does not support StartTLS.
SSL	Is SSL authentication required by your LDAP directory server? If SSL is enabled, you must supply the trusted CA certificates used to validate the server's SSL certificate.
Trusted CA certificates to validate server SSL certificate	The complete path to the file on the LDAP directory server containing the trusted CA certificates, in PEM format, used to verify the LDAP directory server's SSL certificate.
SSL with mutual (or two-way) authentication	You must supply the following information: <ol style="list-style-type: none"> 1 Trusted CA certificates to validate server SSL certificate 2 Trusted CA certificates to validate client SSL certificate 3 Client certificate and (unencrypted) private key.
SSL with client authentication enabled	<ol style="list-style-type: none"> 1 The complete path to the file containing the trusted CA certificates, in PEM format, used to verify the SSL client certificate. 2 The complete path to the file containing the client SSL certificate and its corresponding private key, in PEM format. The client private key must not be encrypted.
Anonymous search to the Directory Information Tree (DIT)	Does the LDAP directory allow anonymous searches to the DIT where user information is stored? Note that this implies that anonymous bind is allowed. For example, does an anonymous user (a user who did not supply a bind DN and password) have read access to the DIT? For most enterprises, anonymous search is not allowed. If anonymous search is disabled, you must supply the bind DN and password of a user who has read access to the DIT.
Bind DN	Required only if anonymous search is disabled. The bind DN for the user who has read access to the DIT.

Prerequisite	Description
Bind password	Required only if anonymous search is disabled. The bind password for the user who has read access to the DIT.
Attribute for unique user name	<p>The attribute for the unique user name.</p> <ul style="list-style-type: none"> • For Active Directory, the default is <code>SAMAccountName</code>. • For Novell eDirectory, the default is <code>cn</code>. • For all other vendors, the default is <code>uid</code>.
Attribute for user display name	<p>The attribute for the user display name.</p> <ul style="list-style-type: none"> • For Active Directory, the default is <code>displayName</code>. • For Novell eDirectory, the default is <code>fullName</code>. • For all other vendors, the default is <code>cn</code>.
Base DN	The base DN, or the portion of the DIT to be considered when searching for users during the user import operation. The LDAP Authentication Configuration tool uses a subtree search; therefore, the search filter is only applicable to users at or below the base DN.
Search Filter Template	<p>The Search Filter Template is used, with optional filter substitution, as the filter in the LDAP search for the user import.</p> <p>Any dollar sign (\$) character in the template is replaced by the filter string specified in the Import Users page of the SA Client. (The default value is an asterisk (*), which matches all entries.)</p> <ul style="list-style-type: none"> • For Active Directory, the default is <code>(&(sAMAccountName=\$)(objectCategory=person)(objectClass=user)(sAMAccountType=805306368))</code>. • For Novell eDirectory, the default is <code>(&(cn=\$)(objectClass=person))</code>. • For all other vendors, the default is <code>uid=\$</code>.

The LDAP Authentication Configuration Process

When you run LDAP Authentication Configuration, you will be prompted depending on whether your LDAP Directory server requires SSL authentication and whether anonymous search is allowed.

Anonymous Search: **No**

SSL: **No**

1. Log in to a server hosting a Slice Component bundle for your SA Core.
2. Log in as the `twist` user:

```
su twist
```

3. Issue the following command:

```
cd /opt/opsware/twist
```

4. Invoke LDAP Authentication Configuration:

```
./ldap_config.sh
```

5. Enter the necessary information. Enter `N` when asked if anonymous search is allowed. Enter `N` when asked if SSL setup is required.
6. After the tool completes, ensure that LDAP authentication configuration is successfully validated and stored.
7. Log on to the Command Center and ensure that external user import works.
8. Ensure that you can log on to the Command Center as an LDAP user.

Note: When running the `ldap_config.sh` script to import ldap users into a Server Automation (SA) core, with a special bind configured, the following message might appear, and the script fails:

```
Error: failed to verify LDAP search configuration. message=null
Failed to verify LDAP search configuration with the specified
LDAP directory server.
Please correct your answers.
```

Additional tests with `ldapsearch` work, as does `ldap_config.sh` with a different base bind.

The error is caused when the `ldap_config.sh` script attempts to resolve a referral to one of the `DomainDnsZones` handling the bind data and encountered a timeout. Unless the script can follow the referral, it cannot validate/populate the ldap entry, resulting in the error messages.

To resolve this issue:

1. Verify that the `DomainDnsZones` are reachable from the core. For example, if you are trying to use a Base bind "`DC=A1,DC=B2,DC=C3,DC=com`", make sure that `DomainDnsZones.A1.B2.C3.com:636` is reachable from the core. If it is not, check if firewalls or routers are functioning correctly.

2. If using SSL with ldap, try running `ldap_config.sh` without SSL. If this works, use the following command to examine the certificate returned by AD:

```
openssl s_client -CAfile /var/-
opt/opsware/crypto/twist/ldapcert.pem -connect DomainDn-
sZones.LA.FRD.DIRECTV.com:636
```

3. If non-SSL works, add the LDAP server certificate into `/var/-
opt/opsware/crypto/twist/ldapcert.pem`.

Anonymous Search: Yes

SSL: No

1. Log in to a server hosting a Slice Component bundle for your SA Core.
2. Log in as the `twist` user:

```
su twist
```
3. Issue the following command:

```
cd /opt/opsware/twist
```
4. Invoke LDAP Authentication Configuration:

```
./ldap_config.sh
```
5. Enter the necessary information. Enter `N` when asked if anonymous search is allowed. Enter `N` when asked if SSL setup is required.
6. After the tool completes, ensure that LDAP authentication configuration is successfully validated and stored.
7. Log on to the Command Center and ensure that external user import works.
8. Ensure that you can log on to the Command Center as an LDAP user.

Anonymous Search: **No**

SSL: **Yes** (SSL server authentication only)

1. Log in to a server hosting a Slice Component bundle for your SA Core.
2. Log in as the `twist` user:

```
su twist
```
3. Issue the following command:

```
cd /opt/opsware/twist
```
4. Invoke LDAP Authentication Configuration:

```
./ldap_config.sh
```
5. Enter `N` when asked if anonymous search is allowed. Enter `Y` when asked if SSL setup is required. Answer `N` when asked whether to use SSL client authentication.
6. After the tool completes, ensure that LDAP authentication configuration is successfully validated and stored.
7. Log on to the Command Center and ensure that external user import works.
8. Ensure that you can log on to the Command Center as an LDAP user.

Anonymous Search: **No**

SSL: **Yes** (SSL mutual authentication required)

1. Log in to a server hosting a Slice Component bundle for your SA Core.
2. Log in as the `twist` user:

```
su twist
```
3. Issue the following command:

```
cd /opt/opsware/twist
```

4. Invoke LDAP Authentication Configuration:

```
./ldap_config.sh
```

5. Enter **N** when asked if anonymous search is allowed. Enter **Y** when asked if SSL setup is required. Enter **Y** when asked whether to use SSL client authentication.
6. After the tool completes, ensure that LDAP authentication configuration is successfully validated and stored.
7. Log on to the Command Center and ensure that external user import works.
8. Ensure that you can log on to the Command Center as an LDAP user.

Anonymous Search: Yes

SSL: Yes (SSL server authentication only)

1. Log in to a server hosting a Slice Component bundle for your SA Core.
2. Log in as the `twist` user:

```
su twist
```

3. Issue the following command:

```
cd /opt/opsware/twist
```

4. Invoke LDAP Authentication Configuration:

```
./ldap_config.sh
```

5. Enter **Y** when asked if anonymous search is allowed. Enter **Y** when asked if SSL setup is required. Enter **N** when asked whether to use SSL client authentication.

Anonymous Search: **Yes**

SSL: **Yes** (SSL mutual authentication required)

1. Log in to a server hosting a Slice Component bundle for your SA Core.
2. Log in as the `twist` user:

```
su twist
```

3. Issue the following command:

```
cd /opt/opsware/twist
```

4. Invoke LDAP Authentication Configuration:

```
./ldap_config.sh
```

5. Enter **Y** when asked if anonymous search is allowed. Enter **Y** when asked if SSL setup is required. Enter **Y** when asked whether to use SSL client authentication.

Note: The values shown as defaults are the values saved during the last LDAP Authentication Configuration Tool session.

Example LDAP Authentication Configuration Session

```
./ldap_config.sh

Retrieving LDAP configuration ...
LDAP Connectivity Configuration
Enter the fully-qualified host name or IP for the LDAP directory
server [sample-centos.example.com] :
Does the LDAP directory server require SSL? [N] :
Enter the port number for the LDAP directory server [8389] :
Does the LDAP directory server support anonymous bind and anonym-
ous read access to the directory information tree? [N] :
Enter the bind distinguished name (DN) of the user who has read
access to the directory information tree (DIT)
[cn=Administrator,cn=users,dc=hyrule,dc=local] :
Do you want to change the bind password for cn=A-
Administrator,cn=users,dc=hyrule,dc=local [N] :

You have entered the following information:
LDAP Directory Server FQHN/IP           : sample-cen-
tos.example.com
LDAP Directory Server Port              : 8389
SSL Enabled?                           : false
Bind DN                                 : cn=A-
Administrator, cn=users,dc=hyrule,dc=local
Bind Password Provided?                 : true

Is this correct? [Y] :

Verifying LDAP directory server connectivity ...
found naming context : DC=hyrule,DC=local
found naming context : CN=Configuration,DC=hyrule,DC=local
found naming context : CN=-
=Schema,CN=Configuration,DC=hyrule,DC=local
found naming context : DC=DomainDnsZones,DC=hyrule,DC=local
```

```
found naming context : DC=ForestDnsZones,DC=hyrule,DC=local
LDAP directory server connectivity successfully verified.
```

LDAP Search Configuration

```
Is the LDAP directory server an Active Directory (AD) directory
server? [Y] :
```

```
Enter the LDAP attribute for the unique username [SamAccountName]
:
```

```
Enter the LDAP attribute for the user's display name [cn] :
```

```
Enter the LDAP search filter template [(&(sAMAccountName=$)
(objectCategory=person)(objectClass=user)
(sAMAccountType=805306368))] :
```

```
Enter the LDAP search base distinguished name (DN). Usually this
is the root naming context. [cn=users,dc=hyrule,dc=local] :
```

You have entered the following information:

```
LDAP Unique Username Attribute : SamAccountName
```

```
LDAP User Display Name Attribute : cn
```

```
LDAP Search Filter Template : (&(sAMAccountName=$)(objectCat-
egory=person)(objectClass=user)
(sAMAccountType=805306368))
```

```
LDAP Search Base Distinguished Name (DN) : cn=use-
ers,dc=hyrule,dc=local
```

```
Is this correct? [Y] :
```

```
Verifying LDAP search configuration ...
```

```
To test LDAP search configuration, you must provide a username of
a LDAP directory user to search.
```

```
LDAP search configuration is successfully verified only if the
given user is successfully returned by the LDAP
directory server.
```

```
Enter a username to search : *
```

You have entered the following information:

Username To Search : *

Is this correct? [Y] :

Resulting LDAP Search Filter : (&(sAMAccountName=*)(objectCategory=person)(objectClass=user)(sAMAccountType=805306368))

Searching LDAP directory server for user * ...

Found 4 users

DN : CN=Administrator,cn=users,dc=hyrule,dc=local

cn : Administrator

SamAccountName : Administrator

DN : CN=Guest,cn=users,dc=hyrule,dc=local

cn : Guest

SamAccountName : Guest

DN : CN=krbtgt,cn=users,dc=hyrule,dc=local

cn : krbtgt

SamAccountName : krbtgt

DN : CN=link,cn=users,dc=hyrule,dc=local

cn : link

SamAccountName : link

Is this correct? [Y] :

LDAP search configuration successfully verified.

Enter the LDAP search filter template to search user groups [(&(cn=\$)(objectCategory=group))] :

Enter the LDAP attribute for the unique user group name [SamAccountName] :

Enter the LDAP attribute in the user group LDAP object class
which contains the DNs of its members [

member] :

You have entered the following information:

LDAP Search User Group Base DN : cn=users,dc=hyrule,dc=local

LDAP Search User Group Search Filter Template : (&(cn=\$)
(objectCategory=group))

LDAP Unique User Group Name Attribute : SamAccountName

LDAP Search User Group Membership Attribute : member

Is this correct? [Y] :

Verifying LDAP user group synchronization configuration ...

Searching LDAP directory server for all users and user groups ...

Searching LDAP directory server for all LDAP users ...

Resulting LDAP Search Filter For All LDAP Users : (&(sAMAc-
countName=*)(objectCategory=person)(object

Class=user)(sAMAccountType=805306368))

Found 4 LDAP users

Parsing search results ...

Searching LDAP directory server for all LDAP user groups ...

Resulting LDAP Search Filter For All LDAP User Groups : (&(cn=*)
(objectCategory=group))

Found 16 LDAP user groups

Parsing search results ...

Do you wish to display detail search result? [N] : y

Parsing search results ...

Denied RODC Password Replication Group: 2 members


```
Administrator : cn=administrator,cn=users,dc=hyrule,dc=local
krbtgt : cn=krbtgt,cn=users,dc=hyrule,dc=local
Allowed RODC Password Replication Group: 0 members
Enterprise Read-only Domain Controllers: 0 members
Group Policy Creator Owners: 1 members
Administrator : cn=administrator,cn=users,dc=hyrule,dc=local
Domain Controllers: 0 members
Cert Publishers: 0 members
Domain Users: 0 members
Enterprise Admins: 1 members
Administrator : cn=administrator,cn=users,dc=hyrule,dc=local
Schema Admins: 1 members
Administrator : cn=administrator,cn=users,dc=hyrule,dc=local
DnsAdmins: 0 members
Read-only Domain Controllers: 0 members
RAS and IAS Servers: 0 members
Domain Guests: 0 members
Domain Admins: 1 members
Administrator : cn=administrator,cn=users,dc=hyrule,dc=local
Domain Computers: 0 members
DnsUpdateProxy: 0 members
Is this correct? [Y] :
LDAP user group synchronization configuration successfully verified.
```

The following properties will be stored into global configuration.

```
aaa.ldap.hostname=gyee-centos.cup.hp.com
aaa.ldap.port=8389
aaa.ldap.ssl=false
aaa.ldap.search.-
binddn=cn=Administrator,cn=users,dc=hyrule,dc=local
aaa.ldap.search.pw=true
```

```
aaa.ldap.search.naming.attribute=SamAccountName
aaa.ldap.search.display.name.attribute=cn
aaa.ldap.search.filter.template=(&(sAMAccountName=*)(objectCategory=person)
(objectClass=user)(sAMAccountType=805306368))
aaa.ldap.search.base.template=cn=users,dc=hyrule,dc=local
aaa.ldap.enable.users.groups.sync=true
aaa.ldap.search.usergroup.naming.attribute=SamAccountName
aaa.ldap.search.usergroup.membership.naming.attribute=member
aaa.ldap.search.user-
group.base.template=cn=users,dc=hyrule,dc=local
aaa.ldap.search.usergroup.filter.template=(&(cn=*)(objectCategory=group))
```

Are you sure? [Y] :

Saving LDAP configuration ...

LDAP configuration successfully saved.

Synchronizing LDAP Users

After you have completed the LDAP Authentication Configuration process, you can use the `ldap_sync.sh` tool to synchronize LDAP users and groups with the SA database from the command line, as described below.

You can also run the LDAP Users & User Groups Synchronization APX from the SA Client to schedule the synchronization process. This program APX (formerly named, "ldap.user_and_user-groups_sync") is listed in the SA Client under **SA Library > By Type > Extensions > Program**.

Note: For instructions on running APXs, see "Run Extensions on Managed Servers" in the *SA User Guide: Server Automation*. This topic is also available in the SA Client help: From the list of Program APXs in the SA Client, click **F1** to open the page help, then click the heading link (Extensions: Properties) to open the how-to topic.

To synchronize LDAP users and user groups using `ldap_sync.sh`:

- 1 On a server hosting a Slice Component bundle for your SA Core, log in as the `twist` user:

```
su twist
```
- 2 Issue the following command:

```
cd /opt/opsware/twist
```
- 3 Invoke LDAP synchronization:

```
./ldap_sync.sh
```

You will see output similar to the following:

```
Retrieving LDAP configuration ...  
Verifying LDAP server connectivity ...
```

```
User Synchronization Phase  
Searching LDAP directory server for all LDAP users ...  
Found 4 LDAP users  
Parsing search results ...  
4 LDAP users do not exist in SA  
Creating them now ...  
Creating user cn=link,cn=users,dc=hyrule,dc=local  
Creating user cn=krbtgt,cn=users,dc=hyrule,dc=local  
Creating user cn=guest,cn=users,dc=hyrule,dc=local  
Creating user cn=administrator,cn=users,dc=hyrule,dc=local
```

```
User Group Synchronization Phase  
Searching LDAP directory server for all LDAP user groups ...  
Found 16 LDAP user groups  
Parsing search results ...  
creating user group Denied RODC Password Replication Group  
creating user group Allowed RODC Password Replication Group  
creating user group Enterprise Read-only Domain Controllers  
creating user group Group Policy Creator Owners  
creating user group Domain Controllers  
creating user group Cert Publishers  
creating user group Domain Users  
creating user group Enterprise Admins  
creating user group Schema Admins  
creating user group DnsAdmins  
creating user group Read-only Domain Controllers  
creating user group RAS and IAS Servers  
creating user group Domain Guests  
creating user group Domain Admins  
creating user group Domain Computers  
creating user group DnsUpdateProxy  
Updating user groups no longer found in LDAP ...
```

LDAP Users & User Groups Sync Results

```
=====
```

```
=  
Number of LDAP Users Found : 4  
Number of LDAP Users Does Not Exist In SA : 4  
Number of LDAP Users Successfully Created in SA : 4  
Number of LDAP Users Failed To Create In SA : 0
```

```
Number of LDAP User Groups Found : 16  
Number of LDAP User Groups Successfully Updated in SA : 0  
Number of LDAP User Groups Successfully Created in SA : 16  
Number of SA User Groups No Longer in LDAP : 0
```

```
Number of SA User Groups Failed To Update : 0  
Number of LDAP User Groups Failed To Process : 0
```

```
Elapsed Time : 00:00:27  
=====
```

LDAP users removed from the LDAP directory will not be removed from SA; however, these user will not be able to log in to SA because their corresponding authentication information has been removed from the LDAP directory.

LDAP user with the same user ID as an existing SA user will be skipped regardless of the user's credential store type. SA will neither create nor update duplicated users.

Additional Steps Required on FIPS-Enabled Cores

The following steps are required if you have a FIPS-enabled core:

- 1 Import the LDAP Server Certification Authorities (CAs) certificates:

```
# /opt/opsware/nss/nssimport.sh cert /tmp/ldap_server.crt  
twist.
```
- 2 For client authentication:
 - a Import the LDAP Client certificates:

```
/opt/opsware/nss/nssimport.sh cert /tmp/client.crt twist
```
 - b Import the LDAP Client private key:

```
/opt/opsware/nss/nssimport.sh key.pem <KeyFilePassword>  
twist.
```

SA Common Access Card (CAC) and Personal Identity Verification (PIV) Smart Card Integration

The Common Access Card (or CAC card) is a smart card about the size of a credit card. It is the standard identification for active-duty military personnel, Selected Reserve, United States Department of Defense (DoD) civilian employees, and eligible contractor personnel. It is also the principal card used to enable physical access to buildings and controlled spaces, and it provides access to defense computer networks and systems. It serves as an identification card under the Geneva Conventions (esp. the Third Geneva Convention). The CAC card meets two-factor authentication standards (something that belongs to the user, and something only known to the user) and standards for digital signature and data encryption technology (authentication, integrity and non-repudiation).

Note: SA/Smart card integration is available only when logging into the SA Client.

Smart Card/SA Integration Authentication Basics

The SA Client can discover the presence of a smart card and provide the user the option to login using regular SA authentication screen or using the new smart card based authentication. The SA Client works with the Card Reader API to access smart card certificates after the user provides the necessary PIN. The smart card's certificate is validated for revocation and unique certificate fields are mapped to an internal SA user account. An SA administrator creates the original mapping of these unique fields.

The information that identifies a user is stored on the smart card in a document called a certificate. This certificate contains an encryption key called a public key. It also contains text fields that identify the user, such as the person's name, usually the first, the last names and the middle initial or perhaps the user's email address within the organization. In order to be able to match a user's smart card authentication information with an existing SA username, the system constructs a username from the text data in the smart card certificate.

A certificate stored on a smart card looks similar to the following:

Certificate:

Data:

Version: 3 (0x2)

Serial Number: 1501 (0x5dd)

Signature Algorithm: ecdsa-with-SHA256

Issuer: C=US, O=Test Certificates 2010, OU=Test CA, CN=Test ECC
P-256 CA

Validity

Not Before: Oct 1 08:30:00 2010 GMT

Not After : Oct 1 08:30:00 2030 GMT

Subject: CN=Test E. Cardholder XV, C=US, O=Test Government, OU=Test
Agency

Subject Public Key Info:

Public Key Algorithm: id-ecPublicKey

EC Public Key:

pub:

04:03:a0:ad:22:46:01:b8:9b:1b:65:b0:94:3f:5e:

...

To derive a username from the certificate, SA uses a pattern specification string set up in the `/etc/opt/opsware/twist/twist.conf` file and a matching and assembly algorithm that constructs the username. The pattern specification might look like this:

```
sc.usernameMakeRule.1=%Subject#CN$1%Subject#CN$2%Subject#CN$3
```

The username creation logic would use the above specification string to create the username:

TestE.Cardholder

Field names from the certificate are specified by using the percent sign (%), the attributes (sub-fields) are specified with a pound sign (#) and positional fields within an attribute are specified using a dollar sign (\$) followed by a number (the position of the field in the text line).

This will be the default pattern supplied with the SA installation. SA administrators must be aware that this pattern may create user names which may NOT be unique and they should plan accordingly.


Note: Do not use Smartcard attributes in the algorithm for username construction.

You must decide on the pattern of creating usernames from smart card certificates prior to installation. It is very important to understand the mechanism, decide on a username creation pattern (you can accept the default pattern or specify a different pattern) and insure that administrators are trained to create smart card user accounts in SA using the correct pattern-based username.

SA Smart Card Integration Architecture

Figure 17 illustrates how CAC/PIV smart card functionality is integrated with SA:

Figure 17. SA/CAC Smart Card Integration Architecture



Setting Up SA/Smart Card Integration

Setting up new users to log in using a CAC smart card is simple:

- Create a new user and specify the credential store as SmartCard.
- When that user logs in to the SA Client, they will swipe their smart card and enter their unique pin number.

Setting Up Smart Card Certificates

- The `/etc/opt/opsware/twist/twist.conf` file must be modified as follows:
- For each signature algorithm, there must be an entry named `sc.sigAlgName.N` where `N` is the number in the series.
- For each algorithm, there must be a path to the certificate file (in `.pem` format) with the name `sc.trustedCertPath.N`.

For example:

```
sc.sigAlgName.0=SHA256withECDSA
sc.trustedCertPath.0=/var/opt/opsware/crypto/twist/smartcard/ECCP256
IssuingCACertificate.pem
sc.sigAlgName.1=SHA384withECDSA
sc.trustedCertPath.1=/var/opt/opsware/crypto/twist/smartcard/ECCP384
IssuingCACertificate.pem
sc.sigAlgName.2=SHA256withRSA
sc.trustedCertPath.2=/var/opt/opsware/crypto/twist/smartcard/RSA2048
IssuingCACertificate.pem
```

The location of the certificate files is optional but it is recommended that the certificate files be stored in the directory:

```
/var/opt/opsware/crypto/twist/smartcard/
```

Setting Up Smart Card Certificates on All Slice Hosts

You must perform the following steps *on each server in the SA Core that hosts a Slice Component bundle*.

1. Create the following folder:

```
mkdir /var/opt/opsware/crypto/twist/smartcard
```

2. For each Slice host, import the users' smart card certificates into the folder you created in Step 1:

```
/var/opt/opsware/crypto/twist/smartcard
```

3. Ensure that the ownership of these certificates is changed to `twist`:

```
chown -R twist:user /var/opt/opsware/crypto/twist/smartcard
```

4. Restart the Web Services Data Access Engine (`twist`) on each Slice host.
5. Set up a user and verify that the user can be authenticated using the smart card.

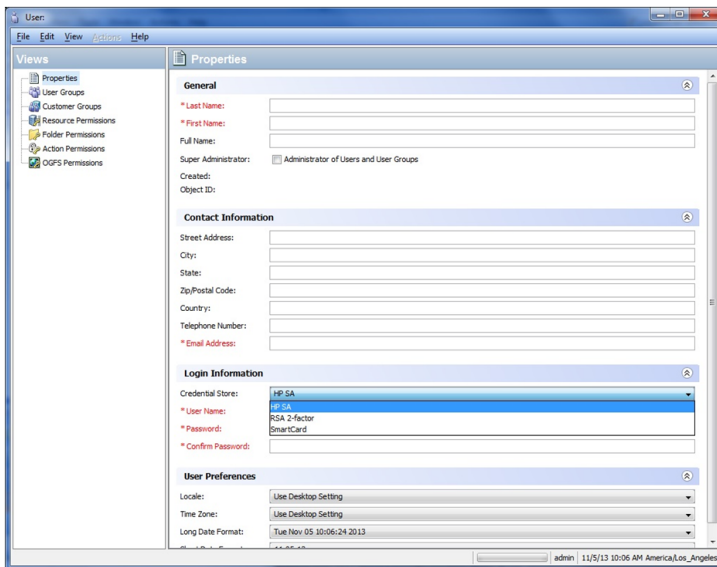
Creating a New Smart Card User

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

1. In the SA Client, select the Administration tab in the navigation pane.
2. Open the Users and Groups node in the navigation pane. This displays the Users node.
3. Select the Users node. This displays all your SA users.
4. Select the **Actions** > **New** menu or select the New User icon. This displays the New User window.

Complete the user information fields as described in [Creating a New User](#), specifying `SmartCard` as the Credential Store.

Note: When you select “SmartCard” as the credential store, the password field is removed from the screen because smart card access is done using smart card encryption techniques and not a preset password.



Note: As described above, the “User Name” field must contain a name which matches the name derived from the user smart card certificate according to the rules described in [Smart Card/SA Integration Authentication Basics](#). The administrator who creates the new smart card user must understand how the username construction pattern rules work so they can enter the text string that matches those rules.

Initial Login to the SA Client as a Smart Card User

When you start the SA Client, you see screens similar to **Figure 18** and **Figure 19**:

Figure 18. Standard SA Client Login Dialog

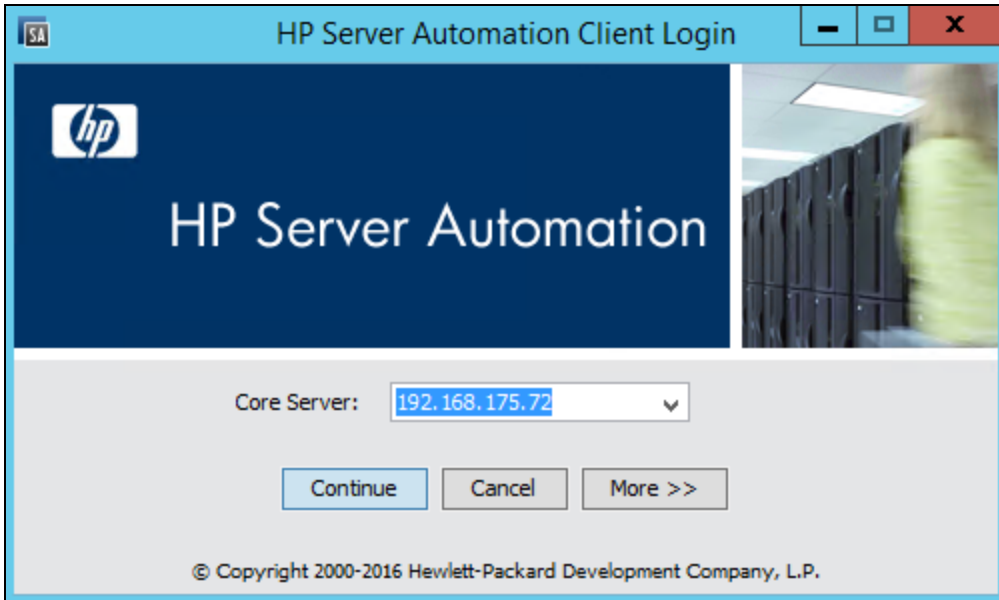
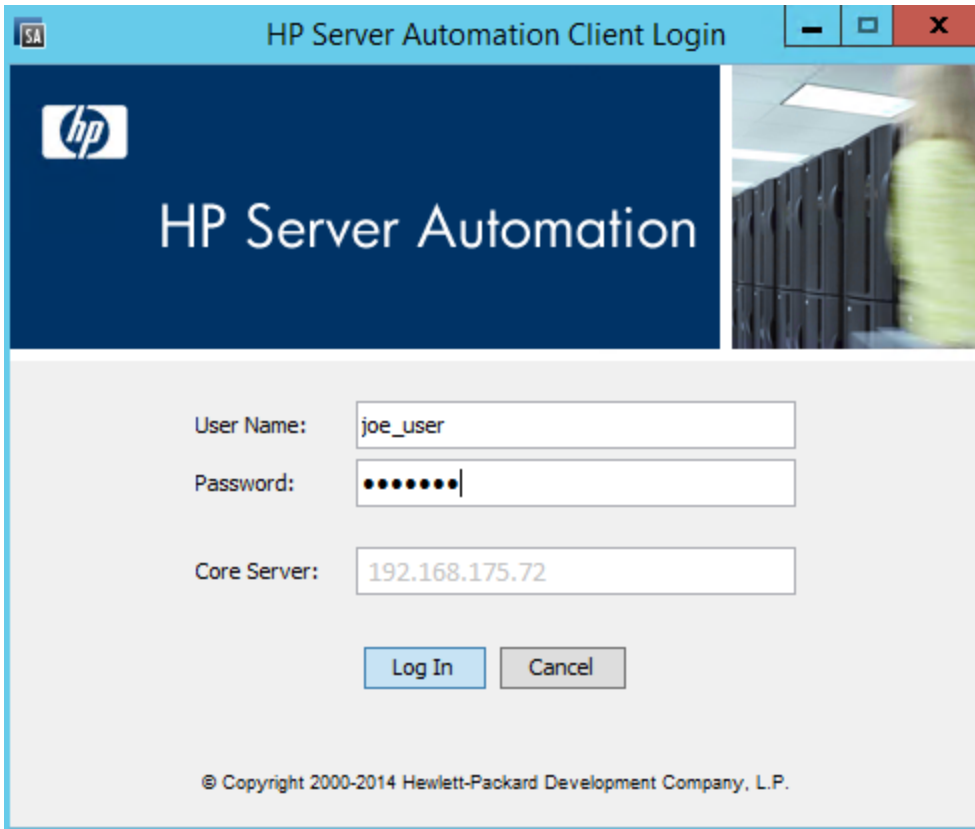
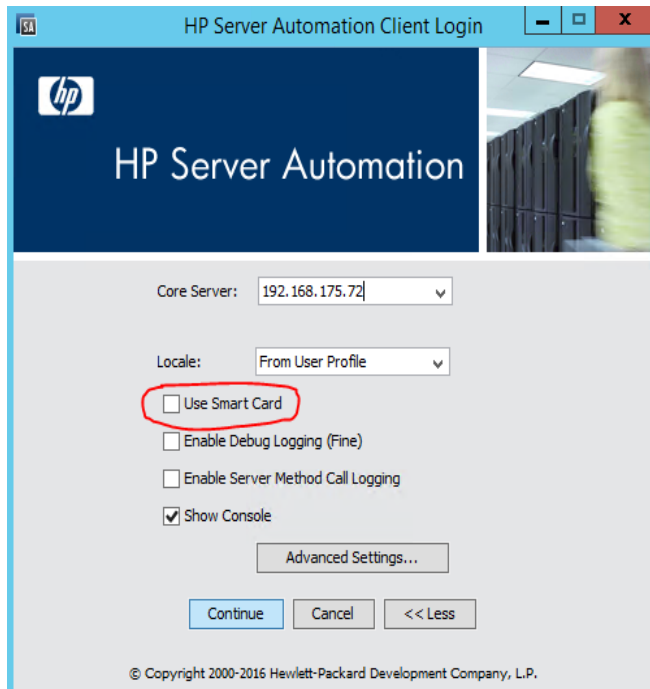


Figure 19. SA Client Username/Password Window



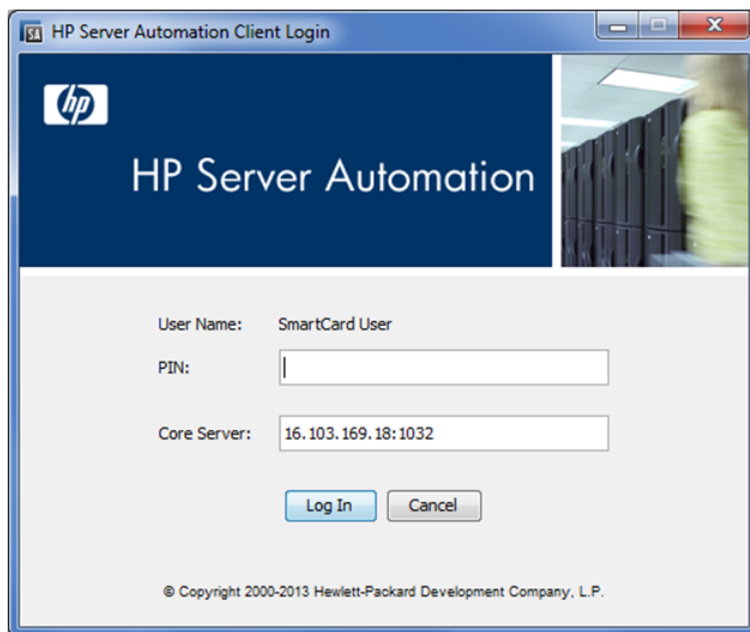
To enable Smart Card login, click on the `More>>` button to access the advanced log in settings. You see a screen similar to **Figure 20**:

Figure 20. Setting the SA Client to Use Smart Card Login



To enable Smart Card log in, select Use Smart Card by checking the box to the left of the option. The log in screen will now look like **Figure 21**:

Figure 21. SA Client Smart Card-Enabled Login Screen



All subsequent logins will display this screen. To revert back to the standard username/password log in, select Advanced Settings and uncheck the Use Smart Card option.

On the Smart Card login screen, the user must be using a PC with an operating smart card reader device. In order for the reader to be usable by SA, ensure that the Windows device is visible in the

“Media” icon application. If the PC used to access SA with the smart card does not have a valid card reader, please contact the IT administrator. To proceed with access to SA, the user must enter the PIN for the smart card and press the “Log In” button.



SA/RSA SecurID® Integration

RSA SecurID® is a two-factor authentication system from RSA Security, Inc. (a division of EMC). Two-factor authentication is based on the concept of *something you know* (a password or PIN) and *something you have* (an authenticator) and provides stronger user authentication than passwords. This section describes how to take advantage of SecurID authentication in your SA system; however, it does not explain how to install, configure, or maintain RSA SecurID.

For detailed information about RSA SecurID, see <http://www.rsa.com>.

This section describes how SA authentication integrates with RSA SecurID. It assumes that you are already using RSA SecurID or will install it. An RSA SecurID server (RSA Authentication Manager or ACE Server) must be installed and fully configured before you can begin using SecurID authentication with SA.

RSA SecurID/SA Integration Overview

SA users are required to authenticate to SA to perform any operations. SecurID integration allows them to use their existing RSA SecurID tokens for authentication. SA authentication can be seamlessly integrated into your existing SecurID environment. As far as the RSA authentication server is concerned, SA (more specifically, the Web Services Data Access Engine server) is just another SecurID agent.

SecurID support is automatic with the installation of an SA Core. Only a few configuration steps are required to enable it:

Note: The first two tasks must be performed on every Web Services Data Access Engine host in your Multimaster Mesh or in SA installations with multiple Web Services Data Access Engines.

- Copying an RSA SecurID configuration file named `sdconf.rec` into a directory on any SA Core servers that host the Web Services Data Access Engine (twist). `sdconf.rec` is located on the RSA Authentication Manager/ACE Server host and contains

required information about the RSA Authentication Manager that must be available to the SA Core.

- Shutting down the Web Services Data Access Engine and restarting after editing the loginModule.conf file to enable SecurID authentication in SA.
- Creating or modifying users in the SA Client to use SecurID authentication.

SA Support for SecurID Authentication Methods

RSA SecurID is based on two-factor authentication, with the SecurID token as the first factor and the Personal Identification Number (PIN) as the second factor.

The SecurID token is the *something you have* and the PIN is the *something you know*. These two factors offer stronger authentication than a user password alone.

SecurID tokens can be either hardware-based (*hardware token* or *hard token*) or software-based (*software token* or *soft token*). The tokens provide a token code which, when combined with a pre-assigned (provisioned) PIN, is called a *passcode*.

Table 12 shows typical authentication methods that are supported by SA/SecurID integration.

Table 12. SecurID Authentication Methods

Authentication Method	Description
Normal Authentication	The most used method. The user's PIN is assigned (<i>provisioned</i>). The passcode is either accepted or rejected.
Next Tokencode Mode (not supported)	This method is used when a user does not enter the passcode correctly. In Next Tokencode Mode, the user must wait for the tokencode to change, and then submit the new tokencode. By default, a user will be put into the Next Tokencode Mode if the incorrect passcode for that user has been submitted three times consecutively.
New PIN Mode (not supported)	This scenario occurs when the user must create a new PIN or modify an existing PIN.

Restrictions

RSA SecurID authentication is not an appropriate method for non-interactive scripts, because the token code changes every 60 seconds and therefore will cause non-interactive scripts to fail. Your options are to rewrite the scripts to be interactive, or avoid using SecurID where such scripts would be affected.

SecurID/SA Integration Platform Requirements

- Solaris
- Linux x86 and x86_64
- RSA ACE Server 6.1 or above.

Configuring SA/SecurID Integration

Support for RSA SecurID authentication is integrated into the SA Core and is installed when the SA Core is installed. However, there are several configuration steps that you must complete to begin using RSA SecurID/SA authentication. The SA Core must also have the IP address of the SecurID authentication server and be able to communicate with it in a secure manner.

Requirement: If you have multiple slices installed in an SA core, the following steps must be performed for each Slice Component bundle host.

Phase 1: The RSA SecurID Authentication Configuration File

1. Contact your RSA SecurID administrator and obtain the file:

```
sdconf.rec
```

2. Copy this file to the following location on all servers in the core that host a Web Services Data Access Engine (twist):

```
/var/opt/opsware/crypto/twist
```

3. Set the file permissions on each server to give the `twist` user ownership of this file and read privileges:

```
chmod 400 /var/opt/opsware/crypto/twist/sdconf.rec
```

```
chown twist /var/opt/opsware/crypto/twist/sdconf.rec
```

4. Ensure that there is no `securid` or `sdstatus.12` file in the `/var/opt/opsware/crypto/twist` directory. If either of these files exist, remove them.

Phase 2: Enable RSA SecurID Authentication in SA

1. By default, RSA SecurID authentication is not enabled. To enable it, on every server in the core that hosts a Web Services Data Access Engine (twist), shut down this component with the following command:

```
/etc/init.d/opsware-sas stop twist
```

2. Locate the file:

```
/etc/opt/opsware/twist/loginModule.conf
```

Edit the file and add the line marked in bold in the example below:

```
TruthLoginModule {  
    com.opsware.login.SecurIDLoginModule sufficient debug=false  
    next_tokencode_mode=false new_pin_mode=false;  
    com.opsware.login.TruthLoginModule sufficient debug=false;  
};
```

3. Restart the Web Services Data Access Engine on all servers with the following command:

```
/etc/init.d/opsware-sas start twist
```
4. If you have multiple Slice Component bundles installed, stop the Command Center (OCC) server and HTTPs proxy on all other Slice Component bundle hosts.
5. At this point only the Command Center for the Slice Component bundle host that is being configured as the RSA server is running. Log into that host's OCC. This will generate the node secret (`securid` file) and the `sdstatus.12` file in the `/var/opt/opsware/crypto/twist` subdirectory as well as register the Slice Component bundle server with ACE.
6. You can now start the OCC and HTTPs proxies on all the other Slice Component bundle hosts in the Core.

Phase 3: Create/Modify SA Users to Use SecurID Authentication

Each user that is to use SecurID Authentication must first exist as an authenticated user in the RSA SecurID authentication server (ACE server) and then must either be created or modified in the SA Client to use SecurID authentication.

In the SA Client, on the user's Profile page, specify that the user's Credential Store should be **RSA 2-factor**.

For detailed information about creating or modifying users, see [Managing Users - SA Client](#).

Troubleshooting

If you receive multiple `Authentication Failed` error messages, first check with your RSA SecurID administrator to insure that the user and passcode is still valid. If you are unable to solve the problem, contact your technical support representative.

User and Security Reports

SA allows you to generate reports that provide a summary of the Client and Feature permissions across servers. These reports are only available when you login to the SA Client as an Administrator. For more information, see the SA Reports Guide.

SA provides following User and Security Reports:

- Client and Feature Permissions
- Customer/Facility Permissions and Device Group Permission Overrides
- User Group Membership
- User Login
- Administrator Actions
- User and Authorizations, By User Group
- User 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

SA Core and Component Security

Introduction to SA Core and Component Security Architecture

- SA can dramatically help improve the security of the typical data center. In particular, SA enables:
- Provisioning security-hardened server operating systems and application software consistently throughout all data centers.
- The introduction of stronger control and accountability across the data center environment; for example, by reducing the number of people who require administrator-level passwords on servers and the creation of digitally signed audit trails of tasks performed on a particular server.
- Automation of the ongoing configuration management challenges of maintaining strong security: identifying servers with missing patches, applying patches consistently, backing up configuration files when they change to enable easy roll-back, and so on.

While the benefits of automating the data center are compelling, organizations need assurance that the automation system itself does not create the potential for new security vulnerabilities. With the ever-increasing sophistication of threats, both from within and external to organizations, it is absolutely mandatory to ensure that your automation software architecture has been designed with security as a primary consideration. SA has been designed with security as a primary consideration.

This section describes how SA uses the most up-to-date security best practices, intended for use in organizations with the most stringent security requirements and with the following design goals:

- **Strict control and accountability:** You can be confident that only authorized administrators can perform management actions, because SA enforces granular role-based access control and generates a digitally signed audit trail of account activity.
- **Secure communication channels throughout the system:** SA is a distributed computing environment in which individual components communicate with each other securely over an IP network. To accomplish this, SA uses SSL/TLS and X.509 certificates to secure the communication between these components.
- **Automated delivery of compliance policies based on industry standards:** SA provides an ongoing stream of immediately actionable compliance policies based

on industry standards. The compliance policies leverage SA's extensive audit and remediation capabilities around granular attributes such as installed patches, installed software, minimum password length, registry key settings, and even individual configuration settings within a file.

Enforcing Strict Control and Accountability

SA provides strong security and accountability, as described in the following sections.

Stronger Controls and Accountability

SA improves security throughout a data center using strong controls and accountability. Using SA, security architects or IT management can control who can perform a particular task on a server. Task control is fine-grained; for example, an administrator can grant comprehensive read-only access with change privileges restricted to patch installation and a specific list of SA Global Shell commands.

SA automatically creates a tamper-proof audit trail that captures details such as which SA user performed a particular management task on a server at a given time. SA's granular role-based access control system is designed around the interaction between users, groups of servers, management tasks, and the SA data model that describes the environment. One immediate security benefit of this powerful access control model is that fewer people need administrator accounts on servers. Instead, they can be given SA user accounts to perform only the management tasks they must perform, a security best practice.

Everyone who logs into SA must have a unique SA user name and password. Administrators can create user names within SA or import them from an external LDAP system. For example, if a company has an existing Microsoft Active Directory implementation, they can synchronize with the directory server to reuse the user accounts that already exist.

When creating user accounts, SA users are assigned to SA groups. Groups are a convenient way of describing what servers users can operate on and what management tasks they can perform on those servers.

Several predefined groups are provided by default in SA. The permissions for these groups can be customized as necessary, and you can create new groups with customized permission levels to satisfy the requirements of any organization. The permissions that you specify for a user group determine what the group's member can do with SA. *Action permissions* specify what actions users can perform; *resource permissions* specify which objects (typically servers) users can perform these actions on. The SA graphical user interface, called the SA Client, as well as the Global Shell interface, are both bound by these task rules, so that users will be able to see and perform only the tasks they are authorized by security administrators to perform.

Security administrators can also control the policy-based software installation environment, which automates the process of installing software and configuring applications on a server. Designated users can model an organization's application software structure in a folder

hierarchy, and set up fine-grained permissions for creation, viewing, modification, and execution. This model provides a clear delineation of specialization, where subject matter experts can implement and adjust policies, and system administrators can manage the servers in their environment by applying software policies to servers.

Note: See [User and User Group Setup and Security](#) user groups and permissions.

Read-only, Digitally Signed Audit Trails

In addition to careful controls of which actions SA users can perform on managed servers, SA automatically maintains a detailed audit trail of events performed by SA users. The audit trail logs details such as the user, the event, the servers acted on, the time the task was performed, the total elapsed time, and any error conditions associated with the task.

The audit trail itself is stored as read-only, digitally signed data in an Oracle database to prevent users from tampering with the data. This audit trail data helps organizations establish strict accountability in their environment—an increasingly urgent topic in the age of Sarbanes-Oxley Act, the Gramm-Leach-Bliley Act (GLB Act), and the Health Information Portability and Accountability Act (HIPAA). Users can select how long the audit trail is stored (the default period is six months), and they can easily create a data warehouse that stores the audit trail (and other SA data) for longer periods of time.

The Audit Trail is housed in the AUDIT_DATA tablespace, and contains the following tables:

AUDIT_OBJTYPE_ATTR

AUDIT_OBJECT_TYPES

AUDIT_OBJECT_COLLECTORS

AUDIT_OBJECT_ATTR

AUDIT_FEATURES

AUDIT_EVENT_OBJECTS

AUDIT_EVENT_DETAIL_VALUES

AUDIT_EVENT_DETAILS

AUDIT_EVENTS

AUDIT_DATA_TYPES

AUDIT_DATA_OBJECTS

AUDIT_DATAOBJ_VALUES

AUDIT_CONFIG_PARAMS

AUDIT_COMPONENTS

AUDIT_ACTIONS

Signed SHA Checksums for Packages in the Software Repository

When SA users upload software to the Software Repository, SA automatically computes an RSA-with-SHA1 signature for the package. To generate the signature, SA uses a combination of the SHA1 checksum calculation, the software package contents, and an internal private RSA key that is known only to the Software Repository. The private key is not modifiable. This prevents users from tampering with the software in the Software Repository. The package and its corresponding digital signature are stored locally at the Software Repository. When SA installs software on a managed server, it validates the RSA key and the SHA1 signature of the software before permitting the download. This helps ensure that the software installed by SA is exactly the same software uploaded into the Software Repository.

Role-Based Authorization

- SA enforces a very granular system of role-based access controls. Security administrators can set up authorization on the following parameters:
- **A facility:** A facility is a collection of servers that are managed by a single SA core. A facility can be all or part of a data center, server room, or computer lab. A facility is the highest level of abstraction in the granular role-based permissioning model.
- **A group of servers (by customer):** Servers are grouped by customers, which can represent any arbitrary group of servers in a single data center. The group might represent a paying customer, a cost center, or servers running a particular business application such as Siebel or the Expense Report application. The software packages managed by SA each belong to a particular customer, although they may also belong to a special account called *Customer Independent*, which means the software is available to provision on any customer's server (for example, patches belong to the customer account *Customer Independent*). This allows security administrators to control the exact set of software packages that may be applied on a particular group of servers.
- **A dynamic group of servers (rules-based):** Security administrators can also create server groups based on *dynamic rules evaluation* (from simple to complex) and grant permissions to all servers belonging to the group. For example, a security administrator can group managed servers that are running the Linux operating system and reside in a particular IP address space, and then assign which SA user groups are authorized to perform management tasks on this server group.
- **Software policy modeling and distribution:** Software policy modeling provides a powerful mechanism to model software using a folder model. Folders provide the ability to define security permissions to control access to their contents across user groups. You can set folder permissions to determine which user groups can view, use, and modify items within a folder.

Audit Logging of User Activities

SA stores audit trails centrally in the Model Repository, where each entry is digitally signed. SA is designed from the ground up with strong cryptographic controls that prevent any undetectable modification to audit logs. Because audit logs are stored centrally, they cannot be deleted from managed servers. In fact, the entire security design of SA is defensive, based on the assumption that an individual managed server being compromised must not endanger the security of the whole system.

Custom Certificate Authority (CA)

You can use your Certificate Authority (CA) certificate to sign all SA certificates. The custom CA is available as an optional post-install/upgrade step in SA. It is used to generate all SA certificates, including intermediate CAs required by SA. The Custom CA will essentially act as the Root CA for SA. All the certificates signed using the Custom CA will inherit attributes like Key Length, Signing Algorithm, Expiry Date, and so on from the Custom CA.

The requirements for custom certificate are as follows:

- The certificate must be a single file with both the certificate and corresponding private key.
- The certificate must be an RSA encrypted certificate. Certificate encrypted by DSA and ECC is not supported.
- The certificate must be encoded in ASCII format. Other formats (such as DER ...) are not supported.
- The certificate must not be signed by a CA, that is, it should not be a server certificate, but a CA certificate.
- The certificate must not have passphrase to protect its private key.
- The certificate must have digest of sha1 or sha256 type.
- The certificate must have key length of either 2048 or 4096 bits.
- The certificate must not have expired.

Securing SA Internal Communications

SA includes several Core Components that communicate with each other over secured communication channels, typically the industry-standard protocols such as HTTPS. These components include:

- The SA user running a secure browser on the user's local desktop or server. The SA browser communicates securely using HTTPS to the SA Command Center. Users provide a user name and password to log in to SA; the credentials are authenticated either within SA or optionally by an external integrated LDAP server.
- SA Server Agents running on the managed servers. The SA Server Agents act both as clients and servers when communicating with SA Core Components. All communication is encrypted,

integrity checked, and authenticated using client certificates that use SSL/TLS. A small number of Core Components can issue commands to the SA Agent over a specific TCP/IP port; the SA Agent can also call back to Core Components, each with its own specified port.

- SA Core Components, which are back-end processes running on a small number of servers. SA Core Components communicate with each other and with the SA Agent, also using strongly authenticated SSL/TLS.


For customers running SA across multiple data centers, communication also occurs between SA cores over a secure channel provided using integrated certified messaging included in SA (SA Bus).

By protecting the communication channel between distributed components, SA prevents intruders from sniffing the network traffic or causing SA to perform unauthorized tasks on SA-managed servers.

Communication Between Components in an SA Core

When an SA component must communicate with another component, it opens a secure (typically SSL/TLS) communication channel using a well-known port. Each SA component has a public-key certificate that is generated when SA is installed. The component uses its public-key certificate when authenticating itself to another component. In this fashion, most interprocess communication is strongly authenticated, encrypted using the strongest ciphers available, and integrity checked.

Figure 22. Component Communication



Communication Between Agents and SA Core Components

The Server Agent participates in the strongly authenticated and encrypted SSL/TLS traffic described above. In addition, when Agents are directed to perform management tasks on a server, the typical flow of control messages (described below) help to ensure that only authorized users are performing those actions. It would be extremely difficult for an intruder to generate a valid command sequence directing the agent to perform an unauthorized task.

The following sequence describes a typical SA management task, namely provisioning software on a managed server. Other operations on managed servers follow the same general protocol:

1. The Data Access Engine opens a communication channel via HTTPS with the SA Server Agent, directing it to perform a management task.
2. The SA Agent calls back to the Data Access Engine to retrieve specifics about the task to perform. To open a communication channel, the Agent must present its public-key certificate, which the SA Core verifies against an internal database mapping the certificate itself to the machine's IP and a unique machine identifier

that SA generates when the agent is installed. This safeguard prevents users from simply copying the digital certificate and corresponding key to another machine in hopes of masquerading as the original managed server.

After successfully opening the communication channel, the SA Agent receives the exact list of software to be installed and removed (as well as any scripts to execute, the order of software installation, and when to reboot during the provisioning process).

3. The SA Agent opens a communication channel to the Software Repository (also via HTTPS) and requests a download of the software it needs to install. Before the Software Repository initiates the download, it recomputes an SHA checksum for the package along with a secret key it knows. Only if the SHA checksum matches the checksum generated when the package was uploaded does the SA Agent receive the software it requested.

Asynchronous, agent-initiated calls to the SA Core provide scalable support for progress reporting and long-running operations, because the SA Core need not manage thousands of synchronous agent operations directly. SA supports these asynchronous calls from the Agent to the Core even in network environments where firewalls prevent Agents from initiating TCP connections, as the SA Gateway infrastructure provides bidirectional tunneling over unidirectional connections.

Other technical details of agent/core communications:

- Connections are SSL v3, mutually authenticated with X.509 certificates (the server checks the client's certificate, and vice versa).
- Private keys for Core and Agent certificates are stored in files that are readable by root only.
- All certificates are generated at installation, are owned by the customer, and are not known to HP.
- Certificates expire 10 years after installation. SA provides a Recertification tool for recertifying Cores and Agents prior to certificate expiration.
- Certificates are signed by SA internal self-signed certificate authorities. To avoid HTTPS security warnings in web browsers, customers may install an externally signed certificate in the SA instance of Apache.

Communication Between SA Cores

If you are running SA across multiple data centers, SA automatically synchronizes relevant data across all SA-managed data centers. Broadly speaking, SA synchronizes two types of data: the SA model of servers (including all hardware, software, and configuration attribute information) and the software packages themselves.

- **Replicating the SA model:** SA uses integrated certified messaging to synchronize the SA model data. SA implements SSL to safeguard the messages flowing across the message bus. The actual messages themselves describe SQL changes that need to be made to the SA database at the receiving end of the communication.

- **Replicating software packages:** SA replicates software packages on demand. That is, they are only copied when they are needed. When the an administrator managing a server in the New Jersey data center directs SA to install a software package that does not exist in New Jersey's Software Repository, SA requests it from another data center. The actual file transfer uses the open source utility rsync, and the communication channel is secured using SSH.

SA Satellite Architecture and Security

An SA satellite, rather than a full SA core, can be installed at secondary locations to enable management of remote servers. Satellites provide the same seamless management of data center servers as an SA core does. The Satellite consists of an SA Gateway and a Software Repository Cache. A Satellite Gateway provides a network connection and bandwidth management to the Satellite. A Satellite can contain multiple Gateways. The Software Repository Cache contains local copies of software packages to be installed on managed servers from the Satellite. Optionally, a Satellite can contain the OS Provisioning Boot Server and Media Server components. A Satellite must be linked to at least one Core, which may be a single core or part of a Multimaster Mesh. Multiple Satellites can be linked to a single core.

The Satellite has the following key capabilities:

- **Automate Regardless of Network Complexity:** Satellites are optimized to work across low-bandwidth connections, through complex, overlapping IP address spaces, and across firewall boundaries.
- **Respond to Network Failures:** SA Satellites implement sophisticated link state routing algorithms that enable dynamic routing around failed network links for redundancy.
- **Ensure Remote Server Security:** Satellites enable IT organizations to proactively ensure remote server security through policy-based patch management, digitally signed and encrypted package installation, and comprehensive audit trails that track complete server change history.

The SA Network: Enabling Risk Mitigation

New vulnerabilities are constantly being reported. The SA Network is a unique service that makes actionable, multi-vendor, prioritized, security alerts available to your SA installation. With the SA Network, you can identify vulnerabilities as soon as you learn about them, and deploy the appropriate fixes without consuming extra resources.

Recognizing that no single standard covers all needs, the SA Network provides a broad collection of compliance policies that are easily customizable and extensible to meet each customer's specific needs.

The SA Network currently focuses on the following three compliance standards:

- **Center for Internet Security (CIS) standards:** A set of standards that detail how to secure a server based on operating system. (<http://www.cisecurity.org/>)
- **Microsoft (MS) Security Guide:** A standard established by Microsoft that details the configuration settings to harden Windows servers. (<http://www.microsoft.com/>)
- **National Security Agency (NSA) Security Configuration Guide (SCG):** A standard established by the United States National Security Agency that details the configuration settings to harden different OSs and applications. (<http://www.nsa.gov/>)

SA Compatibility with Other Security Tools

SA complements many existing security tools such as intrusion detection systems, vulnerability assessment suites, anti-virus scanners, and integrity assurance products. SA can be used to drive change management practices that make these tools an effective safeguard for servers. In particular, SA can be used to install and configure Agents required by these systems consistently, keep configurations (such as the latest anti-virus definition files) up to date, and act on some of the vulnerabilities reported by these systems (such as missing patches or bad configurations).

SA Core Recertification

SA provides a *Core Recertification Tool* that allows you to recertify SA Cores and Agents. The Core Recertification Tool automates and speeds the process of issuing new security certificates.

Note: This tool is separate from and compatible with the existing Agent Recertification tool. For more information, see [Agent Recertification](#).

Carrying out a Core Recertification does not require additional SA downtime. SA services will be fully available during the complete procedure. The following service restarts are required, but can be synchronized with internal maintenance windows:

1. Phases 3 and 7: Automatic restarts for mesh-wide SA gateways.
2. Phases 4, 8, and 12: Automatic restarts for mesh-wide SA Agents.
3. Phases 6: Automatic restarts for primary spin components of each SA facility.
4. Phases 6, 9, and 13: Manual mesh restarts.

