HP Server Automation

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

Software Version: 9.0

Administration Guide

Document Release Date: August 2010
Software Release Date: August 2010

Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notices

Trademark Notices

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

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

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

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

UNIX® is a registered trademark of The Open Group.

Adobe® is a trademark of Adobe Systems Incorporated.

Documentation Updates

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

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

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

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

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

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

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

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

Document Changes

Chapter	Date	Version	Changes
	June 2010	9.0	Document Created
Appendix	August 2010	9.0	New permissions

Support

Visit the HP Software Support Online web site at:

www.hp.com/go/hpsoftwaresupport

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

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

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

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

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

To find more information about access levels, go to:

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

Contents

1 User and User Group Setup and Security

Users, Groups, and Permissions

SA enforces a *role-based security policy* 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.

SA Users and User Groups

When you log in to the SA Client or the SAS Web Client, you are prompted for an SA user name and password. Everyone in your organization who logs on to SA must have a unique SA user name and password.

SA user names are stored in the *Model Repository*. You can create user names by using the SAS Web Client, or you can import them into the Model Repository from an external Lightweight Directory Access Protocol (LDAP) system. SA user names are not case sensitive.

An SA *user group* represents a role, or group of tasks, performed by the members of that group. All users should belong to one or more SA user groups. The tasks that a user is authorized to perform depend on user groups membership.

SA Permissions

The permissions that you specify for a user group determine what an SA user group's members can do.

- *Feature permissions* specify the actions users can perform
- *Resource permissions* specify the objects (typically servers) users can perform these actions on

For example, Jane Doe could belong to a user group called *London Windows Administrators*. This user group has the feature permission to install patches, and the resource permission to Read & Write on the device group named London Windows Servers.

Feature Permissions

With feature permissions, you define the tasks that can be performed by the users of a group. A feature permission is either on or off: The user can either perform a task or not. In the SAS Web Client, you specify feature permissions on the **Features**, **Client Features**, and **Others** tabs of the **Edit Group** page.

Resource Permissions

A *resource* is usually a set of managed servers. A resource permission determines if the users in a user group can view or modify a resource. Resource permissions specify 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 or the SAS Web Client. Users cannot view or modify the resource.

The SAS Web Client organizes resources into the following categories:

- **Customers**: The servers associated with a customer.
- **Facilities**: The servers associated with an SA Facility.
- **Device Groups**: The servers belonging to a specified public device group.

Each of the preceding resource categories corresponds to a tab on the **Edit Group** page of the SAS Web Client.

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

- Application configurations
- 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 Folder Permissions on page 15.)

Server Access and Resource Permissions

Access to a server depends on the server's association with a customer, association with a Facility, and optionally, the server's membership in a public device group. For example, suppose that a server is associated with the `Widget, Inc.` customer, resides in the Fresno facility, and belongs to the Accounting device group. To modify the server, the user group must have the permissions listed in Table 1. (The Read & Write permission for Accounting is required only if user group permissions are specified for public device groups.)

Table 1 Example of Resource Permissions

Resource	**group permission**
Customer: Widget, Inc.	Read & Write
Facility: Fresno	Read & Write
Device Group: Accounting	Read & Write

If the permissions for the customer, facility, or device group do not match, then the most restrictive permissions are enforced. For example, if the permission for the customer is Read & Write, but the permission for the facility is Read, then the Read permission is enforced. 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).

Feature and Resource Permissions Combined

To use a feature on a resource, the user must belong to a group that has the necessary permissions for both the feature and resource. For example, suppose that a server is associated with these resources: the `Widget, Inc.` customer and the Fresno facility. To install a patch on this server, the user must belong to a group with the permissions listed in Table 2.

Table 2 Example of Permissions Resources and Features

Resource or feature	**group permission**
Customer: Widget, Inc.	Read & Write
Facility: Fresno	Read & Write
Feature: Install Patch	Yes

Folder Permissions

Folder permissions control access to the contents of the folder, such as software policies, OS sequences, 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 (grandchildren).

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 children (but not the attributes of the children).
- **Read Objects Within Folder**: View all attributes of the folder's children, open object browsers on folder's children, use folder's children 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, feature and server permissions are required, as well.)

  Selecting this permission automatically adds the List Contents of Folder permission.
- **Write Objects Within Folder**: View, use, and modify the folder's children.

  This permission permits actions such as New Folder and New Software Policy. To perform most actions, client features are required as well.

  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 children.

  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 feature permission. To create scripts, they need the Write Objects Within Folder permission and the appropriate feature 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 children) to another user group.