Major advantages of the Core Recertification Tool are:

- The ability to regenerate all SA certificates before their expiration, which effectively shortens their life span.

- The ability to mitigate certificate compromises.

SA is a closed Public Key Infrastructure (PKI) system that uses X.509 v3 certificates to facilitate authentication, authorization, and secure network communications. An X.509 certificate is a form of identification that binds a specified principal with a public key.

A certificate, along with its corresponding private key, constitutes a digital identity. Like many other forms of identification, a certificate is valid for a finite period of time. X.509 certificate validity period is specified by the `Not Before` and `Not After` date. A given X.509 certificate is considered valid only if the current date is within its validity period. Conversely, a given X.509 certificate is considered invalid if the current date is outside of its validity period. SA does not accept invalid certificates.

SA CAs are automatically generated during bootstrap and subsequently used to issue the rest of the Core Component certificates. SA Agent certificates are issued by the Agent CA during initial Agent registration.

All SA certificates are valid for 10 years by default. There is no way to change the life span of the SA certificates through configuration. The only way to make changes to the SA certificate policies is through customization.

SA uses *class certificates* where all the Core components of a class share one certificate. For example, all the Command Engines share one Command Engine certificate. Compromising one Command Engine certificate means all the Command Engine certificates are compromised. Furthermore, SA does not support *certificate revocation*. The only way to invalidate a compromised Core Component certificate is to recertify the entire Core.

Note: This release of Core Recertification Tool does not support customized Core installations. Any customization that has been done outside the realm of the SA Installer, which requires certain SA certificates and keys to be on a different host or under a different directory, will not be supported by this tool.

SA will warn administrators about upcoming certificate expiration through System Diagnosis on the Data Access Engine. The warning period is configurable (`crypto.expire.warn_days`) with the default being 300 days.

There are two use cases for re-certifying a core; the crypto material is expiring or a security breach has exposed the crypto. In the case of a security breach phases 11 through 13 must be executed.

Agent versus Core Recertification

There is an important distinction between agent and core recertification. Core recertification regenerates the core's certificates and all of the agent certificates on all managed servers. Agent recertification regenerates just the agent certificates on managed servers.

This section describes the full core recertification. For instructions on recertifying just the agent on a managed server, see [Agent Recertification](#).

Adding a New Core or Slice to a Recertified Core Multimaster Mesh

Prior to SA 10.10, the core recertification procedure did not re-sign Model Repository (truth) data and other SA data. During operation, both old/archived and new CAs are loaded to validate the signatures.

From SA 10.10 onwards, core recertification re-signs SA data.

If your mesh was recertified prior to SA 10.1, before adding a new core/slice to this recertified mesh, run re-signing scripts. Contact support to obtain the re-sign scripts and instructions on how to run them. Re-sign scripts might take a long time to finish, depending on the amount of data to re-sign.

Core Recertification Phases

Core Recertification has several phases. Which phases are required depends on your Multimaster configuration.

Table 13 describes the Core Recertification phases:

Table 13. Core Recertification Phases

Phase	Description
1-3	<p>Back up existing crypto material, generates new crypto material, and distributes the new CAs to all the Core Components. These three phases occur sequentially during the first run of the Core Recertification tool. All the existing crypto materials are backup into the <code>crypto.<session number></code> directories. Each Core component has its own backup directory.</p> <p>Create <code>/etc/opt/opsware/crypto/security.conf</code> if it is missing. Update existing <code>/etc/opt/opsware/crypto/security.conf</code>.</p>
4	<p>Distribute the new Agent CAs to all the Agents so that Agents will trust both the new and old Agent CA at the same time. This is to ensure uninterrupted Agent-to-Agent communication.</p> <p>Note: If the <code>agent_recert.using_cdr</code> parameter value is zero (0) in the <code>corerecert.conf</code> file, this phase (phase 4) is skipped. HP recommends that you set the <code>agent_recert.using_cdr</code> parameter to zero (0), as the CDR feature is no longer supported.</p>
6a	<p>Mesh Restart: Restart the Mesh so that it trusts both the new and old CA hierarchies.</p>
6b	<p>Set up the Public Key Infrastructure (PKI) on the primary Spins so that they'll start issuing certificates generated with the new Agent CA.</p>
7	<p>Recertify the Gateways.</p>
8	<p>Recertify the Agents.</p> <p>Note: Make sure all managed servers are functioning and reachable throughout</p>

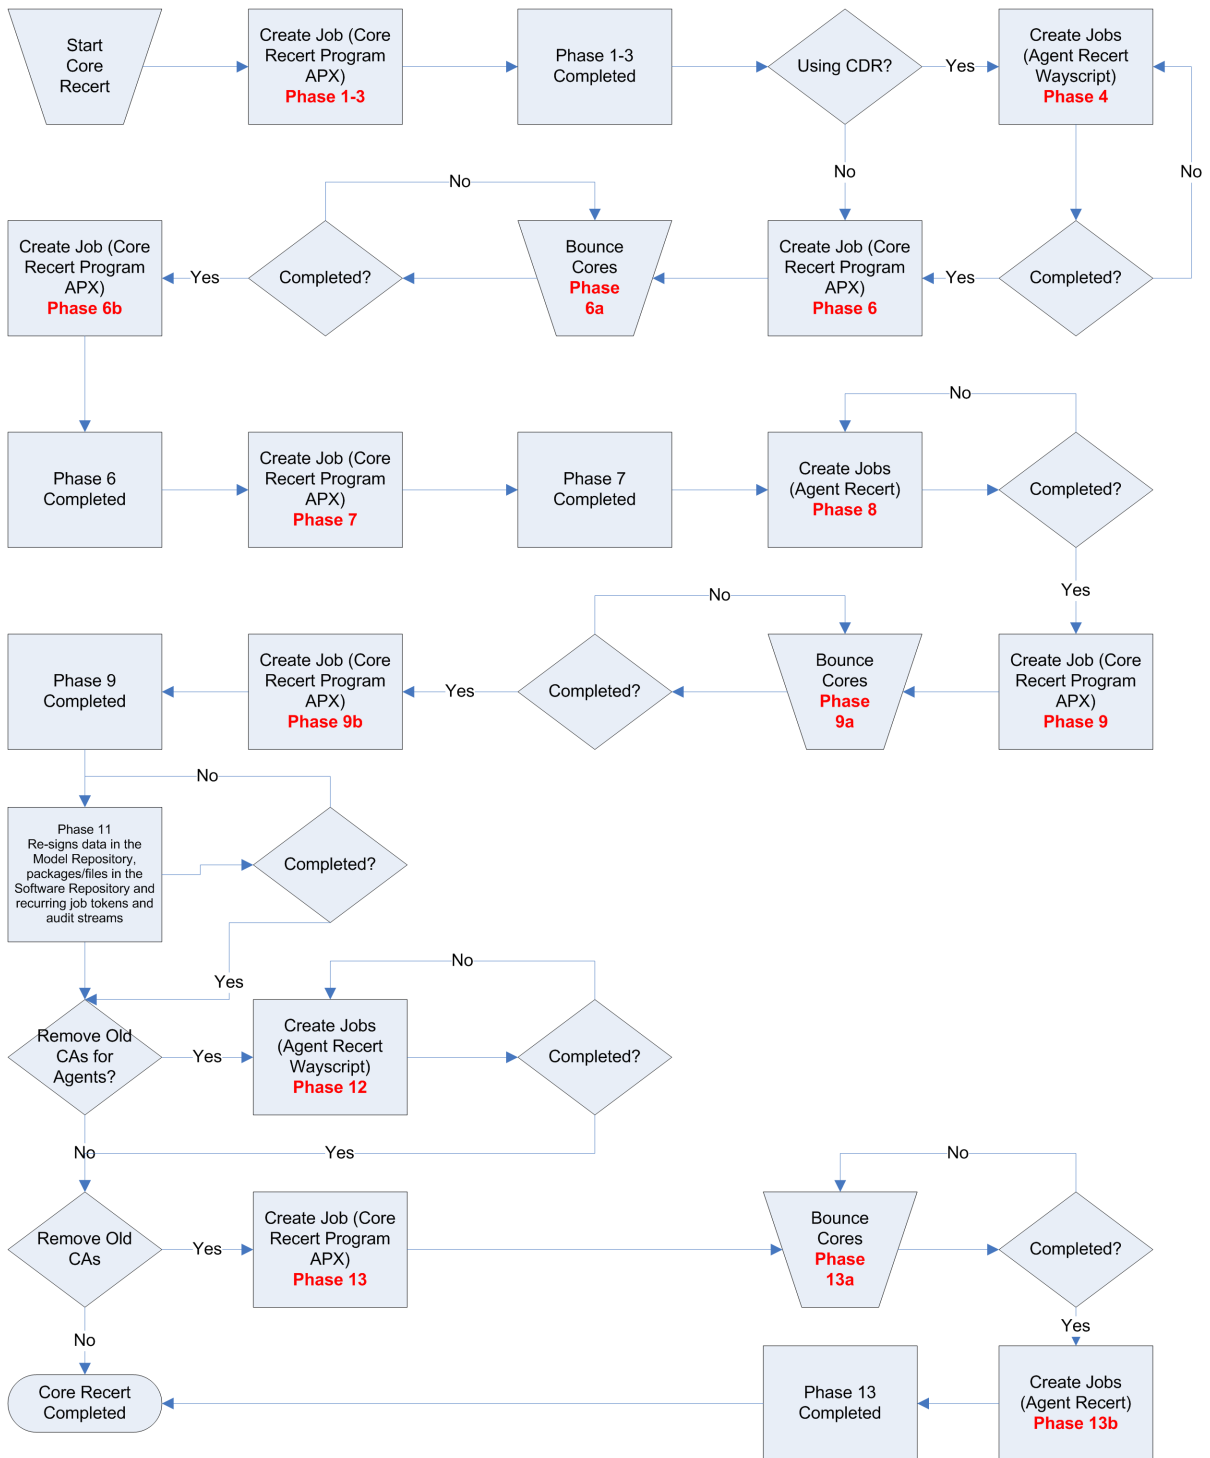
Phase	Description
	phase 8, or the Core will fail to communicate with the servers after the Core Recert process is complete.
9a	Recertify the Core components; issue the command <code>touch /var/opt/opsware/crypto/twist/upgradeInProgress</code> on First Core; Mesh restart; Regenerate Signatures.
9b	Check Mesh Restart status. If the Mesh has successfully restarted, all the Core components are now using the new crypto material while still trusting the old crypto material.
11	[Optional] Re-signs data in the Model Repository, packages/files in the Software Repository and recurring job tokens and audit streams. (Available in SA 10.10 or later.)
12	<p>[Optional] Remove old Agent CAs. Required only when Agent CAs have been compromised or you no longer trust the old CAs.</p> <p>Note: When a managed server that has both an older and a newer CA is not recerted during the Agent Recert phase (Phase 8), that server will not be able to communicate with another managed server that only has an older CA.</p> <p>Note: For Core Recert with custom certificate, HP recommends that you go through phase 12 so the old certificate is removed from the agent trusted CA store, and, therefore, only the customer certificate is used for verification.</p>
13a	<p>[Optional] Remove the old Agent CA hierarchies. Required only when Agent CAs have been compromised or you no longer trust the old CA hierarchies.</p> <p>Note: When a managed server that has both an older and a newer CA is not recerted during the Agent Recert phase (Phase 8), that server will not be able to communicate with another managed server that only has an older CA.</p> <p>Note: For Core Recert with custom certificate, HP recommends that you go through phase 13 so the old Core-component certificate is removed from the trusted CA store, and, therefore, only the customer certificate chain is used for verification</p>
13b	[Optional] Mesh restart. Required only when 13a is also required.

Note: A Mesh Restart means restarting the SA services on all Core and Satellite boxes. The restart has to be performed manually. The following is the startup sequence for multi-host cores:

Model Repository (MR) -> Infrastructure -> Software Repository (SR) -> Slice -> OSProv

Figure 23 shows the flow and phases of the recertification process:

Figure 23. Core Recertification Phases and Flow



Agent Recertification Phases

Three of the phases depicted in **Figure 23** are *Agent Recertification phases*:

- **Phase 4:** Distributing new Agent CA. The purpose of this phase is to ensure continuous Agent-to-Agent communication (recertified Agents communicating with Agents that have yet to be recertified).
- **Phase 8:** Recertify the Agents. This is a *required* phase. The purpose of this phase is to issue new crypto material to the Agents.
- **Phase 12:** Cleanup the old Agent CAs. This phase is *optional*. If you do not wish to trust both the old and new CA hierarchies, you must use this phase to remove the old CAs. Otherwise, you can skip this phase.

Agent Recertification Jobs

Each Agent Recertification phase is accomplished by a recurring job. This job is dictated by the properties shown in **Table 14**, which you must specify in the Core Recertification configuration file:

Table 14. Core Recertification Configuration File: Agent Recertification Properties

Property Name	Req?	Description	Example
agent_recert.all.facilities.delay=<seconds>	No	The delay in seconds for starting the agent recert jobs after entering the agent recert phases. The value must be between 120 and 7200 seconds. This property is optional. The default delay is 3 minutes.	agent_recert.all.facilities.delay=120 The property is available in SA 9.17, 10.03, 10.11,10.22 and later.
agent_recert.all.facilities.start_time=<HH:mm>	No	The start time for the Agent Recertification phase. You may overwrite this value for a given facility by specifying the agent_recert.facility.<facility name>.start property. Start time must be in the following format,	agent_recert.all.facilities.start_time=18:30

Property Name	Req?	Description	Example
		<p>HH:mm, where 00 <= HH < 24 and 00 <= mm < 60.</p> <p>Only the hour and minute components are needed. If the specified time has already passed, the Agent Recertification job will start at the specified time the next day.</p>	
<pre>agent_recert. facility.<facility_ name>.start_time= <HH:mm></pre>	No	<p>If present, the start time of the given facility will be used instead of agent_recert.all.facilities.start_time.</p>	<pre>agent_recert. facility. sacramento.start_ time= 08:00</pre>
<pre>agent_recert.all. facil- ities.- duration=<hours></pre>	Yes	<p>The duration, in hours, for the Agent Recertification job. Duration dictates how long the Agent Recertification job runs before stopping. If the duration has elapsed and the success rate has not been reached, the Agent Recertification job will continue at the next start time. You can overwrite this value for a given facility by specifying the agent_recert.facility.<facility_name>.duration property.</p>	<pre>agent_recert.all. facil- ities.duration=8</pre>

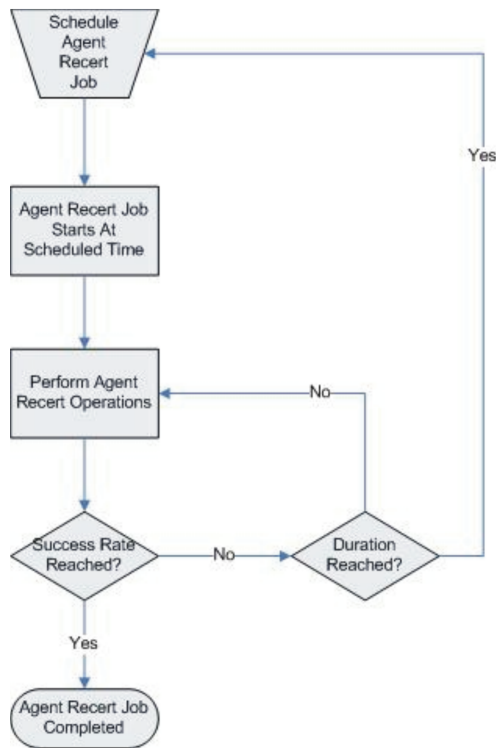
Property Name	Req?	Description	Example
agent_recert.facility.<facility_name>.duration=<hours>	No	Duration must be an integer value between 1 and 24. If present, the duration of the given facility will be used instead of agent_recert.all.facilities.duration	agent_recert.facility.sacramento.duration=10
agent_recert.all.facilities.success_rate=<whole percentage>	Yes	<p>The success rate (in whole percentage) for each facility for the Agent Recertification job. For example, if there are 1000 managed servers in Facility X and the success rate is 98%, the Agent Recertification job will stop if 980 managed servers have been successfully recertified.</p> <p>You can overwrite this value for a given facility by specifying the agent_recert.facility.<facility_name>.success_rate property.</p> <p>Success rate must be an integer value between 1 and 100.</p>	agent_recert.all.facilities.success_rate=100
agent_recert.facility.<facility_name>.success_rate=<whole percentage>	No	If present, the success rate of the given facility will be used instead of agent_recert.all.facilities.success_rate.	agent_recert.facility.sacramento.success_rate=99

Property Name	Req?	Description	Example
<pre>agent_recert.all. facilities.job_ notification=<email addresses></pre>	No	<p>The job notification for the Agent Recertification job. You can overwrite this value for a given facility by specifying the agent_recert.facility.<facility_name>.job_notification property.</p>	<pre>agent_recert.all. facilities.job_ notification= admin@example.com</pre>
<pre>agent_recert. facility.<facility_ name>.job_ notification= <email addresses></pre>	No	<p>If present, the job notification for the given facility will be used instead of agent_recert.all.facilities.job_notification.</p>	<pre>agent_recer- t.facility. sacramento.job_ notification= admin3@example.- com</pre>
<pre>agent_recert.using _ cdr</pre>	No	<p>Indicates Code Deployment & Rollback (CDR) feature is being utilized. Default is 1.</p> <p>Note: HP recommends setting this parameter to zero (0), as the CDR feature is no longer supported.</p>	<pre>agent_recer- t.using _cdr=0</pre>

Agent Recertification Job Flow

Figure 24 shows the Agent Recertification job flow:

Figure 24. Agent Recertification Job Flow



There can be only one Agent Recertification job, scheduled or active, per facility at any given time. An Agent Recertification job will terminate only if:

- The success rate has been achieved
- You explicitly cancel the job
- A fatal error occurs

SA Core Recertification Tool Usage

To run the Core recertification tool, enter the following:

```
/opt/opsware/oi_util/OpwareCertTool/recert_utils/corerecert [--  
phase <phase number>] [--config <complete path to the config  
file>] [--doit]] [-h, --help] [-v, --version] [-s, --status] [-d,  
--debug] [--summary] [--cancel_all_agent_recert_jobs] [--cancel_  
agent_recert_jobs_for_facility <facility name>] [--cancel_all_  
jobs] [--reason <reason for job cancellation>] [--force_resume  
<facility_name>]
```

Arguments to the Core Recertification Tool

Table 15 describes the valid arguments for the Core Recertification tool:

Table 15. Core Recertification Tool Arguments

Argument	Description
-h, --help	Displays help.
--phase	Starts a specified Core Recertification phase. The valid phase numbers are 1, 4, 6, 7, 8, 9, 12, and 13.
--config <config file>	The fully qualified path to the Core Recertification configuration file. The default configuration file is /opt/opsware/oi_util/OpwareCertTool/recert_utils/corecert.conf.
--doit	Reruns or forces a rerun of a given Core Recertification phase. This is useful when certain newly added components have missed the recertification process. It is also used to skip specified phases, such as new Agent CA push or old Agent CA removal.
-v, --version	Prints out the version number of the <code>corecert</code> executable.
-s, --status	Displays the current status of the recertification process.
-d, --debug	Sets Core Recertification to debug mode, debug logs are available in <code>/tmp/recerttool.log</code> .
--summary	Prints out the current status summary, shorter version of <code>--status</code> .
--cancel_all_agent_recert_jobs	Cancels all currently scheduled Agent recertification jobs.
--cancel_agent_recert_jobs_for_facility <facility name>	Cancels the Agent recertification jobs scheduled for a given facility.
--cancel_all_jobs	Cancels all Core and Agent Recertification jobs.
--reason <reason for job cancellation>	Specifies an optional reason for the job cancellation.
--force_resume	Specifies that a new job be automatically scheduled for any facilities with failed agent

Argument	Description
	recertification jobs. Facilities with no failed jobs will be skipped. Alternatively, if you do not specify this parameter, you can resume the job for each facility individually.

The `/tmp/recerttool.log` is *not* cumulative, it is rewritten with each `recerttool` execution. The log contains only the following information: information on starting the background processes for the current phase, parameters that the current phase uses (if applicable), and information on failure to kick off background jobs.

The core recertification background jobs rely on SA's OGS infrastructure. See `/tmp/core_recert.log` (SA 9.1, 10.00 and 10.10) under OGFS of the core used to start the recertification or in `/var/log/opsware/waybot/recert.log` (SA 10.20 and later) for more information.

The agent recertification background jobs are run by the waybot, hence more details can be found in the twist and waybot logs on each core of the mesh.

The Software and Model Repository signature regeneration (Phase 11) will log additional information on the recert's base slice in `UpdateSignatures.log` and `ResignJobTokens.log` under `/opt/opsware/oi_util/OpswareCertTool/recert_utils/`.

Caution: Adding new Core Components during Core Recertification is not recommended. Although adding new Core Components, such as the Slice Component bundle, a Satellite, etc. during Core Recertification is possible under certain circumstances, HP does not recommend doing so unless absolutely necessary. *You must first contact HP Professional Services before adding new Core components while a Core Recertification is in progress.*

Caution: Replacing SA certificates with third-party certificates (not issued by an SA CA) is not supported. During Core Recertification, third-party certificates could be overwritten if they have the same filename as an SA certificate. If you have replaced any SA certificates with certificates issued by a third-party CA, you should contact HP Server Automation Support before performing Core Recertification.

Security Considerations

Consider the following security issues:

Crypto Database File

The SA Core Recertification Tool requires access to the SA crypto database file during recertification.

The SA crypto database consists of the file:

```
/var/opt/opsware/crypto/cadb/realm/opsware-crypto.db.e
```

This file is protected by the crypto material password (`decrypt_passwd`), which was specified during the mesh's First Core installation. During subsequent Core installations, this file is also copied to the new Secondary Core hosts. You must protect this password as compromising the crypto database files means compromising your entire Multimaster Mesh.

The crypto database file is required only during SA installation or upgrade, but it is regenerated during Core Recertification. Therefore, HP strongly recommends that you create procedures that protect the crypto database file. Therefore, before Core Recertification, you must back up this file to a secure location.

During Core Recertification, SA regenerates the crypto database only on the host on which you invoke the Core Recertification Tool. Core Recertification does not copy the newly generated crypto database file to any other hosts in the mesh during recertification. You should also back up this file to a secure location as soon as Core Recertification is complete.

Equally important is to strictly control root access to the Core hosts. Crypto materials (certificates and their corresponding private keys) on the Core hosts are not encrypted. They are protected by the root user account. In other words, these files are protected by the read-only access for the root user. Therefore, having root access to the Core hosts means a user has access to both the crypto material password and the crypto database files, and Core Recertification should only be performed by SA System Administrators, or someone who has legitimate root access to the Core hosts.

Core Recertification Users

There are typically three types of users who will use the SA Core Recertification tool:

- **Core Recertification User:** This user has all the necessary permissions to run the Core Recertification Tool. For all practical purposes, this is the same user as SA System Administrator/Operator.
- **SA Administrator:** Grants or revokes the SA Core Recertification role to the Core Recertification User.
- **SA System Administrator/Operator:** This user is responsible for restarting a given Core. This user has root access to the Core host.

Creating the Core Recertification User

In order to use the Core Recertification tool, you must create a Core Recertification group and user (s) and grant the necessary permissions:

1. As SA Administrator, log on to the SA Command Center.
2. Create a *Core Recertification user group* with the following permissions:
 - Read & Write access to all Facilities
 - Read * Write access to all Customers
 - Read * Write access to all Device Groups
 - Manage Customer
 - Manage Facility
 - Manage Servers and Groups

- Action Permission > Categories (Core Recert > Core Recert)
 - Actions Permission -> Core Recert) -> Agent Recert to Actions
 - Core Recertification (**Client > Core Recert**)
 - Agent Recert (**Client > Agent Recert**)
3. Add the Core Recertification user to the SA System Administrators user group.
 4. List and execute folder permissions on the Library/Tools/Administrative Extensions folder.

Removing a Core Recertification User

To remove a Core Recertification user, perform the following tasks:

1. As SA Administrator, log on to the SA Command Center.
2. Remove the user from the `Core Recertification` user group.

Core Recertification Prerequisites

Before starting Core Recert, you must perform the following tasks:

- Select a new password to protect the crypto materials and decide how that password is to be provided.
- Configure Core Recertification configuration file with the correct values.
- Ensure that all your Cores are up and running.
- Ensure that the Core Recertification tool correctly recognizes your Mesh setup.

Check that all managed servers are reachable by running a Communications test against all managed servers before Core Recert is invoked

Requirements for custom Certificate Authority (CA)

When recertifying an SA Core using a custom CA, make sure that the .pem file you supply meets the following requirements:

- The certificate is single file which contains both the certificate and corresponding private key.
- The certificate is a RSA-encrypted certificate. SA does not support certificates encrypted by DSA and ECC.
- The certificate is encoded in ASCII format. Other formats, such as DER, are not supported.
- The certificate is a CA certificate, and not an end-entity certificate certificate.
- The certificate does not use **passphrase** to protect its private key.
- The certificate uses digest of either **sha1** or **sha256**-type.
- The certificate key length is either 2048 or 4096-bits.
- The certificate is not expired. To avoid frequent recertifications, ensure the certificate is valid for at least ten years.
- The certificate **Subject CN** (Common Name) field is not empty.

SA uses the custom CA to generate some intermediate CA certificates. These intermediate CAs are used for signing all end-entity certificates.

Select a New Password to Protect the Crypto Materials

The crypto database password is required during Core Recertification to protect the newly generated crypto database, the PKCS #12 files, and CA private keys. Core Recertification comprises multiple phases, and most of them require the crypto database password. It is very crucial to protect the crypto database password.

Caution: Some of the Core Recertification tasks are accomplished by Automation Platform Extension (APX) jobs. Therefore, the crypto database password, though obfuscated, may briefly appear in the job parameters or in the job audit logs.

To avoid having the crypto database password appearing in job parameters or audit logs, you may convey the crypto database password using a file by following this procedure:

1. Before invoking the Core Recertification Tool on the Core host, determine the Core host's Server ID. You can obtain the Server ID from either the SA Client or by looking in `/etc/opt/opsware/agent/mid`. You must specify the Server ID value for `base_core_server_ref` in the Core Recertification configuration file.
2. Create a file, `/var/opt/opsware/crypto/cadb/__recert_overwrite__`, which contains the new crypto database password. For example `cadb_password=<new crypto database password>`. Ensure that this file is read-only to the root user.
3. Remove the `/var/opt/opsware/crypto/cadb/__recert_overwrite__` file after Core Recertification has successfully completed.

Because the crypto database password is required in the Core Recertification configuration file, you can specify an invalid password in that file as a security measure.

Core Recertification allows only one password to protect all crypto materials. This includes the crypto database, PKCS #12 files, and all the CA private keys. If you are running a customized version of `OpswareCertTool`, where the crypto materials are protected by multiple passwords and want to continue doing so, *you must contact HP Professional Services before running the Core Recertification Tool.*

Configuring Core Recertification

All Core Recertification properties must be specified in a configuration file. When invoking the Core Recertification Tool, you can specify the location of the configuration file by using the `-config` argument. If the `-config` argument is omitted, the Core Recertification Tool uses the default configuration file located in `/opt/opsware/oi_util/OpswareCertTool/recert_utils/corerecert.conf`.

You can either directly edit the default configuration file or create a new one. Because the configuration file contains sensitive information, it is important that this file be protected accordingly. For example, by ensuring that it is readable and writable only by the root user.

For a core environment upgrade from SA 9.1x or 10.0x to SA 10.2, the core or satellite `/etc/opt/opsware/crypto/security.conf` file is only generated during the Core Recert process.

For core environment upgrade from SA 10.1 to SA 10.2, a fresh install of SA 10.01, or a fresh install of SA 10.2, the `/etc/opt/opsware/crypto/security.conf` file is already generated.

HP does not support a manually created `/edit /etc/opt/opsware/crypto/security.conf` file.

The parameters listed in Table 16 are found in the `corecert.conf` file. Some of these parameters (`fips_enabled` value, key size, signing algorithm, and custom CA), which denote values for the Core, are also found in the `security.conf` file.

Table 16. Core Recertification Configuration File: Properties

Property Name	Req?	Description	Example
Global Properties			
<code>username=<username></code>	Yes	User name of the user who has privilege to perform Core Recertification operations	<code>username=jdoe</code>
<code>password=<password></code>	Yes	Password of the user who has privilege to perform Core Recert operations.	<code>password=dontask</code>
Agent Recertification Properties			
<code>agent_recert.cleanup_old_agent_ca=<0 1></code>	No	Indicates whether to clean up the old Agent CA after Core Recertification. Cleanup of old Agent CA phase is not necessary and can be disabled. The valid values are 1 (true) or 0 (false). Any other value will result in an invalid property error. This is an optional property. Default: 0. Note: If a <code>custom_ca</code> is specified, HP suggests that the <code>agent_recert.cleanup_old_agent_ca</code> parameter should be set to 1, so only the customer certificate is available to be trusted.	<code>agent_recert.cleanup_old_agent_ca=0</code>
<code>agent_recert.all.</code>	Yes	The default start time for the	<code>agent_recert.all.</code>

Property Name	Req?	Description	Example
<p>facilities. start_time= <YYYY:MM:DD:HH:mm></p>		<p>Agent Recertification operation for all facilities.</p> <p>You can override this value for a specified facility (by specifying a default facility start time using the agent_recert.facility.<facilityname>.start_time property).</p> <p>The start time must be in the following format:</p> <p>YYYY:MM:DD:HH:mm, where 2008 <= YYYY <=9999, 0 < MM <= 12, 0 < DD <= 31, 0 <= mm < 12, and 0 <= MM < 60.</p>	<p>facilities.start_time= 2009:02:15:23:00</p>
<p>agent_recert. facility.<facility name>.start_time</p>	No	<p>You can override the default facility start time for a given facility by specifying this property.</p> <p>The start time must be in the following format:</p> <p>YYYY:MM:DD:HH:mm, where 2008 <= YYYY <=9999, 0 < MM <= 12, 0 < DD <= 31, 0 <= mm < 12, and 0 <= MM < 60.</p>	<p>agent_recert.facility. yellow.start_time= 2008:05:01:10:00</p>
<p>agent_recert.all. facilities.duration=<HH></p>	Yes	<p>The default duration, in hours, for the Agent Recertification operation in all facilities.</p> <p>Duration must be an integer value between 1 and 24.</p> <p>You can override the duration for a given facility by specifying the agent_recert.facility.<facility name>.duration property</p>	<p>agent_recert.all. facilities.duration=2</p>
<p>agent_recert. facility.<facility name>.duration=<HH></p>	No	<p>Overrides the default duration for a specific facility.</p>	<p>agent_recert.facility. yellow.duration=10</p>

Property Name	Req?	Description	Example
agent_recert.all.facilities.success_rate=<%>	Yes	The default success rate (in whole percentage) for the Agent Recertification operation in all facilities. You can override this value for a specific facility by specifying the agent_recert.facility.<facility name>.success_rate property	agent_recert.all.facilities.success_rate=50
agent_recert.facility.yellow.success_rate=<%>	No	Overrides the default success rate for a given facility.	agent_recert.facility.yellow.success_rate=98
agent_recert.all.facilities.job_notification=<email_address>	No	The default job email notification for the Agent Recertification operation. You can override the default job email notification for a specific facility by specifying the agent_recert.facility.<facility name>.job_notification property	agent_recert.all.facilities.job_notification=admin@example.com
agent_recert.facility.<facility name>.job_notification=<email_address>	No	Overrides the default job email notification for a specific facility.	agent_recert.facility.yellow.job_notification=sadmin@example.com

Core Recertification Properties

cadb_password=<pswd>	Yes	The password to protect the newly generated crypto database.	cadb_password=crypto123
debug=<0 1>	No	Specifies whether to run the Core Recertification job in debug mode. It can be either 1 (true) or 0 (false). Debug logs are found on the Core machine where the Core Recert is invoked:	debug =1

Property Name	Req?	Description	Example
fips_enablement	No	<p data-bbox="683 323 1057 436">/var/log/opsware/waybot/recert.log. Default: 0.</p> <p data-bbox="683 470 1057 835">Denotes FIPS enablement for mesh and satellite. The default is to use the value in /etc/opt/opsware/crypto/security.conf. If this value is not set or cannot be read, the default is zero (0). If the fips_enablement value is set to 1 (enabled), the signing_algorithm value must be sha1. Values are: 1 (FIPS enabled) and 0 (FIPS disabled).</p> <div data-bbox="683 856 1057 1188" style="background-color: #f0f0f0; padding: 5px;"> <p>Note: SA AGENTS version 10.1 and later are required for FIPS enablement. You can upgrade from SA 9.1x or 10.0x to SA 10.20 and enable support for FIPS if you use the Core Recert process</p> </div> <div data-bbox="683 1209 1057 1398" style="background-color: #f0f0f0; padding: 5px;"> <p>Note: If FIPS is enabled, you must use SHA1, not SHA256, as the hashing algorithm.</p> </div>	fips_enablement=0
base_core_server_ref=<n>	No	Server reference of the host from which you launch Core Recertification.	base_core_server_ref=10010
job_schedule=<YYYY:MM:DD:HH:mm>	No	<p data-bbox="683 1549 1057 1642">Job schedule for the current Core Recertification phase jobs. It must be in the format:</p> <p data-bbox="683 1654 1057 1810">YYYY:MM:DD:HH:mm, where 2008 <= YYYY <=9999, 0 < MM <= 12, 0 < DD <= 31, 0 <= HH < 12, and 0 <= mm < 60.</p> <p data-bbox="683 1835 1057 1864">If this property is not specified,</p>	job_schedule=2009:02:12:23:05

Property Name	Req?	Description	Example
<code>job_schedule.gateway_recert. <facility name>= <YYYY:MM:DD:HH:mm></code>	No	<p>the job starts immediately.</p> <p>Job schedule for the Gateway Recertification phase for a given facility. It must be in the format: YYYY:MM:DD:HH:mm, where 2008 <= YYYY <=9999, 0 < MM <= 12, 0 < DD <= 31, 0 <= HH < 12, and 0 <= mm < 60.</p> <p>If this property is not specified, the <code>job_schedule</code> property for the gateway recertification phase is used.</p>	<code>job_schedule.gateway_recert.<facility name>= 2009:02:12:23:05</code>
<code>keysize</code>	No	<p>The keysize parameter specifies the key length, in bits, for the public key used to verify the certificate. The default is the value in the current SA certificate. If <code>custom_ca</code> is also used, and this value is set, the value must conform to the keysize value in <code>custom_ca</code>. Values are: 2048 and 4096.</p>	<code>keysize=2048</code>
<code>job_notification= <email_address></code>	No	<p>Job notification for all Core Recertification phase jobs.</p> <p>You can override this value for a given phase by specifying the <code>job_notification.<phase_number></code> property</p>	<code>job_notification= admin@example.com></code>
<code>job_notification. <phase_number>= <email_address></code>	No	<p>Job notification for a specified Core Recertification phase.</p>	<code>job_notification.7= saadmin@example.com</code>
<code>job_notification.gateway_recert. <facility name>= <email_address></code>	No	<p>Job notification for the Gateway Recert phase for a given facility.</p>	<code>job_notification.gateway_recert.yellow= admin@acme.com</code>
<code>cleanup_old_opsware_ca=<0 1></code>	No	<p>Specifies whether to clean old SA CA after Core Recert.</p>	<code>cleanup_old_opsware_ca=1</code>

Property Name	Req?	Description	Example
custom_ca	No	<p>SA CA cleanup is not necessary unless the CA has been compromised. In most cases, old SA CA cleanup is not necessary and should be disabled.</p> <p>The valid values are 1 (true) or 0 (false). Any other value will result in an invalid property error.</p> <p>Default: 0 (false)</p> <p>Note: HP suggests that the parameter should be set to 1, so only the customer certificate is available to be trusted.</p>	
		<p>Full path to the valid custom certificate file that conforms to the custom certificate requirements. If the value of this parameter is set to the path of the valid certificate authority, the default behavior is for core recert to use that value to generate all self-signed (customer-specific) certificates used by SA. Core recert uses either the value of the custom_ca parameter or the value of the signing_algorithm parameter. In addition, note the following: The file containing the certificate must also include a concatenated private key. a concatenated private key If fips_enablement is set to 1, custom_ca must have conforming signing_algorithm and the keysize values. If the values conflict, you will see an error message.</p> <p>Note: SA AGENTS version 10.1 and later are required for FIPS enablement. You can upgrade from SA 9.1x or 10.0x to SA 10.20 and enable support for FIPS if you use the Core Recert process.</p> <p>Valid value is full path to custom</p>	<pre>custom_ca=/- /tmp/custom- ca.crt</pre>

Property Name	Req?	Description	Example
signing_algorithm	No	<p>certificate.</p> <p>The signing_algorithm parameter is used to generate the certificate signature when supported keysize values are provided. If you also use custom_ca, and the signing_algorithm value is set, this value must conform to the value in the signing_algorithm in custom_ca. The default is the value in the existing SA certification. Values are: sha1 and sha256. md5 is optionally supported only if the core's existing certificate is md5 base.</p>	signing_algorithm=sha1

Note: During the core recert process, values in the corerecert.conf file and the security.conf file are compared. The security.conf file, generated as part of the core recert process, contains signing_algorithm values and keysize values. If the values in the two files conflict, the process displays a message that asks you if you want to overwrite the values in the security.conf file. If you enter y, SA replaces the values in the security.conf file with the values in the corerecert.conf file. If you do not want to overwrite the values, enter n. The Core Recert process exits and your the current values in the security.conf file remain intact.

Ensure that All Cores are Running/Resolve Conflicts

Before performing Core Recertification, it is strongly recommended that you run System Diagnosis on all Cores to be recertified to ensure that they are running correctly.

You must resolve all transaction conflicts and ensure that there is no transaction backlog in the mesh.

For more information, see [Running a System Diagnosis](#) and [Resolving Mesh Conflicts - SA Client](#).

Ensure That the Core Recertification Tool Correctly Recognizes the Mesh Setup

You must perform the following tasks to ensure that the Multimaster Mesh setup is correctly recognized by the Core Recertification Tool:

1. From the command line, log on to an SA Core host as root user.
2. Run


```
/opt/opsware/oi_util/OpwareCertTool/recert_utils/discover_mesh -p
```
3. Check the output to make sure it reflects your current Mesh setup. If not, contact HP Professional Services before proceeding with Core Recertification.

Recertifying SA Cores

Note: You must clear all backlogs and conflicts on your Multimaster Mesh before you start a core recertification.

Note: Some recertification phases will be performed automatically, while others require multiple runs of the `corerecert` tool.

To recertify SA Cores, perform the following tasks:

1. Ensure that you are classified as a Core Recertification User. If not, see your SA System Administrator.
2. Log on to an SA Core host.
3. Change directory to `/opt/opsware/oi_util/OpawareCertTool/recert_utils/`.
4. Edit:

```
corerecert.conf
```

to ensure that the information is correct for your environment.

5. Run:

```
corerecert --status
```

to ensure Core Recertification is not currently in progress.

6. Run:

```
discover_mesh -p
```

to make sure the Core Recertification Tool can correctly detect your Mesh setup.

7. Run:

```
corerecert --phase 1
```

from the command line to initialize Core Recertification.

Mesh-wide gateways will be automatically restarted.

8. Monitor the progress on screen by running:

```
corerecert --status
```

until it has indicated Phase 4 is in progress.

9. Run:

```
corerecert --phase 4
```

from the command line to start Phase 4, which appends a new Agent CA to all the Agents.

Note: If the `agent_recert.using_cdr` parameter value is 0 in the `corerecert.conf` file, the run phase 4 process is skipped, and the process begins again at the beginning of the next phase.

10. Monitor the progress on screen by running:

```
corerecert --status
```

until all the Agents have successfully had a new agent CA appended.

Note: This step could take days depending on your maintenance windows and the Agent availability. There can be only one scheduled or active Agent Recertification job per facility at any given time. If you encounter any errors during this stage, resolve the errors and go back to step 9. You only need to reschedule the facilities that had errors. You do not need to reschedule the Agent Recert job for the successful facilities.

The recertification will stay in the `agent-recert` phase with a `PHASE_IN_PROGRESS` status until there is a user action. Move on to the next phase when you are satisfied with the success rate.

11. Run:

```
corerecert --phase 6
```

from the command line to start Phase 6 of the core recertification.

12. Monitor the progress on screen by running:

```
corerecert --status
```

until it has indicated `mesh_restart_pending`.

At this point, you must restart the mesh, using the mesh restart instructions and sequences in the SA Administration Guide, SA Maintenance section.

This step could take days depending on your maintenance window. If you encounter any errors during this stage, make sure you resolve the errors and go back to step 11.

13. After manually restarting the mesh successfully, run:

```
corerecert --phase 6
```

from the command line to continue phase 6.

In this step, SA performs two functions:

- Checks to see if the restart took place on the cores.
- Automatically restarts the primary-spin component of each SA facility.

14. Monitor the progress on screen by running:

```
corerecert --status
```

until it indicates that Phase 7 is in progress. If you encounter errors, resolve them and go back to step 13.

15. Run:

```
corerecert --phase 7
```

from the command line to start phase 7.

Note: Mesh-wide gateways will be automatically restarted.

16. Monitor the progress on screen by running:

```
corerecert --status
```

until it indicates that Phase 8 is in progress. If you encounter errors, resolve them and go back to step 15.

17. Run:

```
corerecert --phase 8
```

from the command line to start Phase 8, which recertifies all the Agents.

18. Monitor the progress on screen by running:

```
corerecert --status
```

until all Agents have successfully been recertified.

The recertification will stay in the agent-recert phase with a PHASE_IN_PROGRESS status until there is a user action. Move on to the next phase when you are satisfied with the success rate.

Note: This step could take days depending on a customer's maintenance windows and the agent availability. There can be only one scheduled or active Agent Recertification job per facility at any given time. If you encounter any errors, resolve them and go back to step 17. You only need to reschedule the facilities that had errors, not the Agent Recertification job for the successful facilities.

19. Run:

```
corerecert --phase 9
```

from the command line to start phase 9. The Core Recertification Tool prompts you to confirm that you want to begin phase 9. Press `y` to continue.

20. Monitor the progress on screen by running:

```
corerecert --status
```

until it has indicated `mesh_restart_pending`. If you encounter any errors during this stage, make sure you resolve the errors and go back to step 19.

At this point, ensure that there are no conflicts and no transaction backlogs in the mesh.

21. On the base Slice core server:

a. Issue the following commands:

```
touch /var/opt/opsware/crypto/twist/upgradeInProgress  
/etc/init.d/opsware-sas restart
```

b. Wait till the restart finishes successfully.

At this point, work with your SA System Administrator to restart the rest of the mesh. This step could take days depending on your maintenance window. If you encounter any errors, resolve them and go back to step 19.

22. After the mesh has been successfully restarted, the Recertification User must run:

```
corerecert --phase 9
```

from the command line to continue phase 9.

SA checks to see if the restart took place on the cores.

23. Monitor the progress on screen by running:

```
corerecert --status
```

until it indicates that Phase 11 is in progress. If you encounter any errors, resolve them and go back to step 22.

24. On the base slice core server:

a. Issue the following command:

```
touch /opt/opsware/oi_util/OpswareCertTool/recert_  
utils/TruthResignStatus.txt /opt/opsware/oi_  
util/OpswareCertTool/recert_utils/WordResignStatus.txt
```

b. Run phase 11:

```
corerecert -phase 11
```

from the command line to start Phase 11 which resigns data in model repository, software repository, recurring jobs and audit streams.

25. Monitor the progress on screen by running:

```
corerecert --status
```

until it indicates that Phase 12 is in progress. If you encounter any errors, resolve them and go back to step 24b.

26. If you do not intend to remove the Agent CA, skip to step 28. Otherwise, run:

```
corerecert --phase 12
```

from the command line to start phase 12, which removes the old Agent CA from all the Agents.

Note: At this point, you must restart the mesh, using the mesh restart instructions and sequences in the SA Administration Guide, SA Maintenance section.

27. Monitor the progress on screen by running:

```
corerecert --status
```

until the old Agent CA has removed from all the Agents.

Note: This step could take days depending on customer's maintenance windows and the agent availability. If you encounter any errors during this stage, resolve the errors and go back to step 26. You only need to reschedule the facilities that had errors. You do not need to reschedule the Agent Recertification job for the successful facilities.

Note: For Core Recert with custom certificate, HP recommends that you go through phase 13 so the old Core-component certificate is removed from the trusted CA store, and, therefore, only the customer certificate chain is used for verification.

The recertification will stay in the agent-recert phase with a PHASE_IN_PROGRESS status until there is a user action. Move on to the next phase when you are satisfied with the success rate.

28. Run:

```
corerecert --phase 13
```

from the command line to start phase 13.

A mesh restart is not required in this phase. A restart will remove the old CAs (`cleanup_old_opsware_ca`) in the config file.

29. Monitor the progress on screen by running:

```
corerecert --status
```

until it indicates `mesh_restart_pend` or `core_recert_completed`.

Continue with the remaining instructions only if the status is `mesh_restart_pending`.

Note: At this point, you must restart the mesh, using the mesh restart instructions and sequences in the SA Administration Guide, SA Maintenance section.

Note: This step could take days depending on the customer's maintenance window. If you encounter any errors during this stage, resolve the errors and go back to step 28.

30. After the mesh has been successfully restarted, run:

```
corerecert --phase 13
```

from the command line to continue phase 13.

31. Monitor the progress on screen by running:

```
corerecert --status
```

until it indicates that Core Recertification has completed successfully.

Agent Recertification

This section describes how to recertify the agent on one or more managed servers. You can recertify the agent on one or more servers separately from a full core recertification process. The full core recertification process recertifies the core and all agents. For more information, see [Agent versus Core Recertification](#) and [SA Core Recertification](#).

To recertify the agents on one or more managed servers, perform the following steps:

1. In the SA Client, select the Devices tab.
2. Under the Servers node, select All Managed Servers or Virtual Servers. This displays all the corresponding servers.

Or under Device Groups, select one or more device groups.

3. Select the **Actions** menu, or right-click and select **Run > Agent Recert.**

Or if **Run Extension > Recertify Agent** is not shown, select **Run Extension > Select Extension**. This displays the Select Extension window and lists the available extensions. Select **Recertify Agent** on the Managed Servers in the Select Extension window, then select OK.

This displays the Run Program Extension window showing the servers or device groups you selected.

4. At any time, you can select the Start Job button to accept all the remaining default settings and run the job.
5. Optionally use the Include Devices button to add servers or device groups.
6. Optionally use the Remove button to remove servers or device groups.
7. Select the Next button. This displays the Program screen. Do not make any changes on the Program screen.
8. Select the Next button. This displays the Options screen.
9. On the Options screen, you can change the program timeout value, request detailed information about the job with the `-debug` option, or specify the amount of job output to save.
 1. Program Timeout—Specify the maximum time in minutes you want the agent recertify job to run. If the agent recertify job fails, it will continue running for the specified time period. If after that time period it has not succeeded, it will abort and display an error message.
 2. Usage options—Enter “`-debug`” in the text box if you want additional details about the job to be displayed.
 3. Output Options—Specify what you want done with the program output after the job finishes. If you specify “Discard all program output,” then all the output will be unavailable when you open the completed job.

10. Select the Next button. This displays the Scheduling screen. Specify when you want the job to run.
11. Select the Next button. This displays the Notifications screen.
12. On the Notifications screen, specify the email recipients and whether they should receive email messages if the job fails or succeeds or both.
13. Select the Next button. This displays the Job Status screen.
14. Select the Start Job button. This starts the job and displays the status.
15. Select any server to display details on the status of the job on that server.
16. After the agent recertify job finishes, you can optionally run a communication test on your servers to verify the agents on them.

Multimaster Mesh Administration

This section explains how to administer and maintain a Multimaster Mesh. It does not document how to configure SA for a Multimaster Mesh. For more information about Multimaster architecture and planning for and installing a Multimaster Mesh, see the SA Overview and Architecture Guide and the SA Installation Guide.

Built-In Redundancy of the Multimaster Mesh

Each SA core manages one data center. Each data center is represented as a facility in SA. A multimaster mesh is two or more SA cores managing an equal number of facilities. A multimaster mesh can optionally include one or more SA satellites. An SA satellite is a “mini” SA core that manages a smaller number of servers than a full SA core.

The multimaster mesh configuration of SA is designed for redundancy, reliability, and high availability. A multimaster mesh consists of multiple synchronized cores. All data on each core is synchronized with every other core so that if one core goes down, the other cores handle all requests and jobs.

A multimaster mesh also provides load balancing for better performance.

What Are Multimaster Mesh Conflicts?

In a multimaster mesh (which by definition consists of two or more SA cores), when SA users perform any action on any core, each core forwards the transaction details to all the other cores in the mesh to keep them all synchronized. If two users perform overlapping or conflicting actions on two different cores, when the cores forward the transactions to the other cores, a conflict will occur.

SA can detect these kinds of conflicts, notify you when they occur, and help you resolve them.

The SA core itself cannot resolve the conflicts. SA administrators must use the **Multimaster Tools** in the SA Client to resolve the conflicts at the target databases when they occur to ensure that the transactions are not lost.

1. To view conflicts, see [Viewing the State of the Multimaster Mesh - SA Client](#).
2. To resolve conflicts, see [Resolving Mesh Conflicts - SA Client](#).

3. You can also use the System Diagnosis tools in the SA Client to view information about the health of the multimaster components. For more information, see [Troubleshooting SA - Diagnostic Tests](#).

How SA Handles Mesh Conflicts

Each SA core manages one facility. When an SA core (the source core) sends a transaction to another core (the destination core) and a conflict occurs, SA detects the conflict and the following occurs:

1. The transaction is canceled.
2. All SA database rows affected by the transaction are locked, thereby preventing further changes to those rows.
3. The source core propagates the transaction lock to all other cores in the mesh, thereby locking the rows in all cores.
4. An alert message with the conflict information is emailed to a user-configured mailing list. For more information, see [Multimaster Email Alerts](#).
5. Both the source core and the destination core continue to the next transaction.

If either the source core or the destination core encounters an exception that prevents it from going to the next transaction, it sends an email to the user-configured mailing list describing the problem and shuts down.

To manually resolve conflicts and unlock the database rows, see [Resolving Mesh Conflicts - SA Client](#).

Best Practices for Preventing Mesh Conflicts

This section lists measures you can take to minimize multimaster mesh conflicts.

The probability of multimaster conflicts varies depending on the following factors:

- The number of servers under management—the more servers, the more likely that conflicts can occur.
- The number of cores in the multimaster mesh.
- The number of SA Clients being used by your SA users—the more users making updates, the more opportunities for conflicts.
- The propensity for users to make changes in more than one facility by using different SA Clients.

Users

Your users should be aware of the following:

- Users in multiple facilities are able to modify the same data at the same time, so when possible coordinate updates to avoid conflicts.
- Users should not change data in one facility and immediately make the same change in another facility, because SA automatically propagates changes. Making the same change in multiple facilities will usually result in mesh conflicts.
- A slight time delay occurs before changes that a user makes can propagate to other SA facilities. The length of delay varies depending on a number of factors, including network connectivity and bandwidth. If an update has not yet propagated to all the other Model Repositories in the mesh, wait a reasonable period of time to insure that the transaction has not been delayed before attempting to redo the transaction or perform another update that depends on other recent transactions.

Administrators

Implement the following best practices to reduce the chance of data conflicts:

- Ensure that your network connections are reliable and there is sufficient network bandwidth between facilities in the mesh. The risk of conflicts increases as bandwidth decreases.

See [Network Administration for a Multimaster Mesh](#) for more information.

See the SA Installation Guide for information about network connectivity when running SA in a Multimaster Mesh.

- When possible, partition your data space so that only one user can change the same object in different facilities concurrently.
- Have a user, or a small group of coordinated users, manage a given set of servers. Partitioning the data space ensures accountability of server ownership and prevents users from changing each other's data.

The SA Client facilitates this by allowing you to set permissions by customer, facility, and user group types.

See [Permissions Reference](#) for more information about user groups and SA permissions.

Viewing the State of the Multimaster Mesh - SA Client

The Multimaster Tools show you the status of transactions between each pair of facilities in your SA deployment. They also allow you to resolve any conflicts that occur. You can view details about all the transactions between facilities in the Multimaster Mesh as follows:

1. In the SA Client, select the Administration tab.
2. Under the Multimaster Tools node, select the **State View**. This displays a table showing all your facilities (each facility corresponds to an SA core) and the state of all transactions between each pair of facilities. **Table 17** shows the meanings of the color codes in the state view.

Table 17. Multimaster Transaction State Color Codes

Transaction Color	Transaction State
Blue	Sent - Lists the number of transactions successfully sent to other facilities.
Green	Received - Lists the number of transactions successfully received by the facility.
Purple	Not Sent - One or more transactions in the facility have not yet been sent to the other facilities in the mesh.
Yellow	Not Received - One or more transactions sent from another facility have not yet been received by the facility.
Red	Conflict - One or more conflicts have occurred.

3. To view details about all the conflicting transactions, select the **Conflict View** in the navigation bar. This displays details about each transaction including the following:
 - Transaction—This is a transaction identifier and a link where you can get more detailed information about the conflicting transaction.
 - Action—This describes what the transaction consists of; for example, database updates, inserts, and deletes.
 - Table—This lists the database table affected by the transaction.
 - Count—This lists how many actions were performed on the database elements.
 - User—This lists the SA user who performed the action that resulted in the conflict. Contact this person to verify what they were attempting to do so you can accurately resolve the conflict.
 - Created Time—This is the date and time when the transaction occurred.
 - Source Facility—This is the core from which the transaction was sent.
 - Conflicting Facility—These are the cores where the transaction was received and where the conflict was detected.
4. To view details about a specific transaction conflict, select the Transaction link. This displays details about the selected transaction.
 - Table—This shows the SA database table where the conflict occurred.
 - DB Field—This shows all the SA database field names in the database table where the conflict occurred.

- Facility columns—The remaining columns are for each facility in your SA deployment. Each column lists the values in the corresponding facility. Wherever a conflict occurred, the values are shown in red text.

5. To resolve conflicts, see [Resolving Mesh Conflicts - SA Client](#).

Figure 25 shows the multimaster mesh state view, with no conflicts. All three cores in the multimaster mesh—London, Paris, and Vienna—are up to date. All changes in all cores have been successfully sent to all other cores.

Figure 25. Multimaster Mesh Conflicts, State View—No Conflicts

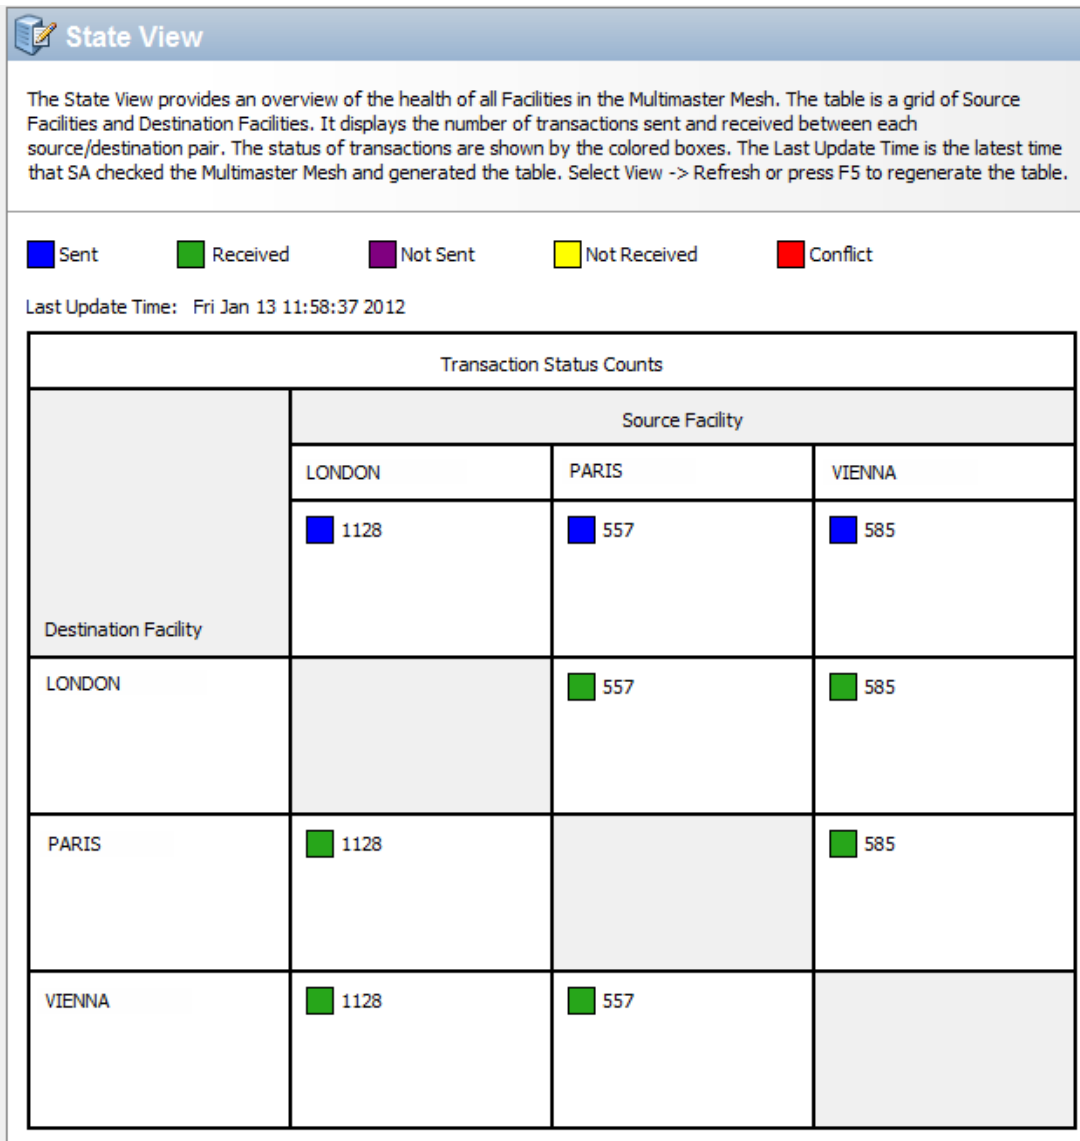


Figure 26 shows the mesh state view with no conflicts, but two changes have been made in two cores and are about to be propagated to the other cores. Two changes have been made to the London core and two changes have been made to the Vienna core.

Figure 26. Multimaster Mesh Conflicts, State View—Changes Waiting to be Sent

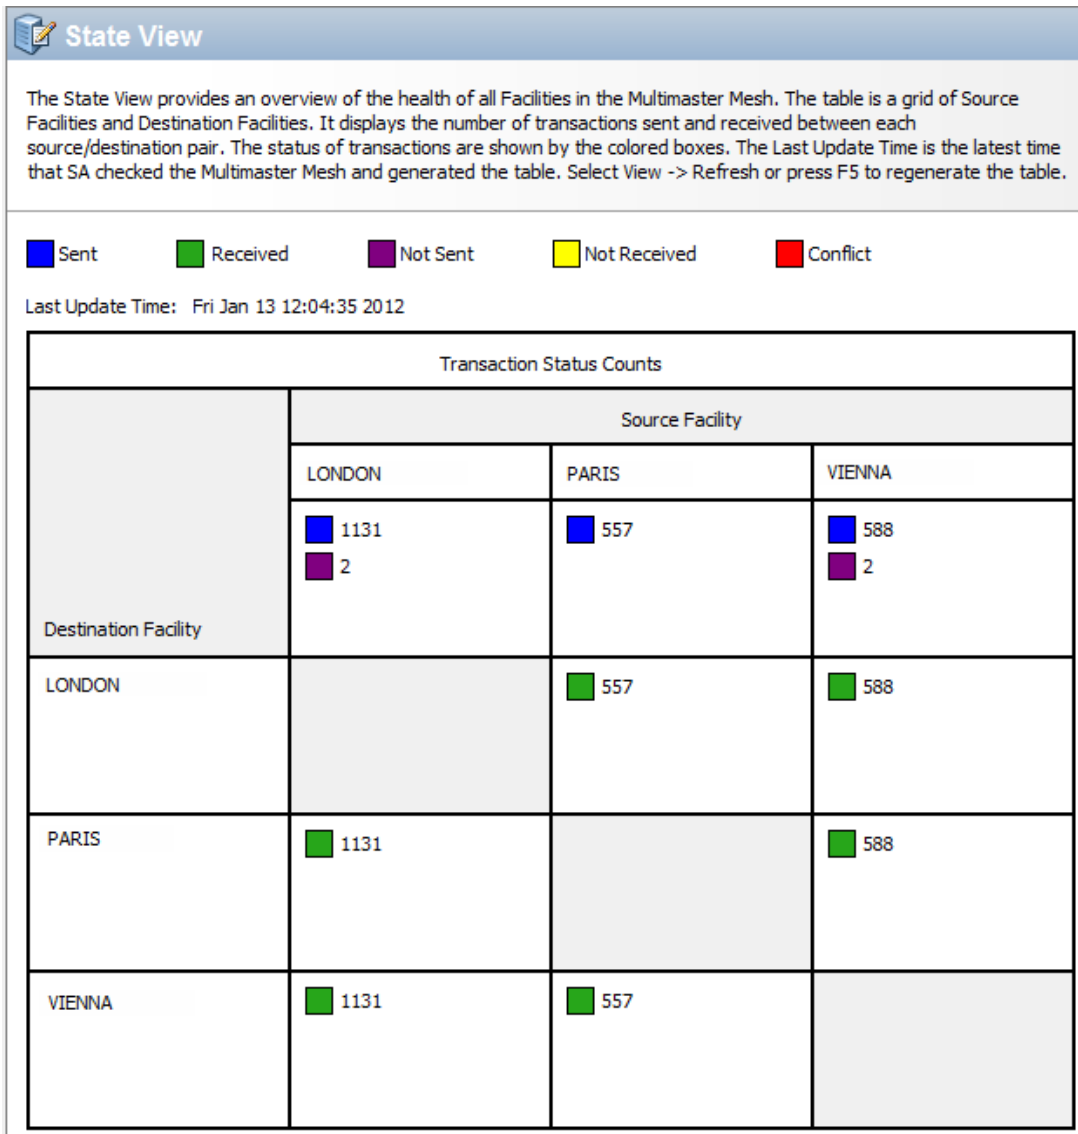



Figure 27 shows the mesh state view with two conflicts, in the London core and in the Vienna core. The London core has a conflict with the Vienna core, and the Vienna core has a conflict with both the London and Paris cores. To resolve conflicts, see [Resolving Mesh Conflicts - SA Client](#).

Figure 27. Multimaster Mesh Conflicts, State View—Two Conflicts

 **State View**

The State View provides an overview of the health of all Facilities in the Multimaster Mesh. The table is a grid of Source Facilities and Destination Facilities. It displays the number of transactions sent and received between each source/destination pair. The status of transactions are shown by the colored boxes. The Last Update Time is the latest time that SA checked the Multimaster Mesh and generated the table. Select View -> Refresh or press F5 to regenerate the table.

■ Sent
 ■ Received
 ■ Not Sent
 ■ Not Received
 ■ Conflict

Last Update Time: Fri Jan 13 12:05:36 2012

Transaction Status Counts			
Destination Facility	Source Facility		
	LONDON	PARIS	VIENNA
Destination Facility	■ 1143 ■ 1	■ 557	■ 590 ■ 1
LONDON		■ 557	■ 590 ■ 1
PARIS	■ 1143		■ 590 ■ 1
VIENNA	■ 1143 ■ 1	■ 557	

Resolving Mesh Conflicts - SA Client

To resolve multimaster mesh conflicts with the SA Client, perform the following steps.

Tip: Before you resolve conflicts, notify the subscribers of the email alert alias. Notifying these users helps to prevent other SA administrators from undoing or affecting each other's conflict resolution efforts. While resolving conflicts, you should resolve the conflict from the SA Client of a single facility. Do not attempt to resolve the same conflict multiple times from the SA Client of different facilities.

Note: If you see a large volume of conflicts that you cannot resolve by using the Multimaster Tools, contact your HP Server Automation Support Representative for assistance in synchronizing databases.

Make sure you have adequate SA permissions to view and resolve conflicts. For more information on permissions, see [Permissions Reference](#).

1. In the SA Client, select the Administration tab.
2. Under the Multimaster Tools node, select the **Conflict View**. This displays details about all the conflicts in the mesh. **Figure 28** shows the Conflict View with two conflicts originating in the London facility and the Vienna facility. For an overview of the conflicts, select the **State View**.

Figure 28. Multimaster Mesh Conflict—Conflict View

Conflict View							
<p>The Conflict View shows all conflicts in the Multimaster Mesh. The table lists each conflict by a transaction ID number, the action that caused the conflict, the database objects affected by the conflict, the user responsible for the conflict, the time the offending action occurred, the source facility that originated the transaction, and the facility where the transaction conflict occurred. The Last Update Time is the latest time that SA checked the Mulitmaster Mesh and generated the table. Select View -> Refresh or press F5 to regenerate the table. To resolve a conflict, click the transaction ID number to show the transaction differences.</p>							
Last Update Time: Fri Jan 13 12:39:26 2012							
Transaction	Action	Table	Count	User	Created Time	Source Facility	Conflicting Facility
7869210001	Insert	DEVICE_CHANGE_LOG	2	TOM	Fri Jan 13 12:0...	LONDON	VIENNA
	Insert	DEVICE_ROLE_CLASSES	1				
	Delete	DEVICE_ROLE_CLASSES	1				
	Update	DEVICE_ROLES	1				
7495990003	Insert	DEVICE_CHANGE_LOG	2	SAL	Fri Jan 13 12:0...	VIENNA	LONDON
	Insert	DEVICE_ROLE_CLASSES	1				PARIS
	Delete	DEVICE_ROLE_CLASSES	1				
	Update	DEVICE_ROLES	1				

3. Optionally press Control-F (the Ctrl and F keys) on your keyboard. This displays the find tool so you can search for a particular conflict. Press the Escape (Esc) key to close the find tool.
4. Examine each conflict, noting the user who performed the action, the source facility, and the conflicting facilities.
5. Select the transaction identifier link from the Transaction column. This displays details about the transaction.
6. Optionally press Control-F (the Ctrl and F keys) on your keyboard. This displays the find tool so you can search the details of a particular conflict. Press the Escape (Esc) key to close the find tool.

7. Examine each conflict, noting the details. You may have to investigate each conflict to determine what the conflict is, what user actions were performed to cause the conflict, who performed the actions, and the intentions of each user.
8. If possible, determine which facility has the correct data, and synchronize from that facility. Synchronizing from a facility copies the data from that facility to all other facilities, thereby resolving the conflict.

If no one facility has the correct data, you can synchronize from one facility, then redo the actions while avoiding the situation that caused the conflict.

You can optionally synchronize each separate database table; however, this method is not recommended unless you have knowledge of the SA database. To synchronize each separate table, select the appropriate buttons labeled Synchronize From This Facility at the bottom of each column, then go to [Select OK in the Mark Conflict Resolved window](#). This removes the conflict..

9. Once you determine which facility has the correct data, select that facility from the drop-down list labeled “Synchronize all objects from” near the top of the window.
10. Select the Synchronize button. This copies the data from the selected facility to all other facilities to resolve the conflict, and displays the Transaction Synchronization Results window.
11. Select OK in the Transaction Synchronization Results window.
12. Select the Mark Resolved button. This displays the Mark Conflict Resolved window, which shows the status of the mesh conflicts that you have resolved.
13. Select OK in the Mark Conflict Resolved window. This removes the conflict.
14. Examine the conflicts in the Conflicts View, and verify that the resolved conflict has been removed.

Advanced Types and Causes of Mesh Conflicts

This section describes some causes and types of multimaster mesh conflicts.

User Overlap Conflicts

Conflicts occur when a user concurrently makes a change using the SA Client in one facility at the same time another user makes a change to the same object in another facility.

For example:

1. Alice removes Node A from a server in the Atlanta facility.
2. Bob removes Node A from the same server in the Boston facility.
3. SA propagates the change from the Atlanta facility to the Boston facility; however, Bob has already removed Node A from the server in the Boston facility. SA

generates a Model Repository Multimaster Component conflict alert, because now it appears that Alice is requesting that a node that does not exist be removed.

4. SA also propagates Bob's update in Step 2 from the Boston facility to the Atlanta facility; however, Alice has already removed Node A from the server in the Atlanta facility. SA generates a second Model Repository Multimaster Component conflict alert.

Conflicts from User Duplication of Actions

Conflicts can also occur when a user, for various reasons, attempts make an update to a Model Repository, does not wait long enough for the update to propagate to the other Model repositories in the Mesh, thinks the update failed, and so attempts to make the update again, thus creating duplicate updates.

For example, this sequence of events could occur:

1. From a server in the Seattle facility, Carol uses the SA command line interface (CLI) to upload the package `carol.conf`.
2. Carol immediately logs in to the SA Client in the Phoenix facility and searches for the package. She does not see the package, because that data has not yet propagated from Seattle to Phoenix. Carol allowed enough time for data propagation between facilities.
3. Carol uploads the package `carol.conf` by using the SA Client in Phoenix.
4. When the data eventually propagates from Seattle, SA generates a conflict because the data already exists in Phoenix.

Conflicts from Out of Order Transactions

Transactions between two facilities usually arrive in the order in which they were sent. However, if a third facility is involved in the transactions, the correct ordering is not guaranteed. For example:

1. A user changes or inserts data at Facility A (Model Repository A).
2. The transaction for that change propagates to Facility B (Model Repository B) and to Facility C (Model Repository C).
3. However, the data is modified again or referenced at Facility B (Model Repository B) and then propagated to Facilities A and C.
4. If the transaction from Facility B (Step 3) reaches Facility C (Model Repository C) before the transaction from Facility A (Step 1), a conflict occurs.

This conflict typically occurs when a user uploads a package using the SA CLI in one facility, and immediately uses the SA Client to add the package to a Software Policy in a different facility.

The occurrence of out of order transactions can be aggravated by concurrent updates in different facilities or problems with inter-facility network connections.

For example:

1. Henry uses the SA CLI on a server in the Denver Facility to upload the package `henry.conf`.
2. SA propagates data about the package to all facilities in the mesh; however, it cannot propagate the data to the Paris Facility because the network connection is down.
3. Henry logs on to a server in the Miami Facility and uses the SA Client to update the description of the package `henry.conf`.
4. SA propagates data about the updated package description to all other facilities in the mesh; however, it cannot propagate the data to the Paris Facility, because the network connection is still down.
5. Network connectivity to the Paris Facility is restored, and the delayed transactions from Steps 2 and 4 are propagated to the Paris Facility.
6. The transaction for the updated package description arrives at the Paris Facility *before* the transaction that uploaded `henry.conf`. Therefore, the Model Repository in the Paris Facility does not contain data about `henry.conf`, so SA generates a conflict alert.
7. The transaction uploading `henry.conf` arrives at the Paris Facility and is processed without error. The package data exists in the Paris Model Repository, but the package description differs from all the other facilities in the mesh.

Database Conflicts

This section provides basic information about identifying the kind of conflicts you may have and the steps you can take to resolve them. See your Oracle database administration documentation for more information about identifying and resolving data and transaction conflicts.

Table 18 shows some types of conflicts:

Table 18. Types of Conflicts

Conflict	Description
Identical data conflict	The Multimaster Tools show a conflicting transaction, but the data is the same between facilities. The data is the same, because users made the same change in different facilities.
Simple transaction conflict	The row exists in all facilities, but some columns have different values or the row does not exist in some facilities (missing objects).
Unique-key constraint conflict	The object does not exist in a facility and cannot be inserted there, because inserting it would violate a unique-key constraint.
Foreign-key constraint conflict	The row does not exist in some facilities and cannot be inserted, because the data contains a foreign key to

Conflict	Description
	another object that also does not exist in that facility.
Linked object conflict	A type of conflict encountered in rare cases. SA includes business logic that links specific related objects in SA, such as a custom attribute name and value, and a customer created in the SA Client (appears in lists) and the associated node for the customer in the node hierarchy. SA ensures that links between related objects are maintained. Resolving a linked object conflict can be complex, because you must attempt to preserve the intent of the transaction that caused the conflict. Contact your HP Server Automation Support Representative to help you resolve linked object conflicts.

Guidelines for Resolving Each Type of Conflict

In general, when you resolve conflicts, apply updates so that the target always reflects the most current data based on the time stamp of the originating changes.

When you cannot follow one of the preceding guidelines, attempt to preserve the intent of the transaction. Contact the users who are generating the transactions and determine what types of changes in the managed environment each user was trying to make.

Identical Data Conflict

All objects in a transaction contain exactly the same data across all facilities. This type of conflict includes the case where the objects do not exist in all facilities.

To resolve an identical data conflict, simply mark the conflict resolved.

Identical Data Conflict (Locked)

All objects in a transaction contain exactly the same data across all facilities, but the objects in the transaction are still locked (marked conflicting).

To resolve this type of conflict, pick an arbitrary facility and synchronize all objects from it. Performing this action unlocks the objects. After synchronizing the data, mark the conflict resolved.

Simple Transaction Conflict

The data is different between facilities or some objects are missing from some facilities. None of the objects depends on the actions of other conflicting transactions. The results of synchronizing the objects does not result in a database foreign-key or unique-key constraint violation.

To resolve a simple transaction conflict, choose the facility that contains the correct data and synchronize from it. How you determine which facility contains the correct data varies depending on the type of transaction:

- If the conflict is the result of two users overriding each other's work, talk to the users and determine which user's change should be correct.
- If the conflict is the result of automated processes overriding each other's data, the most recent change is usually correct.
- If the conflict is the result of out-of-order transactions, the most recent change is usually correct.

After synchronizing the data, mark the conflict resolved.

Unique-Key Constraint Conflict

Resolving these conflicts results in a unique-key constraint violation.

For example, this sequence of events occurs:

1. From the SA Client in the London Facility, John creates Node A1 as a subordinate node of Node A.
2. From the SA Client in the San Francisco Facility, Ann performs the same action. She creates Node A1 as a subordinate node of Node A.
3. Node names must be unique in each branch of the node hierarchy.
4. SA propagates the node changes from the London and San Francisco facilities to the other facilities. Inserting the rows into the Model Repository databases at other facilities causes a unique-key constraint violation and a conflict.

Resolving this conflict by inserting the updates from the London Facility in all facilities would fail with the same unique-key constraint violation.

Perform the following steps to resolve a unique-key constraint conflict:

1. Locate all the involved transactions, and synchronize one transaction from a facility where the object does not exist, thereby deleting it in all facilities.
2. Synchronize the other transaction from a facility where the object exists, thereby inserting the object in all facilities. One of the two uniquely conflicting objects will take the place of the other.

Foreign-Key Constraint Conflict

Resolving these conflicts results in a foreign-key constraint violation.

For example, this sequence of events occurs:

1. Jerry creates Node B in Facility 1.
2. Before that transaction has time to propagate to other facilities, Jerry creates Node C as a subordinate node of Node B.
3. When the first transaction arrives at Facility 2, it generates a conflict for unrelated reasons.

4. When the second transaction arrives at Facility 2, inserting the row for Node C causes a foreign-key constraint conflict, because the parent Node (Node B) does not exist.

Resolving the second conflict first by inserting the update for Node C into all facilities would fail with the same foreign-key constraint violation.

Perform the following steps to resolve a foreign-key constraint conflict:

1. Resolve the conflicting transaction for Node B (the parent Node) by synchronizing the first transaction from the facility where the object exists.
2. Synchronize the second transaction (the Node C update) from the facility where the object exists.

Generally, resolving conflicts in the order in which they were created avoids generating foreign-key constraint conflicts.

Network Administration for a Multimaster Mesh

SA does not require that a Multimaster Mesh configuration meet specific guidelines on network uptime. A Multimaster Mesh configuration can function acceptably in a production environment that experiences temporary inter-facility network outages.

However, as the duration of a network outage increases, the probability of conflicts increases. Extended network outages between facilities can cause the following problems:

- Multimaster messages can fail to propagate between facilities
- The Multimaster Tools can stop functioning
- SA Clients cannot contact the multimaster central Data Access Engine

Production experience for multimaster configurations supports the performance data that **Table 19** shows.

Table 19. Performance Data for Multimaster Configurations

Number of Facilities	Duration Network Outage	Number of Multimaster Conflicts *
8 facilities (SA core installed in each facility)	12 hour outage (1 facility loses network connectivity to the other facilities)	12 to 24 conflicts (average number generated)

* The propensity of users to manage servers in the disconnected facility with SA Clients in other facilities increases the number of conflicts.

Network connectivity issues include SA Bus or multicast routing problems.

Multimaster Email Alerts

When Multimaster conflicts occur or Multimaster components experience problems, SA sends an email to the user-configured Multimaster email alias. You configure this email address when you install SA. If you must change this email address, contact your HP Server Automation Support Representative or see [SA Notification Configuration](#) for more information.

The subject line of the alert email specifies:

- The type of error that occurred when a transaction was being applied to a Model Repository database
- The type of error that caused problems with the Multimaster operation

Contact your HP Server Automation Support Representative for assistance troubleshooting and resolving SA problems that affect the multimaster operation.

Table 20 shows error messages.

Table 20. Multimaster Error Messages

Subject Line	Type of Error	Details
<code>vault.ApplyTransactionError</code>	Multimaster Transaction Conflict	The local database was not successfully updated with the changes from the other database. Each update must affect only one row and not result in any database errors.
<code>vault.configValueMissing</code>	SA Problem	No value was specified for a given configuration parameter. Log into the SA Client and provide the value for this configuration parameter. Contact your HP Server Automation Support Representative for assistance setting SA configuration values.
<code>vault.DatabaseError</code>	Multimaster Transaction Conflict	An error occurred while querying the database for updates to send to other databases or while applying updates from other data-

Subject Line	Type of Error	Details
		bases. Restart the Model Repository Multimaster Component.
<code>vault.InitializationError</code>	SA Problem	<p>An error occurred when the Model Repository Multimaster Component process started. The application returned the message specified. The thread that encountered the error stopped running. This error occurs when running SA in multimaster mode.</p> <p>Resolve the error condition. Restart the Model Repository Multimaster Component.</p>
<code>vault.ParserError</code>	Multimaster Transaction Conflict	<p>An error occurred when parsing the XML representation of the transaction. The application returned the message specified. This error occurs when running SA in multimaster mode.</p> <p>Run the SA Admin Multimaster Tools and verify that the transaction data does not contain special characters that the XML parser might be unable to interpret.</p>
<code>vault.SOAPError</code>	Multimaster Transaction Conflict	<p>An error occurred while using SOAP libraries to marshal or un-marshal transactions into XML. The application returned the message specified. This error occurs when running SA in multimaster mode.</p> <p>Run the SA Admin Multimaster Tools and verify</p>

Subject Line	Type of Error	Details
		that the transaction data does not contain special characters SOAP might be unable to interpret.
<code>vault.UnknownError</code>	SA Problem	The Model Repository Multimaster Component process encountered an unknown error. Contact technical support and provide the database name and SA component's log file.

Facility Administration

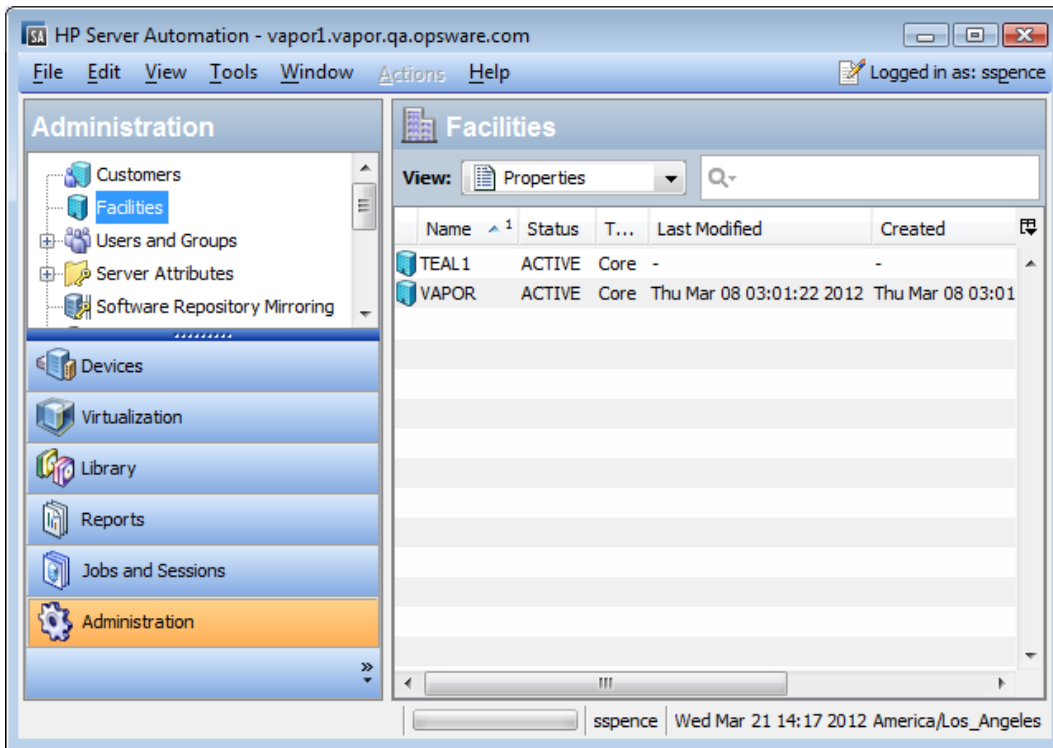
A *Facility* refers to the set of servers that a single SA core or satellite manages. You create a new facility whenever you install an SA core or an SA satellite. A Multimaster Mesh is a primary SA core, one or more secondary SA cores, and zero or more satellites. Whenever you install another SA core or another SA satellite, you create a new facility.

For more information about facilities, cores and satellites and how they fit into the Multimaster Mesh architecture, see the SA Overview and Architecture Guide and the SA Installation Guide.

Viewing Facility Information


You can view information about a facility by selecting the Administration tab in the SA Client, then selecting Facilities. **Figure 29** shows two facilities, Teal1 and Vapor, in the SA Client.

Figure 29. Two Facilities in the SA Client



You can view details about a Facility by opening the facility. **Figure 30** shows details of the Vapor facility, including the facility properties, custom attributes, and realms.

Figure 30. Details of the Facility



Changing the Customers Associated with a Facility

Customers are a way to organize your servers based on the users of your servers. Customers are simply groups of managed servers that provide access control boundaries. You can define as many customers as you need and assign any servers to each customer group. However, you must first associate a customer with one or more facilities before you can place servers from that facility into a customer group. Each server belongs to one and only one facility and each server belongs to one and only one customer (even if it is to the “Not Assigned” customer.)

For more information about customers, see the SA User Guide: Server Automation.

To change the customers associated with a facility, perform the following steps:

- In the SA Client, select the Administration tab.
- Select Facilities in the navigation pane. This displays all your facilities.
- Select the facility you want to change.
- Select the **Actions** menu, or right-click and select the **Open** menu. This displays the facility in a separate window.
- In the facility window, select the Properties view in the navigation pane. This displays information about the facility, including the customers that are associated with the facility.
- To add a new customer, select the “+” icon. This displays the list of existing customers.
- Select one or more customers.
- Click the Select button. This associates the selected customer with the facility.
- To remove a customer, select the customer and select the “-” icon. This removes the customer from the facility.
- Select **File > Revert** to discard your changes.
- Select **File > Save** to save your changes.
- Select **File > Close** to close the facility window.

Adding or Modifying Custom Attributes for a Facility - SA Client

You can create or modify custom attributes for a facility. Custom attributes provide a way for you to store additional information about your servers quickly and easily. Custom attributes are data elements you can create for facilities, servers, and other objects in SA. For more information about custom attributes, see the SA User Guide: Server Automation.

Caution: Be careful when you update or remove existing custom attribute settings, as it can affect or disrupt the operations that depend on custom attributes.

To add, modify, or delete a custom attribute for a facility, perform the following steps:

1. Log into the SA Client.
2. Select the Administration tab.
3. Select Facilities in the navigation pane. This displays all your facilities.

4. Select the facility you want to change.
5. Select the **Actions** menu or right-click and select the **Open** menu. This displays the facility in a separate window.
6. In the facility window, select the Custom Attributes view in the navigation pane. This displays all the custom attributes defined for the facility.
7. To add a new custom attribute, select the “+” icon or the **Actions > Add** menu. Enter the name of the new custom attribute and the value.
8. To modify a custom attribute, select the value field and enter the new value.
9. To delete a custom attribute, select the custom attribute and select the “-” icon or the **Actions > Delete** menu.
10. Select **File > Revert** to discard your changes.
11. Select **File > Save** to save your changes.
12. Select **File > Close** to close the facility window.

Modifying a Facility Name - SA Client

To modify a facility name, you must log into the SA Client with the Manage Facilities permission. The short name of the facility is the internal name that cannot be modified. The display name can be modified.

Perform the following steps to modify a facility’s display name:

1. Log into the SA Client.
2. Select the Administration tab.
3. Select Facilities in the navigation pane. This displays all your facilities.
4. Select the facility you want to change.
5. Select the **Actions** menu, or right-click and select the **Open** menu. This displays the facility in a separate window.
6. In the facility window, select the Properties view in the navigation pane.
7. Enter the new facility name in the Name field.
8. Select **File > Revert** to discard your changes.
9. Select **File > Save** to save your changes.

Satellite Administration

This section describes basic SA Satellite topologies and concepts and the following administrative tasks:

- [Starting/Restarting a Satellite](#)
- [Stopping a Satellite](#)
- [Verifying Satellite Communication with the Primary Core](#)
- [Permissions Required for Managing Satellites](#)
- [Viewing Satellite Information](#)
- [Satellite Monitoring](#)
- [Bandwidth Management of Remote Connections](#)
- [Satellite Software Repository Cache Management](#)
- [Updating Software in the Satellite Software Repository Cache](#)
- [Satellite Software Repository Cache Management](#)
- [SA Satellite Installation and Topologies](#)

Starting/Restarting a Satellite

To start a Satellite, issue the following command:

```
/etc/init.d/opsware-sas start opswgw
```

To restart a Satellite, issue the following command:

```
/etc/init.d/opsware-sas restart opswgw
```

Note: If the Satellite Agent fails to restart (typically due to an NFS error blocking the availability of port 1002, which is required for Satellite Agent communication), restart the Satellite host or temporarily disable the service that is blocking 1002, restart the agent, then restart the blocking service.

Stopping a Satellite

To stop a Satellite, issue the following command:

```
/etc/init.d/opsware-sas stop opswgw
```

Verifying Satellite Communication with the Primary Core

To verify that the Core Management Gateway is communicating with the Satellite, perform the following steps:

1. Log in to the SA Client as a member of a users group that has the Manage Gateway permission.
2. From the Navigation panel, click Administration > Gateway.
3. Verify that the upper left corner of the Manage Gateway page displays a link for the new Satellite.

If the Manage Gateway page does not display the link for the Satellite, you may need to edit the Satellite's properties. The full path name of the properties file follows:

```
/etc/opt/opsware/opswgw/opswgw.properties
```

After modifying the properties file, you must restart the Satellite:

```
/etc/init.d/opsware-sas restart opswgw
```

4. Log in to the SA Client as a member of a users group that has the Read (or Read & Write) permission on the Satellite's facility.
5. From the Navigation panel, click **Devices > All Managed Servers**.
6. Verify that the All Managed Servers page displays the host name of the Satellite server.

For further information, see also "More Troubleshooting Server Communication Tests" in the SA User Guide: Server Automation.

Permissions Required for Managing Satellites

To manage SA gateways, you must have the Manage Gateway permission. By default, this permission is included in the SA System Administrators group. To view facility information, you must have Read (or Read & Write) permission for the specific facility. For more information about user groups and SA permissions, see the [Permissions Reference](#).

Viewing Satellite Information

This section discusses the following topics:

- [Viewing Satellite Facilities and Realms](#)
- [Viewing the Realm of a Satellite Managed Server](#)
- [Viewing and Managing Satellite Gateway Information](#)

Viewing Satellite Facilities and Realms

You can view the core and satellite facilities by selecting the **Administration** tab in the SA Client, then selecting Facilities. Select a facility, then select the Realms view to see the realms associated with the facility, including the bandwidth between realms in the facility. For more information on facilities, see [Facility Administration](#).

Viewing the Realm of a Satellite Managed Server

When installed in a Satellite configuration, SA can manage servers with overlapping IP addresses. This situation can occur when servers are behind NAT devices or firewalls. Servers with overlapping IP addresses must reside in different Realms.

When retrieving a list of servers resulting from a search, you might see multiple servers with the same IP address but in different Realms. You might also see multiple servers with the same IP address when you are planning to run a custom extension and you are prompted to select the servers on which to run the extension.

The Properties view of a server in the SA Client displays additional information that identifies the server corresponding to the IP address.

Viewing and Managing Satellite Gateway Information

To view satellite gateway information, in the SA Client navigation panel, select the Administration tab, then select Gateway. This displays the gateway status, as shown in **Figure 31**. From the list of gateways on the left, select the gateway you want to view. Select the specific gateway information you want to see from the links across the top of the page.

Figure 31. Gateway Status

The screenshot shows the 'Manage Gateway' interface. At the top, it displays gateway details for 'cgw0-C28', including its realm, root, version, and uptime. Below this is a navigation menu with options like Status, Flows, Routing, PathDB, LSDB, Config, History, Isent, Bandwidth, Link Cost, Logging, Process Control, and Page Selection. A sidebar on the left lists gateway instances, with 'cgw0-C28' selected. The main area contains several data tables:

Gateway	Cost	BWLimit Kbits/sec	Send BW Kbits/sec	Recv BW Kbits/sec	Total In Bytes	Total Out Bytes	Payload In Bytes	Payload Out Bytes	Age	Peer
cgw0-C29	1	0	3.21	1.88	382157107	453088297	314805635	305777585	3:5:36.5.45	192.168.196.244-54307
Alice	10	0	1.58	1.23	39021515	56595009	30485960	43893609	3:6:8.8.40	192.168.9.50-41128
cgw0-C28	1	0	1.58	0.00	26460755	62224516	25523838	48562918	3:5:25.80	127.0.0.1-50991

Below the table is a 'Gateway Selection' section with a table of tunnel metrics:

Endpoint	Resolved	Connected	Cost	BWLimit
0:2:128	0:2:1024	0:0:2048	8:39:1024	
0:10:38:57.83	0:2:42.4.43	0:17:43:34.38	3:5:40:20.78	

Further down, there are sections for 'MsgProcessor' and 'DataMover Queue Table'.

Use the gateway status for the following tasks:

- Obtain status information about gateways and the tunnels between gateways. This can be useful for debugging gateways.
- Change the bandwidth limits or tunnel cost between gateway instances.
- Restart Gateway processes.
- Change the logging levels for gateway processes.

Viewing Gateway Diagnostic and Debugging Information

1. In the SA Client, select the Administration tab, then select Gateway.
2. From the list of gateways on the left, select the gateway for which you want to view information. This displays the following Status for the selected gateway:
 - A table of Active Tunnels, including:
 - Tunnel Cost
 - Bandwidth Constraints
 - Bandwidth Estimates
 - Age of the tunnels
 - Information about the internal message queues. Each column in the table for a queue displays data in this format:
 - Number of messages in the queue
 - The message high-water mark for the queue
 - Maximum value configured for the queue

- The last time the message high-water mark was attained for the queue. You can use the time stamp indicating when the message high-water mark was last reached to troubleshoot gateway issues. The time stamp is displayed in the format `DD:HH:mm:ss`.
3. To view the details and statistics for a tunnel between gateways, select the link for the gateway that *terminates* the tunnel, as **Figure 32** shows. This displays the tunnel details and statistics.

Figure 32. Manage Gateway — Status Page

Gateway	Cost	BWLimit Kbits/sec	Send BW Kbits/sec	Recv BW Kbits/sec	Total In Bytes	T
gw1-nat2	1	0	0.00	0.83	686578431	24

4. To view the following pages containing diagnostic information, select one of the following links across the top of the page:
- **Flows** displays information about all open connections for the selected gateway.
 - **Routing** displays the inter-gateway routing table. This table shows which tunnel will be used to reach another gateway in the mesh. The routing table is computed from the data in the path database. The routing computation automatically updates when the link cost for a connection is changed.

Note: When a tunnel collapses, by default, routing information is retained in the routing table for two minutes to provide continuity for the mesh.

- **PathDB - Path Database** displays the route with the lowest cost to all reachable gateways in the mesh. SA determines the lowest cost route to all reachable gateways from the data in the Link State database.
- **LSDB - Link State Database** contains information about the state of all tunnels from the perspective of each gateway instance. The LSDB contains the data for all tunnels and the bandwidth constraint for each tunnel.
- **Config** displays the properties file for the selected gateway, including the path to the properties file on the server running the gateway component. Below the properties values, the page contains crypto file information and the mesh properties database. The **Properties Cache** field is above the

properties values. When you change the bandwidth or link cost for a connection between gateways, the updated value appears in this field if the update was successful.

- **History** displays historical information about the inbound (ingress) and outbound (egress) connections between hosts using the gateway mesh. For example, when host A in Realm A connected to host B in Realm B.

Identifying the Source IP Address and Realm for a Connection

The **Ident** link provides an interface to the real-time connection identification database. If necessary, contact HP Support for additional information about how to run this tool.

1. In the SA Client, select the **Administration** tab, then select **Gateway**.
2. Select the link **Ident**. This displays the real-time connection identification database.
3. In the edit box, enter the protocol and source port for an active connection, separated by a colon; for example, TCP:25679.
4. Select the **Lookup** button. This displays the client Realm and client IP address, which is where the connection came from.

Changing the Bandwidth Usage or Link Cost Between Gateways

The **Edit** link lets you modify the link bandwidth constraint, the link cost, and the load balance rules.

Note: You must apply any bandwidth changes between gateways on core gateways only. Changes made on other gateways will not take effect.

1. In the SA Client, select the **Administration** tab, then select **Gateway**.
2. To specify a bandwidth limit for a connection:
 1. Select the **Edit** link at the top of the page. This displays the Modify Link Bandwidth Constraint control.
 2. Specify two gateway instance names that are connected by a tunnel.
 3. Specify the bandwidth limit you want in kilobits per second (Kbps). Specify zero (0) to remove bandwidth constraints for the connection.
 4. Click **Apply**.
3. To set a link cost for a connection:
 1. Select the **Edit** link at the top of the page. This displays the Modify Link Cost control.
 2. Specify two gateway instance names that are connected by a tunnel.
 3. Specify the cost you want in the **Cost** field.
 4. Click **Apply**.

4. To set the load balance rules for a connection:
 1. Select the **Edit** link at the top of the page. This displays the Modify Load Balance Rules control.
 2. Specify a gateway instance name.
 3. Specify a load balance rule.
 4. Click **Apply**.

Viewing the Gateway Log or Change the Log Level

Note: Changing the logging level to `LOG_DEBUG` or `LOG_TRACE` greatly increases the log output of the gateway and can negatively impact the performance of the gateway.

1. In the SA Client, select the **Administration** tab, then select **Gateway**.
2. Select the **Logging** link at the top of the page. This displays the end of the gateway log file.
3. To change the logging level, select one of `LOG_INFO`, `LOG_DEBUG`, or `LOG_TRACE`.
4. Select **Submit**.

Restarting or Stopping a Gateway Process

1. In the SA Client, select the **Administration** tab, then select **Gateway**.
2. Select the **Process Control** link at the top of the page.
3. To restart the gateway process, click **Restart**.
4. To stop the gateway watchdog and the gateway, click **Shutdown**.

Caution: Stopping a gateway process can cause problems for an SA core. For example, if you stop a core gateway process, you will stop all multimaster traffic to that SA core, and you will be unable to control the gateway from the SA Client.

Requirement: To restart the gateway after stopping it from the SA Client, you must log onto the server running the gateway component and manually restart the process.

Satellite Monitoring

See the following sections in [Overview of SA Monitoring](#):

- [Agent Cache Monitoring](#)
- [Gateway Monitoring](#)

Bandwidth Management of Remote Connections

Bandwidth Management is a measure employed in communication networks to regulate network traffic and minimize network congestion. SA's remote site management model typically uses a Satellite configuration that deploys a remote gateway on every logical location (for example, a branch office) to handle connections to remote servers and manage the network bandwidth of these connections. However, the cost effectiveness of this configuration is significantly reduced for sites that manage only a few servers.

A new SA bandwidth management capability eliminates the need to install a Satellite for remote locations with only a few servers. SA provides the Bandwidth Configuration Management (BCM) tool to control the bandwidth used by Agent or Satellite Gateways when communicating with remote servers.

You can push bandwidth configurations to a peer group by using the BCM tool. After the configuration is pushed to the peers, it is saved to file. During Gateway startup, the configuration is loaded from this file and synchronized with the peers. When a client negotiates a connection through the SA Gateway mesh to connect to a remote TCP service, the client then has a TCP connection to the ingress Gateway. Also, there is a TCP connection leaving the egress Gateway to the remote service.

When the proxy connection through the Gateway mesh is established, the peer addresses of ingress/egress connections are classified, and a runtime queue is created for each classification. At this point, bandwidth throttling is in effect for these connections. The corresponding queue is updated with the bandwidth usage information as the data flows through the connection. The bandwidth usage information is also shared among the peer group so that the local queue can be updated on each gateway cluster. The data can flow through that connection till the maximum bandwidth allowed is reached. Queue bandwidth usage information is reset at a one-second interval.

Note: All Agent Gateways in the same Realm must also be running the same SA version in order to participate in Agent Gateway bandwidth negotiation and communication. Mixed core configurations (core and satellites running a different SA version) is not supported.

The SA Bandwidth Configuration Management Tool

Note: SA BCM is not supported SA Cores/Satellites running Solaris or Red Hat Enterprise Linux 3 x86.

Note: The BCM tool requires that your firewall allows SA network traffic on ports 3001 and 8086. If you plan to use the BCM tool administrative interface, port 8089 must also be open.

This section describes using the BCM tool to create bandwidth management configurations. These configuration can then be automatically synchronized across peer gateways.

Only administrative users who have root access to the gateway host can perform Gateway configuration push operation with the BCM tool.

Note: Although the BCM tool is installed with a default configuration file:

```
/etc/opt/opsware/gateway_name/BWT.conf
```

you should not modify that file directly. Make a copy of the file and edit it to suit your configuration. You can then push the modified configuration file to all the gateway(s) in the realm using the `gwctl -f` command. See [Invoking the Bandwidth Management Configuration Tool](#).

Specified bandwidth configurations are saved to a configuration file. The following is an example of a typical Gateway configuration file:

```
enabled

# Branch offices have only 3M bytes per sec connections, SA
should never use
# more than 512K bytes per sec.
queue branch_office bandwidth 512KB

# Branch offices A and B (non standard addresses)
class 192.168.1.[1-5,10-15,20,30] for branch_office

# Other branch offices
class 192.168.2.0/24 for branch_office
```

Invoking the Bandwidth Management Configuration Tool

You invoke the BCM tool as a command line tool.

On the Satellite whose SA Agent configuration you want to manage, use the following commands:

```
gwctl: [OPTIONS] ...
```

Table 21. Bandwidth Configuration Management Tool Options

Option	Description
-?, --help	Display usage.
-p, --port	When specified with -l. lists the agent gateway proxy port (default 3001). When specified with other options (such -d, -e, -f, -v, -c, -s, etc.), displays the bandwidth throttle configuration port (default 8086).
-l, --list_gws	List all the gateways in this realm.
-f, --conf	Configuration file.
-v, --verify_conf	Verify configuration file and exit; Do not push it to the gateways. Note: This option is used only with the -f <conf_path> option.
-c, --cksum	Display the checksum of the configuration file. Note: This option is used only with the -f <conf_path> option.
-e, --enable_bwt	Enable bandwidth throttling for this realm.
-d, --disable_bwt	Disable bandwidth throttling for this realm.
-r, --request_conf	Request the configuration from the given gateway.
-s, --signature	Request the configuration signatures from the given gateway.
-z, --verbose	Display all messages.

The following are example commands.

To list the gateways in the realm:

```
gwctl -l
```

To specify a different agent gateway port:

```
gwctl --port 2003 -l
```

To verify the configuration file only:

```
gwctl -f myconf.conf -v
```

To push the configuration file to all Agent Gateways in the realm (including localhost):

```
gwctl -f mytconf.conf
```

Enabling/Disabling Remote Connection Bandwidth Management

You must enable or disable remote connection bandwidth management in one of two ways:

- By pushing a bandwidth configuration file containing the `enabled` or `disabled` keyword as the first entry in the file. Each configuration file must contain `enabled` or `disabled` as first line in the file, indicating the status of bandwidth throttling.
- From the command line using `gwctl -e` to enable bandwidth management or `gwctl -d` to disable bandwidth management. The bandwidth management state of `enabled` or `disabled` persists in the bandwidth management configuration file with no version upgrade.

Bandwidth Configuration Grammar

The Context Free Grammar (CFG) of Bandwidth Configuration in EBNF format:

```
config : ((queue | class | version | config_source | config_user | disabled | comment)? '\n')\*

queue : 'queue' queue_name 'bandwidth' d_number bandwidth_spec ('rtt' d_number)? ('parent' queue_name 'borrow')?

queue_name : "[a-zA-Z0-9_]+"

class : 'class' pattern (',' pattern)* 'for' queue_name

pattern : ipv4 | ipv4_cidr

ipv4 : ipv4_address_pattern_element ( '.' ipv4_address_pattern_element )@1:3

ipv4_cidr : d_number ( '.' d_number )@1:3 '/' d_number

ipv4_address_pattern_element : single_number | range | range_class
range_class : '[' (number ('-' number)? ',')+ ']'
```

```
wildcard : '*'

range : '[' number '-' number ']'

single_number : d_number

number : d_number

d_number : "[0-9]+"

x_number : "[a-zA-F0-9]+"

bandwidth_spec : "[GMK]?[bB]"

config_source : 'config-source' ':' "[a-zA-Z0-9.\-]+"

config_user : 'config-user' ':' "[a-zA-Z0-9_!@#$$%^&*
();.`~\-\-]+"

disabled : 'disabled'

comment : '#' "[^\n]*"
```

Satellite Software Repository Cache Management

The largest amount of network traffic in an SA Core occurs between:

- The Software Repository and the Server Agent on a Managed Server during application software or OS patch installations.
- A server being OS Provisioned and the OS Provisioning Media Server that provides the OS media for the provisioning.

When a Satellite is connected by a low-bandwidth network link, performance will be poor during these processes. You can minimize network traffic by creating a copy of the core's Software

Repository contents in the Satellite's Software Repository Cache or installing a local Satellite OS Provisioning Media Server/Boot Server.

Because the Software Repository Cache stores copies of the files in the SA Core's Software Repository (or from another Satellite's Software Repository Cache), SA can supply software requests locally without having the requests pass across the network between the Satellite and the SA Core. Similarly, the OS Provisioning Media Server can supply OS images locally. SA Satellites also support multiple Software Repository Caches per Realm.

The following sections discuss configuring and updating your local Software Repository Cache and, optionally, your OS Provisioning Media and Boot servers.

Availability of Satellite Software Repository Cache Content

The Satellite Software Repository Cache is updated in one of two modes. By default the updates are on demand, and they occur when the agent of a server managed behind this satellite needs to download a package, or manual by the SA Administrator.

When SA is attempting to remediate requested software that is not available locally onto a managed server, the SA Client generates an error and displays a complete list of missing packages to help you identify the packages that need to be copied to the cache. After you have copied the software to the cache, it will continue to be available locally for future installations.

Note: The SA Client does not provide a User Interface to *push* packages to Satellites. However, you can push packages to a Satellite by using the command-line tool `stage_pkg_in_realm`.

This tool is found on the First Core's Model Repository host in:

```
/opt/opsware/mm_wordbot/util/stage_pkg_in_realm.
```

If you use the `checkonly=1` argument in the URL request for the file, the utility requests a file, but the Software Repository will not send the file. If the file is not already cached, the Software Repository Cache will obtain it from the parent Software Repository Cache if the caching policy allows it.

Updating Software in the Satellite Software Repository Cache

To update files in a Satellite's Software Repository Cache, you can configure the cache to update cached copies of files as requests are received (*On-demand Updates*) or to update the cached copy of a file manually (*Manual Updates*):

- **On-demand Update:** The local Software Repository Cache obtains current files as needed from the Software Repository in the SA core.
- **Manual Update:** SA stages the software packages to a Satellite's Software Repository Cache in advance of package installation, so that performance is about the same as if the Managed Server is in the same data center as the core.

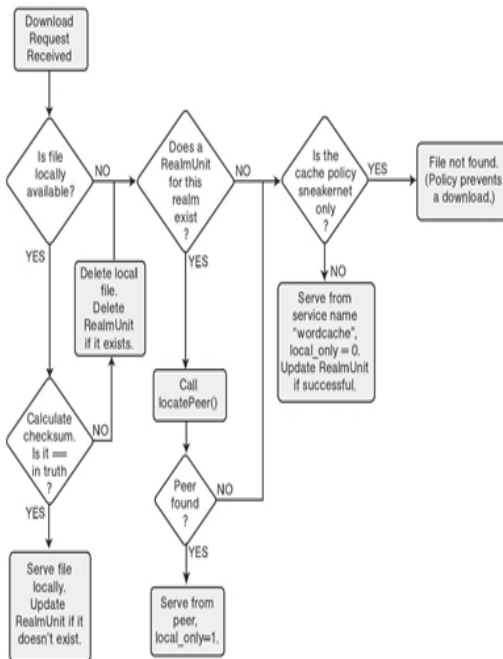
When On-demand update is enabled, if the requested software is already present in the local Software Repository Cache and is current, no action is taken. If the software is not present locally or it is not current, the Software Repository Cache attempts to download the file in the background from the closest upstream Software Repository Cache or from the Core's Software Repository.

If the caching policy is Manual Update and you request an on-demand software update, the Software Repository Cache will raise a `wordbot.unableToCacheFile` exception.

It is always possible to stage a file on a Software Repository Cache regardless of the caching policy. See [Staging Files to a Software Repository Cache](#).

Figure 33 illustrates the logic that the Software Repository Cache uses to update packages in a Satellite.


Figure 33. Software Repository Cache Update Logic



Setting the Software Repository Cache Update Policy

You can specify the Software Repository Cache update policy for each facility by performing the following tasks:

1. Select the **Administration** tab in the SA Client.
2. In the navigation pane, select **System Configuration > Configuration Parameters**. This displays the SA components, facilities, and realms that have system configuration parameters.
3. Select the realm for which you want to set the software repository cache update policy. This displays all the system configurations for that realm.
4. Locate the configuration parameter `word.caching_policy`.

5. Set the value of this parameter to one of the following:
 - Select **Default value: JIT**. This specifies JIT or on-demand update.
 - Select the new value button  and enter the text “SNEAKERNET in the edit field. This specifies manual update.
6. Select the Revert button to discard your changes or the Save button to save your changes.

On-Demand Updates

Enabling on-demand updates allows software to be downloaded to the Satellite Software Repository Cache as soon as that software is requested and when it is not yet locally available. If you have a low-bandwidth network connection, manual updates may be a better solution, as it allows you to pre-download the most commonly requested software into the Software Repository Cache. See [Manual Updates](#).

Each time a Server Agent on a managed server in a Satellite requests software, the local Software Repository Cache checks whether its cached copy of the software is current. If the cached file is not current or is missing, the Software Repository Cache obtains an updated or new local copy of the file from the nearest upstream Software Repository Cache or from the Core’s Software Repository and sends it to the requesting Server Agent.

When configured for on-demand updates, when the Software Repository Cache receives a request for software, it first requests the checksum of the software against the checksum of the Core’s Software Repository to insure that it has the latest copy.

Note: For security purposes, SA caches software checksums for a user-configurable period of time.

If the checksum is the same as the locally stored file, the Software Repository Cache serves the software to the requester. If the checksum does not match or the local file is not present, the Software Repository Cache requests an updated copy of the software from the nearest upstream Software Repository Cache or the Core’s Software Repository.

If network connectivity is lost while the Software Repository Cache is downloading software, the next time a Server Agent requests the same software, the Software Repository Cache will resume the file download from the point at which it stopped.

Manual Updates

For Satellites with low-bandwidth network links, Manual Software Repository Cache updates allow you to *pre-populate* the Software Repository Cache at installation time. You can also configure refreshes for an existing cache. The Software Repository Cache is populated by an out-of-band method, such as by cutting CDs of the required packages and shipping them to the Satellite. To perform manual updates, use the SA DCML Exchange Tool (DET) to copy existing packages from an SA core or use the Staging Utility to perform the update. See [Creating Software Repository Cache Manual Updates](#) and [Staging Files to a Software Repository Cache](#).

When configured for manual updates, a Software Repository Cache does not communicate with upstream Software Repository Caches or the Core's Software Repository until you initiate an update. The Satellite considers its own Software Repository Cache as authoritative.

If the caching policy is manual update and you request an on-demand software update, the Software Repository Cache will raise a `wordbot.unableToCacheFile` exception.

Even if you have configured a Software Repository as on-demand update, you can apply a manual update regardless of its update policy.

Note: When applying manual updates in a Satellite installation with multiple Software Repository Caches, you must apply the update to each Software Repository Cache in the Satellite. Otherwise, when performing operations that retrieve files from the Cache (for example, when installing software on a server in the affected Satellite), you may get the `wordbot.unableToCache file error`.

Emergency Software Repository Cache Updates

You can push Emergency updates manually over the network to Satellites even if the caching policy is manual update. You do not need to reconfigure the Software Repository Cache's caching policy to push emergency updates to a Software Repository Cache. For example, an emergency patch can be staged to a Satellite and applied without waiting for a shipment of CDs to arrive.

Software Repository Cache Size Management

When you apply a manual update to a Software Repository Cache, SA removes files that have not been recently accessed when the cache size limit is exceeded.

The least-recently accessed packages are deleted first.

The Software Repository Cache removes the files the next time it cleans up its cache. By default, the cache is cleaned up every 12 hours. Packages are deleted so that the available disk space stays below the high-water mark.

Requirement: You must have enough disk space to store all necessary packages for the Software Repository Cache to ensure that the Software Repository Cache does not exceed the cache size limit.

Creating Software Repository Cache Manual Updates

To create a manual update, you can use the SA DCML Exchange Tool (DET) to copy existing software from an SA core. You then save an export file you can copy over the network to the Satellite's Software Repository Cache or burn to CD or DVD to be applied later to the cache. You can also use the Staging Utility to upload software. See [Staging Files to a Software Repository Cache](#).

This section discusses the following topics:

- [Creating a Manual Update Using the DCML Exchange Tool \(DET\)](#)
- [Applying a Manual Update to a Software Repository Cache](#)
- [Staging Files to a Software Repository Cache](#)
- [Microsoft Utility Uploads and Manual Updates](#)

Creating a Manual Update Using the DCML Exchange Tool (DET)

You perform this procedure by using the DET. Using the DET, export the software for the Manual Update and export the packages associated with selected software policies.

See the SA Content Utilities Guide for more information about using DET.

To create a manual update, perform the following steps:

1. On the server where you installed the DET component, run the following command to create the following directory:

```
# mkdir /var/tmp/sneakernet
```

2. From the server running the SA Client, copy the following files from the `/var/opt/opsware/crypto/occ` directory:

```
opsware-ca.crt
```

```
spog.pkcs.8
```

to the following directory:

```
/usr/cbt/crypto
```

This is the directory where you installed DET.

3. Create the file, `/usr/cbt/conf/cbt.conf`, so that it contains this content:

```
twist.host=<twist's hostname>
twist.port=1032
twist.protocol=t3s
twist.username=buildmgr
twist.password=buildmgr
twist.certPaths=/usr/cbt/crypto/opsware-ca.crt
spike.username=<your username>
spike.password=<your password>
spike.host=<way's hostname>
way.host=<way's hostname>
spin.host=<spin's hostname>
word.host=<word's hostname>
```

```
ssl.keyPairs=/usr/cbt/crypto/spog.pkcs8  
ssl.trustCerts=/usr/cbt/crypto/opsware-ca.crt
```

4. Create the following DCML Exchange Tool filter file `/usr/cbt/filters/myfilter.rdf` that contains this content:

```
<?xml version="1.0" encoding="UTF-8"?>  
<!DOCTYPE rdf:RDF [  
<!ENTITY filter "http://www.opsware.com/ns/cbt/0.1/filter#">  
>  
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-  
ns#">  
  xmlns="http://www.opsware.com/ns/cbt/0.1/filter#">  
    <ApplicationFilter rdf:ID="a1">  
      <path>/Other Applications</path>  
      <directive rdf:resource="&filter;Descendants" />  
    </ApplicationFilter>  
  </rdf:RDF>
```

In the `<path>` directive of the filter file, replace `/Other Applications` with the path to the node you want to export (all node information about that node, its descendants, and all associated packages will be exported).

This filter will export from the Applications area of the SA Client. If you want to export packages from some other category of software in the SA Client, you need to create a different filter. See the SA Content Utilities Guide for information.

5. On the server where you installed the DET component, run the DCML Exchange Tool by entering the following command:

```
# /usr/cbt/bin/cbt -e /var/tmp/myexport --config /usr/cbt/conf/cbt.conf --filter /usr/cbt/filters/myfilter.rdf
```

The DCML Exchange Tool places the packages associated with the exported nodes in the following directory:

```
/var/tmp/myexport/blob
```

The packages are named `unitid_nnnnnnn.pkg`.

6. Copy all of the `.pkg` files to a directory on the server running the Software Repository Cache, either over the network or by burning the files to a set of CDs or DVDs.

Applying a Manual Update to a Software Repository Cache

To apply a manual update to a Software Repository Cache, run a utility (`import_sneakernet`), which moves or copies the software you want to update into the right location on the Software Repository Cache and registers it with the Model Repository in the SA core.

To apply a manual update to a Software Repository Cache, perform the following steps:

1. Log in as `root` on the server running the Satellite's Software Repository Cache.
2. Copy the export file to a directory on the Software Repository Cache server, mount the CD containing the software export file, or copy the CD contents to a temporary directory.
3. Enter the following command to change directories:

```
# cd /opt/opsware/mm_wordbot/util
```

4. Enter the following command to import the contents of the export file to the Software Repository Cache:

```
# ./import_sneakernet -d dir
```

where *dir* is the CD mount point or the temporary directory containing the export file.

Staging Files to a Software Repository Cache

A Server Agent on a Managed Server can override the caching policy in effect for a Realm. The ability to override the caching policy of a Software Repository Cache allows you to stage software to a cache that is configured to be manual update to resolve the following situations:

- You must circulate an emergency patch, and you do not have time to create a manual update export file and physically visit a Facility to upload the software.
- A necessary patch must be installed during a specified maintenance period, and the period is not long enough to download a patch and install it on all managed servers.
- The utilization of a network link to the Satellite is known to be low at a particular time of day, making that time advantageous for upload.

To force package staging, the Staging Utility provides the argument `override_caching_policy=1`, which is specified in the URL request for the software.