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

Client Feature Permissions and Folders

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

To perform most actions on folders and the items they contain, users need both folder and client feature 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 and the Manage Software Policy permission (Read & Write).

Customer Constraints, Folders, 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 parent folder is Widget and the customer of the package's parent 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 parent folder (Widget) acts as a filter, restricting the objects that can be added to the policy. If you add the Widget customer to the parent 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 parent folder has one or more customers. Software policy actions not listed here, such as New Folder, do not have customer constraints.

- **Add Package**: The customers of the package's parent folder must be a subset of the customers of the software policy's parent folder.
- **Add Application Configuration**: The customers of the application configuration must be a subset of the customers of the software policy's parent folder.
- **Add Software Policy**: If software policy A is added to software policy B, then the customers of A's parent folder must be a subset of the customers of B's parent folder.
- **Attach Software Policy**: The customer of the server being attached must be one of the customers of the software policy's parent folder.
- **Install Software Policy Template**: The customer of the server must be one of the customers of the parent 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.

SA Global File System Permissions

The SA Global File System (OGFS) underlies many SA Client actions, such as browsing managed server file systems and scanning servers for compliance. To perform actions that access the OGFS, you must belong to a user group that has certain OGFS permissions. Table 3 lists the operations you control with OGFS permissions.

Table 3 OGFS Permissions

ogfs permission	**task 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 `rosh` 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 Client, 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 feature 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 `rosh` 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 which managed servers the operation can be applied to. You specify the managed servers by selecting a resource, either a customer, facility, or device group. You also specify the login name of the managed server where the operation runs. (The Launch Global Shell operation is an exception, as explained later in this section.)

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 on, in the SA Client, the SA user jdoe opens a server belonging to the Sunnyvale Servers device group in the Device Explorer. In the Views pane, the string jdoe 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 that the Unix user jdoe has access to.

If you specify root for the login name, make sure that the resource you select 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).

Membership in Multiple Groups

If a user belongs to more than one user group, the user's permissions are derived from the resource and feature permissions 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 feature permissions of all groups that the user belongs to. With a cross product, all feature 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 feature	Atlanta user group permission	Portland user group permission
Resource: Customer Widget, Inc.	Read & Write	None
Resource: Customer Acme Corp.	None	Read & Write
Feature: System Diagnosis	No	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 5. Joe can create packages under the Webster folder because the Sunnyvale group has Read & Write permissions for that folder and for the

Manage Package feature. 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 5 Example of Cumulative Permissions

Resource or feature	Sunnyvale user group permission	dallas user group permission
Resource: Folder Webster	Read & Write	None
Resource: Folder Kiley	None	Read & Write
Feature: Manage Packages	Read & Write	None
Feature: Manage OS Sequences	None	Read & Write

SAS Web Client Restricted Views

The SAS Web Client displays only those features and resources that 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 6. When John logs in, the SAS Web Client displays only the servers for Widget Inc., but not those of Acme Corp. In the navigation panel of the SAS Web Client, the Operating Systems link appears, but not the Scripts link.

Table 6 Example of Permissions and Restricted Views

Resource or feature	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 group that has Read (or Read & Write) permission to both the customer and facility associated with the server. If the server also belongs to a device group with set permissions, then the user group must also have Read (or Read & Write) access to the device group. Otherwise, the user cannot locate the server in the SAS Web Client.

Predefined User Groups

During an SA installation or upgrade, certain predefined user groups are created. You must grant read and/or write permissions to the first Facility and other appropriate permissions to these user groups. Use of the predefined user groups is optional. You can modify the permissions of the predefined user groups and you can also delete or copy these groups to create new groups. Changes or deletions of the predefined user groups are not affected by SA upgrades.

Predefined User Group Feature Permissions

Table 7 shows the Predefined User Groups:

Table 7 Predefined User Group Feature Permissions

User Group	Description
System Administrators	Access to administer the SA application.
Superusers	Complete access to all SA-managed objects and operations.
Viewers	Read-only access to all features.
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.
Hypervisor Managers	Access to create, delete and register VMs.
Virtual Machine Managers	Access to start and stop VMs.
Command Line Administrators	Shell access to servers.
Server Storage Managers	Access to manage server storage.
Users	**Description**
Storage System Managers	Access to manage storage systems.
Storage Fabric Managers	Access to manage storage fabrics.

Pre-Defined User Groups Permissions

SA provides an extended set of role-based, pre-defined user groups. If you plan to use these groups, you must grant the appropriate permissions to the groups. For more information about predefined user groups and permissions, see Appendix A: “Permissions Reference” in this guide.