The Software Repository Cache allows a client to request that it obtain a file but that it not actually send the file to the client. If the file is not already cached, the Software Repository Cache will obtain it from the parent Software Repository Cache if the caching policy allows it. To use this feature, the client includes the argument `checkonly=1` in the URL request for the file.

Running the Staging Utility

To run the staging utility, perform the following steps:

1. On the server running the Software Repository component (part of the Slice Component bundle), verify that the certificate `token.srv` is in your `CRYPTO_PATH`. During installation `token.srv` is copied to:

```
/var/opt/opsware/crypto/gateway/token.srv.
```

2. Log into the server running the Core's Software Repository.
3. Enter the following command to change directories:

```
# cd /opt/opsware/mm_wordbot/util
```

4. To stage the files you want, run the utility `stage_pkg_in_realm`, which has the following syntax:

```
./stage_pkg_in_realm [-h | --help] [-d | --debug]  
[--user <USER>] --pkgid <ID> --realm <REALM> [--gw <IP:PORT>] [-  
-spinurl <URL>] [--wayurl <URL>] [--word <IP:PORT>]
```

To force package staging, the Staging Utility provides the argument `override_caching_policy=1`, which is specified in the URL request for the software. For example:

```
./stage_pkg_in_realm --user admin --pkgid 80002 --realm  
luna  
--gw 192.168.164.131:3001  
Password for admin: <password>  
Package /packages/opsware/Linux/3ES/miniagent is now being  
staged in realm luna
```

Microsoft Utility Uploads and Manual Updates

When you upload new Microsoft patching utilities (described in the SA Installation Guide System Requirements chapter), you should immediately stage those files to all Realms where the Software Repository Cache is configured for manual updates only.

If you do not stage these files to the remote Realms, Server Agents running on Windows servers in those Realms will be unable to download new versions of the utilities and will be unable to register their software packages. It is not necessary to stage packages to Realms where the Software Repository Cache is configured for on-demand updates.

The Software Repository Cache allows a client to request that it obtain a file but that it not actually send the file to the client. If the file is not already cached, the Software Repository Cache will obtain it from the parent Software Repository Cache if the caching policy allows it. To use this feature, the client includes the argument `checkonly=1` in the URL request for the file. See [Running the Staging Utility](#) for information about how to stage files.

SA Satellite Installation and Topologies

A Satellite installation can be a solution for remote sites that do not have a large enough number of potentially Managed Servers to justify a full SA Core installation. A Satellite installation allows

you to install only the minimum necessary Core Components on the Satellite host, which then accesses the Primary (First) Core's database and other services through an SA gateway connection.

A Satellite installation can also relieve bandwidth problems for remote sites that may be connected to a primary Facility through a limited network connection. You can cap a Satellite's use of network bandwidth to a specified bit rate limit. This allows you to insure that Satellite network traffic will not interfere with your other critical systems' network bandwidth requirements on the same pipe.

A Satellite installation typically consists of a *Satellite Gateway* and a *Software Repository Cache* and allows you to fully manage servers at a remote Facility. The Software Repository Cache contains local copies of software packages to be installed on Managed Servers from the Satellite while the Satellite Gateway handles communication with the Primary (First) Core. You can optionally install the *OS Provisioning Boot Server* and *Media Server* on the Satellite host to support Satellite OS Provisioning.

Note: Installing other SA core components on the Satellite host is not supported.

For information about how to install and configure a Satellite, see the SA Installation Guide.

Satellites can be installed using various topologies. For detailed information about Satellite topologies, see the SA Overview and Architecture Guide.

Note: Some advanced topologies require the service of HP Professional Services for installation and upgrade. If the specific installation steps for a topology are not documented, contact HP Technical Support or Professional Services for assistance.

SA Remote Communications Administration

This section describes methods you can use to control SA Gateway Bandwidth use (bandwidth management) and configure software caching for small remote sites with fewer than 50 managed servers without being required to install a full SA Satellite (Managed Server Peer Content Caching):

- [Bandwidth Management of Remote Connections](#)
- [IPv6 in SA](#)
- [SA Managed Server Peer Content Caching](#)
- [Concepts: SA Core Communications Infrastructure](#)

Note: For more information about SA Satellites, Gateways, and Agents, see the SA Overview and Architecture Guide.

Bandwidth Management of Remote Connections

Bandwidth management is a measure employed in communication networks to regulate network traffic and minimize network congestion. SA's remote site management model typically uses a Satellite configuration that deploys a remote gateway on every logical location (for example, a branch office) to handle connections to remote servers and to manage the network bandwidth of these connections. However, the cost effectiveness of this configuration is significantly reduced for sites that manage only a few servers.

A new SA bandwidth management capability eliminates the need to install a Satellite for remote locations with only a few servers. SA provides the BCM tool to control the bandwidth used by Agent or Satellite Gateways when communicating with remote servers.

You can push bandwidth configurations to a peer group by using the BCM tool. After the configuration is pushed to the peers, it is saved to file. During Gateway startup, the configuration is loaded from this file and synchronized with the peers. When a client negotiates a connection through the SA Gateway mesh to connect to a remote TCP service, the client then has a TCP connection to the ingress Gateway. Also, there is a TCP connection leaving the egress Gateway to the remote service.

When the proxy connection through the Gateway mesh is established, the peer addresses of ingress/egress connections are classified, and a runtime queue is created for each classification. At this point, bandwidth throttling is in effect for these connections. The corresponding queue is

updated with the bandwidth usage information as the data flows through the connection. The bandwidth usage information is also shared among the peer group so that the local queue can be updated on each gateway cluster. The data can flow through that connection until the maximum bandwidth allowed is reached. Queue bandwidth usage information is reset at a one-second interval.

Note: All Agent Gateways in the same Realm must also be running the same SA version in order to participate in Agent Gateway bandwidth negotiation and communication. Mixed core configurations (Core and satellites running a different SA version) is not supported.

The SA Bandwidth Configuration Management Tool

Note: SA BCM is not supported SA Cores/Satellites running Solaris or Red Hat Enterprise Linux 3 x86.

Note: The BCM tool requires that your firewall allows SA network traffic on ports 3001 and 8086. If you plan to use the BCM tool administrative interface, port 8089 must also be open.

This section describes using the BCM tool to create bandwidth management configurations. These configuration can then be automatically synchronized across peer gateways.

Only administrative users who have root access to the gateway host can perform Gateway configuration push operation with the BCM tool.

Note: Although the BCM tool is installed with a default configuration file:

```
/etc/opt/opsware/gateway_name/BWT.conf
```

you should not modify that file directly. Make a copy of the file and edit it to suit your configuration. You can then push the modified configuration file to all the gateway(s) in the realm using the `gwctl -f` command. See [Invoking the Bandwidth Management Configuration Tool](#).

Specified bandwidth configurations are saved to a configuration file. The following is an example of a typical Gateway configuration file:

```
enabled  
  
# Branch offices have only 3M bytes per sec connections, SA  
# should never use  
# more than 512K bytes per sec.  
  
queue branch_office bandwidth 512KB
```



```
# Branch offices A and B (non standard addresses)
class 192.168.1.[1-5,10-15,20,30] for branch_office

# Other branch offices
class 192.168.2.0/24 for branch_office
```

Invoking the Bandwidth Management Configuration Tool

You invoke the BCM tool as a command line tool.

On the Satellite whose SA Agent configuration you want to manage, use the following commands:

```
gwctl: [OPTIONS] ...
```

Table 21. Bandwidth Configuration Management Tool Options

Option	Description
-?, --help	Display usage.
-p, --port	When specified with -l. lists the agent gateway proxy port (default 3001). When specified with other options (such -d, -e, -f, -v, -c, -s, etc.), displays the bandwidth throttle configuration port (default 8086).
-l, --list_gws	List all the gateways in this realm.
-f, --conf	Configuration file.
-v, --verify_conf	Verify configuration file and exit; Do not push it to the gateways. Note: This option is used only with the -f <conf_path> option.
-c, --cksum	Display the checksum of the configuration file. Note: This option is used only with the -f <conf_path> option.
-e, --enable_bwt	Enable bandwidth throttling for this realm.
-d, --disable_bwt	Disable bandwidth throttling for this realm.
-r, --request_conf	Request the configuration from the given gateway.
-s, --signature	Request the configuration signatures from the given gateway.
-z, --verbose	Display all messages.

The following are example commands.

To list the gateways in the realm:

```
gwctl -l
```

To specify a different agent gateway port:

```
gwctl --port 2003 -l
```

To verify the configuration file only:

```
gwctl -f myconf.conf -v
```

To push the configuration file to all Agent Gateways in the realm (including localhost):

```
gwctl -f mytconf.conf
```

Enabling/Disabling Remote Connection Bandwidth Management

You must enable or disable remote connection bandwidth management in one of two ways:

- By pushing a bandwidth configuration file containing the `enabled` or `disabled` keyword as the first entry in the file. Each configuration file must contain `enabled` or `disabled` as first line in the file, indicating the status of bandwidth throttling.
- From the command line using `gwctl -e` to enable bandwidth management or `gwctl -d` to disable bandwidth management. The bandwidth management state of enabled or disabled persists in the bandwidth management configuration file with no version upgrade.

Bandwidth Configuration Grammar

The Context Free Grammar (CFG) of Bandwidth Configuration in EBNF format:

```
config : ((queue | class | version | config_source | config_user | disabled | comment)? '\n')\*
```

```
queue : 'queue' queue_name 'bandwidth' d_number bandwidth_spec ('rtt' d_number)? ('parent' queue_name 'borrow')?
```

```
queue_name : "[a-zA-Z0-9_]+"
```

```
class : 'class' pattern (',' pattern)* 'for' queue_name
```

```
pattern : ipv4 | ipv4_cidr
```

```
ipv4 : ipv4_address_pattern_element ('.' ipv4_address_pat-  
tern_element)@1:3
```

```
ipv4_cidr : d_number ('.' d_number)@1:3 '/' d_number
```

```
ipv4_address_pattern_element : single_number | range |  
range_class | wildcard range_class : '[' (number ('-' num-  
ber)? ',')+ ']'
```

```
wildcard : '*'
```

```
range : '[' number '-' number ']'
```

```
single_number : d_number
```

```
number : d_number
```

```
d_number : "[0-9]+"
```

```
x_number : "[a-zA-F0-9]+"
```

```
bandwidth_spec : "[GMK]?[bB]"
```

```
config_source : 'config-source' ':' "[a-zA-Z0-9.\-]+"
```

```
config_user : 'config-user' ':' "[a-zA-Z0-9_!@#$$%^&*  
();.\-~\-\-]+"
```

```
disabled : 'disabled'
```

```
comment : '#' "[^\\n]*"
```

IPv6 in SA

Internet Protocol version 6 (IPv6) is a Layer 3 network protocol in the TCP/IP stack of protocols. IPv6 expands the number of network address bits from 32 bits (in IPv4) to 128 bits. The Internet Engineering Task Force (IETF) designed the IPv6 addressing scheme to provide interoperability with existing IPv4 network architecture and to allow the coexistence of IPv6 networks with existing IPv4 networks (see RFC 4291).

IPv6 solves the IP address shortage problem in IPv4, and it enhances and improves some of the salient features of IPv4. IPv6:

- Enhances routing and addressing capabilities
- Simplifies the IP header
- Supports various types of IP addresses and larger address blocks for use with multicast routing

IPv4/IPv6 Dual-Stack Implementation

The dual-stack protocol implementation in an operating system is a fundamental IPv4-to-IPv6 transition technology. It implements IPv4 and IPv6 protocol stacks independently or in a hybrid form.

Hybrid dual-stack IPv6/IPv4 implementations support a special class of addresses, the IPv4-mapped IPv6 addresses. This address type has its first 80 bits set to 0, the next 16 bits set to 1, and the last 32 bits set to an IPv4 address. These addresses are commonly represented in the standard IPv6 format but have the last 32 bits written in customary IPv4 dot-decimal notation of IPv4; for example: `ffff:192.0.2.128` represents the IPv4 address 192.0.2.128.

SA uses the dual-stack concept for SA core and satellite. Both SA core and satellite require an IPv4 address as well as an IPv6 address; these addresses can be on a single Network Interface Card (NIC) or on two NICs. The reason is that only the SA gateway components are IPv6-enabled, and all other SA core and satellite components are IPv4 only (except for those components accessed directly by agents, such as OGFS, NFS, and Samba).

IPv6 Support in HP SA

When the SA core or satellite is IPv6-enabled, the managed servers will be able to register to the core with their IPv6 addresses and communicate to the core or satellite using the IPv6 protocol. SA core components that directly or indirectly communicate with the managed servers will be able to recognize the IPv6 addresses of managed servers and facilitate IPv6 communication from core to managed servers and vice versa.

Intra-core communication occurs through IPv4.

The agent and satellite gateway uses its IPv6 address to advertise its IPv6 capability but will use its IPv4 address for most of its communication.

SA Agent Installation

The SA Agent is supported in an IPv4, IPv6, and dual-stack network. The management IP for agents on managed servers is chosen at agent installation time by specifying the SA Gateway address as either (a list of) IPv4 or IPv6 address(es). If both IPv4 and IPv6 addresses are passed, they are tried in the same sequence in which they passed, and the first successful connection determines the management IP.

OS Provisioning

OS Provisioning supports IPv4, IPv6, and dual-stack networks.

In IPv6 networks, routing information is configurable only through router advertisement (RA) messages. To use DHCPv6 for addresses and other info, RA must be used for routing configuration.

SA Managed Server Peer Content Caching

In previous SA releases, if you had a smaller site without a sufficient number of managed servers to justify installation of a full SA Core, SA provided the Satellite installation. The Satellite installation allowed you to install only the minimum necessary Core Components on a Satellite host, which then accessed the Primary Core's database and other services through an SA Gateway connection.

SA also provides Managed Server Peer Content Caching, which provides, for facilities with fewer than 50 managed servers, caching of the Software Repository without the need for Satellite components.

Some of the benefits of Managed Server Peer Content Caching are:

- Peer caching uses existing SA managed servers (no additional hardware infrastructure required)
- No SA Satellite installation is required
- No SA Gateway is required
- Peer caching reduces WAN traffic during software staging
- Peer caching allows pre-staging of software packages
- An SA Satellite or Gateway is not required at the remote site
- Software can be manually loaded into the cache

Requirements

Managed Server Peer Content Caching requires:

- A managed server running any SA supported operating system to act as the Peer Cache server.

- Managed Servers must be configured to use peer caching using custom server attributes.

Installing a Peer Cache

1. Decide which managed server(s) will act as a peer cache(s).
2. Upgrade the Agents on those managed servers to SA 9.14 (other managed servers Agents do not need to be upgraded).

Note: Perform the Agent upgrade as described in the “Agent Utilities” appendix of the SA User Guide: Server Automation.

Configuring the Peer Cache and SA Servers

1. Create a custom attribute for each managed server in the branch/remote site.
 1. For example, `peer_cache_dvc_id = 240001`, where 240001 is the device ID of the server acting as a peer cache.
 2. If the branches/remote sites are modeled as device groups, you can apply the custom attributes at the device group level using a script. Managed servers added to the device group later will automatically inherit this custom attribute.
2. Ensure that all managed servers using the peer cache belong to the same customer as the peer cache.
3. (Optional) Create the following custom attributes on the managed server(s) acting as a peer cache(s):
 1. `peer_cache_size = <value in megabytes>`
default: 1TB (but limited to file system size)
 2. `peer_cache_path = <location of file store>`

Note: `sa_cache` is appended to the value you specify for the path. For example, the default for Windows is:

```
\Program Files\Common Files\Opware\sa_cache
```

4. By default, managed servers attempt to connect to the peer cache using the cache's primary IP address. However, you can use a custom attribute to specify a different IP address in the format:
`peer_cache_ip_field = < primary_ip | management_ip | ip:<addr>>`

where:

`primary_ip` - (default) is the IP address of the management interface. This is the locally configured IP address (not NAT translated).

`management_ip` - is the IP address SA uses to communicate with the server. This can be a NAT translated address.

`ip:<addr>` - is used to set an IP address manually (for example, `ip:192.168.2.1`)

See the SA User Guide: Server Automation for more information about configuring the primary IP address and NAT for managed servers.

Remediation with Peer Caching Enabled

You start remediation as described in the SA User Guide: Software Management.

When Managed Server Peer Content Caching is enabled, remediation performs these steps:

1. During the staging phase, managed servers are given the cache IP address (derived from the `peer_cache_dvc_id` custom attribute attached to the server).
2. The managed servers stages packages from the branch/remote site peer cache (see [Retrieve Objects from the Peer Cache](#)).

Retrieve Objects from the Peer Cache

When retrieving objects from the peer cache, SA performs these tasks:

1. The staging code on the managed server is passed on the IP address of the configured peer cache.
2. The staging code makes a secure connection to the Agent port of the peer cache server using the Agent's SA security certificate.
3. The peer cache confirms that the connecting client is configured to use the cache and belongs to the same customer as the peer cache.
4. A request is made to the peer cache to stage a specified unit.
5. The peer cache server responds to the request by sending the unit.
6. During the action phase, the checksum of the object is verified against the checksum of the same object in the Software Repository.

Possible Errors

Step 1: There is no branch cache configured or unable to communicate to the cache agent:

- Staging proceeds across the WAN normally.

Step 3: The client is not authorized to use the peer cache:

1. The cache logs the unauthorized attempt.
2. The cache returns a 403 Forbidden status to the client.
3. Staging proceeds across the WAN normally.

Step 5: The cache does not have the requested object.

1. The cache returns a 503 with a Retry-Later value to the client.
2. The cache requests the object across the WAN from the Software Repository.
3. The client retries the cache after the specified time and retrieves the file.

Step 5: The cache has the requested unit, but the checksum does not match the core checksum:

1. SA treats the file as stale and deletes it when the cache is full.
2. Proceed with Step 5.

Step 5: The software repository does not have the requested object:

1. This situation should be caught during the analysis phase; if not:
2. The cache returns a 404: file not found message.

Viewing the Peer Cache Status Page

1. Install browser certificate: `browser.p12`

`browser.p12` is located in:

```
/var/opt/opsware/crypto/spin/
```

on any Slice Component bundle host. Copy the file to your local machine, and import `browser.p12` into your browser following your browser import certificate instructions.

2. Using your web browser access:

```
https://<peer_cache>:1002/oplets/peer_cache.py
```

Concepts: SA Core Communications Infrastructure

SA is a distributed computing environment in which individual components communicate with each other securely over an IP network. To accomplish this, SA uses SSL/TLS and X.509 certificates to secure the communication between these components.

When an SA Core component must communicate with another component, it opens a secure (typically SSL/TLS) communication channel using a well-known port. Each SA Core component has a public-key certificate, which is generated when SA is installed. The component uses this public-key certificate when authenticating itself to another component. Most interprocess communication is strongly authenticated (encrypted using the strongest ciphers available) and integrity checked.

Communication Between SA Cores

If you are running SA across multiple data centers, SA automatically synchronizes data across all SA-managed data centers. Broadly speaking, SA synchronizes two types of data: the SA model of servers (including all hardware, software, and configuration attribute information) and software packages.

- **Replicating the SA model:** SA uses integrated certified messaging to synchronize the SA model data. SA implements SSL to safeguard the messages flowing across the message bus. These messages describe SQL changes that must be made to the SA database (Model Repository).
- **Replicating software packages:** SA replicates software packages on demand. That is, packages are only copied when needed. For example, when an administrator managing a server in the New Jersey data center directs SA to install a software package that does not exist in New Jersey's Software Repository, SA requests it from another data center.

The actual file transfer uses the open source utility `rsync`, and the communication channel is secured using SSH. The process is similar for Satellites and for peer-cached software repositories.

Figure 34 and **Figure 35** show a two core installation with a Satellite and how the cores' components communicate using Gateways.

Figure 34. Primary SA Core

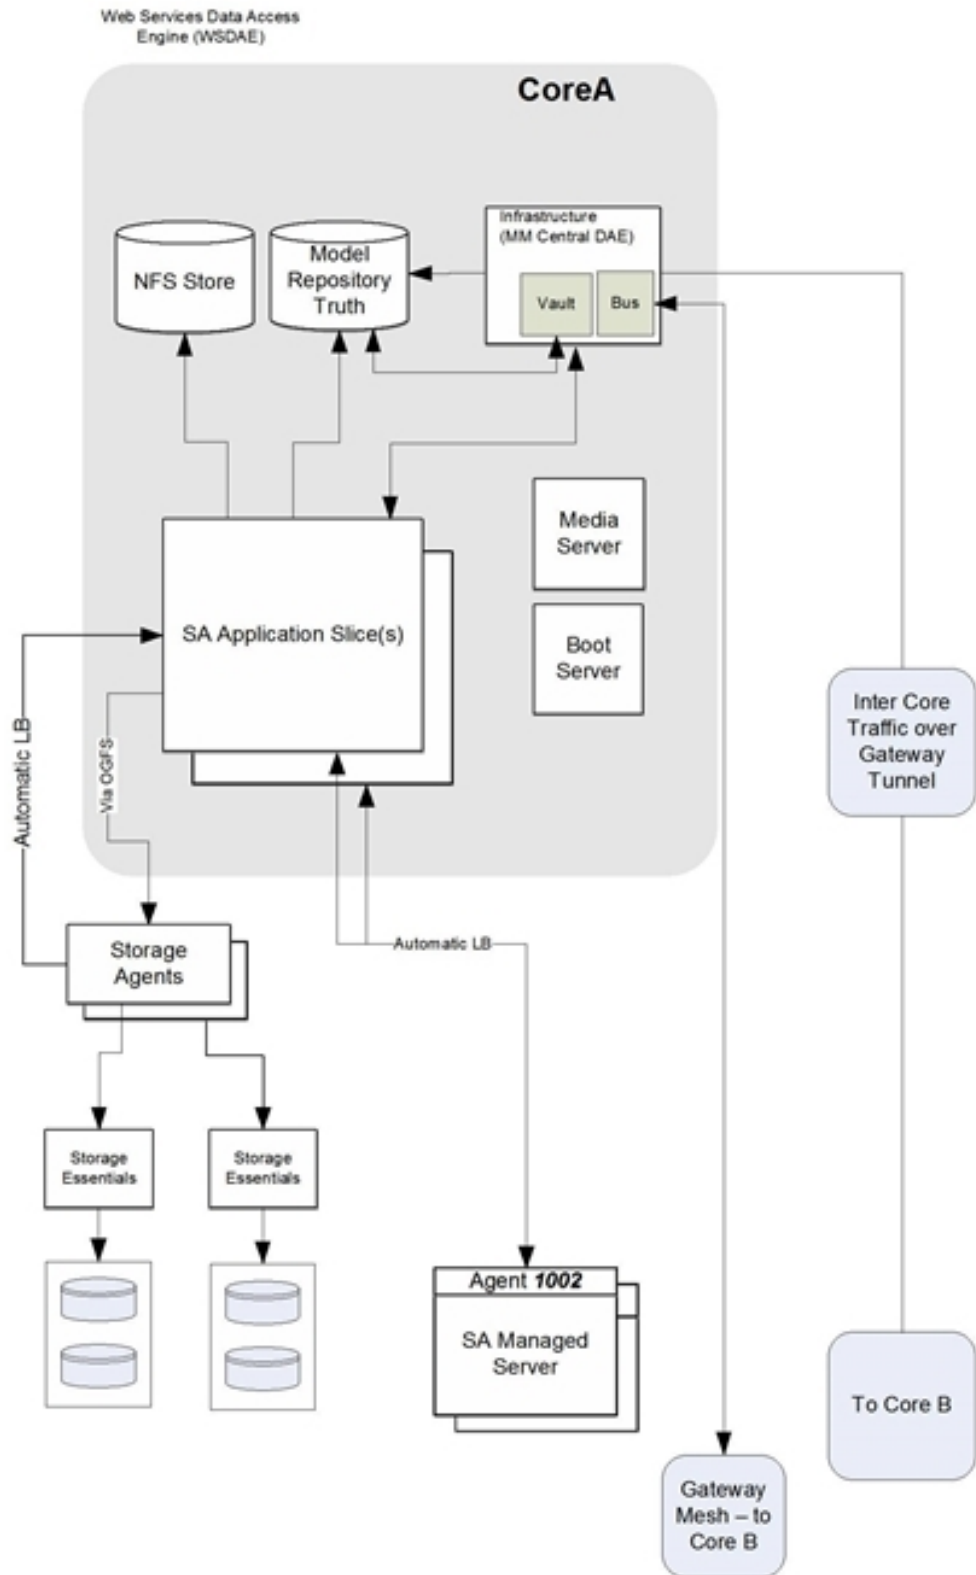
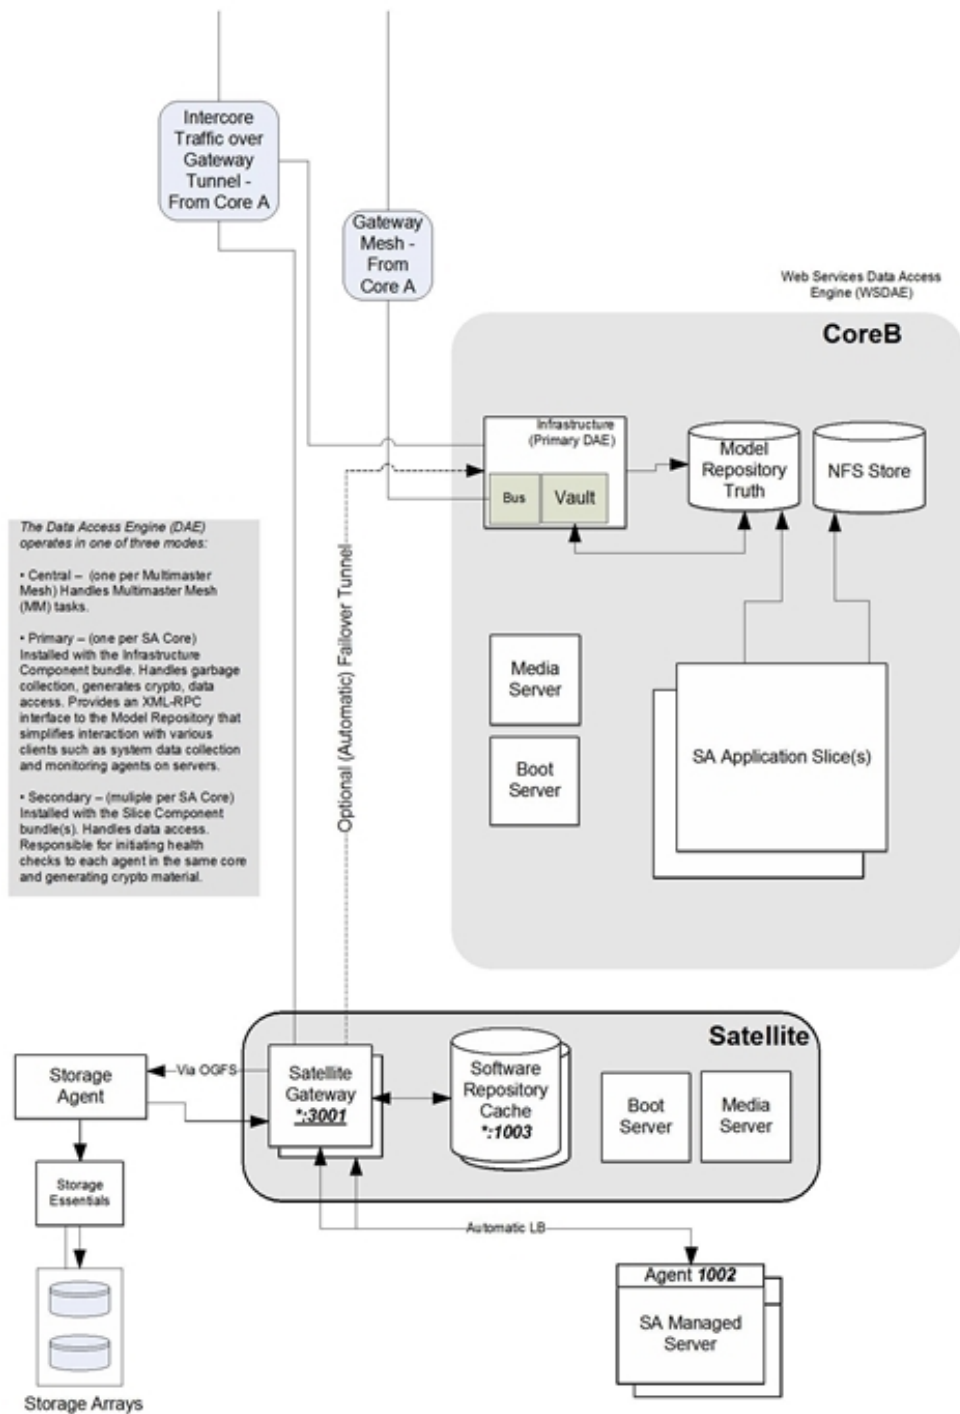


Figure 35. Secondary Core and Satellite



Advanced: Communication Between Agents and SA Core Components

SA Agent installed on managed servers also participate in strongly authenticated and encrypted SSL/TLS traffic. In addition, when Agents are directed to perform management tasks on a server, the typical flow of control messages help to ensure that only authorized users are performing those actions. It would be extremely difficult for an intruder to generate a valid command sequence directing an Agent to perform an unauthorized task.

The following sequence describes a typical SA management task: provisioning software on an SA managed server. Other operations on managed servers follow the same general protocol:

1. The Data Access Engine opens a communication channel through HTTPS with the SA Agent, directing it to perform a management task.
2. The SA Agent calls back to the Data Access Engine to retrieve specifics about the task to perform. To open a communication channel, the Agent must present its public-key certificate, which the SA Core verifies against an internal database mapping the certificate itself to the machine's IP and a unique machine identifier that SA generates when the agent is installed. This safeguard prevents users from simply copying the digital certificate and corresponding key to another machine in hopes of masquerading as the original managed server.

After successfully opening the communication channel, the SA Agent receives the exact list of software to be installed and removed (as well as any scripts to execute, the order of software installation, and when to reboot during the provisioning process).

3. The SA Agent opens a communication channel to the Software Repository (also through HTTPS) and requests a download of the software it needs to install. Before the Software Repository initiates the download, it recomputes an SHA checksum for the package along with a secret key it knows. Only if the SHA checksum matches the checksum generated when the package was uploaded does the SA Agent receive the software it requested, yet another security safeguard.

Asynchronous, agent-initiated calls to the SA Core provide scalable support for progress reporting and long-running operations, as the SA Core need not manage thousands of synchronous agent operations directly. SA supports these asynchronous calls from the Agent to the Core even in network environments where firewalls prevent Agents from initiating TCP connections, since the SA Gateway infrastructure provides bidirectional tunneling over unidirectional connections.

Other technical details of agent/core communications include:

- Connections are SSL v3, mutually authenticated with X.509 certificates (the server checks the client's certificate and vice versa).
- Private keys for Core and Agent certificates are stored in files that are readable by root only.

- All certificates are generated at installation, are owned by the customer, and are not known to HP.
- Certificates expire 10 years after installation. SA provides a Recertification tool for recertifying Cores and Agents prior to certificate expiration.
- Certificates are signed by SA internal self-signed certificate authorities. To avoid HTTPS security warnings in web browsers, customers may install an externally signed certificate in the SA instance of Apache.

This section provides reference information about the parameters in the Gateway Properties file used by the SA Gateway.

SA Gateway Properties File Syntax

The entries in the Gateway Properties file control the operation and configuration of the gateways on the current host.

The SA Gateway Properties file is located in:

```
/var/opt/OPSWgw/gwname/opswgw.properties
```

on each core host.

An SA Gateway properties file can have the following entries:

Note: Do not modify these entries unless you are certain you understand the impact of your change on the core.

Usage: `./opswgw-tc-70 [options]`

`--Gateway name`

(Required) Set the name of the SA Gateway. This name must be unique in a Gateway mesh.

`--Realm realm`

(Required) All Gateways operate in a named Realm. A *Realm* is an SA construct that refers to a set of servers that are serviced by the Gateways in the Realm. Realms can support an IPv4 address space that may overlap with other Realms. Realms are also used to define bandwidth utilization constraints on SA functions.

`--Root true | false`

Specifies that this Gateway will act as a root of the Gateway mesh. All Gateways in a Root Realm must be Root Gateways.

Default: false.

`--Level int`

(Experimental) Routing level for the Gateway. There are eight possible levels, 0 through 7. All Gateways in a realm must have the same level.

Default: 0

`--GWAddress lhost`

Sets the local host address (if you are specifying the value for the Management Gateway, use the IP address only; do not use the hostname. You can, however, use the hostname for other, non-Management Gateways) that this Gateway uses to tell other components how to contact it. This value is used by the core to discover new coreside Gateways. It is also used to communicate the active list of Gateways that are servicing Realm to proxy clients (such as Agents) through the `X-OPSWGWLIST` mime header.

`--Daemon true | false`

Daemonize the process.

Default: false.

`--Watchdog true | false`

Start an internal watchdog process to restart the Gateway in case of a failure or signal. A `SIGTERM` sent to the watchdog will stop the watchdog and Gateway processes.

Default: false.

`--User name`

Change to this user on startup.

`--RunDir path`

Change to this directory on startup.

`--ChangeRoot true | false`

If true chroot into RunDir. This can be used by a helper script to construct a jail.

Default: false

`--PreBind proto:ip:port, ...`

For security reasons, it can be useful to run a Gateway chrooted as a nonprivileged user (only ports above 1024 can be used for any listeners). If you want to use a nonprivileged user *and* a privileged listener port, you can use the `--PreBind` directive to reserve the port while the process is root and before privileges are dropped.

```
--HardExitTimeout seconds
```

The number of seconds after a restart or exit request that the main thread will wait for internal threads and queues to quiesce before performing a hard exit.

```
--LogLevel INFO | DEBUG | TRACE
```

Sets the logging level. Note that `DEBUG` and `TRACE` can produce a large amount of output, which typically is relevant only to developers and can negatively affect performance.

Default: `INFO`.

```
--LogFile file
```

The filename of the SA log file.

```
--LogNum num
```

The number of rolling log files to keep.

```
--LogSize size
```

The size, in bytes, of each log file.

```
--TunnelDst [lip1:]lport1[:cryptol],...
```

If specified, starts a tunnel destination listener. The tunnel listener can listen on multiple ports (a comma-separated list with no spaces). If the port is prefixed with an IP address, the listener will bind only to that IP address. For example: `2001, 10.0.0.2:2001, 2001:/var/foo.pem, 10.0.0.2:2001:/var/foo.pem`

```
--TunnelSrc rhost1:rport1:cost1:bw1[:cryptol],...
```

If specified, creates a tunnel between this Gateway and the Gateway listening at `rhost1:rport1`. The link `cost1` and link bandwidth `bw1` must be set. The `cost` is a 32-bit unsigned int, and bandwidth is in Kbits/sec (K=1024bits). (Additional tunnels are separated by commas.)

Examples: `gw.foo.com:2001:1:0, gw.bar.com:2001:10:256:/var/foo.pem`

```
--ProxyPort [lip1:]lport1,[lip2:]lport2,...
```

The HTTP CONNECT proxy listener port. If more than one proxy listener port is needed, you can add more using a comma-separated list. You can enable interface binding by prepending an IP address to the port.

```
--ForwardTCP [lip:]lport1:realm1:rhost1:rport1,...
```

Creates a static TCP port forward. Forward the local port `lport(x)` to the remote service `rhost(x):rport(x)`, which is in `realm(x)`. A blank `realm` (such as `lport::rhost:rport`) means route to the closest Root Realm.

```
--ForwardTLS [lip:]lport1:realm1:rhost1:rport1, ...
```

Creates a static TCP port forward that specializes in TLS traffic. The TLS session ID is parsed and sent to the egress Gateway for use in load-balancing algorithms. In all other respects, this feature behaves like `ForwardTCP`.

```
--ForwardUDP [lip:]lport1:realm1:rhost1:rport1,...
```

Creates a static UDP port forward. Forward local port `lport(x)` to remote service `rhost(x):rport(x)`, which is in `realm(x)`. A blank `realm` (such as `lport::rhost:rport`) means route to the closest Root Realm. (Note: Some UDP services, such as DHCP, cannot be proxied in this way.)

```
--IdentPort [lip:]lport
```

Starts an IDENT service listening on local port `lport` (optionally bound to the local IP `lip`).

```
--AdminPort [lip:]lport[:cryptol]
```

Starts an administration interface listening on local port `lport`, which is optionally bound to the local IP `lip`. If you use `crypto`, include a `crypto` specification file name.

```
--ConnectionLimit int
```

Specifies the soft memory tuning limit for the maximum number of connections.

```
--OpenTimeout seconds
```

Waits a maximum `seconds` for a remote `CONNECT` call to establish a remote connection.

```
--ConnectTimeout seconds
```


Waits a maximum seconds for a `connect()` to complete. If a timeout occurs, then an HTTP 503 message is returned to the client (via the ingress Gateway). The client will get this message if the `ConnectTimeout` plus the Gateway mesh transit delay is less than the `OpenTimeout`.

`--ReorderTimeout seconds`

In the event of out-of-order messages (for a TCP flow), limits the amount of time (seconds) to wait for messages needed for reassembly to arrive. The most common cause of out-of-order messages is when a transit tunnel fails and a new route is taken mid-flow.

`--TunnelStreamPacketTimeout seconds`

If a portion of a TCP flow cannot be delivered to an endpoint, then tears down the TCP connection after seconds.

`--QueueWaitTimeout seconds`

Specifies the maximum time that a tunnel message can wait at the head of an internal routing queue while waiting for a tunnel to be restored.

`--KeepAliveRate seconds`

Send link keepalive messages once every x seconds on each link.

`--LsaPublishRateMultiple float`

Link State Advertisements (LSA) are published once every $k \cdot M$ seconds, in which M is the number of Gateways in the mesh and k is a floating point constant specified using `--LsaPublishRateMultiple`. For example, if there are 100 Gateways in a mesh and `--LsaPublishRateMultiple` is set to 2.0, then an LSA is published approximately every 200 seconds (due to implementation factors, the actual delay will be somewhere between 190 and 210 seconds).

`--LsaTTLMultiple float`

Sets the TTL for LSA to float multiplied by the `LsaPublishRate`. Example: If `LsaPublishRate` is 10 seconds and `LsaTTLMultiple` is 3, then the TTL for LSA published by this Gateway is set to 30 seconds.

`--MaxRouteAge seconds`

Discards the routes from the routing table that have not been refreshed within seconds.

`--RouteRecalcDutyCycle percentage`

If the time to calculate Dijkstra takes τ seconds, then wait for $\tau \cdot (1/\text{RouteRecalcDutyCycle} - 1)$ seconds until another recalculation can take place.

```
--TunnelTimeoutMultiple float
```

This number, multiplied by the `KeepAliveRate`, gives the maximum time that a tunnel can be idle before it is garbage collected.

```
--DoNotRouteService host1:port1,host2:port2,...
```

Specifies that, when a local client creates a proxy connection to `host:port`, do not route the message; service it locally. Use this property to ensure that certain services are handled locally, in the Gateway's current Realm.

```
--ForceRouteService host1:port1:realm1,host2:port2:realm2,...
```

When a local client creates a proxy connection to `host:port`, forces the message to route to a specified Realm.

```
--HijackService host1:port1,host2:port2,...
```

When the local Gateway sees a connection to `host:port` via a tunnel, and the source Realm is not the local Realm, it must service the connection. If the connection is from the local Realm, the Gateway must allow the message to continue to its destination. You can use this feature to implement transparent caches.

```
--RouteMessages *true | false
```

If specified as `true`, turn on transit routing. If `false`, disable transit routing. If the destination of the message is *not* the local Gateway, then, by default, the message is routed based on the current routing table. If such routing is not desired, set this property to `false`.

```
--EgressFilter proto:dsthost1:dstport1:srchost1:srcrealm1,...
```

When the local Gateway sees a TCP connection attempt to `dsthost:dstport` from `srchost1:srcrealm1`, it must allow the connection. The implied default is to deny all connections. If you want to *allow* all connections, specify the egress filter as `*:*:*:*:*`. It is also common for an egress filter to allow connections only from the Root Realm. This can be expressed by leaving the `srcrealm` blank. Example: `tcp:10.0.0.5:22:172.16.0.5:` allows tcp connections to 10.0.0.5, port 22, from 172.16.0.5 in a Root Realm.

```
--IngressMap ip1:name,ip2:name,...
```

When sending an open message (and the `srcip` is in the ingress map), append (as metadata) the `ip:name` mapping to the open message. This allows a remote egress filter to use the `name` as the `srchost` instead of the `ip`. This feature supports the addition of a server to a farm without the need to individually add the server to many `EgressFilter` entries.

```
--LoadBalanceRule proto:thost:tport:mode:rhost1:rport1:  
rhost2:rport2, ...
```

When receiving a new connection message for `thost:tport`, load balance the connection over real hosts `rhost1:rport1`, `rhost2:rport2` etc. The load balance strategy is defined by `mode`.

There are six load-balancing modes:

STICKY: Send the connection to a working target based on a priority list randomized by a hash of the source IP and source Realm (the hash string can be overridden via the input MIME header `X-OPSW-LBSOURCE`).

LC: Send connection to a working target with the least number of connections.

RR: Send connection to the next working target in a round-robin fashion.

TLS_STICKY: Use an SSLv3/TLSv1.0 session ID to send the connection back to the previous target based on a session ID cache. If the target is in error, or the session ID is missing from the cache, fall back to **STICKY** mode to make a new selection.

TLS_LC: Similar to **TLS_STICKY** mode, but falls back to **LC** mode (least connections).

TLS_RR: Similar to **TLS_STICKY** mode, but falls back to **RR** mode (round-robin). Remember to add an egress filter for `proto:thost:tport`. You do not need to add egress filters for the targets. Non-TLS load balancing modes *can* be used with UDP services.

```
--LoadBalanceRetryWindow seconds
```

If an error occurs when using a load balanced target (such as `rhost1:rport1` above), then the target is marked `inerror`. This property controls how many seconds a Gateway will wait until it retries the target. If the target is missing (such as an `RST` is received upon the connection request), the load balancer will try to find a good target.

The number of seconds a load balanced SSLv3/TLS client can be idle before the `sessionId` association is reaped. This property affects the egress Gateway of a TLS flow.

```
--SessionIdCacheLimit slots
```

A soft limit on the number of SSLv3/TLS session IDs that the cache can hold. If this limit is exceeded, then the garbage collector begins reducing the `SessionIdTimeout` value in order to achieve the cache limit specified by `--SessionIdCacheLimit`.

`--MinIdleTime seconds`

Specifies the minimum number of seconds a connection can be idle during an overload condition before it will be considered for reaping.

`--GCOverloadTrigger float`

Specifies the fraction of `SoftConnectionLimit` at which to start overload protection measures. When the number of open connections reaches this overload trigger point, overload protection starts, reaping the most idle connections over `MinIdleTime`. Overload protection stops when the connection count falls below the overload trigger point.

`--GCCloseOverload true | false`

When a client tries to open a connection after the `ConnectionLimit` has been reached, this property tells the Gateway what to do with the new connection. A value of `true` causes the Gateway to close the new connection. A value of `false` causes the Gateway to park the new connection in the kernel's backlog and to service it after the overload condition subsides. The proper setting is application dependent.

Default: `false`.

`--VerifyRate seconds`

When a connection stops moving data for the specified number of seconds, a connection verify message is sent to the remote Gateway to verify that the connection is still open. This check is repeated periodically and indefinitely when the timeout has expired.

`--OutputQueueSize slots`

Specifies the size of the tunnel output queues. These queues store messages destined for remote Gateways. Each remote Gateway has an output queue. Queues are garbage collected after `MaxQueueIdleTime` is reached.

`--MaxQueueIdleTime seconds`

Specifies the maximum time to keep an idle output queue before garbage collection removes it.

`--TunnelManagementQueueSize slots`

Specifies the size of the queues used to manage tunnel management traffic, such as LSA.

`--TunnelTCPBuffer bytes`

Specifies the size of the TCP SEND and RECV buffer in `bytes`. The operating system must be configured to handle the specified value. You can view the Gateway's log file to see if the specified is denied by the operating system.

```
--DefaultChunkSize bytes
```

Specifies the default (maximum) IO chunk size when encapsulating a TCP stream. This property value can be applied only to links with no bandwidth constraint.

```
--LinkSaturationTime seconds
```

When a link has a bandwidth constraint, the chunk size, `DefaultChunkSize`, is computed based on two parameters. The first is the link's bandwidth constraint. The second is the amount of time that the bandwidth shaper should use the full, real, bandwidth on the link. This parameter controls the duty cycle of the bandwidth shaper. Smaller values give a smoother bandwidth control at the cost of more overhead, because each smaller IO chunk has a header.

```
--TunnelPreLoad slots
```

Specifies the maximum number of output queue slots to use before waiting for the first `Ack` message. This allows for pipelining in Long Fat Pipes. This value is reduced geometrically to one as the number of queue slots diminish.

```
--BandwidthAveWindow samples
```

Specifies the maximum number of IO rate samples for the bandwidth estimation moving window. The samples in this window are averaged to provide a low-pass estimate of the bandwidth in use by a tunnel. This estimate has high frequency components due to the sharp edge of the filter window.

```
--BandwidthFilterPole float
```

Specifies the pole of a discrete-time first-order smoothing filter used to remove the high frequency components of the moving window estimator. Set the value to 0.0 to turn off this filter.

```
--StyleSheet URL
```

Adds a stylesheet link to a URL when rendering the admin UI. This is useful for embedding the admin UI in another web-based UI. In addition to using this property to control the default stylesheet, a dynamic stylesheet override is supported by adding the variable `StyleSheet-t=<url>/style.css` to the admin UI URL.

```
--ValidatePeerCN true | false
```

Specifies whether the peer CN is validated against the peer configuration during a tunnel handshake operation. The peer must be turned off during the installation of an untrusted Gateway.

Default: true.

```
--PropertiesCache file
```

Link cost and bandwidth can be controlled via parametermodify messages over tunnel connections. These realtime adjustments are made to the running process and written to a parameter cache, which will override the properties file or command-line arguments.

```
--PropertiesInclude file
```

Specifies an Include file to load and merge with the current properties. Properties in the include file can override properties from the original Properties File. This property can be specified from the command line. If so, it will override *all* properties, including command-line overrides. It is not recursive and does not support a list.

```
--PropertiesFile file
```

Places all command-line arguments into a properties file within the opswgw name space. Note that the PropertiesFile command-line argument itself *must not* be placed in the properties file within the opswgw name space.

opswgw Command-Line Arguments

All of the parameters in the preceding section can be specified as options for the opswgw command. For example, the opswgw.Gateway foo entry in the Gateway Properties file is equivalent to the following command-line argument:

```
/opt/opsware/opswgw/bin/opswgw --Gateway foo
```

Command-line arguments override corresponding entries in the Gateway Properties file. In addition to the entries listed in the preceding section, the opswgw command can specify a Gateway Properties file as an argument; for example:

```
/opt/opsware/opswgw/bin/opswgw --PropertiesFile filename
```

SA Maintenance

The SA Start/Stop Script

SA provides a multipurpose script for starting, stopping, and getting the status of SA:

```
/etc/init.d/opsware-sas
```

You can use the script to display all SA components installed on a server, to start, stop, or restart all core components, or to start, stop, or restart specific SA components (other than the Oracle database).

For information about starting and stopping the Oracle database, see [Starting the Oracle Database \(Model Repository\)](#).

When running the script on a Core Component host, the script performs the necessary pre-requisite checks for each component installed on the local system.

Note: If an SA Core's components are distributed across multiple servers, the start/stop script cannot interact directly with remote servers to start or stop the remote components. However, the script can connect to the remote servers to determine whether prerequisites are met before starting dependent components locally.

When checking prerequisites for components running on remote servers, the script uses time-out values to allow for different boot times and speed differences among servers. If any of the prerequisite checks fail, the script terminates with an error.

Dependency Checking by the Start/Stop Script

The start/stop script recognizes SA component dependencies and starts SA components in the correct order. The prerequisite checks verify that dependencies are met before the script starts a given component, thus ensuring that the SA components installed across multiple servers start in the correct order.

For example, if the component you are attempting to start requires that another component be running, the script can verify whether:

- The required component's hostname is resolvable
- The host on which the required component is running is listening on a given port

Start/Stop Script Logs

The start/stop script writes to the following logs:

Start/Stop Script Logging

Log	Notes
<code>/var/log/opsware/startup</code>	When the server boots, the script logs the full text (all text sent to <code>stdout</code>) of the start process for all SA components installed on the local system.
<code>stdout</code>	When invoked from the command line, the script displays the full text of the start process for the components.
<code>syslog</code>	When the server boots, the script runs as a background process and sends status messages to the system event logger.

Start/Stop Script Syntax

The SA start/stop script has the following syntax:

```
/etc/init.d/opsware-sas [options] [component1] [component2]...
```

When you specify specific components to start, stop, or restart, those components must be installed on the local system, and you must enter the names exactly as they are displayed by the `list` option. Table 24 lists the options for the SA start/stop script. To see the options of the Health Check Monitor (HCM) also invoked with `opsware-SA`, see **Table 28**.

Table 24. Options for the SA Start/Stop Script

Option	Description
<code>list</code>	Displays all components that are installed on the local system and managed by the script. The script displays the components in the order that they are started.
<code>start</code>	<p>Starts all components installed on the local system in the correct order. When you use the <code>start</code> option to start a specific component, the script performs the necessary prerequisite checks, then starts the component.</p> <p>The <code>start</code> option does not start the Oracle database (Model Repository), which must be up and running before the SA components can be started.</p> <p>Some SA components, such as the Web Services Data Access Engine (<code>twist</code>), can take longer to start. For these components, you can run the script with the <code>start</code> option so that the script runs on the local system as a background process and logs errors and failed checks to the component's log file.</p> <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p>Note: Note: When you use the <code>start</code> option to start multiple components installed on a server, the script will always run the <code>/etc/init.d/opsware-sas</code> command with the <code>startsync</code> option.</p> </div>

Option	Description
<code>startsync</code>	<p>The <code>startsync</code> option starts all components installed on the local system in a synchronous mode.</p> <p>When you use the <code>startsync</code> option, the script runs in the foreground and displays summary messages of its progress to <code>stdout</code>.</p>
<code>restart</code>	<p>Stops and starts all components installed on the local system in a synchronous mode. The script stops all local components in reverse order, then it executes the <code>startsync</code> option to restart the components in the correct order.</p>
<code>stop</code>	<p>Stops all components installed on the local system in the correct order.</p> <p>This option does not stop the Oracle database.</p>

Starting the Oracle Database (Model Repository)

The SA start/stop script cannot start the Oracle database (required for the Model Repository), which must be up and running before the SA components can be started. Before you start the SA components, be sure to start the Oracle listener and database by entering the following command:

```
/etc/init.d/opsware-oracle start
```

Starting a Standalone SA Core

To start a core that has been installed on a single server, perform the following steps:

1. Log in as `root` to the core server.
2. Start the Oracle listener and database for the Model Repository:

```
/etc/init.d/opsware-oracle start
```

3. Start all core components:

```
/etc/init.d/opsware-sas start
```

Starting a Multiple-Server SA Core

SA Core startup order can be affected by several factors. This section describes starting an SA Core in a Multimaster Mesh configuration.

Core Component Hosts Powered Up

If the entire mesh is stopped but the hosts are powered on, the Primary Core must be started first, followed by each secondary core. Each secondary must be started one at a time.

Perform the following steps:

Primary Core

1. If necessary, determine the servers that host the core's components. Log in to the Model Repository host as root and invoke the following command:

```
/etc/init.d/opsware-sas list
```

2. Log in as root to the Primary Core's Model Repository host and start the Oracle listener and database:

```
/etc/init.d/opsware-oracle start
```

3. After the database and listener successfully start, run the SA start script on the following Core Component hosts, one server at a time, in the following order:
 - Infrastructure Component bundle host
 - Slice Component bundle (initial Slice) if not installed on the same host as the Infrastructure Component bundle
 - Subsequent Slice Component bundle hosts
 - OS Provisioning Component bundle host
 - Satellite hosts associated with the core

Invoke the SA start script on each host with this command:

```
/etc/init.d/opsware-sas start
```

The start-up script must complete starting all core components successfully on each host before you invoke it on the next server.

Secondary Core(s)

The start order is the same as above but must be performed after the Primary Core Components have been successfully started. You must start the core components on only one Secondary Core at a time.

Core Component Hosts Powered Down

When the core component hosts are powered off, powering on the hosts also starts SA; therefore, the hosts must be powered on in the following order:

- Infrastructure Component bundle host
- Slice Component bundle (Slice0) if not installed on the same host as the Infrastructure Component bundle
- Additional Slice Component bundle hosts (Slice1 to Slice n), one at a time
- OS Provisioning Component bundle host
- Satellite hosts associated with the core, one at a time

The hosts must be powered up one at a time and the SA Core Components must have successfully started before powering up the next server. You can use the tail command on the the most recently created log file in `/var/opt/opsware/log/startup` to determine the startup status of the components on each host.

Starting Individual SA Core Components

You can specify one or more components to start if those components are running on the local system. You must enter the component names exactly as they are displayed by the `list` option of the `opsware-SA` command.

To start individual components of an SA core, perform the following steps:

1. Log in as `root` to the server that has the component you want to start.
2. (Optional) To list the SA components installed on a server, enter the following command:

```
/etc/init.d/opsware-sas list
```

3. Enter the following command, where *component* is the name as displayed by the `list` option:

```
/etc/init.d/opsware-sas start component
```

For example, if the `list` option displayed `buildmgr`, enter the following command to start the OS Provisioning Build Manager:

```
/etc/init.d/opsware-sas start buildmgr
```

Tip: Alternatively, you can enter the `startsync` option when starting a component on a server. See [Table 24. Options for the SA Start/Stop Script](#) in this chapter for a description of the `startsync` option.

Start Order for Individual SA Core Components

The SA start script starts core components installed on a host in the order listed below. When the script stops components installed on a host, it stops them in reverse of the order in which they were started.

1. `opswgw-mgw`: The SA Primary Core Master Gateway
2. `opswgw-cgws0-<facility>`: The core-side Gateway for the facility in which the core is running
3. `opswgw-cgws`: Other Gateways in the mesh
4. `vaultdaemon`: The Model Repository Multimaster Component
5. `dhcpcd`: A component of the OS Provisioning feature
6. `pxe`: The PXE boot environment
7. `memcached`: An in-memory caching layer that works with the Software Repository Accelerator (`tsunami`) component to support remediation and scalability enhancements for agents that communicate directly with a Linux-based SA Core.
8. `spin`: The Data Access Engine
9. `mm_wordbot`: A component of the Software Repository

10. `tsunami`: Software Repository Accelerator, an object store download accelerator that boosts remediation performance and scalability for any agents that communicate directly with a Linux-based SA Core.
11. `waybot`: The Command Engine
12. `smb`: A component of the OS Provisioning feature
13. `twist`: The Web Services Data Access Engine
14. `buildmgr`: The OS Provisioning Build Manager
15. `opswgw-agw0-<facility>`: The agent-side Gateway for the facility in which the core is running
16. `opswgw-agws`: The Agent Gateways
17. `hub`: A component of the Global File System
18. `sshd`: A component of the Global File System
19. `apxproxy`: The Automation Platform Extension (APX) proxy
20. `spoke`: A component of the Global File System
21. `agentcache`: A component of the Global File System
22. `occ.server`: A component of the SA Client
23. `httpsProxy`: HTTP(S) proxy for core components
24. `da`: The Application Deployment component
25. `opsware-agent`: The Server Agents

Stopping an SA Core with Multiple Hosts

When you shut down a mesh, each core must be stopped in reverse of the start order, and each host within the core must be powered down in reverse start order. Each Secondary Core must be shut down one at a time, followed finally by the Primary Core.

Within each core, primary or secondary. `/etc/init.d/opsware-sas stop` needs to be run on the servers in this order:

- Satellite host(s) associated with the core, one at a time
- OS Provisioning Component bundle host
- Additional Slice Component bundle hosts (Slice1 to Slice n), one at a time
- Slice Component bundle (Slice0) if not installed on the same host as the Infrastructure Component bundle
- Infrastructure Component bundle host
- Database/Model Repository host

To stop the core components on a host, invoke the following command:

```
/etc/init.d/opsware-oracle stop
```

Multiple Data Access Engines

This section discusses the following topics:

- [Overview of Multiple Data Access Engines](#)

- [Reassigning the Data Access Engine to a Secondary Role](#)
- [Designating the Multimaster Central Data Access Engine](#)

Overview of Multiple Data Access Engines

In a core with multiple instances of the Data Access Engine, each instance may be designated in one of the following ways:

1. **Primary Data Access Engine:** Each Facility has only one *primary* Data Access Engine. This Data Access Engine periodically checks the Managed Servers to determine if SA can communicate with them. If a facility has more than one primary Data Access Engine, the competing reachability checks can interfere with each other.
2. **Secondary Data Access Engine:** When a Facility has multiple Data Access Engines installed (for scalability), the non-primary ones are designated as secondary data access engines. The first Data Access Engine installed is designated the Primary or Multimaster Central Data Access Engine. A secondary Data Access Engine does not check managed servers to determine if they are reachable. It only communicates with the Model Repository to write or read data.
3. **Multimaster Central Data Access Engine:** An SA Multimaster Mesh has multiple cores and therefore multiple data access engines. One core's primary data access engine should be designated the *Multimaster Central Data Access Engine*. Although any of the cores may have multiple Data Access Engines, only one per mesh can be the central data access engine.

Reassigning the Data Access Engine to a Secondary Role

This functionality was moved from the SA Web Client to the SA Java Client. Therefore, if you have installed an additional Data Access Engine, you must perform the following steps to reassign the new Data Access Engine to a secondary role:

1. Log into the SA Java Client as a user who belongs to SA Administrators group.
2. Select the Administration tab to display your administration inventory.
3. Select **System Configuration > Service Level Members**. The Service Levels tree is displayed.
4. Choose the Data Access Engine (spin) node. The Managed Servers that are hosting a Data Access Engine will appear in the members table.
5. Select the desired Data Access Engine server and press **Cut**, or choose the Cut action from right-click actions, or from Actions menu.
6. Then select Secondary node from the tree and press **Paste**, or choose the Paste action from right-click actions, or from the Actions menu.
7. A confirmation dialog will appear. Click **Yes**.

8. The new Data Access Engine server will appear in the members table.
9. In a terminal window, log in as root to the server running the additional Data Access Engine, and enter the following command to restart the Data Access Engine:

```
/etc/init.d/opsware-sas restart spin
```

Designating the Multimaster Central Data Access Engine

The HP BSA Installer automatically assigns the multimaster central Data Access Engine.

Caution: In most cases, you should not change the multimaster central Data Access Engine after the installation. Doing so can cause problems when upgrading the SA core to a new version. Before following the steps in this section, contact HP Professional Services.

This functionality is moved from the SA Web Client to the SA Java Client. Therefore, perform the following steps to designate the multimaster central data access engine:

1. Log into the SA Java Client as a user who belongs to the SA System Administrators group.
2. Select the Administration tab to display your administration inventory.
3. Select **System Configuration > Service Level Members**. The Service Levels tree is displayed.
4. Choose the Data Access Engine (spin) node. The Managed Servers that are hosting a Data Access Engine will appear in the members table.
5. Select the desired Data Access Engine server and press **Cut**, or choose the Cut action from right-click actions, or from the Actions menu.
6. Then select Multimaster Central node from tree and press **Paste**, or choose the Paste action from right-click actions, or from the Actions menu.
7. A confirmation dialog will appear. Click **Yes**.
8. The new Data Access Engine server will appear in the members table.
9. In a terminal window, log in as root to the server running the additional Data Access Engine, and enter the following command to restart the Data Access Engine:

```
/etc/init.d/opsware-sas restart spin
```

Scheduling Audit Results and Snapshot Removal



Because audit results and snapshots (results of a snapshot specification) can accumulate over time, especially those that run on a recurring schedule, you can configure your SA core so that after a specified number of days audit results and snapshots will be deleted from the core.

Note that this setting only applies to those audit results and snapshots that have not been archived. Archived results can only be deleted from the SA Client manually.

Additionally, there are two other conditions where an audit result *or a snapshot* will not be deleted by these settings:

- If the snapshot is being used as the target of an audit
- If the audit result *or snapshot is the only result of either an audit or snapshot specification*

To configure audit results and snapshot removal:

1. Select the **Administration** tab in the SA Client.
2. Select System Configuration in the navigation pane. This displays the SA components, facilities, and realms that have system configuration parameters.
3. In the list of SA components, select Data Access Engine. This displays the system configuration parameters for this component.
4. Locate and modify the following system configuration parameters:
 - Locate the `spin.cronbot.delete_audits.cleanup_days` parameter. Enter the new value directly, or select the new value button  and enter the number of days that must elapse before all non-archived audit results will be deleted. If you select **Default value**, no audits will be deleted.
 - Locate the `spin.cronbot.delete_snapshots.cleanup_day` parameter. Enter the new value directly, or select the new value button  and enter the number of days that must elapse before all non-archived snapshots will be deleted. If you select **Default value**, no snapshots will be deleted.
5. Select the Revert button to discard your changes or the Save button to save your changes.

Web Services Data Access Engine Configuration Parameters

This section discusses how to change Web Services Data Access Engine system configuration parameters using the SA Client or by editing the configuration file.

Note: You must restart the Web Services Data Access Engine after changing any system configuration parameters.

Changing a System Configuration Parameter

This section describes how to change some of the system configuration parameters with the SA Client. Other parameters can only be changed by editing a configuration file as described in [Web](#)

Services Data Access Engine Configuration File.

To change a system configuration parameter for the Web Services Data Access Engine in the SA Client, perform the following steps:

1. Select the **Administration** tab in the SA Client.
2. In the navigation pane, select **System Configuration > Configuration Parameters**. This displays the SA components, facilities, and realms that have system configuration parameters.
3. In the list of SA components, select Web Services Data Access Engine. This displays the system configuration parameters for this component.
4. Locate and modify the system configuration parameters you want to change.
5. Select the Revert button to discard your changes or the Save button to save your changes.
6. Restart the Web Services Data Access Engine with the following command:

```
/etc/init.d/opsware-sas restart twist
```

Web Services Data Access Engine Configuration File

The Web Services Data Access Engine configuration file includes properties that affect the server side of the SA Web Services API 2.2. (These properties are not displayed in the SA Client.) The fully qualified name of the configuration file is as follows:

```
/etc/opt/opsware/twist/twist.conf
```

Note: During an upgrade of SA, the `twist.conf` file is replaced, but the `twist_custom.conf` file is preserved. When you upgrade to a new version of SA, to retain the configuration settings, you must edit the `twist_custom.conf` file. The properties in `twist_custom.conf` override those specified in `twist.conf`. The UNIX `twist` user must have write access to the `twist_custom.conf` file.

To change a property defined in the configuration file:

1. Edit the `twist.conf` file with a text editor.
2. Save the changed file.
3. Restart the Web Services Data Access Engine on the server.

Note: You must belong to the Administrators group (`admin`) to modify the `twist.conf` file. Once the file is changed, the Web Services Data Access Engine must be restarted to apply the changes.

The following table lists the properties of the configuration file that affect the SA Web Services API 2.2. Several of these properties are related to the cache (sliding window) of server events. SA maintains a sliding window (with a default size of two hours) of events describing changes to SA objects. This window makes enables software developers to update a client-side cache of objects without having to retrieve all of the objects. For more information, see the API documentation for `EventCacheService`.

Configuration File for SA Web Services API 2.2

Property	Default	Description
<code>twist.webservices.debug.level</code>	1	An integer value that sets the debug level for the SA Web Services API on the server side. Allowed values: 0 - basic info 1 - more detailed information 2 - stack trace 3 - for printing the server event cache entries whenever there is an item added to the cache.
<code>twist.webservices.locale.country</code>	US	The country Internationalization parameter for the Localizer utility. Currently only the US code is supported.
<code>twist.webservices.locale.language</code>	en	Sets the language Internationalization parameter for the Localizer utility. Currently only the en code is supported.
<code>twist.webservices.caching.windowsize</code>	120	In minutes, the size of the sliding window maintaining the server event cache.
<code>twist.webservices.caching.windowslide</code>	15	In minutes, the sliding scope for the window maintaining the server event cache.
<code>twist.webservices.caching.safetybuffer</code>	5	In minutes, the safety buffer for the sliding window maintaining the server

Property	Default	Description
		event cache.
<code>twist.webservices.caching.minwindowsize</code>	30	In minutes, the minimum size of the sliding window that maintains the server event cache.
<code>twist.webservices.caching.maxwindowsize</code>	240	In minutes, the maximum size of the sliding window that maintains the server event cache.

Increasing the Web Services Data Access Engine Maximum Heap Memory Allocation

As data size in a multimaster mesh grows, you may find that you must increase the maximum heap memory allocation for the Web Services Data Access Engine (`twist`). The default value is 1280Mb. To do so, perform the following tasks:

1. Using a text editor, open the file:

```
/etc/opt/opsware/twist/twist_custom.conf
```

2. Modify the following entry to the required allocation:

```
twist.mxMem=<memory size>
```

where memory size corresponds to `-Xmx<memory size>`.

For example:

```
twist.mxMem=2048m
```

would give the Web Services Data Access Engine a maximum of 2048 megabytes of heap memory. This change is preserved even after an upgrade. If you leave this `twist_custom.conf` parameter blank, the default value (1280m) specified in `twist.sh` is used.

Changing Software Repository Mirroring Parameters

Software repository mirroring keeps the software repositories in a multimaster mesh in sync for redundancy and diSAter recovery. This section explains how to change the Software Repository Mirroring configuration parameters. For more information, see [Software Repository Monitoring](#).

Changing a System Configuration Parameter

This section describes how to change some of the system configuration parameters with the SA Client. Other parameters can only be changed by editing a configuration file as described in [Web Services Data Access Engine Configuration File](#).

To change a system configuration parameter for the Web Services Data Access Engine in the SA Client, perform the following steps:

1. Select the **Administration** tab in the SA Client.
2. In the navigation pane, select **System Configuration > Configuration Parameters**. This displays the SA components, facilities, and realms that have system configuration parameters.
3. In the list of SA components, select Web Services Data Access Engine. This displays the system configuration parameters for this component.
4. Locate and modify the system configuration parameters you want to change.
5. Select the Revert button to discard your changes or the Save button to save your changes.
6. Restart the Web Services Data Access Engine with the following command:

```
/etc/init.d/opsware-sas restart twist
```

Software Repository Mirroring Configuration Parameters

You can enable software repository mirroring and set how frequently the mirroring job runs by modifying the following configuration parameters. The software repository mirroring job copies data between software repositories so they are all in sync. For more information, see [Software Repository Monitoring](#).

Software Repository Mirroring Parameters

Parameter	Type	Allowed Values	Default	Description
<code>word.enable_content_mirroring</code>	Boolean Flag	0 or 1	0	Set this value to 1 to enable Software Repository mirroring. Set this value to 0 to disable it.
<code>word.mirror_job_period</code>	Minutes	Any positive integer	60	This value specifies how frequently the Software Repository mirroring job runs.

Monitoring SA Core Components

You will from time-to-time need to monitor the SA internal components for troubleshooting and adjusting component behavior.

Overview of SA Monitoring

SA provides system diagnostic tests in the SA Client to diagnose the functioning of the following SA components:

- Data Access Engine
- Software Repository
- Command Engine
- Web Services Data Access Engine
- Multimaster Infrastructure Components (referred to as the Model Repository Multimaster Component in the SA documentation)

This section provides information for performing basic monitoring of the components listed above and for the following additional SA components:

- Server Agent
- Agent Cache
- SA Client
- Model Repository
- Spoke
- Gateways
- OS Build Manager
- OS Boot Server
- OS Media Server

Use this information when the System Diagnosis tests cannot be used because the SA Client cannot be run or when your managed environment is already set up for automated monitoring. In that case, you can use these commands to automate your system diagnosis and to monitor SA.

This monitoring includes:

- Commands to confirm specific component processes are running, as well as examples of the expected output
- Commands provided by component and by operating system
- Component specific ports, logs, and administrative URLs

Note: The commands shown in this document must be entered all on one line. However, to make sure that the commands and the resulting output are readable, they might have been modified with spaces, blank lines, and line breaks, or backslashes (\) to indicate where a command has been continued on the following line. Also, the output shown is intended as an example only. The output on your servers will be different.

For a description of each of the SA components mentioned in this document, see the SA Overview and Architecture Guide.

Agent Monitoring

A Server Agent is a software module running on each server managed by SA. Whenever a change needs to be made to a managed server, the Server Agent brokers the requests.

For more information about the Server Agent, see the SA User Guide: Server Automation.

To use the SA Client to test an SA Core's communication with a Server Agent running on a managed server, see the following sections in the SA User Guide: Server Automation:

- Agent Reachability Communication Tests
- Communication Test Troubleshooting

Agent Port

The Server Agent uses port 1002.

Monitoring Processes for Agents

On **Windows**, from the **Start** menu, choose **Run**. In the Run dialog, enter `taskmgr`. In the Windows Task Manager dialog, click the Process tab and look for the processes called `watchdog.exe` and `python.exe`.

On **UNIX (Solaris, Linux, AIX, and HP-UX)**, the Server Agent has two running processes.

On **Solaris**, execute the command:

```
# ps -flg `awk -F= '($1=="pgrp") {print $2}' /var/opt/opsware/agent/daemonbot.pid`
```

Running this command should produce output similar to the following:

```
F S  UID  PID  PPID  C  PRI
NI  ADDR  SZ  WCHAN  STIME  TTY  TIME  CMD
8 S  root 9541 9539  0  41  20  ?      1768 ?      Aug
08 ?   1:23 /opt
/opt/opsware/agent/bin/python /opt/opsware/agent/pylibs/shadowbot/daemonbot.pyc --conf /etc/opt/opsware/agent/agent.args
```

```
8 S root 9539 1 0 99 20 ? 398 ? Aug  
08 ? 0:00 /opt
```

```
/opsware/agent/bin/python /opt/  
t/opsware/agent/pylibs/shadowbot/daemonbot.pyc --conf /etc/  
opt/opsware/agent/agent.args
```

On Linux, execute the command:

```
# ps -flg `awk -F= '($1=="pgrp") {print $2}' /var/  
opt/opsware/agent/daemonbot.pid`
```

Running this command should produce output similar to the following output:

```
F S UID PID PPID C PRI  
NI ADDR SZ WCHAN STIME TTY TIME CMD  
1 S root 2538 1 0 85 0 - 3184 wait4 Sep11 ?  
00:00:00  
/opt/opsware/agent/bin/python /opt/  
t/opsware/agent/pylibs/shadowbot/daemonbot.pyc --conf /etc/  
opt/opsware/agent/agent.args  
5 S root 2539 2538 0 75 0 - 30890 schedu Sep11 ?  
00:02:56  
/opt/opsware/agent/bin/python /opt/  
t/opsware/agent/pylibs/shadowbot/daemonbot.pyc --conf /etc/  
opt/opsware/agent/agent.args
```

The daemon monitor is the process with a PPID of 1. The others are server or monitor threads.

On AIX, execute the command:

```
# ps -flg `awk -F= '($1=="pgrp") {print $2}' /var/  
opt/opsware/agent/daemonbot.pid`
```

Running this command should produce output similar to the following output:

```
F S UID PID PPID C PRI  
NI ADDR SZ WCHAN STIME TTY TIME CMD  
40001 A root 110600 168026 0 60 20 2000d018 16208 * Sep 05 -  
7:15 /opt/  
opsware/agent/bin/python /opt/  
t/opsware/agent/pylibs/shadowbot/daemonbot.pyc --conf /etc/  
opt/opsware/agent/agent.args  
40001 A root 168026 1 0 60 20 2000f25c 1352 Sep 05 -  
0:02 /opt/  
opsware/agent/bin/python /opt/  
t/opsware/agent/pylibs/shadowbot/daemonbot.pyc --conf /etc/  
opt/opsware/agent/agent.args
```

On **HP-UX**, execute the command:

```
# ps -flg `awk -F= '($1=="pgrp") {print $2}'  
/var/opt/opsware/agent/daemonbot.pid`
```

Running this command should produce output similar to the following output:

```
F S UID PID PPID C PRI NI ADDR SZ WCHAN STIME TTY  
TIME COMD  
1 R root 10009 1 0 152 20 437eb1c0 266 - Sep 22 ?  
0:00 /opt/  
  
opsware/agent/bin/python /opt-  
t/opsware/agent/pylibs/shadowbot/daemonbot.pyc --conf /etc/-  
opt/opsware/agent/agent.args  
1 R root 10010 10009 0 152 20 434fb440 2190 - Sep 22 ? 3:29  
/opt/  
  
opsware/agent/bin/python /opt-  
t/opsware/agent/pylibs/shadowbot/daemonbot.pyc --conf /etc/-  
opt/opsware/agent/agent.args
```

Agent Logs

The Server Agents create the following log files on managed servers.

Windows:

- %ProgramFiles%Common Files\opsware\log\agent\agent.log*
- %ProgramFiles%Common Files\opsware\log\agent\agent.err*

UNIX:

- /var/log/opsware/agent/agent.log*
- /var/log/opsware/agent/agent.err*

Conditions to monitor in the UNIX logs:

- Strings containing "Traceback"
- Strings containing "OpwareError"

Agent Cache Monitoring

The Agent Cache is a component that serves Server Agent installation files during the Agent deployment process. The Agent Cache component caches the most recent version of the SA Agent. When SA installs the agent on servers in order to manage them, it obtains the agent installation binary file from the Agent Cache component.

Monitoring Processes for the Agent Cache

In all configurations, the Agent Cache component has a single running process.

On **Solaris** or **Linux**, execute the command on the server running the Gateway (in an SA core and an Satellite):

```
# ps auxwww | grep -v grep | grep agentcache
```

Running this command should produce output similar to the following output:

```
root  22288  0.5  0.1  15920  4464  ?  S   19:55   0:08  /opt/opswa-  
re/bin/  
python /opt/opsware/agentcache/AgentCache.pyc -d /var/-  
opt/opsware/agent_installers -p 8081 -b
```

Agent Cache Logs

The Agent Cache logs are in the following files:

- /var/log/opsware/agentcache/agentcache.log
- /var/log/opsware/agentcache/agentcache.err

Conditions to monitor in the logs:

- Strings containing “Error downloading agent”
- Strings containing “Another process is listening on port”

Command Center Monitoring

The Command Center is a web-based user interface to SA. Use the SA Client to access the Command Center.

SA users connect to the Command Center component through an Apache HTTPS Proxy (installed by the HP BSA Installer with the Command Center component).

Command Center Ports

The HTTPS Proxy uses port 443 (HTTPS) and port 80 and directs connections to the Command Center component, which uses port 1031 (the Web Services port).

Monitoring Processes for the Command Center

On Linux, execute the command on the server running the Command Center component:

```
# ps -eaf | grep -v grep | grep java | grep occ
```

Running this command should produce output similar to the following output:


```
occ 17373 1 6 19:46 ? 00:02:35 /opt/opsware/j2sdk1.4.2_10/bin/
```

```
java -server -Xms256m -Xmx384m -XX:NewRatio=3 -Docc.home=/opt/opsware/occ -Docc.cfg.dir=/etc/opt/opsware/occ -Dopsware.deploy.urls=/opt/opsware/occ/deploy/ -Djboss.server.name=occ -Djboss.server.home.dir=/opt/opsware/occ/occ -Djboss.server.
```

Tip: To monitor the Command Center component, you can also set up an automatic monitoring process to send a URL query (using tools such as Wget) to the Command Center URL. If the Command Center component returns its login page, it indicates that both the Apache HTTPS Proxy and Command Center processes are functioning normally.

Command Center Logs

The Command Center does not generate its own logs. The Command Center uses the JBoss server, which writes to the following log files:

- /var/log/opsware/occ/server.log*
- /var/log/opsware/httpsProxy/*log*

Conditions to monitor in the logs:

- java.net.ConnectionException
- java.net.SocketException
- java.lang.NullPointerException

Data Access Engine Monitoring

The Data Access Engine simplifies interaction with various clients in SA, such as the Command Center, system data collection, and monitoring agents on servers.

Data Access Engine Port

The Data Access Engine uses port 1004 (HTTPS) externally and 1007 (the loopback interface) for SA components installed on the same server.

Multimaster Central Data Access Engine Port Forwarding

SQLnet traffic between the Multimaster Central Data Access Engine in a mesh and the Model Repositories in other SA Cores in the mesh is routed over the SA Gateway mesh.

The `tnsnames.ora` file on the server running the Multimaster Central Data Access Engine points to a specified port on each core-side Gateway in the other SA cores. The core-side Gateway in the core running the Multimaster Central Data Access Engine forwards the connection to the

core-side Gateway in each other core, which in turn forwards it to the Model Repositories in the other cores.

The port number on the core-side Gateway is calculated as `20000 + data_center_id`. For example, if the Multimaster Mesh has two facilities, Facility A (facility ID 1) and Facility B (facility ID 2), the Multimaster Central Data Access Engine in Facility A connects to port 20002 on the server running the Gateway to reach the Model Repository in Facility B.

For information about the Multimaster Central Data Access Engine, see [Multiple Data Access Engines](#).

For information about the Gateway mesh topology, see the SA Overview and Architecture Guide.

Monitoring Processes for the Data Access Engine

On **Linux**, execute the command on the server running the Data Access Engine component:

```
# ps auxwww | grep -v grep | grep spin | grep -v java
```

Running this command should produce output similar to the following output:

```
root 30202 0.0 0.0 13592 1500 ? S Sep11 0:01 /opt-  
t/opsware/bin/  
  
python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/-  
opt/opsware/spin/spin.args  
  
root 30204 1.3 0.6 154928 25316 ? S Sep11 411:15 /op-  
t/opsware/  
  
bin/python /opt/opsware/pylibs/shadowbot/daemonbot.pyc  
--conf /etc/opt/opsware/spin/spin.args  
  
root 30256 0.1 0.3 28500 13024 ? S Sep11 50:35 /op-  
t/opsware/  
  
bin/python /opt/opsware/spin/certgenmain.pyc --start  
--conf /etc/opt/opsware/spin/spin.args
```

Data Access Engine URLs

- `https://spin.<data_center>:1004`

To access the Data Access Engine (spin) UI, you need the browser certificate `browser.p12`.

`browser.p12` is located in:

`/var/opt/opsware/crypto/spin/`

on any Slice Component bundle host. Copy the file to your local machine and import `browser.p12` into your browser, following your browser import certificate instructions.

- `https://spin.<data_center-
>:1004/ObjectBrowser.py?cls=Account&id=0`

Accessing the second URL fails when the Model Repository component is not running.

- `https://spin.<data_center>:1004/sys/dbstatus.py`

Accessing this URL shows the database connection status in the HTML page. Your automatic monitoring system can use a regular expression to extract the number of active database connections.

Data Access Engine Logs

The Data Access Engine logs are in the following files:

- `/var/log/opsware/spin/spin.err*` (The main Data Access Engine error file)
- `/var/log/opsware/spin/spin.log*` (The main Data Access Engine log file)
- `/var/log/opsware/spin/spin_db.log`
- `/var/log/opsware/spin/daemonbot.out` (Output from the application server)

In a core with multiple Data Access Engines, each server running a Data Access Engines has a set of these log files.

Web Services Data Access Engine Monitoring

The Web Services Data Access Engine provides increased performance to other SA components.

The Web Services Data Access Engine component is installed as part of the Slice Component bundle.

Web Services Data Access Engine Port

The Web Services Data Access Engine uses port 1032.

The Command Center component communicate with the Web Services Data Access Engine on port 1026 (a private loopback port).

Monitoring Processes for the Web Services Data Access Engine

On **Linux**, execute the command on the server running the Command Center component and on the server running the Slice Component bundle:

```
# ps auxwww | grep -v grep | grep /opt/opsware/twist
```

Running this command should produce output similar to the following output:

```
twist 4039 0.2 11.3 2058528 458816 ? S Sep11 80:51 /opt/opsware/
```

```
j2sdk1.4.2_10/bin/java -server -Xms256m -Xmx1280m -XX:MaxPermSize=192m -Dorg.apache.commons.logging.Log=org.apache.commons.logging.impl.Jdk14Logger .....  
twist 4704 0.0 0.0 4236 1124 ? S Sep11 1:28 /bin/sh /opt/  
opsware/twist/watchdog.sh start 60'  
twist 4743 0.0 0.6 376224 27160 ? S Sep11 18:31 /opt/opsware/  
j2sdk1.4.2_10/bin/java -server -Xms16m -Xmx128m -Dtwist.port=1026 ..... -classpath /opt/opsware/j2sdk1.4.2_10/jre/.....
```

Web Services Data Access Engine URL

`https://occ.<data_center>:1032`

Web Services Data Access Engine Logs

The Web Services Data Access Engine logs are in the following files:

- `/var/log/opsware/twist/stdout.log*`
- `/var/log/opsware/twist/twist.log`
- `/var/log/opsware/twist/access.log`
- `/var/log/opsware/twist/server.log*` (Application level logging)
- `/var/log/opsware/twist/boot.log`
- `/var/log/opsware/twist/watchdog.log`

The `stdout.log` files contain `stdout` and `stderr` and logs the output of any `System.out.println()`, `System.err.println()` and `e.printStackTrace()` messages; however, only some of the exceptions will show up in these logs. The number of files and the size of each file can be configured via `twist.conf`. Additional logs are created when the specified maximum file size is reached. The `stdout.log` is the most recent, and `stdout.-log.1` through `stdout.log.5` are progressively older files. The file is also rotated on startup.

The `twist.log` file contains WebLogic-specific messages and WebLogic level exceptions. These files are rotated on startup. Monitor the `twist.log` files for exceptions that indicate when the Web Services Data Access Engine (Twist) component failed to start correctly. If errors are encountered during Model Repository (Truth) connection setup, errors are logged in the `twist.log` files; for example, you might see the following error message:

```
####<Oct 14, 2006 1:37:43 AM UTC> <Error> <JDBC> <localhost.localdomain> <twist> <main> <<WLS Kernel>> <> <BEA-001150>  
<Connection Pool "TruthPool" deployment failed with the following  
error:
```

```
<Specific message, such as Oracle error codes and tracebacks>
```

The `access.log` file contains access information in common log format. These files are rotated when the file reaches 5MB in size.

The `server.log` files contain application level exceptions and debug messages generated from the Web Services Data Access Engine. The `server.log` files will also contain errors resulting from Model Repository (Truth) connection setup problems. The debug messages are controlled by the log level set at the package or class level in the `twist.conf` file. The number of files and the size of each file can both be configured via `twist.conf`. The `server.log.0` is always the current file, while `server.log.9` is the oldest.

The `boot.log` file contains information on the initial stdout and stderr messages generated when the Web Services Data Access Engine starts. In addition, the `boot.log` file contains the output from `Kill -QUIT` commands.

The `watchdog.log` file records the status of the Web Services Data Access Engine once every minute.

Command Engine Monitoring

The Command Engine is the means by which distributed programs such as Server Agents run across many servers. Command Engine scripts are written in Python and run on the Command Engine server. Command Engine scripts can issue commands to Server Agents. These calls are delivered in a secure manner and are auditable by using data stored in the Model Repository.

Command Engine Port

The Command Engine uses port 1018.

Monitoring Processes for the Command Engine

On **Linux**, execute the command on the server running the Command Engine component:

```
# ps auxwww | egrep '(COMMAND$|waybot)' | grep -v grep
```

Running this command should produce output similar to the following output:

```
USER  PID  %CPU  %MEM  VSZ   RSS  TTY   STAT   START   TIME  COMMAND
root  412  0.0   0.0  13600  1472  ?    S      Sep11   0:00  /opt/opsware/bin/python /opt/opsware/pylibs/shadowbot/daemonbot.pyc
--conf /etc/opt/opsware/waybot/waybot.args
```

On Linux servers running kernel 2.4 or later, the Command Engine has one process.

Command Engine Logs

The Command Engine logs are in the following files:

- `/var/log/opsware/waybot/waybot.err*`
- `/var/log/opsware/waybot/waybot.log*`
- `/var/log/opsware/waybot/daemonbot.out*`

Software Repository Monitoring

The Software Repository, a component of the SA core, is where all software managed by SA is stored. The Software Repository is part of the SA Library. Each core has one or more software repositories. This section describes how to monitor the software repositories in your cores.

Software repository mirroring keeps the software repositories in a multimaster mesh in sync for redundancy and disaster recovery. For example, if you upload a software package to one core in the mesh, the software repository mirroring job will replicate that package to all the other software repositories in the mesh.

To enable or disable software repository mirroring or to change how frequently the software repository mirroring job runs, see [Changing Software Repository Mirroring Parameters](#).

Software Repository Ports

The Software Repository uses the following ports:

- 1003 (Encrypted)
- 1006 (Clear text)
- 1005 (Replicator administrative user interface)
- 5679 (Multimaster Software Repository)

Monitoring Processes for the Software Repository - Linux

To check the software repository processes on Linux, run the following command on the server running the Software Repository component:

```
#ps auxwww | grep -v grep | grep mm_wordbot
```

This command produces output similar to the following:

```
root 31006 0.0 0.0 13612 1492 ? S Sep11 0:00 /opt/
opsware/bin/
python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/-
opt/opsware/mm_wordbot/mm_wordbot.args
root 31007 0.0 0.1 103548 7688 ? S Sep11 7:33 /op-
t/opsware/bin/
python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/-
opt/opsware/mm_wordbot/mm_wordbot.args
root 31092 0.0 0.0 13608 1480 ? S Sep11 0:00 /op-
t/opsware/bin/
python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/-
opt/opsware/mm_wordbot/mm_wordbot-clear.args
```

```
root 31093 0.0 0.1 70172 6424 ? S Sep11 2:11 /opt/
t/opsware/bin/

python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/-
opt/opsware/mm_wordbot/mm_wordbot-clear.args
```

On Linux, the Software Repository has multiple running processes (most are threads), which are for the encrypted Software Repository and for the clear text Software Repository.

Software Repository Logs

The logs for the Software Repository are in the following files:

- /var/log/opsware/mm_wordbot/wordbot.err*
- /var/log/opsware/mm_wordbot/wordbot.log*
- /var/log/opsware/mm_wordbot-clear/wordbot-clear.err*
- /var/log/opsware/mm_wordbot-clear/wordbot-clear.log*

Software Repository Mirroring - SA Client

Software repository mirroring keeps all your software repositories in sync for redundancy and diSAter recovery. If one software repository fails, the other software repositories can continue servicing requests for software. To enable software repository mirroring, see [Changing Software Repository Mirroring Parameters](#).

If you have Software Repository mirroring enabled, you can view and monitor the status of software repository mirroring as follows:

1. Log in to the SA Client as a user with the Multimaster Tools permissions. For more information on permissions, see [Permissions Reference](#).
 2. Select the Administration tab.
 3. Select Software Repository Mirroring in the navigation panel. This displays the status of software repository mirroring in your multimaster mesh. The information displayed includes:
 - **Number of Files in the Mesh:** This is the total number of files in each fully synced software repository.
 - **Total Disk Space Used:** This is the approximate total disk space required by a fully synchronized software repository.
 - **Status:** Shows which software repositories have all needed files (green), which need files (yellow), and which have mirroring disabled (grey).
- **Green:** All needed files are present in the facility's software repository. The number of missing files is zero.

- **Yellow:** One or more files are missing from the facility's software repository and need to be updated. These facilities will be updated when the mirroring job next runs. The mirroring job runs periodically as defined by the mirroring job run period.
 - **Grey:** Software repository mirroring is disabled in the facility.
-
- **Facility:** Shows the SA facility in which the software repository is running.
 - **Files:** The number of files currently in the host's Software Repository.
 - **Size:** The approximate total disk space currently used by the Software Repository files.
 - **Missing:** The number of files that need to be mirrored by the facility's Software Repository but that have not yet been replicated.

To change how frequently the software repository mirroring job runs, see [Changing Software Repository Mirroring Parameters](#).

Figure 36 shows the Software Repository Mirroring status with three SA cores named Bangalore, London, and New York. A software package was uploaded to the London core. The yellow status indicators show that Bangalore and New York cores are out of sync—the software package has not been replicated to those two cores yet.

Figure 36. Software Repository Mirroring Status—Out of Sync

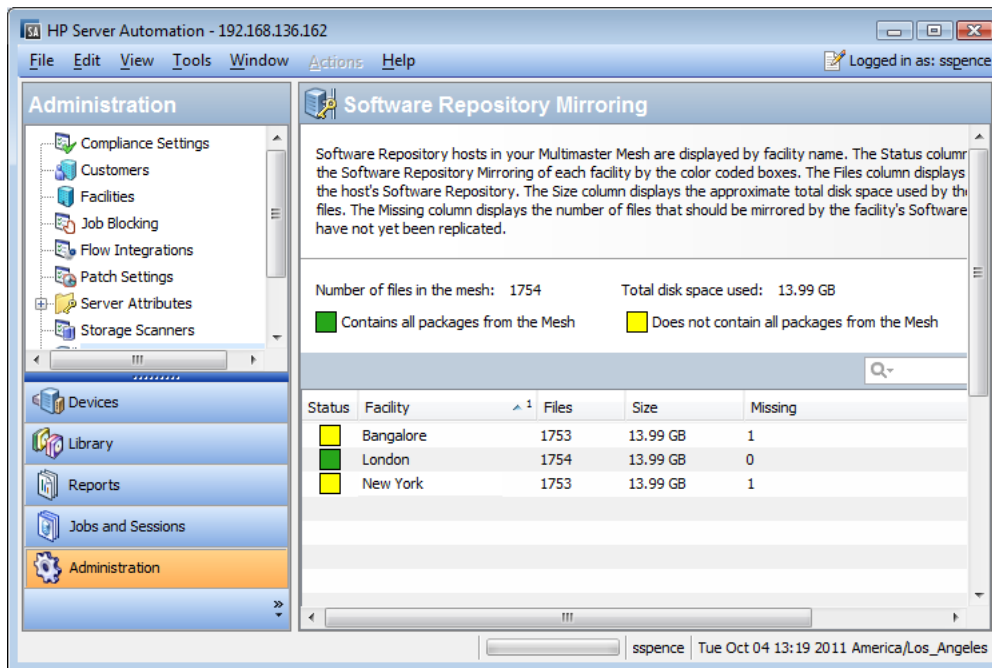
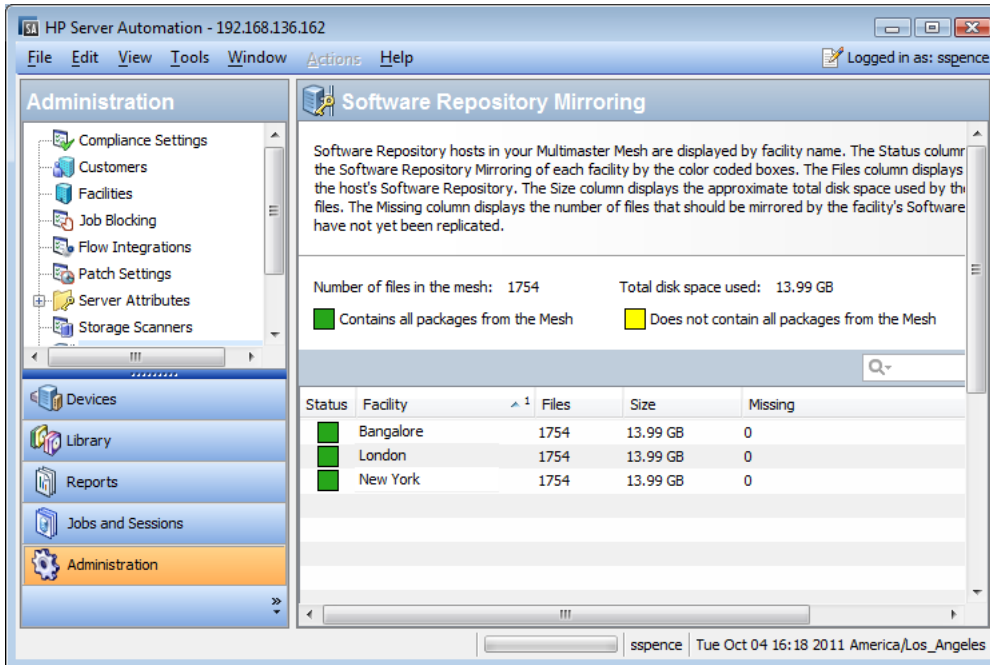


Figure 37 shows the Software Repository Mirroring state after the mirroring job has run and replicated the software package to all cores. The green status indicators show that all cores are in sync.

Figure 37. Software Repository Mirroring Status - In Sync



Model Repository Monitoring

The Model Repository is an Oracle database that contains essential information necessary to build, operate, and maintain a list of all managed servers, their hardware, their configuration, the operating system, and all other applications.

For more information about the Model Repository, including detailed information about monitoring the Model Repository, see “Appendix A: Oracle Setup for the Model Repository” in the SA Installation Guide.

Model Repository Port

The default port for the Model Repository is 1521; however, this might have been modified by the database administrator who installed it.

Monitoring Processes for the Model Repository

Monitor the Oracle Database process. If the process is not found, the database has failed or was not started.

On Linux, execute the command on the server running Oracle:

```
# ps -fu oracle | grep pmon
```

Running this command should produce output similar to the following:

```
oracle 2112 1 0 21:22 ? 00:00:00 ora_pmon_truth
```

(The process name might include the database SID, truth, as shown in this example.)

If the process is not found, the listener has failed or was not started.

On Linux, use this command to monitor the Oracle Listener process:

```
# ps -fu oracle | grep tnslnsr
```

Running this command should produce output similar to the following:

```
oracle 2021 1 0 21:22 ? 00:00:01  
/u01/app/oracle/product/11.2.0/db_2/bin/tnslnsr LISTENER -  
inherit
```

Model Repository Logs

Log files for the Model Repository are produced by the Oracle database, and their location is specific to your installation.

By default, SA uses a directory for each SID (in this case truth) for the Model Repository logs. (This could be different based on how Oracle was installed.)

```
/u01/app/oracle/admin/truth/bdump/alter_truth.log
```

Conditions to monitor:

Not all errors indicate a problem with the database. Some errors might be caused by an application.

In these examples, there is a problem if the command has output.

```
grep ORA- /u01/app/oracle/admin/truth/bdump/alter_truth.log  
ORA-00600: internal error code, arguments: [729], [480],  
[space leak], [], [], [], [], []  
ORA-07445: exception encountered: core dump [lxmcpn()+0]  
[SIGSEGV] [Address not mapped to object] ...
```

Table Space Usage

Tablespace usage should be monitored against a threshold, usually increasing in severity (for example., over 80% is a warning, over 90% is an error, over 95% is a critical error).

There are several ways to monitor tablespace usage. For a SQL query that you can run to check for sufficient free disk space in the tablespaces, see “Appendix A: Oracle Setup for the Model Repository” in the SA Installation Guide. The SQL query provided in the installation guide must be executed as a privileged database user.

Multimaster Conflicts

The number of conflicting transactions in any Model Repository can be found by running the following SQL query as any SA database user.

```
select count(*) from transaction_conflicts where resolved = 'N';
```

Multimaster conflicts should be monitored in stages, with increasing numbers of conflicts resulting in increasing levels of escalation. The values used for the stages depend on patterns of use.

The SA administrator should record the number of conflicts for some period of time (perhaps a week) and use that information to determine the level of alert raised by the monitoring system.

Model Repository Multimaster Component Monitoring

The Model Repository Multimaster Component is a Java program responsible for keeping multiple Model Repositories synchronized and propagating changes for the originating Model Repository to all other Model Repository databases.

Model Repository Multimaster Component Port

The Model Repository Multimaster Component uses port 5678.

Monitoring Processes for the Model Repository Multimaster Component

On **Linux**, execute the command on the server where you installed the Infrastructure Component bundle:

```
# ps auxwww | grep -v grep | grep vault | grep -v twist
```

Running this command produces output similar to the following:

```
root 28662 0.0 0.0 2284 532 ? S Sep27 0:00 /opt/
opsware//bin/
python /opt/opsware//pylibs/shadowbot/etc/daemonizer.pyc
--runpath /var/opt/opsware/vault --cmd /opt/opsware/j2sdk1.4.2_
10/bin/java -classpath /op-
t/opsware/vault/classes:/opt/opsware/vault ..... -ms120m -
mx1024m
-DCONF=/etc/opt/opsware/vault/
-DHOSTNAME=m234.dev.opsware.com com.loudcloud.vault.Vault

root 28663 0.0 6.3 1285800 130896 ? S Sep27 5:32 /op-
t/opsware/
j2sdk1.4.2_10/bin/java -classpath /op-
t/opsware/vault/classes:/opt/opsware/vault ..... -ms120m -
mx1024m
-DCONF=/etc/opt/opsware/vault/
-DHOSTNAME=m234.dev.opsware.com com.loudcloud.vault.Vault
```

Model Repository Multimaster Component Logs

The Model Repository Multimaster Component logs are in the following files:

- `/var/log/opsware/vault/vault.n.log`

To configure the log file name, log file size, or logging level, perform the following steps.

1. Select the **Administration** tab in the SA Client.
2. In the navigation pane, select **System Configuration > Configuration Parameters**. This displays the SA components, facilities, and realms that have system configuration parameters.
3. In the list of SA components, select Model Repository, Multimaster Component. This displays the system configurations for that component.
4. Locate and modify the `log`, `logLevel` or `logsize` configuration parameters, as needed.
5. Select the Revert button to discard your changes or the Save button to save your changes.

Global File System Monitoring

The Global Shell feature is installed as part of any Slice Component bundle. It dynamically constructs the Global File System (OGFS) virtual file system.

The Global Shell can connect to an Server Agent to open a UNIX shell or a Windows Remote Desktop connection on a managed server.

For information about using the Global Shell, see the Global Shell chapter and appendices in the SA User Guide: Server Automation.

The Global File System component consists of the following programs:

- **Hub**: A Java program that interacts with other Core Components and Agents on Managed Servers (through the Agent Proxy) to compose the file system view.
- **Adapter**: On Linux, a C program that transports file system requests and replies between the FUSE (a module in the kernel) and the Hub and uses the FUSE userspace library to communicate with the FUSE kernel module.
- **Agent Proxy**: A Python program that provides the Hub with SSL connectivity to Agents running on managed servers.
- **FUSE (Linux Only)**: A file system in Userspace (FUSE) (software governed by the GNU GPL license) that provides in-kernel dispatch of file system requests into the Adapter.

The process group ID file for the Hub is located in the following directory:

- `/var/opt/opsware/hub/hub.pgrp`

All Global File System programs (Hub, Adapter, Agent Proxy, and their log rotators) run in this process group.

Monitoring Process for the Global File System

On Solaris, execute the command on the server(s) running the Slice Component bundle:

```
# ptree $(ps -g $(cat /var/opt/opsware/hub/hub.pgrp) -o pid=)
```

Running this command produces output similar to the following:

```
7594 /opt/opsware/bin/python /opt/opsware/hub/bin/rotator.py  
/opt/
```

```
opsware/j2sdk1.4.2.....
```

```
7598 /opt/opsware/j2sdk1.4.2_10/bin/java -server -Xms64m -  
Xmx1024m
```

```
-Dhub.kernel=SunO.....
```

```
7613 /opt/opsware/bin/python /op-  
t/opsware/adapter/SunOS/bin/rotator.py
```

```
/opt/opsware/.....
```

```
7617 /opt/opsware/ogfsutils/bin/python2.4 /op-  
t/opsware/adapter/
```

```
SunOS/lib/adapter.py.....
```

```
7618 /opt/opsware/adapter/SunOS/bin/mount -o hostpath=  
/hostpath,nosuid /dev/ogdrv /v.....
```

```
7619 /opt/opsware/bin/python /op-  
t/opsware/agentproxy/bin/rotator.pyc
```

```
/opt/opsware/bi.....
```

```
7625 /opt/opsware/bin/python /opt/opsware/agentproxy/lib/  
main.pyc.....
```

On Solaris, the OGFS (specifically, the programs Hub, Adapter, and Agent Proxy) has seven running processes.

On Linux, execute the following command on the server running the Slice Component bundle.

```
# ps u -g $(cat /var/opt/opsware/hub/hub.pgrp)
```

Running this command produces output similar to the following:

```
USER PID %CPU %MEM VSZ RSS TTY STAT START TIME COMMAND
```

```
root 8862 0.0 0.0 2436 1356 ? S Sep29 0:00 /op-  
t/opsware/bin/python /opt/opsware/hub/bin/rotator.py /op-  
t/opsware/j2sdk1.4.2_10/b.....
```

```
root 8868 0.1 1.8 1256536 76672 ? S Sep29 35:51 /opt/
t/opsware/j2sdk1.4.2_
10/bin/java -server -Xms64m -Xmx1024m -Dhub.kernel=Linux -
Dh.....
root 8906 0.0 0.0 2412 1304 ? S Sep29 0:28 /op-
t/opsware/bin/python /opt/
opsware/adapter/bin/adapter.....
root 8908 0.0 0.0 13088 684 ? S Sep29 0:10 /op-
t/opsware/adapter/Linux/
bin/adapter.bin /var/opt/opsware/ogfs/mnt/ogfs -f -o none.....
root 8913 0.0 0.0 2308 1132 ? S Sep29 0:00 /op-
t/opsware/bin/python /opt/
opsware/agentproxy/bin/rotator.pyc /opt/opsware/bin/pyt.....
root 8923 0.0 0.1 153120 6544 ? S Sep29 5:56 /op-
t/opsware/bin/python
/opt/opsware/agentproxy/lib/main.pyc.....
```

On Linux, OGFS (specifically, the programs Hub, Adapter, and Agent Proxy) has six running processes.

The Global File System also supports a `status` option to the `init` script for both Linux and Solaris.

On Linux or Solaris, execute the following command on the server running the Slice Component bundle to run this `status` option:

```
# /etc/opt/opsware/startup/hub status
```

Running this command produces output similar to the following:

```
Testing for presence of Hub process group file (/var/-
opt/opsware/hub/hub.pgrp) ... OK
Testing that processes are running in Hub process group (8862)
... OK
Testing that OGFS is mounted ... OK
Testing that the OGFS authenticate file is present ... OK
OGFS is running
```

Global File System Logs

The Hub logs are in the following files:

- /var/log/opsware/hub/hub.log*
- /var/log/opsware/hub/hub.out*

Conditions to monitor in the Hub logs:

- Strings containing ““Can’t establish twist connection””

The Adapter logs are in the following files:

- /var/log/opsware/adapter/adapter.err*

The Agent Proxy logs are in the following files:

- /var/log/opsware/agentproxy/agentproxy.err*

Monitoring Processes for FUSE (Linux Only)

On Linux, execute the command on the server running the Slice Component bundle:

```
# lsmod | grep -v grep | grep fuse
```

Running this command produces output similar to the following:

```
fuse          31196 2
```

FUSE logs messages in the following file:

- /var/log/messages

Monitoring Processes for the SunOS Kernel Module

On Solaris, the OGFS functionality relies on the SunOS kernel module.

Execute the command on the server running the Slice Component bundle:

```
# modinfo | grep -i opsware
```

Running this command produces output similar to the following:

```
137 1322cd8 43a9 272 1 ogdrv (Opware GFS driver v1.13)  
138 13ac227 338df 18 1 ogfs (Opware Global Filesystem v1.14)
```

The Global File System logs messages related to SunOS kernel module in the following file:

- /var/adm/messages

Spoke Monitoring

The Spoke is the back-end component of the SA Client. The Spoke, a Java RMI server, provides access to the files in the OGFS and provides access to run commands inside an OGFS session.

Spoke Ports

The Spoke uses port 8020.

Monitoring Processes for the Spoke

On Linux, execute the command on the server running the Slice Component bundle:

```
# ps -ef | grep -v grep | grep spoke
```

Running this command produces output similar to the following:

```
root 29191 1 0 Aug28 ? 01:12:11 /opt/opsware/j2sdk1.4.2_
10/bin/
java -server -Xms32m -Xmx256m -Dbea.home=/opt/opsware/spoke/etc -
Dspoke.home=/opt/opsware/spoke
-Dspoke.cryptodir=/var/opt/opsware/crypto/spoke
-Dspoke.logdir=/var/log/opsware/spoke
-Djava.util.logging.config.file=/opt/opsware/spoke/etc/logg
```

On Linux, the Spoke component has a single, running Java process.

Spoke Logs

The Spoke logs are in the following files:

- /var/log/opsware/spoke/spoke-*.log
- /var/log/opsware/spoke/stdout.log

Gateway Monitoring

SA Management and Core Gateways allow an SA Core to manage servers that are behind one or more NAT devices or firewalls. Connectivity between gateways is maintained by routing messages over persistent TCP tunnels between the gateway instances.

For information about configuring the Gateways, the SA Overview and Architecture Guide.

For information about maintaining Satellite Gateways, see [Satellite Administration](#).

Gateway Ports

By default, the Gateway uses the following ports:

- 2001—Management Gateway Listener Port
- 2001—Slice Component Core Gateway Listener Port
- 3001—Agent Gateway Port
- 3001—Satellite Gateway Port

Monitoring Processes for the Gateway

In all configurations, the Gateway component has two running process—the Gateway process itself and its watchdog process.

On Solaris or Linux, execute the commands on the server running the Gateway component:

```
# ps -eaf | grep -v grep | grep opswgw | grep cgw
```

Running this command produces output similar to the following:

```
root 17092 1 0 Sep21 ? 00:00:00 [opswgw-watchdog-2.1.1:
cgw0-C43]
--PropertiesFile /etc/opt/opsware/opswgw-cgw0-C43/opswg-
w.properties --BinPath /opt/opsware/opswgw/bin/opswgw
root 17094 17092 0 Sep21 ? 02:23:21 [opswgw-gateway-2.1.1:
cgw0-
C43] --PropertiesFile /etc/opt/opsware/opswgw-cgw0-C43/opswg-
w.properties --BinPath /opt/opsware/opswgw/bin/opswgw --Child
true
# ps -eaf | grep -v grep | grep opswgw | grep agw
```

Running this command produces output similar to the following:

```
root 17207 1 0 Sep21 ? 00:00:00 [opswgw-watchdog-2.1.1:
agw0-C43]
--PropertiesFile /etc/opt/opsware/opswgw-agw0-C43/opswg-
w.properties --BinPath /opt/opsware/opswgw/bin/opswgw
root 17208 17207 0 Sep21 ? 01:18:54 [opswgw-gateway-2.1.1:
agw0-
C43] --PropertiesFile /etc/opt/opsware/opswgw-agw0-C43/opswg-
w.properties --BinPath /opt/opsware/opswgw/bin/opswgw --Child
true
```

In a Satellite facility on Solaris or Linux, execute the command on the server running the Satellite Gateway component:

```
# ps -eaf | grep -v grep | grep opswgw | grep <gateway-name>
```

Where <gateway-name> in this example is Sat1.

Running this command produces output similar to the following:

```
root 17092 1 0 Sep21 ? 00:00:00 [opswgw-watchdog-2.1.1:
Sat1]
--PropertiesFile /etc/opt/opsware/opswgw-Sat1/opswgw.properties -
-BinPath /opt/opsware/opswgw/bin/opswgw
root 17094 17092 0 Sep21 ? 02:23:21 [opswgw-gateway-2.1.1:
Sat1]
--PropertiesFile /etc/opt/opsware/opswgw-Sat1/opswgw.properties -
-BinPath /opt/opsware/opswgw/bin/opswgw --Child true
```

Gateway URL

Log into the SA Client UI and select **Gateway** under **Administration** in the navigation panel.

`https://occ.<data_center>/com.opsware.occ.gwadmin/index.jsp`

Gateway Logs

The Gateway logs are in the following files:

- `/var/log/opsware/gateway-name/opswgw.log*`

Conditions to monitor in the logs:

- Strings containing “ERROR”
- Strings containing “FATAL” (indicates that the process will end soon)

OS Build Manager Monitoring

The OS Build Manager component facilitates communications between OS Build Agents and the Command Engine. It accepts OS provisioning commands from the Command Engine, and it provides a runtime environment for the platform-specific build scripts to perform the OS provisioning procedures.

OS Build Manager Ports

The OS Build Manager uses the following ports:

- 1012 (HTTPS)
- 1017 (SA Build Agent)

Monitoring Processes for the OS Build Manager

In all configurations, the OS Build Manager component has a single running process.

On Linux, execute the command on the server running the OS Build Manager component:

```
# ps -eaf | grep -v grep | grep buildmgr
```

Running this command produces output similar to the following:

```
root 2174 1 0 Sep27 ? 00:13:54 /opt/opsware/j2sdk1.4.2_10/bin/  
java -Xmx256m -Dbuildmgr -  
Djava.security.properties=/opt/opsware/buildmgr/etc/java.secur  
ity -DDEBUG -DDEBUG_VERBOSE=1 -DLOG_OPTIONS=tTN -DLOG_FILE_  
THRESHOLD=10485760 -DLOG_FILE_RETAIN_COUNT=7 -DLOG_  
CLASSES=com.opsware.buildmgr.OutputStreamLo
```

OS Build Manager URL

`https://buildmgr.<data_center>:1012`

The OS Build Manager UI is read-only and port 1012 for the UI is configurable.

OS Build Manager Logs

The OS Build Manager logs are in the following files:

- `/var/log/opsware/buildmgr/buildmgr.log` (Build Agent activities, OS provisioning activities)
- `/var/log/opsware/buildmgr/*.request.log` (Web Server log; one file per day; 90 logs maximum)
- `/var/log/opsware/buildmgr/console.log`
- `/var/log/opsware/buildmgr/servers/<IP_address or machine_ID or MAC_address>` (A per connection log)

Conditions to monitor in the logs: the string “Traceback”

OS Boot Server Monitoring

The OS Boot Server, part of the OS Provisioning feature, supports network booting of Sun and x86 systems with inetboot and PXE, respectively. The process used to provide this support is the Internet Software Consortium DHCP server.

These applications are installed by the SA Installer but are not specific to SA. Monitor them by using standard system administration best practices for these applications.

OS Boot Server Ports

The OS Boot Server uses the following ports:

- 67 (UDP) (DHCP service)
- 69 (UDP) (TFTP service)

OS Boot Server Logs

The OS Boot Server does not generate its own logs. The OS Boot Server uses these services: TFTP with INETD, NFS server, and ISC DHCPD. All of these services log with syslog. Consult your vendor documentation for more information. See also the `syslog.conf` file that was used to configure the OS Boot Server to determine how the logging has been configured for this component.

OS Media Server Monitoring

The OS Media Server, part of the OS Provisioning feature, is responsible for providing network access to the vendor-supplied media used during OS provisioning. The processes used to provide this support include the Samba SMB server and Sun Solaris NFS.

These applications are installed by the HP BSA Installer but are not specific to SA. Specifically, SA provides a Samba package for Linux and Solaris that customers can use to install the OS Media Server. NFS services are provided by the operating system. Using the HP BSA Installer to install the OS Media Server configures NFS on Linux and Solaris.

Monitor the Samba SMB server and Sun Solaris NFS applications by using standard system administration best practices for these applications.

OS Media Server Ports

The OS Media Server uses the following ports:

- The portmapper used by NFS is port 111.
- Samba SMB uses ports 137, 138, 139, and 445.

OS Media Server Logs

The OS Media Server logs are in the following files:

- `/var/log/opsware/samba/log.smbd`
- `/var/log/opsware/samba/log.nmbd`

Solaris and Linux OS provisioning use of vendor-provided services such as NFSD. These services typically log through syslog. Consult your vendor documentation for more information on these log files.

Troubleshooting SA - Diagnostic Tests

This section describes:

- The **Core Health Check Monitor** that checks the health of individual SA components. See [Core Health Check Monitor \(HCM\)](#).
- The **System Diagnosis** tool that checks the overall health of the SA core. See [Running a System Diagnosis](#).

You can use these tools to diagnose the following types of problems you may encounter while maintaining SA:

- **Operational problems:** processes failing or becoming unresponsive (for example, the Data Access Engine, Command Engine, or Software Repository)
- **SA Core Component Failure:** which causes other components to fail.

The following examples describe the effects of some core component failures:

- If the Data Access Engine fails, the SA Client, the Command Engine, and the Software Repository components will fail.
- If the Software Repository fails to contact the Data Access Engine, downloads from the Software Repository are impossible.
- If the Model Repository fails, the Data Access Engine fails.
- If the Software Repository has neither a functioning DNS, nor a properly-configured /etc/hosts file, it fails to contact the Data Access Engine.
- If unreachable servers exist in the managed environment, communication is disrupted.

Note: System diagnosis can only be run on one facility at a time.

SA Core Component Internal Names

For legacy reasons, certain SA Core Components are referred to in this documentation using internal naming. **Table 27** shows the internal and external names of SA components.

Table 27. Internal and External Component Names

Internal Name	External Name
agentcache	A component of the Global File System
buildmgr	OS Provisioning Build Manager
hub	A component of the Global File System
mm_wordbot	A component of the Software Repository
occ	SA Command Center
opswgw-agw0	Agent Gateway
opswgw-mgws0	Master Gateway
spin	Data Access Engine
spoke	A component of the Global File System
truth	Model Repository
twist	Web Services Data Access Engine
vault/vaultdaemon	Model Repository Multimaster Component
way/waybot	Command Engine
word	Software Repository

Core Health Check Monitor (HCM)

The Health Check Monitor (HCM) includes a suite of tests to check the status of an SA core. The scripts in the HCM are installed by the SA Installer. There is some functional overlap between HCM and the System Diagnosis Tool described in [System Diagnostic Tests](#).

HCM provides two types of tests:

- **Local Tests:** Validate the health of a core on a component-by-component basis.
- **Global Tests:** Validate the health of a core on a holistic basis.

Overview of HCM Local Tests

The HCM local tests validate *individual core components*. The local tests reside on the same server as the components they validate. Run local tests by running the SA Start script (`/etc/init.d/opsware-sas`) and specifying a test mode argument and optional component names.

The test mode specifies the set of tests to run (you cannot specify individual tests.) Each test is run only once, even if you specify multiple components that require the same test. The test results are displayed on `stdout`.

Note: You cannot run the Health Check Monitor from a Satellite host.

Syntax of the Script for HCM Local Tests

HCM local tests use the following syntax:

```
/etc/init.d/opsware-sas <mode> [<component>[<component>...]]  
[<name>=<value>[<name>=<value>]...]
```

Running HCM Local Tests

To run the local tests, perform the following steps:

1. Log on as `root` to the server running the SA core components that you want to test.
2. Run the SA start up script using the `status` argument or specify the `mode` (test category) argument and one or more components (see the next section for the command options). For example, the following verifies that the Web Services Data Access Engine is available.

```
/etc/init.d/opsware-sas status twist
```

Table 28 describes the HCM command-line arguments. For a description of the `opsware-SA` options for starting and stopping a core, see [Table 24. Options for the SA Start/Stop Script](#).