Super Administrators

A super administrator is an SA user who manages SA's security structure. Super administrators create users and groups, specify permissions for groups, and assign users to 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 SAS Web Client as a user with super administrator privileges.

The HP BSA 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 afterwards.

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

Customer Administrators

The super administrator delegates 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 cannot create users. The user groups that a customer administrator can manage depend on the relationships between several objects, including customer permissions and customer groups.

For example, suppose that the user Joe Smith is a customer administrator who can manage the user group named Sunnyvale Admins. The users who belong to the Sunnyvale Admins group are responsible for managing servers owned by the Widget, Inc. customer. Figure 1 shows the relationships required between various objects. In order for Joe Smith to manage the Sunnyvale Admins user group, the following relationships must exist:

- The R & W customer permission Widget, Inc. is assigned to the Sunnyvale Admins.
- The Widget, Inc. customer belongs to the customer group named Widget Admins.
- The user Joe Smith is assigned to the Widget Admins customer group.

For instructions on setting up the relationships shown in Figure 1, see Delegating User Group Management to a Customer Administrator on page 34.

Figure 1 Relationships Required for a Customer Administrator

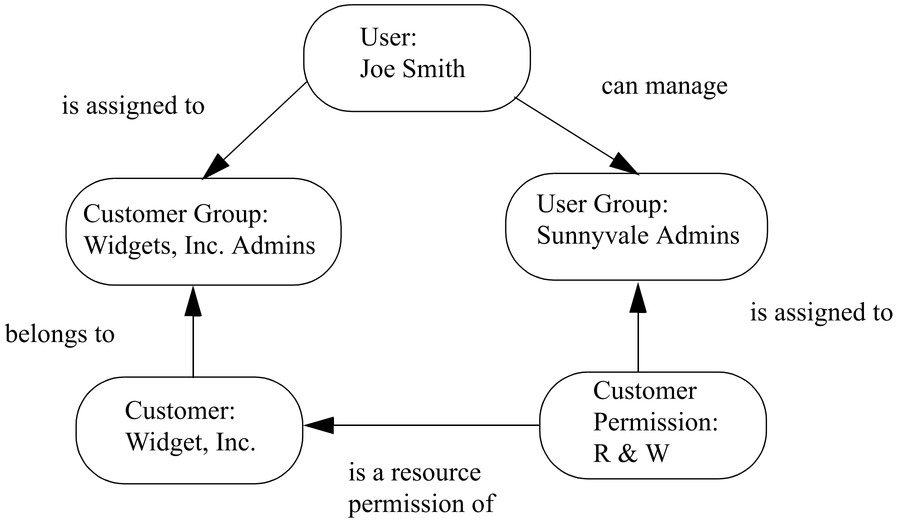

Process Overview for Security Administration

The person responsible for the security of SA creates and maintains users, groups, and permissions. This person must be able to log in to the SAS Web Client as a user who is a super administrator. The default super administrator is the `admin` user.

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 on page 34.
3. Note the facility that the managed servers belong to.

   A facility is an SA object that 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 a default facility for the core.
4. Associate customers with managed servers.

   In SA, a customer is an object that represents a business organization, such as a division or a corporation. Typically, a server is associated with a customer because it runs applications for that customer.
5. (Optional) Create device groups and assign servers to the groups.

   For more information, see the “Device Groups” section of 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.
7. Create the user groups.

   For instructions, see Creating a User Group on page 26.
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 instructions, see Setting the Customer Permissions on page 27 and the adjacent sections.
9. (Optional) Delegate the management of user groups to other users.

   For instructions, see Delegating User Group Management to a Customer Administrator on page 34.
10. Set the feature (action) permissions on the user groups.

    To determine which feature permissions are required to perform a specific task, see the tables in Permissions Required for the SA Client on page 207. For example, if you have a user group named Policy Setters, see Software Management Permissions Required for User Actions on page 227.

    For instructions, see Setting the SA Client Features Permissions on page 30 and the adjacent sections.

11 Set the OGFS permissions on the user groups.

OGFS permissions are required for some actions. The OGFS permissions are included in the tables in Permissions Required for the SA Client on page 207.

For instructions, see Adding OGFS Permissions on page 32.

12 Create the folder hierarchy in the Library of the SA Client.

For information on folders, see the "Software Management Setup" chapter of the *SA Policy Setter Guide*.

13 Set the folder permissions.

Again, see Permissions Required for the SA Client on page 207. In general, you need read permission on a folder to use its contents in an operation, write permission to create folder contents, and execute permission to run scripts that reside in a folder.