Table 28. Options for the HCM Local Test Script

Option	Description
<code>mode</code>	The set of tests to run. The <code>mode</code> can be one of the following strings: <code>status</code> : Runs tests that verify the availability of the specified components. For example, the tests verify that the components are listening on the correct ports and responding to basic queries. <code>verify_post</code> : Same as <code>status</code> . <code>verify_pre</code> : Runs tests that validate the conditions necessary for the specified components to operate. <code>verify_functionality</code> : Runs tests that are similar to the tests run by the <code>status</code> mode; however, they might take longer to run. Therefore, you might choose to skip these tests to save time. <code>health</code> : Runs the tests of the <code>status</code> , <code>verify_pre</code> , and <code>verify_functionality</code> modes and provides an overview of the overall state of the specified components.
<code>component</code>	The internal name of the core component. If this option is not spe-

Option	Description
	<p>cified, then all components are validated. To view the internal names of the components installed on the local server, enter the following command:</p> <pre data-bbox="553 443 1052 474">/etc/init.d/opsware-sas list</pre>
name=value	<p>Options that control how the tests are run. Allowed values:</p> <p><code>terse=[true false]</code>: If <code>true</code>, summarizes the results of all successful tests for each component in a single <code>SUCCESS</code> message; however, the results of failed tests are displayed individually. By default, this option is set to <code>false</code>. (This option is passed to the individual tests.)</p> <p><code>parsable=[true false]</code>: If <code>true</code>, summarizes the results from all tests for each component with a single <code>SUCCESS</code> or <code>FAILURE</code> message. By default, this option is set to <code>false</code>. (This option is passed to the individual tests.)</p> <p><code>verify_filter=<regex></code>: Runs only the tests whose file names match the regular expression you enter. For example, specifying <code>verify_filter="OPSW"</code> runs only tests with file names that contain the string <code>OPSW</code>, such as <code>100_OPswcheck_host_spin.sh</code>. By default, this option is not defined. (This option is not passed to the individual tests.)</p> <p>If a given test is a symbolic link to another file, the filter will be evaluated against the target of the symbolic link, not the name of the symbolic link. If the test is a symbolic link, <code>verify_filter</code> uses the file name of the file it is pointing to for comparisons.</p>

Note: You can find a list of the internal name used for certain Core Components and their standard names in [SA Core Component Internal Names](#).

Overview of HCM Global Tests

A *global* HCM test checks an entire SA Core. Run these tests by executing the `run_all_probe-s.sh` script on the following hosts:

- **Sliced configuration**—the server hosting the core’s Management Gateway and/or Infrastructure Component (in a Typical Install, the Management Gateway is installed on the server that hosts the Infrastructure Component).
- **Non-sliced configuration**—the server hosting the Primary Model repository Multimaster Component for the core being validated.

Test results are displayed on `stdout`. The global tests cannot check the status of other cores in a multimaster mesh.

In a multiserver core, the global tests connect to the other core servers using SSH. All connections are made as `root`, or a `non-root` user with `sudo` permission. Authentication is performed by specifying the user password or the key file on the command line. If both are specified, then the user password is used. One of these authentication methods must be specified unless the server is the local host. Passwordless login with keyfile is supported *only* for the root user.

Running HCM Global Tests

To run the HCM global tests, perform the following steps:

1. Log in as `root` to the server that hosts the Model Repository Multimaster Component and/or the Infrastructure Component.
2. Execute the `run_all_probes.sh` script with the `run` option (see the following section for details on the options). For example, to check the table space usage in the Oracle database of the Model Repository, enter the following command:

```
/opt/opsware/oi_util/bin/run_all_probes.sh run \  
check_database_tables
```

Syntax of the Script for HCM Global Tests

The script that runs HCM global tests has the following syntax:

```
/opt/opsware/oi_util/bin/run_all_probes.sh run|list  
[<test> [<test>...]  
[hosts=" [<user>@]<system>[:<password>] [[<user@>]<system>[:<password>]]..."  
[keyfile=<keyfiletype>:<keyfile>[:<passphrase>]]
```

Table 29 describes the options for this syntax.

Table 29. Options for the HCM Global Test Script

Option	Description
<code>list</code>	Lists the available tests.
<code>run</code>	Runs the specified tests.
<code>test</code>	The name of the test to run. If no tests are specified, all tests are run. When shipped, the script includes the following tests: <ul style="list-style-type: none">• <code>check_opsware_services</code>: Runs the local tests on all specified servers by running the following command remotely on

Option	Description
	<p>each core server: <code>/etc/init.d/opsware-sas health</code></p> <ul style="list-style-type: none"> • <code>check_MM_state</code>: For a multimaster source core, checks the multimaster state of the core. • <code>check_time</code>: In a multiserver core, verifies that the system clocks are synchronized across core servers. • <code>check_opsware_version</code>: Validates that the versions of all the components in the core are the same version. • <code>check_database_tables</code>: Validates that the Model Repository tablespace usage is within acceptable limits. For more information on table spaces, see “Oracle Setup for Model Repository” in the SA Installation Guide. • <code>check_OS_resources</code>: Validates whether the virtual memory and disk space on SA partitions is within acceptable thresholds. • <code>check_fully_functional</code>: Validates full functionality of all SA components. For an alternative way to run System Diagnostics Comprehensive tests from the SA Client, see System Diagnostic Tests.
system	Name of a reachable SA core system.
keyfiletype	<p>Specifies the type of key file to use. Allowed values are:</p> <ul style="list-style-type: none"> • <code>rsa_key_file</code> • <code>dsa_key_file</code>.
keyfile	Specifies the file containing the current server’s SSH private key. Passwordless login with keyfile is supported <i>only</i> for the root user.
passphrase	Specifies the <code>passphrase</code> that was used to encrypt the SSH private key.
user	Optional user to access the remote system. The user needs to have sudo permission. Default is "root".
password	Optional <user> password for <system>.

Setting Up Passwordless SSH for Global Tests

The global tests access remote servers in a core through the SSH daemon. These tests require you to supply `user` passwords or to use SSH public/private keys. Passwordless login with keyfile is supported *only* for the root user.

To set up authentication using public/private keys generated by `ssh-keygen`, perform the following steps:

1. Run the following commands on the trusted server and accept the defaults. The commands are different for Linux and Solaris.

Linux:

```
cd /root/.ssh  
ssh-keygen -t dsa
```

Solaris:

```
cd /.ssh  
ssh-keygen -t dsa
```

2. Update the client server by copying the `id_dsa.pub` file to the client server's `.ssh` directory and then renaming it to `authorized_keys`. Here are some example commands for Linux and Solaris:

Linux:

```
scp id_dsa.pub <host>:/.ssh/authorized_keys  
/root/.ssh/authorized_keys
```

Solaris:

```
scp id_dsa.pub <host>:/.ssh/authorized_keys  
/.ssh/authorized_keys
```

3. Verify the trusted server. Run the following command to validate that the trusted server can connect to the client server without a password:

```
ssh -l root <host>
```

Extending the Health Check Monitor

This section is intended for advanced system administrators with experience in UNIX shell programming and SA administration.

The HCM is implemented as a series of UNIX shell scripts that perform local or global tests on the core servers. The scripts conform to specific naming conventions and reside in predefined directories. You can extend the HCM by writing your own scripts and copying them to the correct directories under `/opt/opsware/oi_util`.

Requirements for Extensions to HCM Local Tests

An HCM local test is a script that is run by the `/etc/init.d/opsware-sas` script (see [Running HCM Local Tests](#)). A local test script must meet the following requirements:

- **UNIX Shell Script:** It is a UNIX shell script that runs as `root`.
- **Component Server:** The script resides and runs on the server of the component validated by the script. For example, if the script validates the Data Access Engine (spin), it resides on the server that runs the Data Access Engine.
- **Executable:** The script is an executable file (`chmod u+x`).
- **File Name:** The file name of the script has the following syntax:

```
<int><test>.sh
```

In this syntax, `int` is an integer that specifies the test execution order and `test` is the name of the test. Note that the HCM scripts provided with SA contain `OPSW` in the script file name; for example, `100_OPSPWportping.sh`.

- **Directory:** The script resides in the following directory:

```
/opt/opsware/oi_util/local_probes/<component>/[verify_pre | verify_post | verify_functionality]/
```

In this path, `component` is the internal name of the core component, such as `spin` or `twist`. The directories beneath the `component` directory match the category of the test. For example, if the test performs a runtime validation on a core component, the script resides in the `verify_functionality` subdirectory. For details, see [Categories and Local Test Directories](#).

The directories beneath the `component` directory map to the `mode` options of the `/etc/init.d/opsware-sas` command. For example, if you save a script in the `verify_pre` subdirectory, the script is executed when you run `opsware-SA` with the `verify_pre` option. If you specify the `health` option of `opsware-SA`, the scripts in all three directories are executed. **Table 30** describes the mapping between the directory names and the mode options.

Table 30. Modes of `opsware-SA` and the Subdirectories of Local Test Scripts

Mode Option of Command Line	Subdirectory of Scripts Run for This Option
health	verify_pre verify_post verify_functionality
status	verify_post
verify_functionality	verify_functionality
verify_post	verify_post
verify_pre	verify_pre

- **Exit Code:** The script returns an exit code of zero to indicate success or nonzero for failure. The `/etc/init.d/opsware-sas` command uses the exit code to determine the status for the test.
- **Results Displayed:** The script displays test results on `stdout`.

- **Local Preamble Script:** The test script runs the `local_probe_preamble.sh` script, as shown by [HCM Local Test Example](#). The `local_probe_preamble.sh` script contains a superset of the libraries and shell variables used by the `/etc/init.d/opsware-sas` command.

The `local_probe_preamble.sh` script performs the following tasks:

- Sets shell variables used by the local tests. For example, it sets `$PYTHON` (which points to the Python interpreter) and `$UTILS_DIR` (which points to the directory of utilities available to the tests).
- Parses the command line, evaluates all `name=value` pairs, and sets shell variables. For example, if you specify `timeout=60` on the command line when running `/etc/init.d/opsware-sas`, the `local_probe_preamble.sh` script sets the variable `$timeout` to the value `60`.
- Provides access to useful functions such as `retry`, which executes a command multiple times until it succeeds or exceeds the specified timeout.
- **Shell Variables:** The test script takes into account the variables specified by the `name=value` options on the command line. For a list of predefined names, see the `name=value` option in [Table 28. Options for the HCM Local Test Script](#).

Categories and Local Test Directories

The `/opt/opsware/oi_util` directory has the following subdirectories.

local_probes/<component>/verify_pre

This directory includes prerequisite tests for each component. These tests validate that the necessary conditions exist for the component to operate. For example, the directory `twist/verify_pre` contains the test script `10check_localhost_spin.sh` because the Data Access Engine component must be available for the Web Services Data Access Engine component to function.

local_probes/<component>/verify_post

This directory includes validation tests for each component. These tests verify that a given component is available. For example, the directory `spin/verify_post` contains the test script `10check_primary_spin.sh` to validate that the Data Access Engine component is listening on port 1004 and responds to basic queries.

local_probes/<component>/verify_functionality

This directory includes runtime validation tests for each component. These tests verify that a component is fully operational. They are similar to `verify_post` tests; however, they might take longer to run. You might choose to skip these tests to save time.

Directory Layout for HCM Local Tests

The following directory layout shows where the local tests reside:

```
/opt/opsware/oi_util/  
|  
|_lib  
| |_local_probe_preamble.sh  
|  
|_local_probes  
|  
|_COMMON  
| |_<test>  
| |_ ...  
|  
|_<component>  
| |  
| |_verify_pre  
| | |_ <int><test> (can be symlink to ../../COMMON/<test>)  
| | |_ ...  
| |  
| |_verify_post  
| | |_ <int><test> (can be symlink to ../../COMMON/<test>)  
| | |_ ...  
| |  
| |_verify_functionality  
| |_<int><test> (can be symlink to ../../COMMON/<test>)  
| |_...  
|  
|_<component>  
...  
...
```

HCM Local Test Example

The following script verifies that the `cron` utility is running on the local server:

```
#!/bin/sh
# Verify that cron is running
# Read in our libraries / standard variable settings and parse
# the command line.
/opt/opsware/oi_util/lib/local_probe_preamble.sh
printf "Verify \"cron\" is running:"
process_running=`ps -eo fname | egrep '^cron$' | head -1`
if [ -z "$process_running" ]; then
echo "FAILURE (cron does not exist in the process table)"
exit 1
else
echo "SUCCESS"
exit 0
fi
```

Requirements for Extensions to HCM Global Tests

An HCM global test is a script invoked by the `run_global_probes.sh` command (see [Running HCM Global Tests](#)). A global test script must meet the following requirements:

- **UNIX Shell Script:** It is a UNIX shell script that runs as `root`.
- **Model Repository Server:** The script resides on the Model Repository Server, but it can run remotely on any core server.
- **Executable:** The script is an executable file (`chmod u+x`).
- **File Name:** The file name of the script has the following syntax:

```
<int><test>.sh[.remote]
```

In this syntax, `int` is an integer that specifies the test execution order and `test` is the name of the test specified on the command line. Note that the HCM scripts provided with SA contain `OPSW` in the script file name; for example, `300_OPswcheck_time.sh`.

- **Remote Execution:** If the test script runs on a core server other than those described in [Overview of HCM Global Tests](#), then the file name must have the `.remote` extension. When you execute `run_all_probes.sh` and specify such a test, the script is automatically copied to all specified servers and executed remotely with the SSH protocol.

The `.remote` file name extension is not required for tests that run on the same server as the Model Repository. Multimaster Component (in non-sliced installations) or the Management Gateway/Infrastructure Component (in Sliced installations). Examples of these tests are the checks for Model Repository integrity and multimaster conflicts. If the script

does not have the `.remote` extension and it needs to communicate with remote servers, the script must use SSH. The global preamble script includes helper functions for handling remote communications with SSH.

- **Directory:** The script resides in the following directory:

```
/opt/opsware/oi_util/global_probes/[verify_pre | verify_post  
]/
```

For details, see [HCM Global Test Directories](#).

- **Exit Code:** The script returns an exit code of zero to indicate success or nonzero for failure. The `run_global_probes.sh` command uses the exit code to determine the status for the test.
- **Results Displayed:** The script displays test results on `stdout`.
- **Global Preamble Script:** The test script runs the `global_probe_preamble.sh` script, as shown by [HCM Global Test Example](#). The `global_probe_preamble.sh` script contains a superset of the libraries and shell variables used by the HCM global tests.

The `global_probe_preamble.sh` script performs the following tasks:

- Sets shell variables used by the tests.
- Parses the command line and evaluates all `name=value` pairs, setting them as shell variables. For example, if you specify `hosts="sys1:pw1 sys2:pw2"` on the command line with `run_all_probes.sh`, the `global_probe_preamble.sh` script sets the variable `$hosts` to the value `"user1@sys1:pw1 user2@sys2:pw2"`.
- Provides access to the following functions:
 - `copy_and_run_on_multiple_hosts`: Copies and executes a shell script on multiple remote servers.
 - `copy_from_remote`: Copies a file from a remote server.
 - `copy_to_remote`: Copies a file to a remote server.
 - `run_on_multiple_hosts`: Runs an existing command on multiple servers.
 - `run_on_single_host`: Runs an existing command on a single server.
- **Shell Variables:** The test script takes into account the shell variables specified by the `name=value` options on the command line.
- **Authentication:** The script sets up authentication or public/private key generation. See [Setting Up Passwordless SSH for Global Tests](#).

HCM Global Test Example

The following script checks the free disk space of the file systems used by SA. This script runs on the core servers specified by the `hosts` option of the `run_all_probes.sh` command:

```
# Check for freespace percentage on Opsware SA filesystems  
# Read in our libraries, standard variable settings, and parse
```



```
# the command line.
/opt/opsware/oi_util/lib/global_probe_preamble.sh
MAX_PERCENTAGE=80
for filesystem in /opt/opsware /var/opt/opsware \
/var/log/opsware; do
# The leading and trailing spaces in the following printf
# are to improve readability.
printf " Checking $filesystem: "
percent_free=`df -k $filesystem 2> /dev/null | \
grep -v Filesystem | \
awk '{print $5}' | \
sed 's/%//'\`
if [ $percent_free -ge $MAX_PERCENTAGE ] ; then
echo "FAILURE (percent freespace > $MAX_PERCENTAGE)"
exit_code=1
else
echo "SUCCESS"
exit_code=0
fi
done
exit $exit_code
```

Directory Layout for HCM Global Tests

The following directory layout shows where the global tests reside:

```
/opt/opsware/oi_util/
|_bin
| |_run_all_probes.sh
| |_remote_host.py
| |_<support_utility>
| |_...
| |_lib
| |_global_probe_preamble
```

```
|  
|_global_probes  
|  
|_verify_pre  
| |_<int><probe>.remote  
|  
|_verify_post  
| |_int<probe>[.remote]  
|_ ...
```

HCM Global Test Directories

The `/opt/opsware/oi_util` directory has the following subdirectories:

global_probes/verify_pre

This directory includes tests that determine whether the specified servers are core servers. When a global test in this category determines that a server is not running an SA component or the server is unreachable, no further tests are run against that server.

Only tests with a `.remote` extension are allowed under the `verify_pre` directory.

global_probes/verify_post

This directory includes tests to determine the state of a specific aspect of the entire core. For example, the directory includes the `600_OPswcheck_OS_resources.sh.remote` script, which checks resources such as virtual memory and disk space.

Running a System Diagnosis

This describes how to run a set of system diagnosis. For details on each individual diagnostic test, see [System Diagnostic Tests](#).

To run system diagnostic tests, you must have the System Diagnosis action permission. For more information on permissions, see [Permissions Reference](#).

Before running the diagnostic tests, it is recommended that you run the Health Check Monitor first. For instructions, see [Core Health Check Monitor \(HCM\)](#), [Running HCM Local Tests](#), and [Running HCM Global Tests](#).

To run system diagnosis tests, perform the following steps:

1. In the SA Client, select the Administration tab in the navigation pane.
2. Select the Facilities node in the navigation pane. This displays all your SA facilities.
3. Select the facility where you want to run the diagnostics test.
4. Select the **Actions** menu or right-click and select **Run System Diagnosis**. This displays the Run Program Extensions window showing the System Diagnostics extension.
5. **Program Properties:** Select Next to display the Options window.
6. **Options:** Set the following options, then select Next. Or to accept the remaining defaults and run the tests, select Start Job.
 1. Verify or change the facility on which you want to run the diagnostic tests.
 2. Select the tests you want to run. For details on the tests, see [System Diagnostic Tests](#).
 3. Verify or set the job time out. If the job does not complete in the specified time, it will be aborted.
7. **Scheduling:** Select when you want the system diagnostics job to run, then select Next.
8. **Notifications:** Enter email addresses to receive notifications when the job finishes. Select the type of notifications you want. Optionally enter a ticket identifier to be associated with the job, then select Next.
9. **Job Status:** Select the Start Job or Schedule Job button. This runs the job or schedules the job to be run in the future and displays the Job ID number in the window banner. You can use the Job ID number to look up the job under the Jobs and Sessions tab.

When the job runs, it runs the diagnostic tests and displays the results.

10. Select any line in the job status to see the details of each diagnostic test that ran.
11. Press Ctrl-F to display the search bar.
12. Select Export All Results to create a file containing the results for further analysis. You can save the results as a zip file, a text file, or a comma-separated value file.

For details on each diagnostic test, see [System Diagnostic Tests](#).

System Diagnostic Tests

The System Diagnosis tool checks the functioning of the SA core components and the ability of managed servers to interact with the SA core. You can troubleshoot most of the errors that occur within the SA core with the SA diagnosis tool.

The System Diagnosis tool tests the SA core components first, and then, optionally, tests any servers in the managed environment that you specify. The System Diagnosis tool performs intensive testing of core components' functionality:

- **Standalone Tests:** Test as much of the functionality of a component as possible without the use of other SA components. Standalone Tests verify base level functionality and a component's ability to respond to XML-RPC calls.
- **Comprehensive Tests:** Test the full functionality of all core components.

Upon completion of Comprehensive Tests, the System Diagnosis tool displays the success or failure of each test, the test results, and error information for any tests that failed.

The core components are not tested in a specific order; however, the tests generally occur in this order:

- Component Standalone Tests
- Component Comprehensive Tests

Core Components Tested by the System Diagnosis Tool

The component tests simulate all the component functionality. In addition to errors, the tests verify that each component is functioning within certain conditions (for example, whether database connections are near maximum on the Data Access Engine).

The System Diagnosis tool tests the following components:

- Model Repository
- Data Access Engine
- Software Repository (and Word Store)
- Command Engine
- Server Agents on SA Core servers
- OS Build Manager
- Model Repository Multimaster Component
- Web Services Data Access Engine

Data Access Engine Tests

The following section describes the tests that occur during Data Access Engine diagnostic tests.

Standalone Tests

- Check for the current Data Access Engine version.
- Check for the current Model Repository database version.
- Verify that all Oracle objects are valid.
- Obtain a Device object.
- Obtain a MegaDevice object.
- Verifies advanced query functioning.
- Verify a Device object.
- Obtain the list of facilities.
- Obtain the names of the Data Access Engine cronbot jobs.
- Check whether the usage of database connections is below the acceptable level.

- Check whether any database connection has been open more than 600 seconds.
- Check whether the Data Access Engine and Model Repository are in the same facility.
- Verify that all Model Repository garbage-collectors are running when the Model Repository is running in multimaster mode.
- If the Data Access Engine is configured as the central multimaster Data Access Engine:
 - Check whether multimaster transactions are being published.
 - Check whether multimaster transactions are showing up at remote facilities.
 - Check for multimaster transaction conflicts.

Comprehensive Tests

- Test connectivity to the Model Repository on the configured port.
- Test connectivity to the Command Engine on the configured port.
- Test connectivity to the Software Repository on the configured port.

Errors Caused By Additional Database Privileges

If an additional privilege (permission) has been made manually to the Oracle database (Model Repository), the following error message might appear:

```
Test Results: The following tables differ between the Data Access Engine and the Model Repository: facilities.
```

To fix this problem, revoke the database grant. For instructions, see “Troubleshooting System Diagnosis Errors” in the SA Installation Guide.

Software Repository Tests

The following section describes the tests that occur during Software Repository diagnostic tests.

Standalone Tests

None.

Comprehensive Tests

- Test whether a file that is not a package can be uploaded to the Software Repository process that serves encrypted files. This test verifies whether the file is present in the Software Repository file system and that the file size matches the source.
- Verify that a file can be downloaded from the Software Repository.

- Verify whether the Software Repository process that serves unencrypted files is running and serving files.
- Try to download a file without encryption.
- Verify that a package can be uploaded to the Software Repository and that the package is registered with the Model Repository.
- Verify that a package can be deleted from the Software Repository and removed from the Model Repository.

Web Services Data Access Tests

The following section describes the tests that occur during Web Services Data Access diagnostic tests.

Standalone Tests

- Connect to the Web Services Data Access Engine and retrieve its version information.

Comprehensive Tests

- Connect to the Web Services Data Access Engine.
- Read a server record from the Model Repository and thereby check connectivity to the Model Repository.

Command Engine Tests

The following section describes the tests that occur during Command Engine diagnostic tests.

Standalone Tests

- Check the state machine.
- Check session tables.
- Check lock-down status.
- Check for signature failures.
- Check command and service tables.
- Check the facility cache.

Comprehensive Tests

- Check Data Access Engine connectivity.
- Check security signatures.
- Check lock operation.
- Run an internal script.
- Run an external script.

Model Repository Multimaster Component Tests

The following section describes the tests that occur during Model Repository Multimaster Component diagnostic tests.

Standalone Tests

- Check the ledger state by examining the ledger file.
- Report the total number of messages sent, number of messages still in the ledger file (for example, not confirmed by all listeners), and the sequence number of the last message confirmed by each listener.
- Check the sender health by examining the state of the Outbound Model Repository Multimaster Component.
- Check the receiver health by examining the state of the Inbound Model Repository Multimaster Component.

Comprehensive Tests

None.

Troubleshooting SA - Log Files

SA components record events in log files. One of the most valuable tools for troubleshooting SA problems is these component log files. Understanding SA components and how they log information can help you troubleshoot and resolve problems quickly. When you file a support request, HP Support may request you to send one or more log files or session data files.

This section describes log files, where they are located, and how you can use them for troubleshooting. It also describes how to create a session data file.

For a list of SA internal component names, see [SA Core Component Internal Names](#).

Viewing Log Files

To view a log file in a terminal window, log into the server running the component and use a command-line utility such as `more`, `less`, `grep`, or `vi`. See the following sections for locations of specific SA component log files.

Note: The log file for a component resides on the server where the component is installed.

Where Log Files Are Stored

Most SA log files are stored in `/var/log/opsware`. However, some components either log to their own directories (such as Oracle) or use syslog (such as NFS and DHCPD). **Table 31** lists SA components and their log directories. This information can help you determine which components or log files may be helpful in troubleshooting your particular problem.

Table 31. SA Log Files

Product Area	SA Component	Log File Directory
Database	Model Repository (truth or Oracle database)	Various directories under <code>/u01/app/oracle</code> , or as configured
Data Access, API	Data Access Engine (spin)	<code>/var/log/opsware/spin</code>
	Web Services Data	<code>/var/log/opsware/twist</code>

Product Area	SA Component	Log File Directory
	Access Engine (twist)	
Object Storage	Software Repository (word / wordcache)	/var/log/opsware/mm_wordbot
	Tsunami	/var/log/opsware/tsunami
	Memcached	/var/log/opsware/memcached
Job & Session Management	Command Engine (way)	/var/log/opsware/waybot
Global Shell, APX	Global File System, OGFS (hub)	/var/log/opsware/hub
	Global File System, OGFS (spoke)	/var/log/opsware/spoke
	APX Proxy	/var/log/opsware/apxproxy
	Other	/var/log/opsware/adapter /var/log/opsware/ogfs /var/log/opsware/agentproxy /var/log (opswsshd)
Mesh Communication	Agent Gateway	/var/log/opsware/opswgw-agwsN-FACILITY
	Core Gateway	/var/log/opsware/opswgw-cgwsN-FACILITY
	Management Gateway	/var/log/opsware/opswgw-mgwsN-FACILITY
Front-End	SA Web Client (occ)	/var/log/opsware/occ
	HTTPS Proxy	/var/log/opsware/httpsProxy
Mesh Replication	Model Repository Multimaster Component (vault/OMB)	/var/log/opsware/vault
OS Provisioning	Build Manager	/var/log/opsware/buildmgr
	DHCPD	/var/log, or as configured by syslog
	Samba	/var/log/samba
	NFS	/var/log, or as configured by syslog

Product Area	SA Component	Log File Directory
Agent Deployment	Agent Cache	/var/log/opsware/agentcache
Startup	SA Init Scripts	/var/log/opsware/startup
SA Agent	SA Agent	/var/log/opsware/agent

Product Areas and Related Component Log Files

Understanding the functional purpose of each component listed in **Table 31** can help you determine which components and logs to start with when troubleshooting. In many cases, the problem context including error messages or tracebacks can give you an idea of which logs to examine.

For example, when troubleshooting agent communication problems, a key step is to realize that one or more gateways are involved in all mesh communications and that if a gateway is down or not functioning properly, mesh communication will be impacted.

Table 32 lists SA product areas and log files to check when troubleshooting.

Table 32. Product Areas and Related Component Log Files

Product Area	Database Logs	Data Access Logs	Object Storage Logs	Job Mgmt Logs	Global Shell Logs	Mesh Comm Logs	Agent Logs
Agent Deployment	X	X	X		X	X	X
Audit and Compliance	X	X	X	X	X	X	X
Remediation for Software Management	X	X	X	X		X	X
Patching	X	X	X	X		X	X
Run Scripts	X	X		X	X	X	X
Application Configuration	X	X		X		X	X
OS Provisioning	X	X		X	X	X	X
Global Shell,	X	X			X	X	X

Product Area	Data-base Logs	Data Access Logs	Object Storage Logs	Job Mgmt Logs	Global Shell Logs	Mesh Comm Logs	Agent Logs
APX							
Ad hoc Device Management	X	X			X	X	X

About Log File Sizes

The default for the maximum log file size is 10 MB. When the specified maximum file size is reached, additional log files are created.

If you raise the log level for any components, the log files typically will grow significantly faster than the default log level. It is very important that you only raise the log level for a short period of time, long enough to gather log information about the problem you are troubleshooting, and then set the debug level back to the default value.

About Component Log Levels

By default, most SA components are configured to log-only errors and warnings. Temporarily raising the log level on individual components can reveal more detailed messages and help you understand what is going wrong with a particular component.

Raising the log level may cause additional overhead and performance loss, so do not keep the logging level raised for an extended period of time. Raise it only when actively diagnosing a problem, then restore it when you are finished.

Before changing log levels, save the original log level for easier reversion when you are finished. Back up the original configuration file prior to editing it, then restore it when you are finished.

Log levels typically follow a common format for naming:

- Trace
- Debug
- Info
- Warn or Warning
- Error
- Fatal
- Finest

Log-level naming can vary from component to component, but it mostly follows the standardized naming practices.

Changing Component Log Levels

This section discusses how to change logging levels for the various SA components that support it. Because multiple component instances may exist in a mesh, it may be necessary to perform these steps on multiple servers, such as SA slices or SA satellites.

Boot Server Logs

The Boot Server does not generate its own logs. The Boot Server uses these services: TFTP with INETD, NFS server, and ISC DHCPD. All of these services log with `syslog`. Consult your vendor documentation for more information. See also the `syslog.conf` file that was used to configure the Boot Server to determine how the logging has been configured for this component.

Build Manager Logs

These logs are in the following file:

```
/var/log/opsware/buildmgr/buildmgr.log
```

Command Engine Logs

These logs are in the following files:

```
/var/log/opsware/waybot/waybot.err*  
/var/log/opsware/waybot/waybot.log*
```

Changing Log Levels

To change the log level for the Command Engine, edit the file `/etc/opt/opsware/waybot/waybot.args` and add the following line with the desired log level:

```
loglevel: DEBUG
```

You must restart the Command Engine for this change to take effect.

Data Access Engine Logs

These logs are in the following files:

```
/var/log/opsware/spin/spin.err*  
/var/log/opsware/spin/spin.log*
```

Note: In a core with multiple Data Access Engines, each server running an engine has a set of these log files.

Media Server Logs

These logs are in the following files:

```
/var/log/opsware/samba/log.smbd  
/var/log/opsware/samba/log.nmbd
```

Solaris and Linux OS provisioning use of vendor-provided services such as NFS. These services typically log through `syslog`. Consult your vendor documentation for more information on these log files.

Model Repository Logs

The Model Repository is an Oracle database. The location logs the database is specific to your installation. For more information, see the Monitoring Oracle Log Files section in the SA Installation Guide.

Model Repository Multimaster Component Logs

These logs are in the following files:

```
/var/log/opsware/vault/err*  
/var/log/opsware/vault/vault.n.log
```

Changing Logging

To configure the log file name, log file size, or logging level for the Model Repository Multimaster component, in the SA Client select the Administration tab, select System Configuration in the navigation panel, then select the Model Repository Multimaster Component. This displays the log file, log level, and log size system configuration parameters available for the model repository multimaster component. After setting the desired values, select the Revert button to discard your changes or the Save button to save your changes.

Alternatively, to change the log level for the Model Repository Multimaster component, edit the file `/etc/opt/opsware/vault/logging.properties` and change the following line.

```
.level=INFO
```

The default log level value is INFO.

You must restart the Model Repository Multimaster Component for this change to take effect. For instructions, see [Starting Individual SA Core Components](#).

Agents Logs

The Agents create the following log files on managed servers:

UNIX:

```
/var/log/opsware/agent/agent.log*  
/var/log/opsware/agent/agent.err*
```

Windows:

```
%ProgramFiles%Common Files\opsware\log\agent\agent.log*  
%ProgramFiles%Common Files\opsware\log\agent\agent.err*
```

SA Client Logs

The SA Client does not generate its own logs. The SA Client uses the JBoss server, which writes to the following log files:

```
/var/log/opsware/occ/server.log*  
/var/log/opsware/httpsProxy/*log*
```

Changing Log Levels

To change the log level for the SA Client, edit the `/opt/opsware/occ/occ/conf/log4j.xml` file and change the `org.jboss.logging.XLevel` attribute value for the desired namespace. The default value is INFO.

You must restart the SA Client for this change to take effect.

Software Repository Logs

These logs are in the following files:

```
/var/log/opsware/mm_wordbot/wordbot.err*  
/var/log/opsware/mm_wordbot/wordbot.log*
```

Changing Log Levels

To change the log level for the Software Repository, edit the file `/etc/opt/opsware/mm_wordbot/mm_wordbot.args` and change the following property to the desired log level:

```
logLevel: logging.Level.INFO
```

For example, to set logging to debug, set this value to the following:

```
logLevel: logging.Level.DEBUG
```

You must restart the Software Repository for this change to take effect. For instructions, see [Starting Individual SA Core Components](#).

Web Services Data Access Engine Logs

The Web Services Data Access Engine contains the following log files:

```
/var/log/opsware/twist/stdout.log*  
/var/log/opsware/twist/twist.log  
/var/log/opsware/twist/access.log  
/var/log/opsware/twist/server.log*  
/var/log/opsware/twist/boot.log  
/var/log/opsware/twist/watchdog.log
```

The `stdout.log` file contains debug output and logging of every exception that the server generates. The file does not conform to a specific format. * indicates the files are `log.1`, `log.2`, `log.3`, and so forth. The number of files and the size of each file can both be configured using `twist.conf`. Additional logs are created when the specified maximum file size is reached. The `stdout.log` is the most recent, and `stdout.log.1` through 5 are progressively older files. The file is also

rotated on startup. This file also contains the output of any `System.out.println()`, `System.err.println()`, and `e.printStackTrace()` statements.

The `twist.log` file contains JBoss-specific error or informational messages and Weblogic specific messages. These files are rotated on startup.

The `access.log` file contains access information in common log format. These files are rotated when the file reaches 5MB in size.

The `server.log` file contains debug messages generated from the Web Services Data Access Engine. The debug messages are controlled by the log level set at the package or class level in the `twist.conf` file. * indicates the files are `log.1`, `log.2`, `log.3`, and so forth. The number of files and the size of each file can both be configured via `twist.conf`. The `server.log.0` is always the current file, while `server.log.9` is the oldest.

The `boot.log` file contains information on the initial `stdout` and `stderr` messages generated when the Web Services Data Access engine starts. In addition, the `boot.log` file contains the output from `Kill -QUIT` commands.

The `watchdog.log` file records the status of the Web Services Data Access Engine once every minute.

Changing Log Levels

To change the log level for the Web Services Data Access Engine edit the file `/etc/opt/opsware/twist/twist.conf`. Change the log level from `WARNING` to `FINEST` or another value for the default log level or for another logger namespace you are interested in. There are multiple namespaces in this file. You can change the log level for all namespaces or for individual namespaces.

Gateway Logs

These logs are in the following files:

```
/var/log/opsware/<gateway-name>/opswgw.log*
```

where `<gateway-name>` is the directory of a specific gateway component.

Changing Log Levels

To change the log level for any of the gateway components, create or edit the file `/etc/opt/opsware/<gateway-name>/opswgw.custom` and set the log level in the following line:

```
opswgw.LogLevel=INFO
```

You must restart the gateway after changing the log level. For instructions, see [Restarting or Stopping a Gateway Process](#).

Global File System Logs

The OGFS log files are in the following files:

```
/var/log/opsware/hub/OPSWhub.log*  
/var/log/opsware/ogfs/ogsh.err*  
/var/log/opsware/adapters/adapters.err*  
/var/log/opsware/agentcache/agentcache.log  
/var/log/opsware/spoke/spoke-*.log  
/var/log/opsware/spoke/stdout.log
```

Changing Log Levels - OGFS Hub Component

To change the log level for the hub component of the OGFS, perform the following steps:

1. Log in to the global shell (OGSH) as an administrative user. For instructions, see the [SA User Guide: Server Automation](#).
2. To determine the current log level, examine the file `/opsw/sys/hub/loglevel`. For example, run the following OGSH command:

```
more /opsw/sys/hub/loglevel
```

3. To change the log level, enter the following OGSH commands:

```
echo "MESSAGE ON" > /opsw/sys/hub/loglevel  
echo "LEVEL FINE" > /opsw/sys/hub/loglevel
```

The default values are “MESSAGE OFF” and “LEVEL INFO.”

Changing Log Levels - OGFS Spoke Component

To change the log level for the OGFS Spoke component, edit the file `/etc/opt/opsware/spoke/spoke_custom.conf`. Modify or add the following to this file and set the desired log level:

```
.level=INFO
```

You must restart the OGFS spoke component after changing the log level. For instructions, see [Starting Individual SA Core Components](#).

HTTPS Server Proxy Logs

These logs are found in:

```
/cust/apache/servers/https-Proxy/logs
```

Note: The log file `ssl_request_log` can grow quite large and should be inspected if you are concerned about disk space availability.

APX Proxy Logs

The APX proxy log files are in `/var/log/opsware/apxproxy/`.

Changing Log Levels

To change the log level for the APX proxy component, create or edit the file `/etc/opt/opsware/apxproxy/apxProxyOverrides.conf`. Add or modify the following lines and set the desired log level:

```
.level = INFO  
com.opsware.level=INFO  
com.opsware.apxproxy.level=CONFIG
```

You must restart APX proxy after changing the log level. For instructions, see [Starting Individual SA Core Components](#).

The possible values for these properties are listed in the file `/etc/opt/opsware/apxproxy/apxProxy.conf`.

SSHD Logs

The SSHD log files are in the location configured by syslog, typically `/var/log`.

Changing Log Levels

To change the log level for the SSHD component, edit the file `/etc/opt/opsware/sshd/sshd_conf`. Modify the following and set the desired log level:

```
LogLevel INFO
```

You must restart SSHD after changing the log level. For instructions, see [Starting Individual SA Core Components](#).

Global Shell Audit Logs

When a user accesses or modifies a managed server with the Global Shell feature, SA records the event in an audit log. The Global Shell audit logs contain information about the following events:

- Logins and logouts with Global Shell and Remote Terminal sessions
- The commands entered in Global Shell and Remote Terminal sessions
- File system operations (such as create and remove) on managed servers
- Commands and scripts that run on managed servers through the Remote Shell (`rosh`)

Note: The Global Shell audit logs are on the server where the OGFS is installed.

To view a log file, open a terminal window, log into the server running the OGFS, and use a command-line utility such as `more`, `grep`, or `tail`. For an example that uses the `tail` command, see [Example of Monitoring Global Shell Audit Logs](#).

The Global Shell audit logs are made up of three sets of logs files:

- Shell event logs
- Shell stream logs
- Shell script logs

Shell Event Logs

The shell event logs contain information about operations that users have performed on managed servers with the Global Shell. These logs are in the following directory (where *ogfs-host* is the name of the server running the OGFS):

```
/var/opt/opsware/ogfs/mnt/audit/event/ogfs-host
```

The log file name has the following syntax (where *n* is the log rotation number):

```
audit.log.n
```

For each event, SA writes a single line to an event log file. Each line in the log file contains the following information about the event:

- Unique ID of the event
- Unique ID of the parent event
- Date of the operation
- ID of the SA user who performed the operation
- Name of the SA user who performed the operation
- Name of the component that generated the audit event
- Version of the SA component that generated the audit event
- Name of the SA feature which generated the audit event
- Name of the operation (action)
- Verbosity level
- Exit status of the event
- ID of the managed server
- Name of the managed server
- Details of the event

The following example shows a single line in an audit event log file:

```
jdoue@m185:051202182224813:13 jdoue@m185:051202182224790:12  
2006/01/28-12:40:19.622 User.Id=2610003 User.Name=jdoue  
Hub:1.1 GlobalShell AgentRunTrustedScript 1 OK  
Device.Id=10003 Device.Name=m192.dev.opsware.com  
ConnectMethod=PUSH RemotePath= RemoteUser=root  
ScriptName=__global__.sc_snapshot.sh  
ScriptVersion=30b.2.1572 ChangeTime=1128971572  
RemoteErrorName=
```

In this example, the first field is the ID of the event:

```
jdoue@m185:051202182224813:13
```

This ID field has the following syntax:

```
opsware-user@ogfs-host: YYYYMMDDHHmmssSSS: n
```

The *n* at the end of the ID field is a sequence number of the audit event generated in a session. The ID field matches the name of a shell stream log file.

Shell Stream Logs

The shell stream logs contain the `stdout` of scripts that are run from the Global Shell. These logs are in the following directory (where *ogfs-host* is the name of the server running the OGFS):

```
/var/opt/opsware/ogfs/mnt/audit/streams/ogfs-host
```

The log file name has the following syntax:

```
opsware-user@ogfs-host: YYYYMMDDHHmmssSSS: n
```

The log file name matches the ID field in the shell event log. A header line in the log file contains the file name, character set, version, and SA user name. If the `stdout` of the script contains control characters, the shell stream log will contain the same control characters.

Shell Script Logs

The shell script logs contain the contents of scripts that are run from the Global Shell. These logs are in the following directory (where *ogfs-host* is the name of the server running the OGFS):

```
/var/opt/opsware/ogfs/mnt/audit/scripts/ogfs-host
```

The log file name is a hash string based on the script contents; for example:

```
23f1d546cc657137fa012f78d0adfd56095c3b5
```

A header line in the log file contains the file name, character set, version, and SA user name.

Example of Monitoring Global Shell Audit Logs

The following example monitors the commands entered by an end user who logs into a managed server with a Remote Terminal session:

1. In a terminal window, as `root`, log into the core server running the OGFS. The following steps refer to this window as the “auditing window.”
2. In the auditing window, go to the `audit/event` directory:

```
cd /var/opt/opsware/ogfs/mnt/audit/event/ogfs-host
```

3. In the SA Client, open a Remote Terminal to a UNIX managed server.
4. In the auditing window, examine the last line in the `audit.log` file:

```
tail -1 audit.log.n
```

For example, the following entry from the `audit.log` file indicates that the SA user `jdoue` opened a Remote Terminal to the host (`Device.Name`) `toro.example.com`. The event ID is `jdoue@m235:060413184452579:59`.

```
jdoue@m235:060413184452595:60 jdoue@m235:060413184452579:59  
2006/04/13-18:44:52.728 User.Id=6220044 User.Name=jdoue
```

```
Hub:1.1 GlobalShellAgentLogin 1 OK Device.Id=840044  
Device.Name=toro.example.com ConnectMethod=JUMP RemotePath=  
RemoteUser=root
```

5. In the auditing window, go to the `audit/streams` directory:

```
cd /var/opt/opsware/ogfs/mnt/audit/streams/ogfs-host
```

6. In the auditing window, use the `tail -f` command to monitor the file that corresponds to the Remote Terminal session. The file name is the same as the event ID. For example, if the event ID is `jdoh@m235:060413184452579:59`, then you would enter the following command:

```
tail -f jdoh*59
```

7. In the Remote Terminal window, enter some UNIX commands such as `pwd` and `ls`.
8. Watch the auditing window. The commands (and their output) from the Remote Terminal session are written to the file in the `audit/streams` directory.

Digital Signatures in the Global Shell Audit Logs

The shell stream and script log files contain digital signatures and fingerprints, which are generated with the RSA-SHA1 algorithm. To verify the signature and fingerprint of a log file, open a terminal window, log into the OGFS, and enter the following command:

```
/opt/opsware/agentproxy/bin/auditverify stream_file_name \  
rsa_key_path
```

This is an example in `bash`:

```
STREAMDIR=/var/opt/opsware/ogfs/mnt/audit/streams/acct.opsw.com  
STREAMFILE=jdoh@somehost:051210003000111:61  
RSAKEYPATH=/var/opt/opsware/crypto/waybot/waybot.srv  
  
/opt/opsware/agentproxy/bin/auditverify $STREAMDIR/$STREAMFILE \  
$RSAKEYPATH
```

If the log file has not been modified, `auditverify` displays the following message:

```
[AuditVerify]: Verification Result: Valid Signature
```

By default, the logs are signed with the private key in the following file:

```
/var/opt/opsware/crypto/agent/agent.srv
```

To change the key file used for signing, modify the `audit.signature.key_path` system configuration parameter as described in [Configuring the Global Shell Audit Logs](#).

Storage Management for the Global Shell Audit Logs

By periodically removing the shell stream and script log files, SA prevents these files from filling up the available disk space. SA provides system configuration parameters that determine when

the log files are removed. These parameters enable you to specify the removal of the log files based on the age (`archive_days`) of the files or the amount of disk space (`archive_size`) used by the files.

The following parameters specify the age of the files to remove:

`audit.stream.archive_days`

`audit.script.archive_days`

The following parameters specify the amount of disk space that the files can occupy before they are removed:

`audit.stream.archive_size`

`audit.script.archive_size`

For details on these parameters, see **Table 33**. For instructions on modifying these system configurations, see [Configuring the Global Shell Audit Logs](#).

Table 33. Parameters for Global Shell Audit Log Configuration

Parameter	Description	Default Value
<code>audit.script.archive_days</code>	<p>Audit script files older than this value (in days) are deleted. 0 means files are never deleted.</p> <p>Note: Using a 0 value or very high number results in OGSB connection issues, due to the high number of large files created under <code>/var/opt/opsware/mnt/audit/streams</code>.</p> <p>If this occurs, renaming the subdirectories under <i>streams</i> will temporarily resolve the issue.</p>	90
<code>audit.script.archive_size</code>	<p>Maximum amount of disk space (in MB) used by all audit script files. Older files are removed first. Zero (0) means no maximum.</p>	1000

Parameter	Description	Default Value
	<p>Note: Using a 0 value or very high number results in OGSN connection issues, due to the high number of large files created under <code>/var/opt/opsware/mnt/audit/streams</code>.</p> <p>If this occurs, renaming the subdirectories under <code>streams</code> will temporarily resolve the issue.</p>	
<code>audit.signature.algorithm</code>	Signature algorithm to use when signing audit streams.	RSA-SHA1
<code>audit.signature.key_path</code>	Location of the private key used when signing audit streams.	<code>/var/opt/opsware/crypto/waybot/waybot.srv</code>
<code>audit.stream.archive_days</code>	Audit stream files older than this value (in days) are deleted. 0 means files are never deleted.	10
<code>audit.stream.archive_size</code>	Maximum amount of disk space (in MB) used by all audit stream files. Older files are removed first. 0 means no maximum.	1000
<code>audit.stream.-file_keep</code>	Maximum number of rotated audit stream files.	50
<code>audit.stream.-file_size</code>	Maximum file size for audit streams. Specified in MB. The largest allowed value is 50MB.	10

Configuring the Global Shell Audit Logs

You can change some system configuration parameters for the global shell audit logs such as the maximum log file size. For a list of the parameters you can change, see **Table 33**. To configure

the parameters, perform the following steps:

1. Select the **Administration** tab in the SA Client.
2. Select System Configuration in the navigation pane. This displays the SA components, facilities and realms that have system configuration parameters.
3. In the list of SA components, select Hub. This displays the system configuration parameters for this component.
4. Locate and modify the system configuration parameters you want to change, as listed in **Table 33**.
5. Select the Revert button to discard your changes or the Save button to save your changes.

Extracting Session Data

SA saves context and other information about jobs, also known as “way sessions” or simply “sessions.” By default, this session data is kept for seven days before being garbage-collected to reuse space. This data can be useful for troubleshooting job and session problems. You also may want to save valid session data for comparison with problematic cases.

You can use the `dump_session` tool to extract and save this information. The `dump_session` tool generates a tarball file containing the session data in a file named `Session<job_ID>.pkl.gz`.

This section describes the `dump_session` tool and how to use it to extract session data.

To capture session data for an SA job, perform the following steps:

1. Determine the numeric job ID of the problematic job or command. For jobs, select the Jobs and Sessions tab in the SA Client and locate the desired job. The job ID is listed in the Job ID column.
2. Log into the SA core server.
3. Run the `dump_session` tool, and provide the job ID as the first argument. For example:

```
# /opt/opsware/bin/dump_session <job_ID>
```
4. Save the session output, which is a tarball in the current working directory named `Session<ID>.pkl.gz`.
5. If requested by HP Support, attach the tarball to the support incident for the problem.

Listing Recent Sessions

You can list the most recent set of jobs by running `dump_session` with the `-l` option and specifying the number of jobs you want to see. For example, the following command lists the most recent 25 jobs:

```
# /opt/opsware/bin/dump_session -l 25
```

The default number of jobs listed with `-l` is ten.

The following is sample output for five sessions:

```
# /opt/opsware/bin/dump_session -l 5
Session ID | Start Date | Session Desc
26000001 | 20100902T12:00:01 | 'Automated Communications Test for
core 1'
25980001 | 20100902T15:00:00 | 'opsware.patch_compliance'
26030001 | 20100902T17:51:57 | 'Communication Test'
25990001 | 20100903T00:00:00 | 'Automated Hypervisor Scan for
core: 1'
26010001 | 20100903T00:00:01 | 'Automated Communications Test for
core 1'
```

Sample Output

The following shows a sample `dump_session` command and sample output for SA job ID 1870001:

```
# /opt/opsware/bin/dump_session 1870001
Dumping session to 'Session1870001.pkl.gz'
Session:1870001
MegaServiceInstance:20001
WayScriptVersion:1830001
SecurityUser:60001
Realm:0
Device:10001
WayScript:1830001
```

dump_session Command Reference

This describes the `dump_session` command syntax and options. The `dump_session` command is at `/opt/opsware/bin/dump_session`. It extracts and formats SA sessions and related commands from the SA database.

Syntax

```
dump_session [<session_id> ...] [<session_file> ...] [-h] [-l <num>] [-d<num>]
```

Options

Table 34 lists the options to the `dump_session` command.

Table 34. dump_session Options

Option	Description
<code><session_id></code>	Specifies one or more SA job IDs. Information about these jobs will be copied from the SA database to a gzipped, multi-pickle file named “<session_id>.pkl.gz” in the current working directory.
<code><session_file></code>	Specifies one or more previously saved <session_id>.pkl.gz files. These files will be processed and converted into a static HTML directory structure resembling the waybot’s backend web UI.
<code>-h</code>	Displays help information.
<code>-l <num></code>	Displays to stdout the last <num> number of SA jobs executed on each core in the mesh. If <num> is omitted, then 10 is assumed. <num> can only be omitted when -l is the last argument on the command line.
<code>-d<num></code>	Sets the debug level to the specified number.

SA Notification Configuration

This section describes user-definable configuration parameters that allow you to modify contact information in the SA Client help, configuring a core mail server, setting core email alert addresses, and so on.

Configuration parameters are typically specified during the SA Core installation interview process. For more information, see the *SA Installation Guide*.

Caution: There are many default values for the various system configuration parameters that should not be changed unless expressly directed to do so by your technical support representative or consultant.

Note: Server Agents read system configuration values at installation time only. If you change any configuration values, all Agents' configurations must be updated manually. Contact your HP Server Automation Support Representative for help making these changes or in making any other changes in SA System Configuration.

Configuring SA Administrator Contact Information in SA Help

To configure SA administrator contact information that appears on the Server Automation Help page, perform the following tasks:

1. Log on as root to the server running the Core's Command Center (OCC).
2. Change to the following directory:

```
/etc/opt/opsware/occ
```
3. Open the `psrvr.properties` file in a text editor.
4. Change the values in the following fields to specify contact information in the SA Client Help:

```
pref.occ.support.href
```

```
pref.occ.support.tex
```
5. Save the file and exit the editor.


- Restart the OCC by entering the following command:

```
/etc/init.d/opsware-SAS restart occ.server
```


Configuring the Mail Server for a Facility

SA core components use the system configuration parameter `opsware.mailserver` to determine the address of the mail server to use for email notifications. By default, the value of `opsware.mailserver` is `smtp`, which is used if no value is specified. Most systems can use this value successfully.

However, if you need to specify a different value for `opsware.mailserver`, perform the following steps:

- Select the **Administration** tab in the SA Client.
- In the navigation pane, select **System Configuration > Configuration Parameters**. This displays the SA components, facilities, and realms that have system configuration parameters.
- In the list of SA components, select a facility. This displays the system configuration parameters for the facility.
- Locate the parameter `opsware.mailserver`.
- In the value column, enter the new value directly, or select the new value button  and enter the host name of your mail server.
- Select the Revert button to discard your changes or the Save button to save your changes.

Configuring the Command Engine Notification Email

- Select the **Administration** tab in the SA Client.
- In the navigation pane, select **System Configuration > Configuration Parameters**. This displays the SA components, facilities, and realms that have system configuration parameters.
- In the list of SA components, select Command Engine. This displays the system configuration parameters for this component.
- Locate the parameter `way.notification.email.fromAddr`.
- In the value column, enter the new value directly, or select the new value button  and enter the “from” email address for the email messages that will be sent by the Command Engine to notify users about scheduled jobs.
- Select the Revert button to discard your changes or the Save button to save your changes.

- Restart the Command Engine component with the following command:

```
/etc/init.d/opsware-sas restart occ.server
```

- If SA is running in multimaster mode, restart the Model Repository Multimaster Component.

When restarting multiple SA components, you must restart them in the correct order. See [Starting a Standalone SA Core](#).

Configuring Email Alert Addresses for an SA Core

Requirement: Server agents read system configuration values at installation time only. If you change any configuration values, all agents' configurations must be updated manually. Contact HP SA support representative for help making these changes or in making any other changes in SA system configurations.

Perform the following tasks to configure email alert addresses. SA core installation uses the default value (`EMAIL_ADDR`) for these parameters.

- Select the **Administration** tab in the SA Client.
- In the navigation pane, select **System Configuration > Configuration Parameters**. This displays the SA components, facilities, and realms that have system configuration parameters.
- In the list of SA components, select SA Agent. This displays the system configuration parameters for this component.
- Locate and modify the following parameters, as needed:
 - In the parameter, `CronbotMailAlertsEnabled`, specify the value 1 to enable cronbot email alerts. To disable cronbot email alerts, specify the value 0.
 - In the parameter, `CronbotAlertAddress`, enter the email address that the Server Agent should use to alert the recipient about failed scheduled jobs.
- Select the Revert button to discard your changes or the Save button to save your changes.

Configuring Email Alert Addresses for a Multimaster Mesh

Perform the following tasks to configure email alert addresses for Multimaster alerts. An SA core installation uses the default value `EMAIL_ADDR` for these parameters.

1. Select the **Administration** tab in the SA Client.
2. In the navigation pane, select **System Configuration > Configuration Parameters**. This displays the SA components, facilities and realms that have system configuration parameters.
3. In the list of SA components, select Model Repository, Multimaster Component. This displays the system configuration parameters for this component.
4. Locate and modify the following parameters, as needed.
 - In the field, `sendMMErrorsTo`, enter the email address to which multimaster conflicts will be sent.
 - In the field, `sendMMErrorsFrom`, enter the email address that SA will use as the “from” address for Multimaster conflicts alert emails.
5. Select the Revert button to discard your changes or the Save button to save your changes.

Restart the Model Repository Multimaster Component in all SA cores in the Multimaster Mesh. See [Starting Individual SA Core Components](#).

Global Shell: Windows Subauthentication Package

Under Microsoft® Windows, a program (service or application) cannot obtain a handle to a login session for a user account without supplying the password for that user account. Without both the user name and password, a running program cannot impersonate or act as a user other than the user in whose identity the program is currently running.

This restriction also applies to SA Agents. The SA Agent is installed to run in the LocalSystem security context. The LocalSystem logon session is a special, trusted, and privileged security context that is created at boot time on every Windows server that is running Windows Server 2003, 2008, and 2012 operating systems. However, if the SA Agent needs to run a child process in the security context of another user (such as `<DOMAIN>\<username>`), it requires the password for that user account. The user name, password, and child program name are all passed to the Win32 API `LogonUser()`.

The SA Agent performs actions on a managed server on behalf of the SA Global Shell feature. An SA user can perform registry read operations, file creation, and browsing operations on a managed server by using the Global Shell feature and the SA Agent. If an SA user wants to perform the operation as a LocalSystem user, the SA Agent only needs to create a subprocess running in the same security context of the Agent itself. If an SA user wants to perform a Global Shell operation as a non-LocalSystem user, the Agent cannot use the Win32 API `LogonUser()` because it requires the user account password. See the SA User Guide: Server Automation for more information about Global Shell operations.

Microsoft Windows Authentication Process

Microsoft Windows authentication is a process that verifies whether a user is authorized to access a system. During this verification process, the user provides a password that is cryptographically hashed. This hashed value is then compared with a stored value.

Windows provides a subsystem that supports different forms of authentication. This subsystem is called the Microsoft® Windows Local Security Authority Subsystem (LSAS) and takes the form of a process running the `LSAs.exe` application on a Windows server.

The design of LSAS allows Windows to support multiple authentication packages. These authentication packages verify a password, a Kerberos token, a thumbprint, a retina pattern, and so on.

In a standard Windows NT4 installation, LSAS has a single authentication package that is called `MSV1_0`. `MSV1_0` is the authentication package that implements NT4 domain authentication. Any time you log in to a Windows NT4 server, providing a user name, password, and domain name, or

any time you mount a share on a Windows NT4 server, you are interacting with the MSV1_0 authentication package. On a Windows 2000 server, the set of standard authentication packages consists of MSV1_0 and Kerberos. Depending on the domain configuration, any login attempt will have the user interacting with one of these authentication packages. MSV1_0 and Kerberos are also available as authentication packages on Windows Server 2003, 2008, and 2012.

Microsoft Windows Subauthentication Package

All of the main Microsoft Windows authentication packages support delegation of the credential check to code that is known as a subauthentication package. A subauthentication package is a DLL that supplements or replaces part of the authentication and validation criteria used by the main authentication package.

The MSV1_0 authentication package can (on the request of a client) defer the verification of user name and password to a previously registered subauthentication package. By default, MSV1_0 use its own internal user name and password checking software. It is only when a Windows client (such as the SA Agent) requests a specific subauthentication module that MSV1_0 delegates to the identified module.