For instructions, see Setting Folder Permissions on page 31.

14 (Optional) Delegate the management of folder permissions to other user groups, individual users, or customer administrators.

For instructions, see Setting Folder Permissions on page 31.

15 Create new users in SA or import existing users from an external LDAP.

For instructions, see Creating a User on page 24 or External LDAP Directory Service with SA on page 38.

16 Assign users to the appropriate groups.

For instructions, see Assigning a User to a Group on page 27.

Private User Group

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 feature and resource permissions to the private user group. The permissions that you specify for a private user group determines what the user can do with SA. Feature permissions specify what actions the user can perform; resource permissions indicate which objects (typically servers) the user can perform the actions on. OGFS permissions cannot be assigned to a private 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 feature 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 feature permissions of all groups that the user belongs to.

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

To access a private user group see Setting Private User Group Permissions on page 33. After accessing a private user group, the SA administrator may assign the following permissions:

1 Set the resource permissions on the private user group.

These permissions specify read and write access to servers associated with facilities, customers, and device groups. Resource permissions control which servers the user can access.

For instructions, See “Setting the Customer Permissions” on page 27.” and See “Setting the Facility Permissions” on page 28.

2 Set the feature (action) permissions on the private user group.

To determine which feature permissions are required to perform a specific task, See “Setting the SA Client Features Permissions” on page 30.” and See “Setting the General Feature Permissions” on page 30.

3 Set the folder permissions on the default home folder of the user. When an SA Administrator creates a new user, a private user group is also created for the user in the following location: `/Home/user_name`. By default, the user has Read and Write permissions to this folder and the SA administrator has List and Edit permissions to this folder.

For instructions, See “Setting Folder Permissions” on page 31.

Managing Users and User Groups

To manage users, you must log in to the SAS Web Client as a super administrator (`admin`).

User Management

You can create, modify, delete, suspend, users and view user permissions.

Creating a User

You can create SA users with the SAS Web Client, or you can import users from an external LDAP directory. See External LDAP Directory Service with SA on page 38 in this chapter for more information.

To create a user with the SAS Web Client, perform the following steps:

1 From the navigation panel, select Administration ➤ Users & Groups.

The Users tab appears. (See Figure 2.)

Figure 2 Users Tab

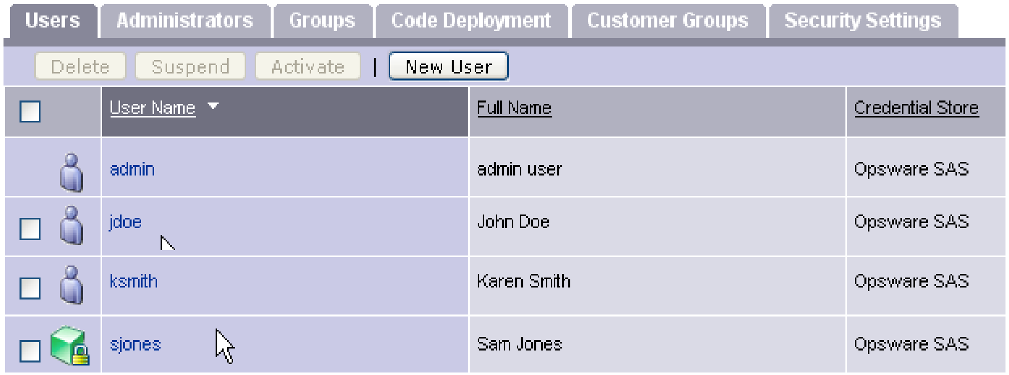

□	User Name ▾	Full Name	Credential Store
	admin	admin user	Opsware SAS
□	jdoe	John Doe	Opsware SAS
□	ksmith	Karen Smith	Opsware SAS
□	sjones	Sam Jones	Opsware SAS

2 Click **New User**.

3 On the Profile Editor page, fill in the required fields, which are labelled in bold font.

The Login User Name may be different than the first, last, and full names. The Login User Name is not case sensitive and cannot be changed after the user is created.

4 (Optional) If both SA and HP Network Automation (NA) are installed, and you want the user to authenticate with NA, then select NA for the Credential Store.

The Credential Store field can be either SA (the default), Network Automation (NA), External, or RSA 2-factor. The value NA specifies that the user was configured on a TACACS+/RADIUS server connected to NA, *not* a native NA user. The value External indicates that the user was imported from an external LDAP directory. The value RSA 2-factor specifies that the user was configured on an RSA server connected to SA. You can change the user password in the SAS Web Client only if the Credential Store is SA.