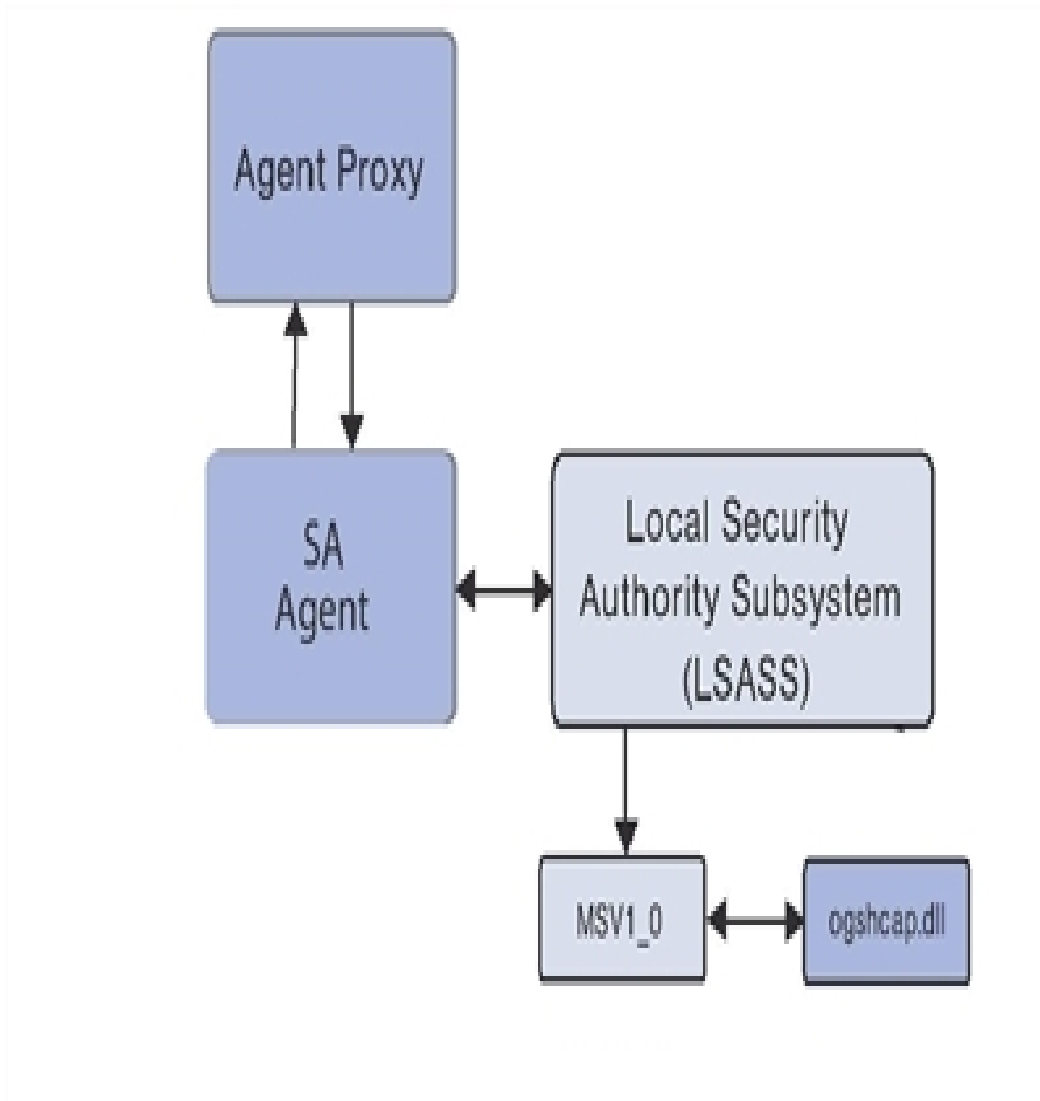
SA Subauthentication Package

SA provides an MSV1_0 subauthentication package that is requested by the SA Agent when the Agent is authenticating a user on whose behalf a Global Shell operation (such as a child process) must be run. This subauthentication package is a DLL known as *ogshcap.dll* (where *ogshcap* represents the Global Shell Custom Authentication and Subauthentication Package).

The *ogshcap.dll* file is passed the credentials that are supplied to Windows by the client application. This DLL is used on all supported Windows operating systems (Windows Server 2003, 2008 and 2012) and is used in an identical way on each operating system.

Figure 38 illustrates the subauthentication process in SA.

Figure 38. Subauthentication Process Flow



In the case of the SA Agent, the Agent passes a NULL password along with the user name when it calls a special Windows API to request subauthentication by the SA subauthentication package (ogshcap.dll). The Windows API then calls the MSV1_0 authentication package which, in turn, passes the credentials, including the NULL password to the requested subauthentication package.

The SA subauthentication package performs checks to verify that the user account is not locked out or disabled, and that the calling client is the SA Agent. The DLL ignores the password field, which is empty (NULL). After its verification steps are passed, the DLL returns a success status to MSV1_0, which creates a login session that is then passed to LSAS. In turn, LSAS passes a handle to this login session to the SA Agent. This handle to a login session is then passed by the SA Agent to a call to the Win32 API `CreateProcessSAUser()` to run the child process in the identity of the non-LocalSystem user.

After Windows has been requested to perform a single subauthentication operation using the `ogshcap.dll` file, Windows opens this file and keeps it open until the server next reboots. This means that the `ogshcap.dll` cannot be deleted before the next reboot, nor can it be overwritten during an Agent installation or upgrade without a reboot.

Note: For all Windows operating systems, the user name of the security principal being authenticated must be a member of the Administrators group on the local server or of the Domain Admins group of the Primary Domain of which the server is a member.

SA Agent Installation Changes

During an SA Agent installation on all Windows operating systems, a new Windows registry value is created (if it does not already exist) as the following registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\MSV1_0
```

The new registry value is of type `REG_SZ` and contains:

- **Name:** `Auth155`
- **Value:** `ogshcap`

The SA Agent Installer contains the `ogshcap.dll` file. During an Agent installation, the `ogshcap.dll` file is copied to the following source location:

```
%SystemDrive%\Program Files\Opware\bin\ogshcap.dll
```

After this DLL file is created at this location, the Agent Installer tries to copy it to the following destination location:

```
%SystemRoot%\system32\ogshcap.dll
```

If no such file currently exists at the destination location, the copy succeeds. If the copy fails because the file is open and is in use, the Agent Installer calculates a cryptographic hash of both source and destination files. If the source and destination files are different by hash, the Agent Installer calls the `Win32 API MoveFileEx()`, which creates a Windows-internal registry key. This registry key informs Windows that it must replace the destination file with the source file at the next reboot.

If the hash for one or both DLL files cannot be successfully calculated, the Agent Installer assumes that the replacement of the DLL is warranted. For example, if the Microsoft cryptographic modules cannot be loaded by the Agent Installer, the hash cannot be calculated. The Agent Installer then assumes that the DLL must be replaced.

A post-install reboot can be initiated after the Agent installation by specifying the installer option (`--reboot`) on the Agent Installer command line.

Note: When a post-install reboot is required to get the latest version of the DLL, the reboot performs a move operation in which the DLL in the source location is moved to the destination location. Therefore, the source DLL file overwrites the destination DLL.

If the existing `ogshcap.dll` on the operating system must be replaced and a reboot is required to accomplish this, the Agent Installer will not (by default) initiate the reboot. A reboot occurs only if the person performing the installation specifies `--reboot` as a command-line option.

The `--reboot` option is accepted by the Agent Installer on all operating systems; however, it is performed only on Windows operating systems. For example, if the `--reboot` option is specified during an Agent installation on a Linux 7.2 operating system, a reboot will not be performed by the Agent Installer. In comparison, if the `--reboot` option is specified during an Agent installation on a Windows 2000 operating system, a reboot will be performed by the Agent Installer.

If the hashes have been calculated and the source and destination files are verified as identical, no attempt to overwrite the opened `ogshcap.dll` is made.

The Agent always performs the first-time installation of the `ogshcap.dll` or the analysis of whether an existing DLL should be overwritten with the version of the DLL that is in the Agent Installer payload. In this case, there is no way to prevent installation of this DLL by the Agent Installer.

If the Agent Installer indicates that a reboot is required and the reboot does not occur after the Agent installation, the SA Agent will be using the out-of-date version of the DLL until the reboot occurs. This means that any bug fixes or modified functionality that are in the new DLL will not be used by the SA Agent until the reboot. However, Windows authentication, on behalf of the SA Agent by the old DLL, will still successfully occur, even while the DLL is marked for replacement by the newer DLL.

The following sample Agent Installer log is from an installation of the `ogshcap.dll`. In this case, the existing DLL on the operating system does not need to be replaced.

```
[08/Jun/2005 20:59:18] [INFO] Install CAP file if differing check-
sum between new and existing file.
[08/Jun/2005 20:59:18] [TRACE] NeedToReplaceOGSHCAPDLL()
[08/Jun/2005 20:59:18] [INFO] Testing CAP file existence:
C:\WINDOWS\system32\ogshcap.dll
[08/Jun/2005 20:59:18] [INFO] C:\WINDOWS\system32\ogshcap.dll CAP
file exists
[08/Jun/2005 20:59:18] [TRACE] GenerateKeyToFile()
[08/Jun/2005 20:59:18] [TRACE] Successfully called CreateFile
(C:\Program
Files\Common Files\Opware\cogbot\hmac.key)
[08/Jun/2005 20:59:18] [TRACE] Key file already exists
```

```
[08/Jun/2005 20:59:18] [TRACE] C:\Program Files\Common
Files\Opware\cogbot\hmac.key size: 36 bytes
[08/Jun/2005 20:59:18] [TRACE] Successfully called CloseHandle
(C:\Program
Files\Common Files\Opware\cogbot\hmac.key)
[08/Jun/2005 20:59:18] [TRACE] GenerateKeyToFile() = 1
[08/Jun/2005 20:59:18] [INFO] Calculate MAC for File:
C:\WINDOWS\system32\ogshcap.dll
[08/Jun/2005 20:59:18] [TRACE] C:\WINDOWS\system32\ogshcap.dll
size: 40960 bytes
[08/Jun/2005 20:59:18] [TRACE] C:\Program Files\Common
Files\Opware\cogbot\hmac.key size: 36 bytes
[08/Jun/2005 20:59:18] [TRACE] Successfully called CreateFileMap-
ping() for
C:\WINDOWS\system32\ogshcap.dll
[08/Jun/2005 20:59:18] [TRACE] Successfully called CreateFileMap-
ping() for
C:\Program Files\Common Files\Opware\cogbot\hmac.key
[08/Jun/2005 20:59:18] [TRACE] CalculateMAC()
[08/Jun/2005 20:59:18] [TRACE] PrintHexBytes()
[08/Jun/2005 20:59:18] [TRACE] HMAC for C:\WINDOWS\sys-
tem32\ogshcap.dll: 0x02
0x95 0x2B 0x03 0x51 0x02 0x9F 0x6D 0x58 0xF6 0xF1 0x5E 0x1C 0xFC
0x2A 0x72 0x5D
0x7E 0x5F 0xDA
[08/Jun/2005 20:59:18] [TRACE] CalculateMACFromFile() = 1
[08/Jun/2005 20:59:18] [INFO] Calculate MAC for File: C:\Program
Files\Opware\bin\ogshcap.dll
[08/Jun/2005 20:59:18] [TRACE] C:\Program Files\Opware\agent\bin\ogshcap.dll size:
40960 bytes
[08/Jun/2005 20:59:18] [TRACE] C:\Program Files\Common
Files\Opware\cogbot\hmac.key size: 36 bytes
```

```
[08/Jun/2005 20:59:18] [TRACE] Successfully called CreateFileMap-
ping() for
C:\Program Files\Opware\agent\bin\ogshcap.dll
[08/Jun/2005 20:59:18] [TRACE] Successfully called CreateFileMap-
ping() for
C:\Program Files\Common Files\Opware\cogbot\hmac.key
[08/Jun/2005 20:59:18] [TRACE] CalculateMAC()
[08/Jun/2005 20:59:18] [TRACE] PrintHexBytes()
[08/Jun/2005 20:59:18] [TRACE] HMAC for C:\Program
Files\Opware\agent\bin\ogshcap.dll: 0x02 0x95 0x2B 0x03 0x51
0x02 0x9F 0x6D 0x58
0xF6 0xF1 0x5E 0x1C 0xFC 0x2A 0x72 0x5D 0x7E 0x5F 0xDA
[08/Jun/2005 20:59:18] [TRACE] CalculateMACFromFile() = 1
[08/Jun/2005 20:59:18] [INFO] C:\WINDOWS\system32\ogshcap.dll CAP
file does not
need to be replaced
[08/Jun/2005 20:59:18] [TRACE] NeedToReplaceOGSHCAPDLL() = 0
[08/Jun/2005 20:59:18] [TRACE] UpdateCAPRegistrySetting()
[08/Jun/2005 20:59:18] [INFO] Update SubAuthentication Package
Registry key
[08/Jun/2005 20:59:18] [TRACE] Successfully opened registry key
SYSTEM\CurrentControlSet\Control\Lsa\MSV1_0.
[08/Jun/2005 20:59:18] [TRACE] Successfully found registry value:
'Auth255' at
this key, retrieved value 'ogshcap' (8) bytes.
[08/Jun/2005 20:59:18] [TRACE] Existing registry value matches
expected value:
'ogshcap'
[08/Jun/2005 20:59:18] [TRACE] UpdateCAPRegistrySetting() = 1
[08/Jun/2005 20:59:18] [INFO] UpdateCapRegistrySetting() was suc-
cessful
[08/Jun/2005 20:59:18] [TRACE] Win32InstallN() = 1
[08/Jun/2005 20:59:18] [INFO] Installation completed suc-
cessfully.
```

```
[08/Jun/2005 20:59:18] [INFO] An Agent install time reboot is NOT needed.
```

SA Agent Uninstallation Changes

During an SA Agent uninstallation, the Windows uninstaller tries to remove the following file:

```
%SystemRoot%\system32\ogshcap.dll
```

If the removal fails (because the file is open and is in use by Windows), the uninstaller calls `MoveFileEx()`, instructing Windows to remove the file during the next reboot. The uninstaller will prompt the user whether it should initiate a reboot immediately, if the attempt to remove the file fails.

The uninstaller also removes the special subauthentication registry key value created at Agent install time. See [SA Agent Uninstallation Changes](#) for more information.

Appendix A Permissions Reference

This appendix lists the permissions required to perform tasks with SA. For more information on permissions, see [User and User Group Setup and Security](#).

Server Objects Permission

Table 35 specifies the permissions required for server objects such as Registered Software, Internet Information Server, Local Security Settings, Runtime State, Users and Groups, and .Net Framework Configuration.

Table 35: Server Object Permissions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Browse Server Objects	Manage Server Modules: Read & Write Allow Execute Server Modules: Yes	N/A	N/A
Add to Library (From the Server Browser)	Manage Server Modules: Read & Write Allow Execute Server Modules: Yes Manage Package: Read and Write		Write
Add to Software Policy	Manage Server Modules: Read and Write Allow Execute Server Modules: Yes Manage Package: Read and Write Manage Software Policy: Read & Write	N/A	Write

Server Property and Reboot Permissions

Table 36 specifies the permissions required by users to modify server properties, reboot servers, and deactivate SA agents. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

Table 36. Server Property and Reboot Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)
Deactivate SA Agent	Deactivate: Yes	Read & Write
Modify Property: Server Name or Description	N/A	Read & Write
Reboot Server	Reboot Server: Yes	Read & Write

Device Group Permissions

To use device groups in the SA Client, you must have the permissions described in **Table 37**. For a list of tasks that require the Model Public Device Group permission, see **Table 45**.

Table 37. Device Groups Action Permissions

User Action	Action Permission
Creating a public static device group	Manage Public Device Group: Yes
Creating a public dynamic device group	Manage Public Device Group: Yes
Adding a server to a public static device group	Manage Public Device Group: Yes
Adding a server to a public dynamic device group	Manage Public Device Group: Yes
Removing a server from a public static device group	Manage Public Device Group: Yes
Removing servers from a public dynamic device group	Manage Public Device Group: Yes
Moving a public device group	Manage Public Device Group: Yes
Duplicating a public device group	Manage Public Device Group: Yes
Deleting a public device group	Manage Public Device Group: Yes
Adding devices to a device group being used as an Access Control Group	Manage Public Device Group and Super Administrator

Server Agent Deployment Permissions

To install a server agent on servers using the SA Client, you must have the permissions described in **Table 38**.

Table 38. Agent Action Permissions

User Action	Action Permission
Install the SA agent on servers	Allow Install Agent: Yes
Scan the network for agentless servers	Allow Scan Network: Yes
View servers running agents and device groups	Managed Servers and Groups: Yes
Modify facilities	Facilities: Yes
Resource	Facilities: Read and Write to the facilities to scan for servers and install agents. Read and Write to the customers who will be assigned servers.

Note: If you are installing an SA Agent in a custom location, make sure you do *not* disable the symlinks permissions, which are already set by default.

Virtualization Service Management Permissions

To manage virtualization services (VSs), virtual machines (VMs), and VM templates, you must have the action permissions listed in **Table 39**.

If a user does not have a particular action permission (the permission is set to No), the corresponding menu item will not appear in the SA Client Actions menu.

Table 39. Virtualization Action Permissions

Action Permission	Description
View Virtualization Inventory	Also requires the permission Managed Servers and Groups = Yes. Allows you to view virtualization inventory (across supported technologies) and perform the “Reload Data” operation to view the most up-to-date virtualization information. If this permission is set to No, the Virtualization

Action Permission	Description
Manage VM Lifecycle: Clone VM	tab in the SA Client and the Oracle Solaris Zones view are not displayed. Clone virtual machines and perform compatibility checks. “Customize Guest OS” is also required for guest customization.
Manage VM Lifecycle: Create VM	Create VMs and perform compatibility checks. When running the OS Build Plan from a Create VM job, also required are the permissions listed for “Run OS Build Plan” listed in Table 42 .
Manage VM Lifecycle: Customize Guest OS	Allows OS guest customization during “Clone VM” or “Deploy VM from VM Template.”
Manage VM Lifecycle: Delete VM	Delete VMs.
Manage VM Lifecycle: Deploy VM from VM Template	Deploy VMs from VM templates and perform compatibility checks. “Customize Guest OS” is also required for guest customization.
Manage VM Lifecycle: Migrate VM	Migrate virtual machines (host only, storage only, or both host and storage) and perform compatibility checks.
Manage VM Lifecycle: Modify VM	Modify configuration of VMs.
Manage VM Power State	Ability to perform power control operations for VMs (for example, power on, power off, pause, suspend, reset, restart guest, and shutdown guest).
Manage VM Templates: Convert VM to VM Template	Convert VMs to VM templates.
Manage VM Templates: Delete VM Template	Delete VM templates.
Administer Virtualization Services	Register, modify and remove virtualization services.
Add Host to Virtualization Service	Add hypervisors to a virtualization service so that they can be managed.

Virtualization Container Permissions and Server Resource Permissions

In addition to action permissions, virtualization container permissions are required to perform all virtualization actions. Virtualization container permissions give you access to virtualization

containers such as datacenters, hypervisors, host groups, clusters, resource pools, folders, projects, and their children.

The access-control list (ACL) inheritance rule defines what user groups are automatically granted access to any newly added or discovered virtualization containers, based on what ACLs the user group has for the parent container.

Permission options are **L** (List), **READ**, **WRITE**, **X** (execute), and **PM** (edit permissions). If you want the setting for groups with X or PM to inherit ACLs, then use “X,PM.” The path to the rule is located here: Administration/System Configuration/Server Automation/Web Services Data Access Engine/Twist.v12n.inventory.inheritance.acl.

The PM option, which is the default, is the most strict option and is good for use with multi-tenant control. PM requires that a user with Edit permissions (generally a virtualization administrator) manually assign access to other groups. Only user groups that already have PM for the parent of the newly added or discovered container gets access.

The List option is the most permissive. If the user group has List permissions for the parent container, the group is automatically added to the new container with the group’s same permissions. For example, Group A has List and Read permissions, and Group B has List, Read, Write, and Execute permissions, for Datacenter 1. A new cluster is added under Datacenter 1. Group A now has List and Read permissions for the new cluster, and Group B has List, Read, Write, and Execute for the new cluster.

In addition to action permissions and virtualization container permissions, server resource permissions are required on servers running in a Virtualization Service. Server resource permissions are granted through facilities, customers, and device groups.

For more information about virtualization permissions and server resource permissions, see the SA User Guide: Virtualization Management.

Where **Table 39** lists just the action permissions, **Table 40** lists the user tasks you can perform and the complete set of action permissions, virtualization container permissions, server resource permissions, and in some cases folder permissions required to perform each user action.

Virtualization Tasks and Required Permissions

Table 40 lists the permissions required to perform each task on the virtualization inventory. The tasks in this table are used with VMware vCenter, Microsoft SCVMM. For more information on these tasks, see the SA User Guide: Virtualization Management.

Table 40. Virtualization Tasks and Required Permissions for vCenter and SCVMM

User Action	Required Action Permissions	Required Virtualization Container Permissions	Required Server Resource Permissions (Facility, Customer, Device Group)
View Virtualization tab in SA Cli-	View Virtualization Inventory: Yes	VS: List And	VS server: Read

User Action	Required Action Permissions	Required Virtualization Container Permissions	Required Server Resource Permissions (Facility, Customer, Device Group)
ent	Managed Servers and Groups: Yes	Separate permissions on each container under the VS Datacenter: Read (for access to the underlying datastores) On the parent container of VMs and templates: Read	
Add VS	Administer Virtualization Services: Yes View Virtualization Inventory: Yes Managed Servers and Groups: Yes	None needed.	VS server: Read
Edit VS, Remove VS	Administer Virtualization Services: Yes View Virtualization Inventory: Yes Managed Servers and Groups: Yes	VS: Write	VS server: Read
Reload Data for the VS or a container under the VS	View Virtualization Inventory: Yes Managed Servers and Groups: Yes	VS or container under the VS: Read	None needed

User Action	Required Action Permissions	Required Virtualization Container Permissions	Required Server Resource Permissions (Facility, Customer, Device Group)
Add Host to Virtualization Service	Add Host to Virtualization Service: Yes View Virtualization Inventory: Yes Managed Servers and Groups: Yes	Container where the hypervisor is being added: Write Or VS container if no container is specified: Write	Server (hypervisor) being added: Read
VM Power Controls - Start, Stop, Reset, Restart Guest, Shutdown Guest, Suspend, and Pause VM	View Virtualization Inventory: Yes Manage VM Power State: Yes Managed Servers and Groups: Yes	Container where the VM resides: Read	
Create VM	View Virtualization Inventory: Yes Manage VM Lifecycle: Create VM: Yes Managed Servers and Groups: Yes Allow Execute OS Build Plan: Yes, if specifying an OSBP. Manage Package: Read, for non-PXE Create VM with OSBP.	Destination container (hypervisor, cluster, or resource pool) where the VM will reside: Write Folder in the vCenter VS inventory where the VM will reside: Write	Server.write for the newly created VM Note - Execute permission is also required on the SA Library folder containing the selected OS Build Plan. For non-PXE Create VM with OSBP: Read on the <code>Opaware/Tools/OS Provisioning/WinPE</code> folder (Windows) Read on the <code>Opaware/Tools/OS Provisioning</code> folder (Linux).

User Action	Required Action Permissions	Required Virtualization Container Permissions	Required Server Resource Permissions (Facility, Customer, Device Group)
Modify VM	View Virtualization Inventory: Yes Manage VM Lifecycle: Modify VM: Yes Managed Servers and Groups: Yes	Container where the VM resides: Write And Hypervisor container the VM is on (vCenter only): List	VM server: Write
Migrate VM	View Virtualization Inventory: Yes Manage VM Lifecycle: Migrate VM: Yes Managed Servers and Groups: Yes	Container where the VM resides: Write Additional: To migrate storage - Hypervisor: List To migrate host or host and storage - destination container (hypervisor, cluster, or resource pool) where the VM will reside: Write	VM server: Read
Clone VM (vCenter only)	View Virtualization Inventory: Yes Manage VM Lifecycle: Clone VM: Yes Managed Servers and Groups: Yes	Container where the VM resides: Read Destination container (hypervisor, cluster, or resource pool) where the new VM will reside: Write Folder in the vCenter VS	Source VM server: Read New VM server: Write

User Action	Required Action Permissions	Required Virtualization Container Permissions inventory where the new VM will reside: Write	Required Server Resource Permissions (Facility, Customer, Device Group)
Customize Guest OS - When performed as part of a Clone VM operation or a Deploy VM from VM Template operation	<p>Same as Clone VM when performed as part of a clone VM operation.</p> <p>Same as Deploy VM from VM Template when performed as part of a deploy VM operation.</p> <p>Manage VM Lifecycle: Customize Guest OS: Yes</p> <p>Allow Execute OS Build Plan: Yes</p>	<p>Same as Clone VM when performed as part of a clone VM operation.</p> <p>Same as Deploy VM from VM Template when performed as part of a deploy VM operation.</p>	<p>Same as Clone VM when performed as part of a clone VM operation.</p> <p>Same as Deploy VM from VM Template when performed as part of a deploy VM operation.</p> <p>For Linux customization, Execute on the <code>Opaware/Tools/Build Plans/Virtualization/Guest Customization/Linux</code> folder.</p> <p>For Windows customization, Execute on the <code>Opaware/Tools/Build Plans/Virtualization/Guest Customization/Windows</code> folder.</p>
Delete VM	<p>View Virtualization Inventory: Yes</p> <p>Manage VM Lifecycle: Delete VM: Yes</p> <p>Managed Servers and Groups: Yes</p>	Container where the VM resides: Write	VM server: Write
Deploy VM from VM Template	<p>View Virtualization Inventory: Yes</p> <p>Manage VM Lifecycle: Deploy VM from VM</p>	<p>Container where the VM template resides: Execute</p> <p>Destination container (hypervisor, cluster, or</p>	<p>VM template server: Read</p> <p>New VM server: Write</p>

User Action	Required Action Permissions	Required Virtualization Container Permissions	Required Server Resource Permissions (Facility, Customer, Device Group)
	<p>Template: Yes</p> <p>Managed Servers and Groups: Yes</p>	<p>resource pool) where the new VM will reside: Write</p> <p>Folder in the vCenter VS inventory where the new VM will reside: Write</p>	
Convert VM to VM Template	<p>View Virtualization Inventory: Yes</p> <p>Manage VM Templates: Convert VM to VM Template: Yes</p> <p>Managed Servers and Groups: Yes</p>	<p>Container where the VM resides: Write</p> <p>VM Templates folder in SCVMM Library: Write</p>	VM server: Read
Delete VM Template	<p>View Virtualization Inventory: Yes</p> <p>Manage VM Templates: Delete VM Template: Yes</p> <p>Managed Servers and Groups: Yes</p>	<p>Container where the VM template resides: Write</p>	VM server: Write
Merge Servers	<p>View Virtualization Inventory: Yes (in order to merge a Virtualization server with</p>	<p>Container where the VM or Template resides: Write</p> <p>or</p> <p>Hypervisor:</p>	Server.write for both servers to merge

User Action	Required Action Permissions	Required Virtualization Container Permissions	Required Server Resource Permissions (Facility, Customer, Device Group)
	another server) Merge Servers: Yes Managed Servers and Groups: Yes	Write	

Solaris Virtualization Permissions

Table 41 lists the permissions required for managing Oracle Solaris Zones. For more information, see the SA User Guide: Virtualization Management.

Table 41. Solaris Virtualization Permissions

User Action	Required Action Permissions	Required Server Resource Permissions (Facility, Customer, Device Group)
Create Zone	Manage VM Lifecycle: Create VM View Virtualization Inventory: Yes Managed Servers and Groups: Yes	Hypervisor server: Read Customer the new VM is assigned to: Write
Reload Data	View Virtualization Inventory: Yes Managed Servers and Groups: Yes	Hypervisor server: Read VM server: Read
Modify	Manage VM Lifecycle: Modify VM View Virtualization Inventory: Yes Managed Servers and Groups: Yes	Hypervisor server: Read VM server: Write
Remove	Manage VM Lifecycle: Delete VM View Virtualization Inventory: Yes Managed Servers and Groups: Yes	Hypervisor server: Read VM server: Read
Start, Stop	Manage VM Power State: Yes View Virtualization Inventory: Yes Managed Servers and Groups: Yes	Hypervisor server: Read VM server: Write

OS Provisioning Permissions

This section describes the permissions required for OS provisioning. For security administrators, **Table 42** answers this question: To perform a particular action, what permissions does a user need?

In **Table 42**, the Server Permission column is for the servers referenced by the OS sequence or installation profile. Server permissions are specified by the Customer, Facility, and Device Groups permissions in the SA Client. To create and save an OS sequence in a folder, you will need write permissions to the folder.

Table 42. OS Provisioning Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
OS Build Plan			
Create OS Build Plan	Manage OS Build Plan: Read & Write	None	Write
View OS Build Plan	Manage OS Build Plan: Read	None	Read
Edit OS Build Plan	Manage OS Build Plan: Read & Write	None	Write
Delete OS Build Plan	Manage OS Build Plan: Read & Write	None	Write
Add Device Group to OS Build Plan	Any of the permission combinations below is valid: 1) Manage Servers and Groups + Manage OS Build Plan: Read & Write, or 2) Manage Public Device Group (in Client Features tab, Servers	None	Folder containing the OS Build Plan: Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
	section) + Manage OS Build Plan: Read & Write, or 3) Manage Public Device Groups (SA Client) (from Others tab, Servers and Device Group Permission section) + Manage OS Build Plan: Read & Write		
Add OGFS Script to OS Build Plan	Manage OGFS Script: Read + Manage OS Build Plan: Read & Write	None	Folder containing the OGFS Script: Read + Folder containing the OS Build Plan: Write
Add Server Script to OS Build Plan	Manage Server Script: Read + Manage OS Build Plan: Read & Write	None	Folder containing the Server Script: Read + Folder containing the OS Build Plan: Write
Add ZIP Package to OS Build Plan	Manage Package: Read + Manage OS Build Plan: Read & Write	None	Folder containing the package: Read + Folder containing the OS Build Plan: Write
Attach Software Policy to OS Build Plan	Manage Software Policy: Read + Man-	None	Folder containing the Software Policy: Read + Folder containing the OS Build Plan: Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
	age OS Build Plan: Read & Write		
Attach Windows Patch Policy to OS Build Plan	Manage Windows Patch: Policy + Manage OS Build Plan: Read & Write	None	Folder containing the OS Build Plan: Write
Run OS Build Plan (from server or from OS Build Plan node)	Managed Servers and Groups + Manage OS Build Plan: Allow Execute OS Build Plan: Yes	Read & Write	Folder containing the OS Build Plan: Execute
Run OS Build Plan (for VMware ESXi 4.1)	Manage Servers and Groups + Manage OS Build Plan: Read + Allow Execute OS Build Plan: Yes + Allow Manage Server + View Virtual Servers + Manage Virtual Servers	Read & Write	Folder (/Opware /Tools/OS Provisioning) contains the Run OS Build Plan web extension: Execute + Folder containing the OS Build Plan: Execute + List and Execute folder permission on /Opware/Tools/Virtualization Programs/Hypervisor Scanner folder
OS Sequence			
Create OS Sequence	Manage OS Sequence: Read & Write	Note: To create an OS Sequence	Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
	+ Operating Systems + Wizard: Prepare OS	<p>using an OS Installation Profile that is assigned to a customer, a user must have at least Read permission to the customer</p> <p>Note: To create an OS Sequence using a Customer Independent OS Installation Profile, no Customer permission is required.</p>	
View OS Sequence	Manage OS Sequence: Read	None	Read
Edit OS Sequence	Manage OS Sequence: Read & Write	None	Write
Delete OS Sequence	Manage OS Sequence: Read & Write	None	Write
Run OS Sequence (From server or from OS)	Manage OS Sequence: Read and	Read & Write	Read

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
sequences)	Allow Execute OS Sequence: Yes		
View unprovisioned servers	SA Client permission: Server Pool	Read	N/A
Attach Software Policy	Manage Software Policy: Read + Manage OS Sequence: Read & Write	NA	Folder containing the Software Policy: Read + Folder containing the OS Sequence: Write
Attach Windows Patch Policy	Manage Windows Patch: Policy + Manage OS Sequence: Read & Write	NA	Folder containing the OS Sequence: Write
Attach Solaris Patch Policy	Manage Software Policy: Read + Manage OS Sequence: Read & Write	NA	Folder containing the Solaris Patch Policy: Read + Folder containing the OS Sequence: Write
OS Installation Profile			
Create, edit, delete OS installation profile	Operating System + Wizard: Prepare OS	Note: To create an OS Sequence using an OS Installation Profile that is assigned to a customer, the customer	N/A

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
		must have read & write permission. Note: To create an OS Sequence using a Customer Independent OS Installation Profile, no Customer permission is required.	
Unprovisioned Server List			
View servers in the unprovisioned server list	Server Pool	N/A	N/A
Manage Boot Clients			
Execute Managed Boot Clients Web Application	Allow Configuration of Network Booting + Managed Server and Groups + Manage Customers + Server Pool	Read/Write to the Facility and Customer + Read/Write to customer Not Assigned	List and Execute on the <code>/Opware/Tools/OS Provisioning/Manage Boot Clients</code> folder

Table 43 lists the actions that users can perform for each OS Provisioning permission. **Table 43** has the same data as **Table 42**, but is sorted by action permission.

For security administrators, **Table 43** answers this question: If a user is granted a particular action permission, what actions can the user perform?

Table 43. User Actions Allowed in the SA Client by OS Provisioning Permissions

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder
Manage OS Sequence: Read	View OS sequence	Read	Read
Manage OS Sequence: Read & Write + Operating System + Wizard: Prepare OS	Create OS sequence	Read	Write
Allow Execute OS Sequence: Yes	Run OS sequence	Write	Read
Manage OS Sequence: Read Allow execute OS Sequence: Yes	Run OS sequence	Write	Read
Manage OS Sequence: Read Allow Execute OS Sequence: No	View OS sequence	Read	Read
Manage OS Sequence: Write Allow Execute OS Sequence: Yes	Run OS sequence Edit OS sequence	Write	Write
Manage OS Sequence: Write Allow Execute OS Sequence: No	Edit OS sequence	Read	Write
Operating System+ Wizard: Prepare OS	Create, edit, delete OS installation profile	Read & Write, N/A, N/A	N/A
Server Pool	View servers in the unprovisioned server list	Read	N/A

Manage Boot Clients Permissions

The following section describes the permissions required to use the Manage Boot Clients (MBC) Utility for OS Provisioning.

Table 44. Manage Boot Client Utility Permissions

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder
Allow Execute OS Build Plan	Run OS Build Plan	Write	Read
Allow Execute OS Sequence	Run OS Sequence	Write	Read
Manage Server and Groups	Manage Server and Groups	Write	Read
Manage Customers	Create, edit Customers	Write	Read
Server Pool	Access Server Pool	Write	Read
Read & Write permission to customer Not Assigned	Access to servers assigned to customer Not Assigned	Write	Read
Allow Configuration of Network Booting	Configuration of Network Booting	Write	Read

Software Management Permissions

Table 45 specifies the Software Management permissions required by users to perform specific actions in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

If a customer is assigned to a folder, then customer constraints might limit the objects that can be associated with a software policy contained in the folder. For a list of tasks affected by these constraints, see [Folders, Customer Constraints, and Software Policies](#).

To install software, you must belong to a user group that has the install software permissions. This user group must also have folder permissions for the software you want to install.

Table 45. Software Management Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Software Policy			
Create Software Policy	Manage Software Policy: Read & Write	N/A	Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Delete Software Policy	Manage Software Policy: Read & Write	N/A	Write
Open Software Policy (View)	Manage Software Policy: Read	N/A	Read
Edit Software Policy Properties	Manage Software Policy: Read & Write	N/A	Write
Add Packages	Manage Software Policy: Read & Write Manage Packages: Read	N/A	Folder containing the software policy: Write
Add RPM Packages	Manage Software Policy: Read & Write Manage Packages: Read	N/A	Folder containing the software policy: Write
Add Patches	Manage Software Policy: Read & Write Manage Patches: Read	N/A	Folder containing the software policy: Write
Add Application Configurations	Manage Software Policy: Read & Write Manage Application Configuration: Read	N/A	Folder containing the software policy: Write
Add Scripts	Manage Software Policy: Read & Write Manage Server Scripts: Read	N/A	Folder containing the software policy: Write
Add Server Objects	Manage Software Policy: Read & Write Manage Packages: Read	N/A	Folder containing the software policy: Write
Add Software Policies	Manage Software Policy: Read & Write	N/A	Folder containing the software policy: Write
Remove Packages	Manage Software Policy: Read & Write	N/A	Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Remove RPM Packages	Manage Software Policy: Read & Write	N/A	Write
Remove Patches	Manage Software Policy: Read & Write	N/A	Write
Remove Application Configurations	Manage Software Policy: Read & Write	N/A	Write
Remove Software Policies	Manage Software Policy: Read & Write	N/A	Write
Remove Scripts	Manage Software Policy: Read & Write	N/A	Write
Remove Server Objects	Manage Software Policy: Read & Write	N/A	Write
Install/ Uninstall Software	Manage Software Policy: Read Allow Attach/Detach Software Policy: Yes Allow Install/Uninstall Software: Yes Model Public Device Groups: Yes (Required if you remediate a public device group)	Read & Write	Read
Attach Software Policy	Manage Software Policy: Read Allow Attach/Detach Software Policy: Yes Model Public Device Groups: Yes (This permission is required if you are attaching the software policy to a public device group)	Read & Write	Read
Detach Software Policy	Manage Software Policy: Read	Read & Write	Read

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
	Allow Attach/Detach Software Policy: Yes Model Public Device Groups: Yes (This permission is required if you are attaching the software policy to a public device group)		
Remediate	Manage Software Policy: Read Allow Remediate Servers: Yes Model Public Device Groups: Yes (Required if you remediate a public device group)	Read & Write	Read
Run ISM Control	Manage Software Policy: Read Allow Run ISM Control: Yes Model Public Device Groups: Yes (Required if you run ISM Control on a public device group)	Read & Write	Read
Duplicate Zip Package	Manage Software Policy: Read & Write	N/A	Write
Edit ZIP Installation Directory	Manage Software Policy: Read & Write	N/A	Write
Scan Software Compliance	N/A	Read	N/A
Rename Software Policy	Manage Software Policy: Read & Write	N/A	Write
Cut Software Policy	Manage Software Policy: Read & Write	N/A	Write
Copy Software Policy	Manage Software Policy:	N/A	Read

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
	Read		
Paste Software Policy	Manage Software Policy: Read & Write	N/A	Source Folder: Read (for copy and paste) Source Folder: Write (for cut and paste) Destination Folder: Write
Move Software Policy	Manage Software Policy: Read & Write	N/A	Source Folder: Write Destination Folder: Write
Folder			
Create Folder	N/A	N/A	Write
Delete Folder	N/A	N/A	Write
Open Folder	N/A	N/A	Read
View Folder Properties	N/A	N/A	Read
Edit Folder Properties	N/A	N/A	Write
Manage Folder Per- missions	N/A	N/A	Edit Folder Per- missions
Cut Folder	N/A	N/A	Write
Copy Folder	N/A	N/A	Read
Paste Folder	N/A	N/A	Source Folder: Read (for copy and paste) Source Folder: Write (for cut and paste)

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
			Destination Folder: Write
Move Folder	N/A	N/A	Source Folder: Write Destination Folder: Write
Rename Folder	N/A	N/A	Write
Package			
Import Package	Manage Package: Read & Write	N/A	Write
Export Package	Manage Package: Read	N/A	Read
Open Package (View)	Manage Package: Read	N/A	Read
Edit Package Properties	Manage Package: Read & Write	N/A	Read
Delete Package	Manage Package: Read & Write	N/A	Write
Rename Package	Manage Package: Read & Write	N/A	Write
Cut Package	Manage Package: Read & Write	N/A	Write
Paste Package	Manage Package: Read & Write	N/A	Source Folder: Read (for copy and paste) Source Folder: Write (for cut and paste) Destination Folder: Write
Move Package	Manage Package: Read & Write	N/A	Source Folder: Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
			Destination Folder: Write

Table 46 lists the actions that users can perform for each Software Management permission. **Table 46** has the same data as **Table 45**, but is sorted by action permission. For security administrators, **Table 46** answers this question: If a user is granted a particular action permission, what actions can the user perform?

Table 46. User Actions Allowed by Software Management Permissions

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Manage Software Policy: Read & Write	Create Software Policy	N/A	Write
	Delete Software Policy	N/A	Write
	Edit Software Policy	N/A	Write
	Rename Software Policy	N/A	Write
	Cut Software Policy	N/A	Write
	Paste Software Policy	N/A	Write
	Move Software Policy	N/A	Write
	Remove Packages	N/A	Write
	Remove Patches	N/A	Write
	Remove Application Configurations	N/A	Write
	Remove Scripts	N/A	Write
	Remove Server Objects	N/A	Write
	Remove Software Policy	N/A	Write
	Duplicate ZIP packages	N/A	Write

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Manage Software Policy: Read	Open Software Policy (View)	N/A	Read
	Copy Software Policy Properties	N/A	Read
Manage Software Policy: Read & Write And Manage Package: Read	Add Packages Add RPM Packages	N/A	Folder containing the software policy: Write Folder containing the package: Read
Manage Software Policy: Read & Write And Manage Patches: Read	Add Patches	N/A	Folder containing the software policy: Write Folder containing the patch: Read
Manage Software Policy: Read & Write And Manage Application Configuration: Read	Add Application Configurations	N/A	Folder containing the software policy: Write Folder containing the application configuration: Read
Manage Software Policy: Read & Write	Add Software Policies	N/A	Folder containing the software policy: Write Folder containing the software policy to be added to another software policy: Read
Manage Software Policy: Read & Write And Manage Server Scripts: Read	Add Scripts	N/A	Folder containing the software policy: Write Folder containing the scripts: Read

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Manage Software Policy: Read & Write And Manage Packages: Read	Add Server Objects	N/A	Folder containing the software policy: Write Folder containing the server objects: Read
Manage Software Policy: Read & Write	Remove Packages	N/A	Write
	Remove RPM Packages	N/A	Write
	Remove Patches	N/A	Write
	Remove Application Configurations	N/A	Write
	Remove Scripts	N/A	Write
	Remove Server Objects	N/A	Write
	Remove Software Policies	N/A	Write
Manage Software Policy: Read And Allow Attach/Detach Software Policy: Yes And	Attach Software Policy	Read & Write	Read
Model Public Device Groups: Yes (Required if you are attaching the software policy to a public device group)	Detach Software Policy	Read & Write	Read
Manage Software Policy: Read And Allow Remediate Servers: Yes	Remediate	Read & Write	Read

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
<p>And</p> <p>Model Public Device Groups: Yes (Required if you remediate a public device group)</p>			
<p>Manage Software Policy: Read</p> <p>And</p> <p>Allow Attach/Detach Software Policy: Yes</p> <p>And</p> <p>Allow Install/Uninstall Software: Yes</p> <p>And</p> <p>Model Public Device Groups: Yes (Required if you remediate a public device group)</p>	<p>Install/ Uninstall Software</p>	<p>Read & Write</p>	<p>Read</p>
<p>Manage Software Policy: Read</p> <p>And</p> <p>Allow Run ISM Control: Yes</p> <p>And</p> <p>Model Public Device Groups: Yes (Required if you run ISM Control on a public device group)</p>	<p>Run ISM Control</p>	<p>Read & Write</p>	<p>Read</p>

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Manage Package: Read & Write	Import Package	N/A	Write
	Delete Package	N/A	Write
	Rename Package	N/A	Write
	Cut Package	N/A	Write
	Paste Package	N/A	Write
	Move Package	N/A	Write
Manage Package: Read & Write	Edit Package Properties	N/A	Read
Manage Package: Read	Export Package	N/A	Read
	Open Package (View)	N/A	Read

Chef Cookbook Management Permissions

This section specifies the Chef Cookbook Management permissions required by users to perform specific actions in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

Note: In addition to the action permissions listed, every user action also requires the Managed Servers and Groups permission.

Permissions for Running a Chef Recipe from a Cookbook with No Dependencies

The following permissions are required in order to run a Chef Recipe from a cookbook with no dependencies:

- These **Action Permissions** control the Chef tasks you can perform.

Permission	Setting	Task Enabled
Run Chef Recipes	Yes	The ability to start or schedule a specific Run Chef Recipe job.
Manage Package	Read (or stronger)	The ability to use Cookbooks (which is a type of SA package) in Run Chef Recipe jobs.

The user running the Run Chef Recipe job must belong to a user group with the *Run Chef Recipes* and *Manage package* permissions.

- **Folder Permissions** control the access to the SA Library folder where the cookbook resides.

The user running the Run Chef Recipe job must belong to a user group with *Read* permission on the folder where the cookbook resides.

- **Resource Permissions** control the access of the current user to the managed servers in SA.

The user running the Run Chef Recipe job must belong to a user group with *Read&Write* permission on the server's facility, customer, and at least one of its Device Groups.

For more information about setting resource permissions, see [About Resource Permissions](#).

- **Customer Constraints on Folders** determine which servers can be the target of a Run Chef Recipe job. As each server is assigned to a *Customer*, the customer constraints of the cookbook folder must include the Customer of the target server.

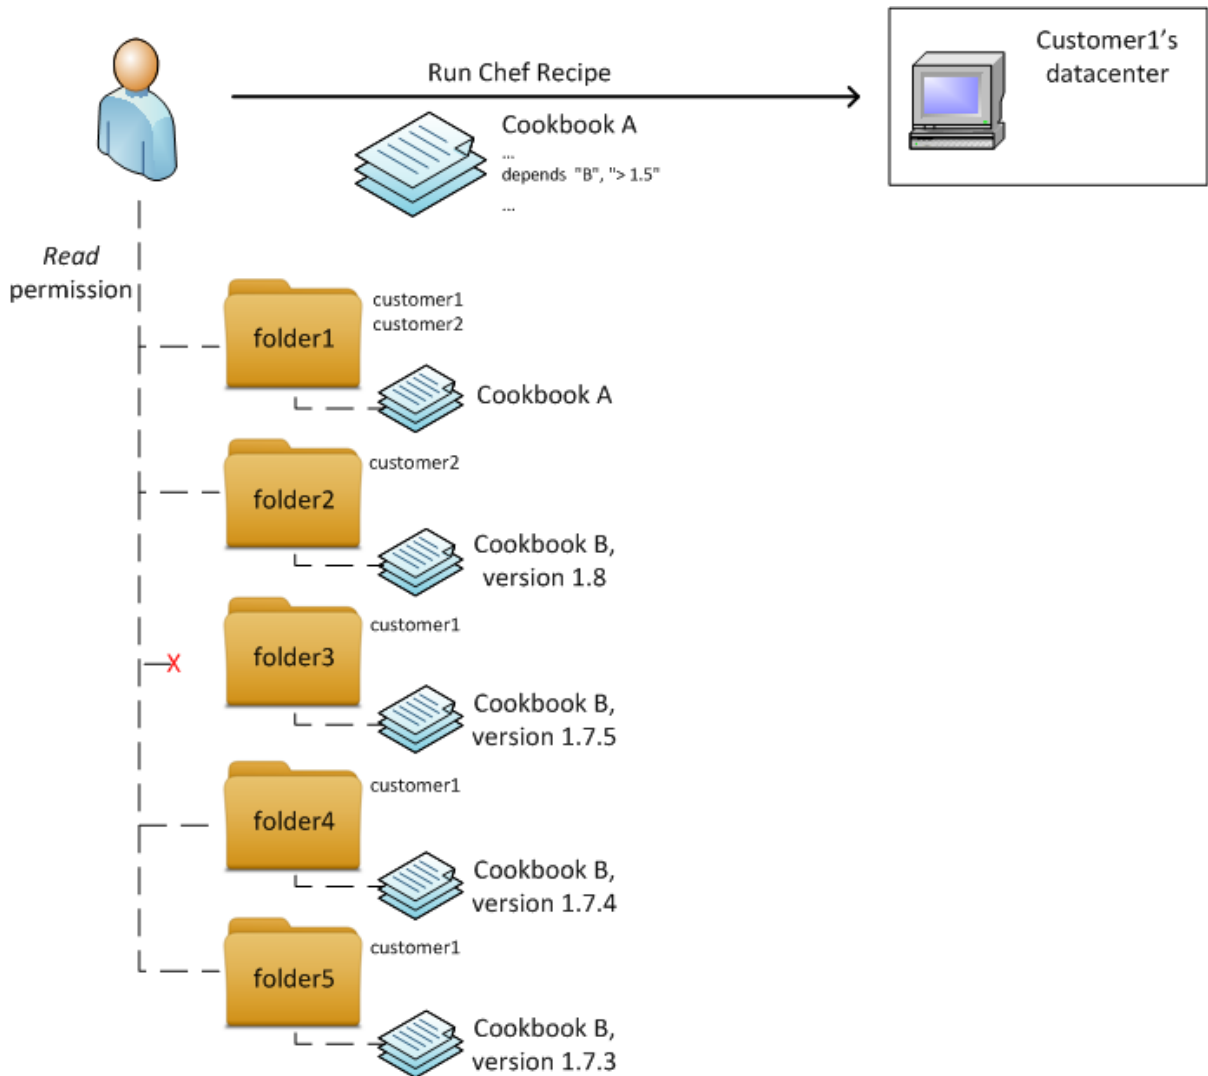
Alternatively, you can ignore folder customer permissions entirely by assigning the *Customer Independent* customer to the cookbook folders.

For more information about setting folder permissions, see [About Folder Permissions](#).

Permission Management for Cookbooks with Dependencies

The dependencies of a cookbook must satisfy the same permission requirements as the main cookbook: Read folder permissions and the proper folder customer constraints. If multiple versions of the dependent cookbooks exist, SA will use the newest version of the dependent cookbooks for which the entire dependency graph satisfies all required permissions.

Example: In the following setup, when the user tries to run a recipe from cookbook A, SA will resolve its dependency on cookbook B to version 1.7.4.

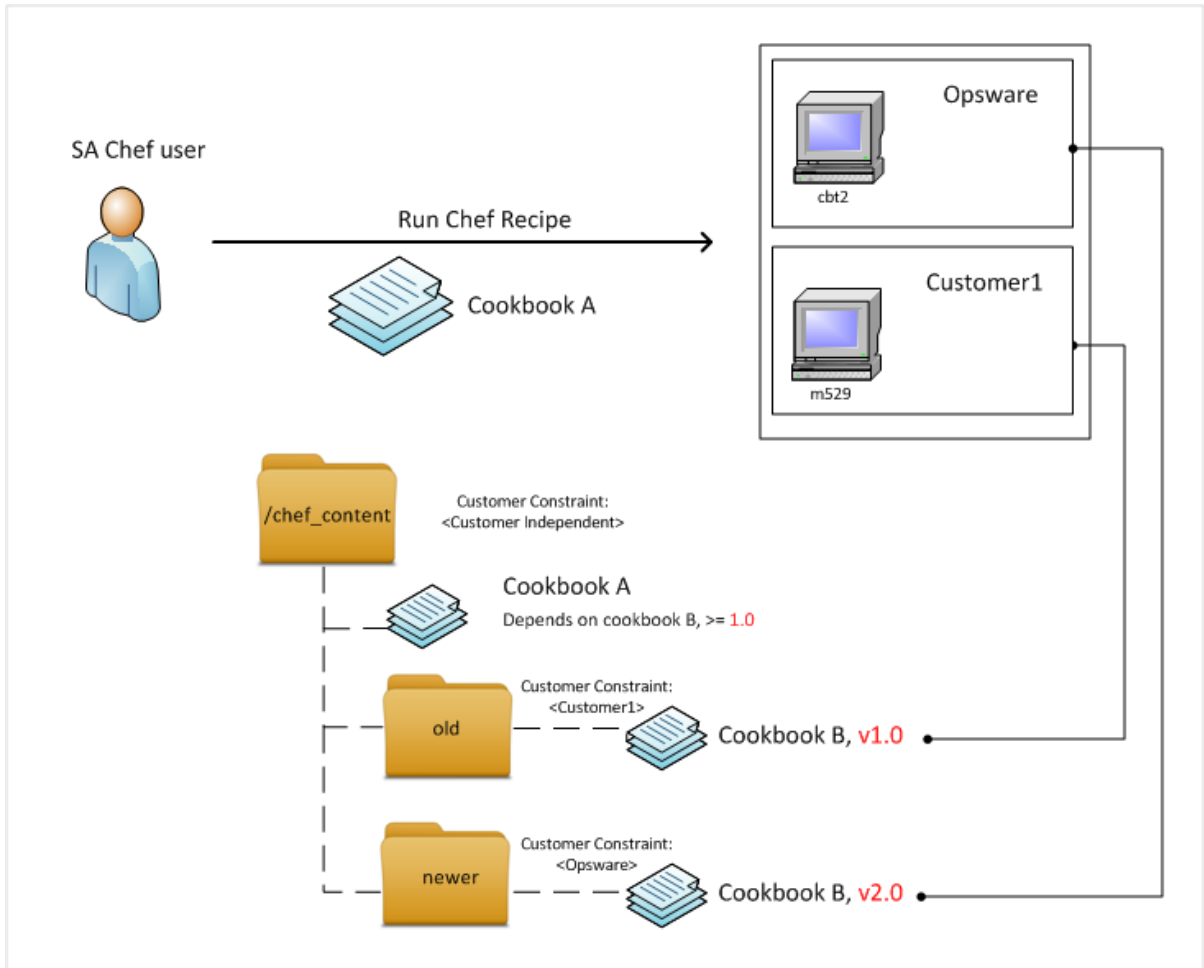


More in-depth, version 1.8 of cookbook B cannot be used because folder2 is not associated to customer1 (the customer of the targeted server). Version 1.7.5 of cookbook B can't be used because the user doesn't have any permissions on folder3. Versions 1.7.4 and 1.7.3 are both accessible and SA will choose the higher version, therefore 1.7.4.

Multi-tenancy

Customer constraints on folders provide the mechanism to support multi-tenancy, which allows you to apply different content to different customers.

In the example below, applying cookbook A to a group of two managed servers (cbt2 and m529) will result in applying version 1.0 of cookbook B to server m529 and version 2.0 of cookbook B to server cbt2.



Application Configuration Management Permissions

[Application Configuration Management Permissions Required for User Actions](#) specifies the permissions required by users to perform specific actions with application configurations in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

Note: In addition to the action permissions listed in [Application Configuration Management Permissions Required for User Actions](#), every user action also requires the Managed Servers and Groups permission.

In [Application Configuration Management Permissions Required for User Actions](#), the Server Permission column is for the servers referenced by the application configuration or configuration template. Server permissions are specified by the Customer, Facility, and Device Groups

permissions in the SA Client. In [Application Configuration Management Permissions Required for User Actions](#), the Folder Permission column is for the folders in the SA Library that contain the application configurations and configuration templates.

To perform an action, the user requires several permissions. For example, to attach an application configuration to a server, the user must have the following permissions:

- Manage Application Configurations: Read
- Manage Configuration Templates: Read
- Manage Installed Configuration and Backups on Servers: Read & Write
- Managed Servers and Groups
- Read & Write permissions to the facility, device group, and customer of the server
- Read permission for the folder in the SA library that contains the application configuration or template

Application Configuration Management Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
Application Configuration			
Create Application Configuration	Manage Application Configurations: Read & Write and Manage Configuration Templates: Read	None	Read & Write
View Application Configuration	Manage Application Configurations: Read & Write and Manage Configuration Templates: Read	None	Read
Edit Application Configuration	Manage Application Configurations: Read & Write and Manage Configuration Templates: Read	None	Read & Write
Delete Application	Manage Application	None	Read & Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
Configuration	Configurations: Read & Write and Manage Configuration Templates: Read		
Specify Template Order	Manage Application Configurations: Read & Write and Manage Configuration Templates: Read	None	Read & Write
Attach Application Configuration to Server	Manage Application Configurations: Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read & Write	Read & Write	Read
Attach Application Configuration to Device Group	Manage Application Configurations: Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read & Write and Manage Public Device Group: Yes and Model Public Device Group: Yes	Read & Write	Read
Set Application Configuration Values on Server	Manage Application Configurations:	Read & Write	Read

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
	Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read & Write		
Push Application Configuration to Server	Manage Application Configurations: Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read & Write	Read & Write	Read
Schedule Application Configuration Push	Manage Application Configurations: Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read & Write	Read & Write	Read
Scan Configuration Compliance	Allow Configuration Compliance Scan: Yes and Manage Application Configurations: Read and Manage Configuration Templates: Read	Read	Read
Schedule Application	Allow Configuration	Read	Read

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
Configuration Audit	Compliance Scan: Yes and Manage Application Configurations: Read and Manage Configuration Templates: Read		
Roll Back (Revert) Application Configuration Push	Manage Application Configurations: Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read & Write	Read & Write	Read
Application Configuration Templates			
Create Application Configuration Template	Manage Configuration Templates: Read & Write	None	Read & Write
View Application Configuration Template	Manage Configuration Templates: Read & Write	None	Read
Edit Application Configuration Template	Manage Configuration Templates: Read & Write	None	Read & Write
Delete Application Configuration Template	Manage Configuration Templates: Read & Write	None	Read & Write
Load (Import) Application Configuration Template	Manage Application Configurations: Read & Write and Manage Configuration Templates:	None	Read & Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
	Read & Write		
Set Application Configuration Template to Run as Script	Manage Configuration Templates: Read & Write	None	Read & Write
Compare Two Application Configuration Templates	Manage Configuration Templates: Read	None	Read
Compare Application Configuration Template Against Actual Configuration File (Preview)	Manage Application Configurations: Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read	Read	Read

[User Actions Allowed by Application Configuration Management Permissions](#) lists the actions that users can perform with application configurations for each permission. [User Actions Allowed by Application Configuration Management Permissions](#) has the same data as [Application Configuration Management Permissions Required for User Actions](#), but is sorted by permission. Although not indicated in [User Actions Allowed by Application Configuration Management Permissions](#), the Managed Servers and Groups permission is required for all OS provisioning actions.

For security administrators, [User Actions Allowed by Application Configuration Management Permissions](#) answers this question: If a user is granted a particular permission, what actions can the user perform?

User Actions Allowed by Application Configuration Management Permissions

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
Allow Configuration Compliance Scan: Yes and Manage Application Configurations: Read and Manage Configuration Templates: Read	Scan Configuration Compliance	Read	Read
	Schedule Application Configuration Audit	Read	Read
Manage Application Configurations: Read & Write and Manage Configuration Templates: Read	Create Application Configuration	None	Read & Write
	Delete Application Configuration	None	Read & Write
	Edit Application Configuration	None	Read & Write
	Specify Template Order	None	Read & Write
	View Application Configuration	None	Read
Manage Application Configurations: Read & Write and Manage Configuration Templates: Read & Write	Load (Import) Application Configuration Template	None	Read & Write
Manage Application Configurations: Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read	Compare Application Configuration Template Against Actual Configuration File (Preview)	Read	Read

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
Manage Application Configurations: Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read & Write	Attach Application Configuration to Server	Read & Write	Read
	Push Application Configuration to Server	Read & Write	Read
	Roll Back (Revert) Application Configuration Push	Read & Write	Read
	Schedule Application Configuration Push	Read & Write	Read
	Set Application Configuration Values on Server	Read & Write	Read
Manage Application Configurations: Read and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read & Write and Manage Public Device Group: Yes and Model Public Device Group: Yes	Attach Application Configuration to Device Group	Read & Write	Read
Manage Configuration Templates: Read	Compare Two Application Configuration Templates	None	Read
Manage Configuration Templates: Read & Write	Create Application Configuration Template	None	Read & Write
	Delete Application Configuration Template	None	Read & Write
	Edit Application	None	Read & Write

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
	Configuration Template		
Manage Configuration Templates: Read & Write (cont.)	Set Application Configuration Template to Run as Script	None	Read & Write
	View Application Configuration Template	None	Read

Patch Management for Windows Permissions

Table 49 specifies the Windows Patch Management permissions required by users to perform specific actions in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

Note: In addition to the permissions listed in **Table 49**, every user action also requires the Managed Servers and Groups permission.

In **Table 49**, most of the entries in the User Action column correspond to menu items in the SA Client. In addition to action permissions, server permissions are required on the managed servers affected by the patching operation.

Note: If either Allow Install Patch or Allow Uninstall Patch permission is set to Yes, then the Manage Patch and the Manage Windows Patch Policies permissions are automatically set to Read.

Table 49. Windows Patch Management Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)
Patches		

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)
Install Patch (Available)	Allow Install Patch: Yes Manage Patch: Read	Read & Write
Uninstall Patch (Available)	Allow Uninstall Patch: Yes and Manage Patch: Read	Read & Write
Install Patch (Limited Availability)	Allow Install Patch: Yes Manage Patch: Read & Write	Read & Write
Uninstall Patch (Limited Availability)	Allow Uninstall Patch: Yes and Manage Patch: Read & Write	Read & Write
Open Patch (View Patch)	Manage Patch: Read	N/A
Change Patch Properties	Manage Patch: Read & Write	N/A
Import Patch	Manage Patch: Read & Write and Package	N/A
Import Patch Database	Manage Patch: Read & Write	N/A
Export Patch	Manage Patch: Read and Package	N/A
Export Patch	or Allow Install Patch: Yes and Package: Yes	N/A
Export Patch	or Allow Uninstall Patch: Yes and Package	N/A
Export Patch	or Manage Policy: Read and Package	N/A
Delete Patch	Manage Patch: Read & Write	N/A
Patch Policies and Exceptions		
Remediate Policy	Allow Install Patch: Yes	Read & Write
Open Patch Policy (View)	Manage Windows Patch Policy: Read	N/A
Add Patch to Patch Policy	Manage Patch: Read and Manage Windows Patch Policy: Read & Write	N/A

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)
Remove Patch from Patch Policy	Manage Windows Patch Policy: Read & Write	N/A
Set Exception	Allow Install Patch: Yes	Read & Write
Set Exception	or Allow Uninstall Patch: Yes	Read & Write
Copy Exception	Allow Install Patch: Yes	Read & Write
Copy Exception	or Allow Uninstall Patch: Yes	Read & Write
Attach Patch Policy to Server (or Device Group)	Manage Windows Patch Policy: Read	Read & Write
Detach Patch Policy from Server (or Device Group)	Manage Windows Patch Policy: Read	Read & Write
Create Patch Policy	Manage Windows Patch Policy: Read & Write	N/A
Delete Patch Policy	Manage Windows Patch Policy: Read & Write	N/A
Change Patch Policy Properties	Manage Windows Patch Policy: Read & Write	N/A
Patch Compliance Rules		
Edit Patch Products (Patch Configuration window)	Manage Patch Compliance Rules: Yes	N/A
Scan Patch Compliance	Manage Windows Patch Policy: Read	N/A
Schedule a Patch Policy Scan	Manage Patch Compliance Rules: Yes	N/A
Change Default Patch Availability	Manage Patch Compliance Rules: Yes	N/A
Change Patch Policy Compliance Rules	Manage Patch Compliance Rules: Yes	N/A
View Patch Policy Compliance Rules	Manage Windows Patch Policy: Yes	N/A

Table 50 lists the actions that users can perform for each Patch Management permission. **Table 50** has the same data as **Table 49**, but is sorted by action permission. Although it is not indicated in **Table 50**, the Managed Servers and Groups permission is required for all Patch Management actions.

For security administrators, **Table 50** answers this question: If a user is granted a particular action permission, what actions can the user perform?

Table 50. User Actions Allowed by Windows Patch Management Permissions

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)
Allow Install Patch: Yes	Copy Exception	Read & Write
	Remediate Policy	Read & Write
	Set Exception	Read & Write
Allow Install Patch: Yes and Manage Patch: Read	Install Patch (Available)	Read & Write
	Uninstall Patch (Available)	Read & Write
Allow Install Patch: Yes and Manage Patch: Read & Write	Install Patch (Limited Availability)	Read & Write
	Uninstall Patch (Limited Availability)	Read & Write
Allow Install Patch: Yes and Package: Yes	Export Patch	N/A
Allow Uninstall Patch: Yes	Copy Exception	Read & Write
	Set Exception	Read & Write
Allow Uninstall Patch: Yes and Package	Export Patch	N/A
Allow Uninstall Patch: Yes and Manage Patch: Read	Uninstall Patch	Read & Write
Manage Patch Compliance Rules: Yes	Change Default Patch Availability	N/A
	Change Patch Policy Compliance Rules	N/A
	Edit Patch Products (Patch Configuration window)	N/A
	Schedule a Patch Policy Scan	N/A
Manage Windows Patch Policy: Read	Attach Patch Policy to Server (or Device Group)	Read & Write
	Detach Patch Policy from Server (or	Read & Write

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)
	Device Group)	
	Open Patch Policy (View)	N/A
Manage Windows Patch Policy: Read & Write	Change Patch Policy Properties	N/A
	Create Patch Policy	N/A
	Delete Patch Policy	N/A
	Remove Patch from Patch Policy	N/A
Manage Windows Patch Policy: Yes	View Patch Policy Compliance Rules	N/A
Manage Patch: Read	Open Patch (View Patch)	N/A
	Scan Patch Compliance	
Manage Patch: Read & Write	Change Patch Properties	N/A
	Delete Patch	N/A
	Import Patch Database	N/A
Manage Patch: Read & Write and Package	Import Patch	N/A
Manage Patch: Read and Manage Windows Patch Policy: Read & Write	Add Patch to Patch Policy	N/A
Manage Patch: Read and Package	Export Patch	N/A
Manage Policy: Read and Package	Export Patch	N/A

Patch Management for Ubuntu Permissions

In Ubuntu Patch Management, all user roles are combined, which means that a single user can perform all patch management actions. Ubuntu out-of-the-box settings give the user the following User Group roles:

- Patch Policy Setter
- Patch Deployer
- Software Policy Setter
- Policy Deployer

In addition, the conditions listed as follows must be met:

- To configure Ubuntu patch policies:
 - The user must belong to both Patch Policy Setters and Software Policy Setters User Groups.
 - The user must have Read & Write resource permissions for the Customers to which the server belongs.
 - The Datacenter must be added for both of the above groups.
- To deploy Ubuntu Patch Policies:
 - The user must belong to both Patch Deployers and Software Deployers user groups.
 - The user must have Read & Write resource permissions for the Customers to which the server belongs.
 - The Datacenter must be added for both of the above groups.
- To attach Ubuntu Patch Policies to a Ubuntu server:
 - The user must to have Read & Write permissions on the folder where the target patch policies reside.
 - To import a Debian package, the user must have Read & Write resource permission on the Opsware Customer.

Note: See [Patch Management for Windows Permissions](#) for the standard patching action permissions.

For users in User Group Roles at the facility where the server is managed to have the correct permissions to use Ubuntu patching, they must have the folder permissions shown in **Table 51**.

Table 51. Folder Permissions for Ubuntu User Group Roles

Folder	User Group Role	Permission
/Opsware	Patch Policy Setter	Read & Write
/Opsware	Software Policy Setter	Read & Write
/Opsware	Patch Policy Deployer	Read
/Opsware	Software Policy Deployer	Read
/Opsware	Superuser	Read & Write
/Opsware	Opsware System Administrator	Read & Write

Folder	User Group Role	Permission
/Opware/Patching/Tools	Patch Policy Setter	Read, List, Execute
/Opware/Patching/Tools	Software Policy Setter	Read, List, Execute
/Opware/Patching/Tools	Patch Policy Deployer	Read, List, Execute
/Opware/Patching/Tools	Software Policy Deployer	Read, List, Execute
/Opware/Patching/Tools	Superuser	Read, List, Execute
/Opware/Patching/Tools	Opware System Administrator	Read, List, Execute
/Opware/Patching/Tools	Command-Line Administrator	Read, List, Execute

Patch Management for Solaris Permissions

This section describes permissions for managing patches on Solaris systems. For patch information on other UNIX systems, see [Patch Management for Other UNIX Permissions](#). For permissions on Solaris patch policies, see [Solaris Patch Policy Management Permissions](#).

Table 52 specifies the Patch Management permissions required by users to perform specific actions in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

Note: In addition to the permissions listed in **Table 52**, every user action also requires the Managed Servers and Groups permission.

In **Table 52**, most of the entries in the User Action column correspond to menu items in the SA Client. In addition to action permissions, server permissions are required on the managed servers affected by the patching operation.

Note: If either Allow Install Patch or Allow Uninstall Patch permission is set to Yes, then the Manage Patch and the Manage Windows Patch Policy permissions are automatically set to Read. If you plan to use Solaris patch policies, you should also set Manage Software Policy to Read or Read and Write. For more information, see [Solaris Patch Policy Management Permissions](#).

Table 52. Solaris Patch Management Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)
Patches		
Install Patch (Available)	Allow Install Patch: Yes Manage Patch: Read	Read & Write
Uninstall Patch (Available)	Allow Uninstall Patch: Yes Manage Patch: Read	Read & Write
Install Patch (Limited Availability)	Allow Install Patch: Yes Manage Patch: Read & Write	Read & Write
Uninstall Patch (Limited Availability)	Allow Uninstall Patch: Yes Manage Patch: Read & Write	Read & Write
Open Patch (View Patch)	Manage Patch: Read	N/A
Change Patch Properties	Manage Patch: Read & Write	N/A
Import Patch	Manage Patch: Read & Write	N/A
Export Patch	Manage Patch: Read Allow Install Patch: Yes (optional) Allow Uninstall Patch: Yes (optional) Manage Software Policy: Read (optional)	N/A
Delete Patch	Manage Patch: Read & Write	N/A

Table 53 lists the actions that users can perform for each Solaris Patch Management permission. **Table 53** has the same data as **Table 52**, but is sorted by action permission. Although it is not indicated in **Table 53**, the Managed Servers and Groups permission is required for all Patch Management actions.

For security administrators, **Table 53** answers this question: If a user is granted a particular action permission, what actions can the user perform?

Table 53. User Actions Allowed by Solaris Patch Management Permissions

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)
Allow Install Patch: Yes	Remediate Policy	Read & Write
Allow Install Patch: Yes Manage Patch: Read	Install Patch (Available)	Read & Write
	Uninstall Patch (Available)	Read & Write
Allow Install Patch: Yes Manage Patch: Read & Write	Install Patch (Limited Availability)	Read & Write
	Uninstall Patch (Limited Availability)	Read & Write
Allow Install Patch: Yes (Also sets Manage Patch: Read)	Export Patch	N/A
Allow Uninstall Patch: Yes (Also sets Manage Patch: Read)	Export Patch	N/A
Allow Uninstall Patch: Yes (Also sets Manage Patch: Read)	Uninstall Patch	Read & Write
Manage Patch: Read	Open Patch (View Patch)	N/A
	Export Patch	N/A
Manage Patch: Read & Write	Change Patch Properties	N/A
	Delete Patch	N/A
	Import Patch	N/A

Solaris Patch Policy Management Permissions

Table 54 specifies the Solaris Patch Policy Management permissions required by users to perform specific actions in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

If a customer is assigned to a folder, then customer constraints might limit the objects that can be associated with a Solaris patch policy contained in the folder. For a list of tasks affected by these constraints, see [Folders, Customer Constraints, and Software Policies](#).

Table 54. Solaris Patch Policy Management Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Solaris Patch Policy			
Create Solaris Patch Policy	Manage Software Policy: Read & Write	N/A	Write
Delete Solaris Patch Policy	Manage Software Policy: Read & Write	N/A	Write
Open Solaris Patch Policy (View)	Manage Software Policy: Read	N/A	Read
Edit Solaris Patch Policy Properties	Manage Software Policy: Read & Write	N/A	Write
Add Patches	Manage Software Policy: Read & Write Manage Patches: Read	N/A	Folder containing the software policy: Write
Add Scripts	Manage Software Policy: Read & Write Manage Server Scripts: Read	N/A	Folder containing the software policy: Write
Remove Patches	Manage Software Policy: Read & Write	N/A	Write
Remove Scripts	Manage Software Policy: Read & Write	N/A	Write
Attach Solaris Patch Policy	Manage Software Policy: Read Allow Attach/Detach Software Policy: Yes Model Public Device Groups: Yes (This permission is required if you are attaching the Solaris patch policy to a public device group.)	Read & Write	Read
Detach Solaris Patch Policy	Manage Software Policy: Read	Read & Write	Read

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
	Allow Attach/Detach Software Policy: Yes Model Public Device Groups: Yes (This permission is required if you are attaching the Solaris patch policy to a public device group.)		
Remediate	Manage Software Policy: Read Allow Remediate Servers: Yes Model Public Device Groups: Yes (Required if you remediate a public device group.)	Read & Write	Read
Scan Solaris Patch Compliance	N/A	Read	N/A
Rename Solaris Patch Policy	Manage Software Policy: Read & Write	N/A	Write
Cut Solaris Patch Policy	Manage Software Policy: Read & Write	N/A	Write
Copy Solaris Patch Policy	Manage Software Policy: Read	N/A	Read
Paste Solaris Patch Policy	Manage Software Policy: Read & Write	N/A	Source Folder: Read (for copy and paste) Source Folder: Write (for cut and paste) Destination Folder: Write
Move Solaris Patch Policy	Manage Software Policy: Read & Write	N/A	Source Folder: Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
			Destination Folder: Write

Patch Management for Other UNIX Permissions

This section describes permissions for managing patches on UNIX systems other than Solaris. For Solaris information, see [Patch Management for Solaris Permissions](#). You can use software policies with UNIX patches. For more information, see [Software Management Permissions](#).

Table 55 specifies the Patch Management permissions required by users to perform specific actions in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

Note: In addition to the permissions listed in **Table 55**, every user action also requires the Managed Servers and Groups permission.

In **Table 55**, most of the entries in the User Action column correspond to menu items in the SA Client. In addition to action permissions, server permissions are required on the managed servers affected by the patching operation.

Note: If either Allow Install Patch or Allow Uninstall Patch permission is set to Yes, then the Manage Patch and the Manage Windows Patch Policy permissions are automatically set to Read. If you plan to use policies, you should also set Manage Software Policy to Read or Read and Write.

Table 55. UNIX Patch Management Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)
Patches		
Install Patch (Available)	Allow Install Patch: Yes	Read & Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)
	Manage Patch: Read	
Uninstall Patch (Available)	Allow Uninstall Patch: Yes and Manage Patch: Read	Read & Write
Install Patch (Limited Availability)	Allow Install Patch: Yes Manage Patch: Read & Write	Read & Write
Uninstall Patch (Limited Availability)	Allow Uninstall Patch: Yes and Manage Patch: Read & Write	Read & Write
Open Patch (View Patch)	Manage Patch: Read	N/A
Change Patch Properties	Manage Patch: Read & Write	N/A
Export Patch	Manage Patch: Read and Package	N/A
Export Patch	or Allow Install Patch: Yes and Package: Yes	N/A
Export Patch	or Allow Uninstall Patch: Yes and Package	N/A
Export Patch	or Manage Policy: Read and Package	N/A
Delete Patch	Manage Patch: Read & Write	N/A

Table 56 lists the actions that users can perform for each Patch Management permission. **Table 56** has the same data as **Table 55**, but is sorted by action permission. Although it is not indicated in **Table 56**, the Managed Servers and Groups permission is required for all Patch Management actions.

For security administrators, **Table 56** answers this question: If a user is granted a particular action permission, what actions can the user perform?

Table 56. User Actions Allowed by UNIX Patch Management Permissions

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)
Allow Install Patch: Yes	Copy Exception	Read & Write
	Remediate Policy	Read & Write
	Set Exception	Read & Write
Allow Install Patch: Yes and Manage Patch: Read	Install Patch (Available)	Read & Write
	Uninstall Patch (Available)	Read & Write
Allow Install Patch: Yes and Manage Patch: Read & Write	Install Patch (Limited Availability)	Read & Write
	Uninstall Patch (Limited Availability)	Read & Write
Allow Install Patch: Yes and Package: Yes	Export Patch	N/A
Allow Uninstall Patch: Yes	Copy Exception	Read & Write
	Set Exception	Read & Write
Allow Uninstall Patch: Yes and Package	Export Patch	N/A
Manage Patch: Read	Open Patch (View Patch)	N/A
Manage Patch: Read & Write	Change Patch Properties	N/A
	Delete Patch	N/A
	Import Patch Database	N/A
Manage Patch: Read & Write and Package	Import Patch	N/A
Manage Patch: Read and Manage Policy: Read & Write	Add Patch to Policy	N/A
Manage Patch: Read and Package	Export Patch	N/A
Manage Policy: Read and Package	Export Patch	N/A

Audit and Remediation Permissions

Table 57 specifies the Audit and Remediation permissions required by users to perform specific actions in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

Note: In addition to the permissions listed in **Table 57**, every user action also requires the Managed Servers and Groups permission.

Server Permissions for Audit and Remediation

Audit and Remediation actions require both action and server permissions. For example, the Create Audit action requires the action permission “Manage Audit: Read & Write” and the Managed Servers and Groups permission. This action also needs Read permission on the server referenced by the Audit. In **Table 57**, the Server Permission column is for the servers referenced by the Audit or Snapshot Specification—depending on the action. Server permissions are specified by the customer, facility, and device groups permissions in the SA Client.

If an Audit and Remediation object (such as a snapshot specification) references multiple servers, at a minimum, Read permission is required for all servers referenced. Otherwise, the object cannot be viewed or modified.

Audit and Remediation objects are not directly associated with customers and facilities. Customer and facility permissions do control access to servers that are referenced by Audit and Remediation objects, such as snapshot specifications and audits.

“Allow Create Task Specific Policy Permission” for Audit and Remediation

As a best practice, do *not* enable this permission—do *not* set this permission to “Yes.” By default, this permission is disabled—it is already set to “No.” It is recommended that you create audit rules in an audit policy and then, subsequently, link audit tasks and snapshot specifications to that audit policy.

OGFS Permissions for Audit and Remediation

For the actions that access a managed server’s file system, the OGFS Read Server File System permission is required. For example, the Read Server File System permission is required to create a snapshot specification with rules that include the files of a managed server. Such rules include Application Configurations, Custom Scripts, COM+ objects, File System, IIS Metabase entries, and Windows Registry.

Other types of selection criteria require the corresponding OGFS permissions:

- Read Server Registry
- Read COM+ Database

- Read IIS Metabase

Audit and Remediation User Action Permissions

The following table lists typical Audit and Remediation user actions and the permissions required to perform them.

Table 57. Audit and Remediation Permissions Required for User Actions

User Action	Action Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Snapshot Specification			
View contents of Snapshot Specification	Manage Snapshot Specification: Read & Write	N/A	Read & Write
Schedule and run a Snapshot Specification	Manage Snapshot Specification: Read & Write	N/A	Read & Write
Create Snapshot Specification	Manage Snapshot Specification: Read & Write	N/A	Read & Write
Create Application Configuration Rule	Manage Snapshot Specification: Read & Write	Write Server File System	Read & Write
Create COM+ Rule	Manage Snapshot Specification: Read & Write	Read COM+ Database	Read & Write
Create Custom Script Rule	Manage Snapshot Specification: Read & Write Allow Create Custom Script Policy Rules: Yes.	Write Server File System	Read & Write
Create Files	Manage Snapshot Specification: Read & Write	Write Server File System	Read & Write
Create IIS Metabase Rule	Manage Snapshot Specification: Read & Write	Read IIS Metabase	Read & Write
Create Registry Rule	Manage Snapshot Specification: Read & Write	Read Server Registry	Read & Write
Link Audit Policy into Snapshot Specification	Manage Snapshot Specification: Read & Write Manage Audit Policy: Read	N/A	Read & Write

User Action	Action Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
	Library Folder: Read		
Import Audit Policy into Snapshot Specification	Manage Snapshot Specification: Read & Write Manage Audit Policy: Read Library Folder: Read	N/A	Read & Write
Save As Audit Policy	Manage Snapshot Specification: Read & Write Manage Audit Policy: Read & Write Library Folder: Read & Write	N/A	Read & Write
Snapshots			
View, list contents of a Snapshot	Manage Snapshot: Read Manage Snapshot Specification: Read	N/A	Read
Create Audit from Snapshot	Manage Snapshot: Read Manage Snapshot Specification: Read Manage Audit: Read	N/A	Read
View Archived Snapshot	Manage Snapshot: Read	N/A	Read
Create Audit from archived Snapshot	Manage Snapshot: Read Manage Audit: Read	N/A	Read
Delete Snapshot results	Manage Snapshot: Read & Write	N/A	Read & Write
Detach Snapshot from a server	Allow General Snapshot Management: Yes Manage Snapshot: Read & Write Manage Snapshot Specification: Read	N/A	Read

User Action	Action Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Remediate Snapshot results	Manage Snapshot: Read Manage Snapshot Specification: Read Allow Remediate Audit/Snapshot Results: Yes	N/A	Read & Write
Remediate Snapshot Results: Application Configuration	Manage Snapshot: Read Allow Remediate Audit/Snapshot Results: Yes Manage Snapshot Specification: Read	Write Server File System	Read & Write
Remediate Snapshot Results: COM+	Manage Snapshot: Read Allow Remediate Audit/Snapshot Results: Yes Manage Snapshot Specification: Read	Read COM+ Database	Read & Write
Remediate Snapshot Results: Custom Scripts	Manage Snapshot: Read Allow Remediate Audit/Snapshot Results: Yes Manage Snapshot Specification: Read	Write Server File System	Read & Write
Remediate Snapshot Results: File System	Manage Snapshot: Read Allow Remediate Audit/Snapshot Results: Yes Manage Snapshot Specification: Read	Write Server File System	Read & Write
Remediate Snapshot Results: Metabase	Manage Snapshot: Read Allow Remediate Audit/Snapshot Results: Yes Manage Snapshot Specification: Read	Read IIS Metabase	Read & Write

User Action	Action Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Remediate Snapshot Results: Registry	Manage Snapshot: Read Allow Remediate Audit/Snapshot Results: Yes Manage Snapshot Specification: Read	Read Server Registry	Read & Write
Audits			
View an Audit	Manage Audit: Read	N/A	Read & Write
Run an Audit	Manage Audit : Read	N/A	Read & Write
Schedule an Audit	Manage Audit : Read	N/A	Read & Write
Create an Audit	Manage Audit: Read & Write	N/A	Read
Create Application Configuration Rule	Manage Audit: Read & Write	Write Server File System	Read & Write
Create COM+ Rule	Manage Audit: Read & Write	Read COM+ Database	Read & Write
Create Custom Script Rule	Manage Audit: Read & Write Allow Create Custom Script Policy Rules: Yes	Write Server File System	Read & Write
Create Discovered Software Rule	Manage Audit: Read & Write Manage Server Modules: Read	N/A	Read & Write
Create Files Rule	Manage Audit: Read & Write	Write Server File System	Read & Write
Create Hardware Rule	Manage Audit: Read & Write	N/A	Read & Write
Create IIS Metabase Rule	Manage Audit: Read & Write	Read IIS Metabase	Read & Write
Create Internet Information Server Rule	Manage Audit: Read & Write	N/A	Read & Write
Create Registered Software	Manage Audit: Read & Write	N/A	Read & Write

User Action	Action Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Rule	Manage Server Modules: Read		
Create Software Rule	Manage Audit: Read & Write	N/A	Read & Write
Create Storage Rule	Manage Audit: Read & Write Manage Server Modules: Read	N/A	Read & Write
Create Weblogic Rule	Manage Audit: Read & Write Manage Server Modules: Read	N/A	Read & Write
Create .NET Framework Configurations Rule	Manage Audit: Read & Write Manage Server Modules: Read	N/A	Read & Write
Create Windows Registry Rule	Manage Audit: Read & Write	Read Server Registry	Read & Write
Create Windows Services Rule	Manage Audit: Read & Write	N/A	Read & Write
Create Windows/UNIX Users and Groups Rule	Manage Audit: Read & Write Manage Server Modules: Read	N/A	Read & Write
Link an Audit Policy into an Audit	Manage Audit: Read & Write Manage Audit Policy: Read SA Client Library Folder: Read	N/A	Read & Write
Import an Audit Policy into an Audit	Manage Audit: Read & Write Manage Audit Policy: Read Library Folder: Read	N/A	Read & Write
Save as Audit Policy	Manage Audit: Read & Write Manage Audit Policy: Read & write	N/A	Read & Write

User Action	Action Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
	Library Folder: Read & Write		
Audit Results			
View Audit Results	Manage Audit Results: Read Manage Audit: Read	N/A	Read
View Archived Audit Results	Manage Audit: Read	N/A	Read
Delete Audit Results	Manage Audit Results: Read & Write	N/A	Read & Write
Remediate Audit Results	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes	N/A	Read & Write
Remediate Audit Results: Application Configuration	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes	Write Server File System	Read & Write
Remediate Audit Results: COM+	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes	Read COM+ Database	Read & Write
Remediate Audit Results: Custom Script Rule	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes	Write Server File System	Read & Write
Remediate Audit Results:	Manage Audit: Read	N/A	Read & Write

User Action	Action Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Discovered Software	Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes Manage Server Module: Read Allow Execute Server Modules: Yes		
Remediate Audit Results: Files	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes	Write Server File System	Read & Write
Remediate Audit Results: IIS Metabase	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes	Read IIS Metabase	Read & Write
Remediate Audit Results: Remediate Internet Information Server	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes	Read IIS Metabase	Read & Write
Remediate Audit Results: Remediate Discovered Software	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes Manage Server Module: Read Allow Execute Server Modules: Yes	N/A	Read & Write

User Action	Action Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Remediate Audit Results: Remediate Software	Manage Audit: Read Manage Audit Results: Read & Write	N/A	Read & Write
Remediate Audit Results: Remediate Storage	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes Manage Server Module: Read Allow Execute Server Modules: Yes	N/A	Read & Write
Remediate Audit Results: Remediate Weblogic	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes Manage Server Module: Read Allow Execute Server Modules: Yes	N/A	Read & Write
Remediate Audit Results: Remediate Windows .NET Framework Configurations	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes Manage Server Module: Read Allow Execute Server Modules: Yes	N/A	Read & Write
Remediate Audit Results: Windows Registry	Manage Audit: Read Manage Audit Results: Read & Write	Read Server Registry	Read & Write

User Action	Action Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
	Allow Remediate Audit/Snapshot Results: Yes		
Remediate Audit Results: Windows Services	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes	N/A	Read & Write
Remediate Audit Results: Remediate Windows/UNIX Users and Groups	Manage Audit: Read Manage Audit Results: Read & Write Allow Remediate Audit/Snapshot Results: Yes Manage Server Module: Read Allow Execute Server Modules: Yes	N/A	Read & Write

Table 58 lists the actions that users can perform for each Audit and Remediation permission. **Table 58** has the same data as **Table 57**, but is sorted by action permission. Although it is not indicated in **Table 58**, the Managed Servers and Groups permission is required for all Audit and Remediation actions.

For security administrators, **Table 58** answers this question: If a user is granted a particular action Audit and Remediation permission, what actions can the user perform?

Table 58. User Actions Allowed by Audit and Remediation Permissions

Action Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Allow Create Custom Script Rule Policy: No and	View Custom Script Rule: Audit	N/A	Read

Action Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Manage Audit: Read			
Allow Create Custom Script Rule Policy: Yes and Manage Audit: Read & Write	Create Custom Script Rule: Audit	Write Server File System	Read & Write
Allow Create Custom Script Rule Policy: No and Manage Snapshot: Read & Write	View Custom Script Rule: Snapshot	N/A	Read
Allow Create Custom Script Rule Policy: Yes and Manage Snapshot: Read & Write	Create Custom Script Rule: Snapshot	Write Server File System	Read & Write
Allow General Snapshot Management: Yes	Detach Snapshot from a server	N/A	Read
Manage Snapshot Specification: Read & Write and Allow Remediate Audit/Snapshot Results: No and Manage Audit or Manage Snapshot: Read	View Audit or Snapshot, No Remediation	N/A	Read
Manage Snapshot Specification: Read and Allow Remediate Audit/Snapshot Results: Yes	Remediate Audit/Snapshot Results	N/A	Read & Write

Action Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
and Manage Audit or Manage Snapshot: Read & Write			
Manage Snapshot Specification: Read and Allow Remediate Audit/Snapshot Results: Yes and Manage Audit or Manage Snapshot Results: Read & Write	Remediate Application Configuration Rule	Write Server File System	Read & Write
	Remediate COM+ Rule	Read COM+ Database	Read & Write
	Remediate Custom Script Rule Registry Rule	Write Server File System	Read & Write
	Remediate File System Rule	Read IIS Metabase	Read & Write
	Remediate IIS Metabase Rule	Read Server Registry	Read & Write
	Remediate Windows Registry Rule	Write Server File System	Read & Write
Manage Audit: Read	View, schedule, run Audit	N/A	Read
	View, schedule, run Audit with custom scripts in it	N/A	Read & Write

Action Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Manage Audit: Read & Write	Create, edit, delete Audit	N/A	Read & Write
	Save Audit as Audit Policy	N/A	Read & Write
	Link Audit Policy into Audit	N/A	Read & Write
	Create Application Configuration Rule	Write Server File System	Read & Write
	Create COM+ Rule	Read COM+ Database	Read & Write
	Create File System Rule	Write Server File System	Read & Write
	Create IIS Metabase Rule	Read IIS Metabase	Read & Write
	Create Window Registry Rule	Read Server Registry	Read & Write
Manage Audit: Read & Write and Allow Create Custom Script Policy Rules: Yes	Create Custom Scripts Rule	Write Server File System	Read & Write
Manage Audit: Read & Write and Manage Server Module: Read	Create the following Audit Rules: <ul style="list-style-type: none"> • Discovered Software • Registered Software • Storage • Weblogic • Windows .NET Framework Configurations • Windows Users and Groups 	N/A	Read & Write
Manage Audit Results: Read	View Audit Results	N/A	Read

Action Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Manage Audit Results: Read & Write	Delete Audit Results	N/A	Read & Write
Manage Snapshot Specification: Read & Write	View, schedule, run Snapshot Specification	N/A	Read
	View, schedule, run Snapshot Specification with custom scripts in it	N/A	Read & Write
Manage Snapshot Specification: Read & Write	Create, edit, and delete Snapshot Specification	N/A	
	Save Snapshot Specification as Audit Policy (This action requires REad & Write for the library folder where policy lives.)	N/A	
	Link Audit Policy Into Audit	N/A	Read & Write
	Create Application Configuration Rule	Write Server File System	Read & Write
	Create COM+ Rule	Read COM+ Database	Read & Write
	Create Discovered Software		
	Create File System Rule	Write Server File System	Read & Write
	Create IIS Metabase Rule	Read IIS Metabase	Read & Write
Manage Snapshot Specification: Read & Write and Manage Server Module: Read	Create the following Snapshot Rules:	N/A	Read & Write
	<ul style="list-style-type: none"> • Discovered Software • Registered Software 		

Action Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
	<ul style="list-style-type: none"> • Storage • Weblogic • Windows .NET Framework Configurations • Windows Users and Groups 		
Manage Snapshot Specification: Read & Write and Create Custom Script Policy Rule	Create Custom Rule for Snapshot Specification	Write Server File System	Read & Write
Manage Snapshot: Read	View contents of Snapshot	N/A	Read
Manage Snapshot: Read & Write	Delete Snapshot results	N/A	Read & Write
Manage Audit Policy: Read	View contents of Audits and Snapshot Specifications	N/A	Read
Manage Audit Policy: Read & Write	Create, edit Audit Policy	N/A	Read & Write
	Create Application Configuration Rule	Write Server File System	Read & Write
	Create COM+ Rule	Read COM+ Database	Read & Write
	Create File System Rule	Write Server File System	Read & Write
	Create IIS Metabase Rule	Read IIS Metabase	Read & Write

Action Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
	Create Windows Registry Rule	Read Server Registry	Read & Write
Manage Audit Policy: Read & Write Manage Server Module: Read	Create the following Snapshot Rules: <ul style="list-style-type: none"> • Discovered Software • Registered Software • Storage • Weblogic • Windows .NET Framework Configurations • Windows Users and Groups 	N/A	Read & Write
Manage Audit Policy: Read & Write and Allow Create Custom Script Policy Rule	Create Custom Script Rule	Write Server File System	Read & Write

Compliance View Permissions

The following section describes the Compliance View permissions required by users to perform specific actions in the SA Client. For security administrators, the following table answers this question: To perform a particular action, what permissions does a user need?

Table 59. Compliance View Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)
Audit		
View Details	Manage Audit Result: Read	Read
Run Audit	Manage Audit: Read Manage Audit Result: Read & Write	Read & Write
Remediate	Allow Remediate Audit/Snapshot Result: Yes For other permissions needed to remediate for specific audit rules, see Audit and Remediation User Action Permissions and Table 58. User Actions Allowed by Audit and Remediation Permissions .	Read & Write
Software		
Remediate	Manage Software Policy: Read Allow Remediate Servers: Yes	Read & Write
Scan Device	Manage Software Policy: Read Or Allow Attach/Detach Software Policy: Yes Or Allow Install/Uninstall Software: Yes Or Allow Remediate Servers: Yes	Read & Write
Patch		
Remediate	Manage Patch Policy: Read Install Patch: Yes	Read & Write
Scan Device	Manage Patch: Read Or Manage Patch Policy: Read	Read & Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)
	Or Allow Install Patch: Yes Or Allow Uninstall Patch: Yes Or Allow Install/Uninstall Software Or Allow Remediate Servers	
App Config		
Viewing Details	Manage Application Configurations: Read	Read
Scan Device	Allow Configuration Compliance Scan: Yes	Read
Specific App Config Remediation	See Application Configuration Management Permissions for permissions required for remediating application configurations.	Read & Write

Job Permissions

To manage jobs in the SA Client, you must have the permissions described in **Table 60**. When you select the Edit or Cancel Any Job permission, the View All Jobs permission is automatically selected.

To view any job in the SA Client, you must have permissions to run or execute the job. For example, if you had the permissions for an action such as Manage Application Configurations set to Read, but did not have Write permissions for this action, you would not be able to see any Application Configuration Push jobs in the SA Client.

Table 60. Job Management Permissions

User Action	Action Permission
Enable Approval Integration	Manage Approval Integration

User Action	Action Permission
Set Job Types Requiring Approval	Manage Approval Integration
Invoke JobService API Methods to Manage Blocked (Pending Approval) Jobs (This action is performed by customized software on the backend, not by end-users logged onto the SA Client.)	Edit or Cancel Any Job View All Jobs
End (Cancel) Job	Edit or Cancel Any Job View All Jobs
Delete Schedule	Edit or Cancel Any Job View All Jobs

Script Execution Permissions

Table 61 specifies the Script Execution permissions required by users to perform specific actions in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

If a customer is assigned to a folder, then customer constraints might limit the objects that can be associated with a software policy contained in the folder. For a list of tasks affected by these constraints, see [Folders, Customer Constraints, and Software Policies](#).

Table 61. Script Execution Permissions Required for User Actions

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Creating a Non Super User Server Script	Manage Server Script: Read & Write	N/A	Write
Creating a Super User Server Script	Manage Server Script: Read & Write Allow Control of Super User Server Scripts: Yes	N/A	Write
Creating an OGFS Script	Manage OGFS Script: Read & Write	N/A	Write
Opening (Viewing all script properties)	Manage Server Script: Read	N/A	Execute

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
except script contents) a Non Super User Server Script			
Opening (Viewing all script properties including script contents) a Non Super User Server Script	Manage Server Script: Read	N/A	
Opening (Viewing all script properties except script contents) a Super User Server Script	Manage Server Script: Read Allow Control of Super User Server Scripts: Yes	N/A	
Opening (Viewing all script properties including script contents) a Super User Server Script	Manage Server Script: Read Allow Control of Super User Server Scripts: Yes	N/A	
Opening (Viewing all script properties except script contents) an OGFS Script	Manage OGFS Script: Read	N/A	Execute
Opening (Viewing all script properties including script contents) an OGFS Script	Manage OGFS Script: Read	N/A	Read
Editing Non Super User Server Script Properties	Manage Server Script: Read & Write Note: The Allow Control of Super User Server Scripts: Yes permission is required to edit the script property, "Can Run as Super User".	N/A	Write
Editing a Super User Server Script	Manage Server Script: Read and Write	N/A	Write

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
	Allow Control of Super User Server Scripts: Yes		
Editing OGFS Script Properties	Manage OGFSr Script: Read & Write	N/A	Write
Locating Server Script in Folders	Manage Server Script: Read	N/A	Read
Locating OGFS Script in Folders	Manage OGFS Script: Read	N/A	Read
Exporting a Server Script	Manage Server Script: Read	N/A	Read
Exporting an OGFS Script	Manage OGFS Script: Read	N/A	Read
Renaming a Server Script	Manage Server Script: Read & Write	N/A	Write
Renaming a Super User Server Script	Manage Server Script: Read & Write Allow Control of Super User Server Scripts: Yes	N/A	Write
Renaming an OGFS Script	Manage OGFS Script: Read & Write	N/A	Write
Deleting a Server Script	Manage Server Script: Read & Write	N/A	Write
Deleting a Super User Server Script	Manage Server Script: Read & Write Allow Control of Super User Server Scripts: Yes	N/A	Write
Deleting an OGFS Script	Manage OGFS Script: Read & Write	N/A	Write
Running Server Script as Super User	Managed Servers and Groups: Yes	Read and Write	Execute

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Running Server Script as a Super User (by copying the script contents from another script)	Manage Server Script: Read Run Ad-Hoc Scripts: Yes Run Ad-Hoc Scripts and Source Visible Server Scripts as Super User: Yes Managed Servers and Groups: Yes	Read and Write	Read
Running Server Script as a specified user	Managed Servers and Groups: Yes	Read and Write	Execute
Running Server Script as a specified user (by copying the script contents from another script)	Manage Server Script: Read Run Ad-Hoc Scripts: Yes Managed Servers and Groups: Yes	Read and Write	Read
Running Ad-Hoc Scripts	Run Ad-Hoc Scripts: Yes Managed Servers and Groups: Yes	Read and Write	N/A
Running Ad-Hoc Scripts as Super User	Run Ad-Hoc Scripts: Yes Run Ad-Hoc Scripts and Source Visible Server Scripts as Super User: Yes Managed Servers and Groups: Yes	Read and Write	N/A
Running OGFS Scripts	N/A	N/A	Execute

Table 62 lists the actions that users can perform for each Script Execution permission. **Table 62** has the same data as **Table 61**, but is sorted by action permission. For security administrators, **Table 62** answers this question: If a user is granted a particular action permission, what actions can the user perform?

Table 62. User Actions Allowed by Script Execution Permissions

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Manage Server Script: Read & Write	Creating a Non Super User Server Script	N/A	Write
	Editing Non Super User Server Script Properties	N/A	Write
	Deleting a Non Super User Server Script	N/A	Write
	Renaming a Non Super User Server Script	N/A	Write
Manage Server Script: Read	Opening (Viewing all script properties including script contents) a Non Super User Server Script Opening (Viewing all script properties including script contents) a Super User Server Script	N/A	Read
	Locating Server Script in Folders	N/A	Read
	Exporting Server Scripts	N/A	Read
Manage Server Script: Read	Opening (Viewing all script properties excluding script contents) a Non Super User Server Script Opening (Viewing all script properties excluding script contents) a Super User Server Script		Execute
Manage Server Script: Read & Write And Allow Control of Super	Creating a Super User Server Script	N/A	Write

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
User Server Scripts: Yes			
	Editing Super User Server Script Properties Editing Non Super User Server Script Properties	N/A	Write
	Renaming a Super User Server Script Renaming a Non Super User Server Script	N/A	Write
	Deleting a Super User Server Script Deleting a Non Super User Server Script	N/A	Write
Manage OGFS: Read & Write	Creating an OGFS Script	N/A	Write
	Editing OGFS Script Properties	N/A	Write
	Deleting an OGFS Script	N/A	Write
	Renaming an OGFS Script	N/A	Write
Manage OGFS Script: Read	Opening (Viewing all the OGFS Script Properties, including script contents) an OGFS Script	N/A	Read
	Locating OGFS in Folders	N/A	Read
	Exporting OGFS Scripts	N/A	Read
Manage OGFS Script: Read	Opening (Viewing all the OGFS Script Properties, excluding script contents) an OGFS Script	N/A	Execute

Action Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Run Ad-Hoc Scripts	Running Ad-Hoc scripts	Read and Write	N/A
Run Ad-Hoc Scripts and Source Visible Server Scripts as Super User	Running Ad-Hoc scripts as Super User Running any visible Saved Server Scripts as Super User. This also applies to Server Scripts that are not designated as Super User Scripts.	Read and Write	N/A
N/A	Running Non Super User Server Script	Read and Write	Execute
N/A	Running Private Scripts	Read and Write	Execute (on Home folder)
N/A	Running OGFS Scripts	N/A	Execute

The following table lists the script execution permissions required for running scripts using a software policy.

Table 63. Script Execution Permissions Required for Software Management

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Adding a Server Script to a software policy	Manage Server Scripts: Read	N/A	Read
Adding a Server Script to the Options step in the Remediate window	N/A	N/A	Execute
Adding a Server Script to the Options step in the Remediate window	Manage Server Scripts: Read Run Ad-Hoc Scripts: Yes	N/A	Read

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
(Copying the script contents)			
Adding a Super User Server Script to the Options step in the Remediate window	Manage Server Scripts: Read Run Ad-Hoc Scripts: Yes Run Ad-Hoc Scripts and Source Visible Server Scripts as Super User: Yes	N/A	Read
Specifying an Ad-Hoc Script to the Options step in the Remediate window	Run Ad-Hoc Scripts: Yes	N/A	N/A
Specifying an Super User Ad-Hoc Script to the Options step in the Remediate window	Run Ad-Hoc Scripts: Yes Run Ad-Hoc Scripts and Source Visible Server Scripts as Super User: Yes	N/A	N/A
Adding a Server Script to the Options step in the Install Software window	N/A	N/A	Execute
Adding a Server Script to the Options step in the Install Software window (Copying the script contents)	Manage Server Scripts: Read Run Ad-Hoc Scripts: Yes	N/A	Read
Adding a Super User Server Script to the Options step in the Install Software window	Manage Server Scripts: Read Run Ad-Hoc Scripts: Yes Run Ad-Hoc Scripts and Source Visible Server Scripts as Super User: Yes	N/A	Read
Specifying an Ad-Hoc Script to the Options step in the Install Software window	Run Ad-Hoc Scripts: Yes	N/A	N/A

User Action	Action Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Specifying an Super User Ad-Hoc Script to the Options step in the Install Software window	Run Ad-Hoc Scripts: Yes Run Ad-Hoc Scripts and Source Visible Server Scripts as Super User: Yes	N/A	N/A

Flow Permissions - HP Operations Orchestration

The following permissions are required to administer flows or to run flows in SA:

Table 64. Flow-Related Permissions

User Action	Permission
Configure SA-00 integration	Administer Flow Integrations
Run flows in the SA Client as an SA user	Run Flow

Service Automation Visualizer Permissions

Table 65 specifies the Service Automation Visualizer (SAV) permissions required to perform specific actions in the SA Client. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

In **Table 65**, most of the entries in the User Action column correspond to menu items in the SA Client. In addition to action permissions, server read permissions are required on the managed servers affected by the analyze operation, such as permissions to open a Remote Terminal or a Remote Desktop Client, open the Device Explorer, and open a Global Shell session from the Service Automation Visualizer.

Note: SAV permissions required to scan a server are the same for both physical servers and virtual servers.

For complete information, see the SA User Guide: Service Automation Visualizer.

Table 65. SAV Permissions Required for User Actions

User Action	Action Permission	Source Server Permission (Customer, Facility)	Folder Permission
SAV-Only Operations			
Launch the Service Automation Visualizer	Allow Analyze: Yes	Read	N/A
Generate a scan or refresh Snapshot— regular or virtual servers	Allow Analyze: Yes	Read	N/A
Create a Snapshot or edit a scheduled Snapshot	Allow Analyze: Yes Manage Business Applications: Read & Write	Read	N/A
Start, stop, pause, restart virtual server inside of SAV (pause VM for VMware only—cannot pause a Solaris local zone)	Administer Virtual Server: Yes	Read	N/A
SA Client Operations			
Run script (as a non-Super User)	Run Ad-hoc Scripts: Yes	Read and Write	N/A
Run script (as a Super User)	Run Ad Hoc & Source Visible Server Scripts As Super User: Yes	Read and Write	N/A
Execute OGFS script	Manage OGFS Scripts: Yes	Read and Write	N/A
Storage Operations (SE-enabled core)			
Viewing SAN arrays or NAS filer data, including relationships.	View Storage Systems: Yes	Read	N/A
Viewing any SAN switch data, including relationships	View Storage Systems: Yes	Read	N/A

User Action	Action Permission	Source Server Permission (Customer, Facility)	Folder Permission
SA Client Folder Operations			
Open a Business Application from a folder	N/A	N/A	Read Objects Within Folder
Create a Business Application and save to a folder	Manage Business Applications: Yes	N/A	Write Objects Within Folder
Rename a Business Application inside a folder	N/A	N/A	Write Objects Within Folder
Delete a Business Application from a folder	N/A	N/A	Write Objects Within Folder
Cut, copy, or paste a Business Application from a folder	N/A	N/A	Write Objects Within Folder

Note: In order to save a Business Application to a user's own home directory in the Library, for example, `/home/username`, this user's private user group will also need to have the Manage Business Applications permission set to Yes. For more information, see the User Group and Setup chapter in the SA Administration Guide.

Viewing Storage in SAV 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.

Storage Visibility and Automation Permissions

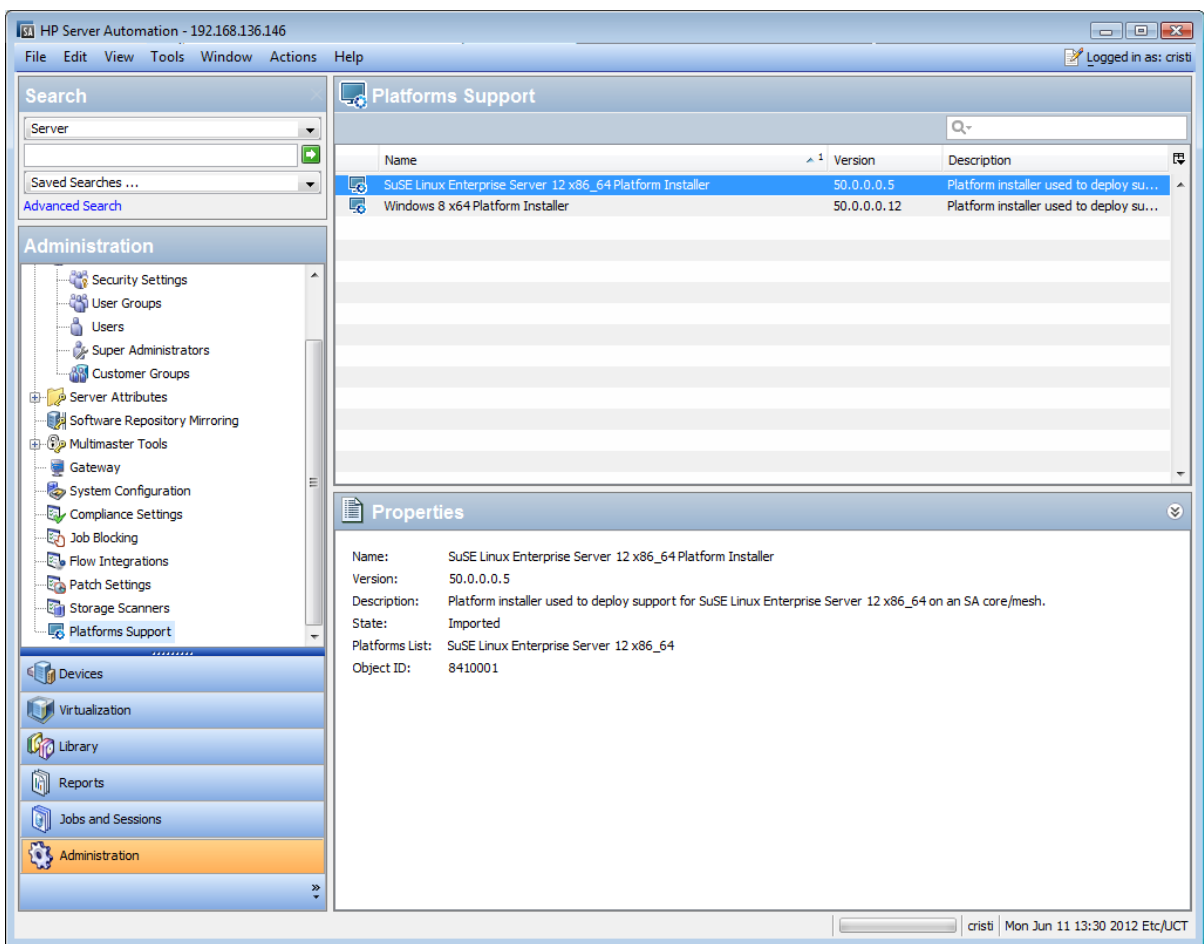
You must have certain permissions to perform actions with Storage Visibility and Automation. See the *Storage Visibility and Automation Installation & Administration Guide* for a description of these permissions.

Appendix B Managed Platform Support

Managed platform support provides a simplified way to add platforms to the SA. Managed platform support allows you to perform changes on the entire SA Core automatically and reduce the need to restart core components.

For each new managed platform, a program APX called the Platform Installer will be made available through ITOM Marketplace. The Platform Installer will perform necessary operations on the SA Core to add support for each new platform. **Figure 39** shows contents of the entire new-platform package:

Figure 39. Managed Platform Support New Platform Package



This appendix describes how to import the new platform package and deploy the new platform on the SA Core.

Note: For product support and compatibility information, see the support matrix for the relevant product release. You can download the *HP Server Automation Support and Compatibility Matrix* for this release from the HP Software Support Online Product Manuals website: <http://h20230.www2.hp.com/selfsolve/manuals>.

Importing the New Platform Package

You can download platform packages from ITOM Marketplace individually and import them to an SA Core.

1. Enter the following URL, which takes you to the ITOM Marketplace portal:
<https://marketplace.saas.hpe.com/itom/content/managed-platform-content-server-automation-2>
2. A list of installers appears. Download one installer on the SA Core file system.
3. Because the installer is an APX, import it to the SA Core using the following command:

```
/opt/opsware/bin/apxtool import <Platform Installer File-name>
```
4. Run the platform installer.

Note: Importing a platform installer on an SA Core does not deploy support for the new platform automatically. The platform installer must be run by an SA user to implement updated information and changes. The next section describes deploying support for the newly installed platform.

Deploying Support for the New Platform

This section describes actions you must take to deploy the newly imported platform.

Required Manage Platforms Permission

To see the SA Client platform support feature and its list of platform installers and to run one of the installers, the SA User Group must have the Manage Platforms permission.


To assign this permission to an SA User Group:

1. Open the user group in SA Client, and go to the Action Permissions node.
2. In the right panel, search for Manage Platforms under the System Administration category.
3. Set Manage Platforms to **Yes**, and save.

Using the Platform Installer

After you have the Managed Platforms permission, the Platform Support entry under the Administration tab is visible in the SA Client (see **Figure 40**).

Figure 40. Platforms Support Window



This window lists the platform installers imported on the SA Core. Each installer has the following attributes:

- Name
- Description
- Version
- List of platforms that it will deploy
- State

Status of the platform installer, as follows:

- **NOT RUN**—Installer was imported on the SA Core but was not yet run, so support for that OS is not available.
- **FAILED**—Installer was imported on the SA Core and run, but it failed. In this case, support for the new OS is only partially deployed, and the new OS cannot be used until the installer is run successfully.
- **INSTALLED**—Installer was imported on the SA Core and run successfully. Support for the new OS was deployed successfully as well, and the new platform can be used by SA.
- **UNKNOWN**—Installer status could not be determined.

Running a Platform Installer

To run an installer, you can:

- Right-click on the installer, and choose **Run...**, or
- Select one installer and choose **Actions > Run...** from the main menu.

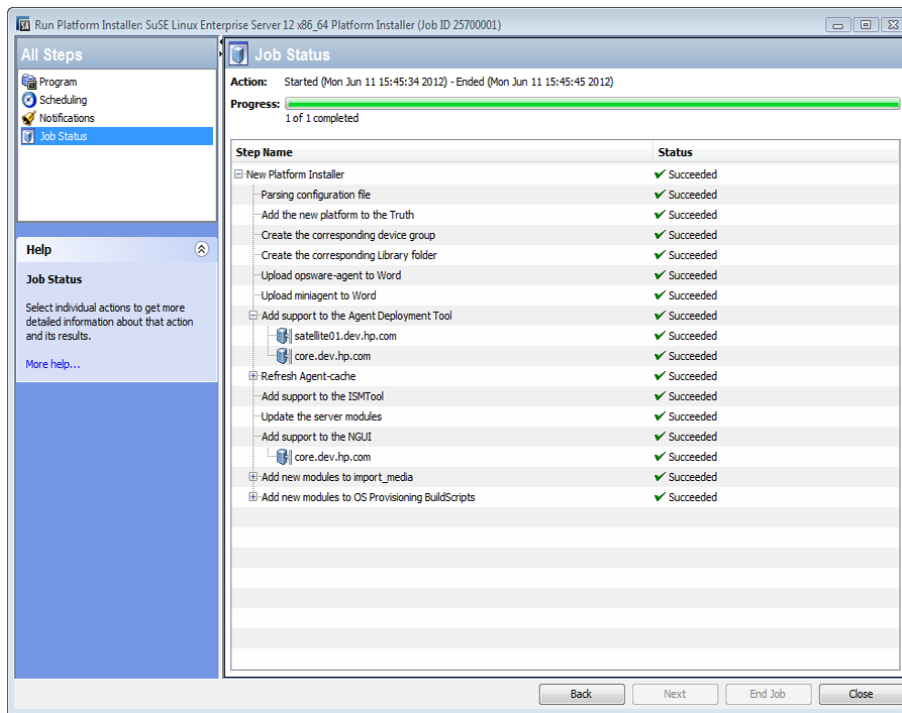
The Run Platform Installer job window will appear. This window provides the option to schedule the task to run at a specified time with no recurrence and to set up email notifications.

After the job is started, the installer determines what changes must occur in your mesh and creates a series of steps (see **Figure 41**).

Some steps can be executed only once (such as 'Add the new platform to the Truth'), and other steps must be run on multiple machines in your mesh/core configuration (such as 'Add support to the Agent Deployment Tool').

- By selecting each step, you will be able to see the captured **stdout** and **stderr** files. If a step must be run on multiple machines, its corresponding node in the job result window will have one child for each machine.
- By selecting that child node, you will be able to see the **stdout** and **stderr** files that resulted by running the step on that particular machine.

Figure 41. Run Platform Installer Job Status Window



Deleting a Platform Installer

To delete an installer, you can:

- Right-click on the installer, and choose **Delete**, or
- Select one installer and choose **Actions > Delete** from the main menu.

Deleting an installer does not mean removing the support for that OS if it was deployed on the SA Core. So, after importing and running a platform installer, you can safely delete it without losing the support for the new OS in SA.

Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on SA Administration Guide (Server Automation 10.23)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to hpe_sa_docs@hpe.com.

We appreciate your feedback!