5 (Optional) Assign the user to one or more of the groups listed at the bottom of the page.

You can also assign the user to a group at a later time. If a user does not belong to a group, the user cannot view servers or perform tasks with the SA Client.

6 Click **Save** to create the user.

Editing User Profile Information

Each SA user can edit the profile information for his or her own login user. If you log in as a super administrator (`admin`), you may view or edit the information of any SA user. To do so, perform the following steps:

1 From the navigation panel, select Administration ➤ Users & Groups.

2 On the Users tab, select an entry in the User Name column.

3 In the Profile Editor, modify the information as appropriate.

4 Click **Save**.

Viewing a User's Permissions

You do not assign permissions directly to a user. Instead, you set the permissions on a user group and then assign a user to a group. To view the permissions of a user, perform the following steps:

1 From the navigation panel, select Administration ➤ Users & Groups.

2 On the Users tab, select an entry in the User Name column.

3 If the user belongs to more than one group, on the Edit User page, select a user group in the "View as" field. The permissions displayed depend on the user group you select.

4 View the permissions on the Resource Privileges and Action Privileges tabs.

Deleting a User

When you delete a user, the user's login and logout history is permanently stored, and the user is unassigned from user groups. After a user is deleted, you can create another user with the same name.

To delete an SA user, perform the following steps:

1 From the navigation panel, select Administration ➤ Users & Groups.

2 On the Users tab, select the check box next to the user to be deleted.

3 Click **Delete**.

Suspending a User

A suspended user cannot log in to SA, but has not been deleted from the Model Repository. A suspended user is indicated by the lock icon on the Users tab of the SAS Web Client. A user can be suspended in the following ways:

- **Login Failure**: If you specify 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 on page 36.
- **Account Inactivity**: If you specify 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**: To suspend the user's account immediately, go to the Users tab, select the user's check box, and click **Suspend**.

To activate a suspended user, go to the Users tab, select the user's check box, and click **Activate**.

User Group Management Tasks

You can create, copy and delete user groups and assign users to user groups.

Figure 3 User Group Page

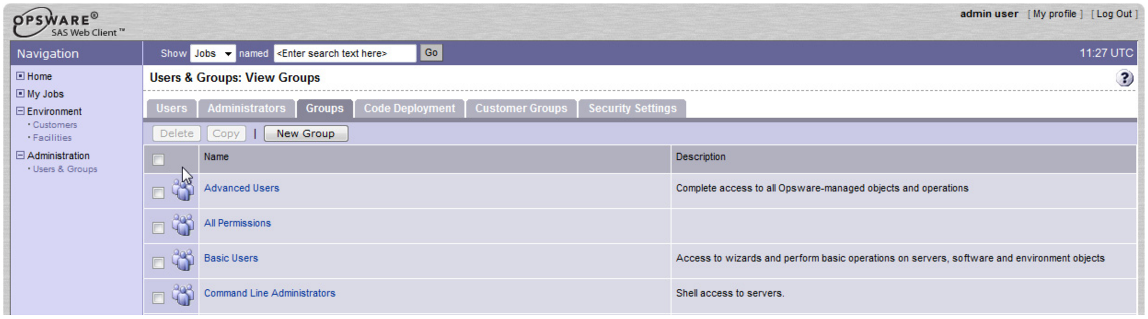

Users & Groups: View Groups

□	Name	Description
□	Advanced Users	Complete access to all Opsware-managed objects and operations
□	All Permissions	
□	Basic Users	Access to wizards and perform basic operations on servers, software and environment objects
□	Command Line Administrators	Shell access to servers.

Creating a User Group

To create an SA user group, perform the following steps:

1 From the navigation panel, select **Administration ➤ Users & Groups**.

2 On the Groups tab, click **New Group**.

3 On the New Group page, enter a role in the Group name field.

4 At this point, you can select the check boxes under the Feature column to assign permissions to the group. The New Group page, however, does not display all available permissions.

5 Click **Save**.

Copying a User Group

To copy (clone) an existing user group to create a new group with the same configuration, perform the following tasks:

1. Select the checkbox next to the user group you want to copy (clone).
2. Click the Copy button.
3. The Copy User Group page appears. Specify a name for the new user group and, optionally, a description.
4. Press the Copy button. When the copy completes, the new group appears in the User Groups list.

Deleting a User Group

To delete an existing user group, perform the following tasks:

1. Select the checkbox next to the user group(s) you want to delete.
2. Click the Delete button.

Assigning a User to a Group

You should 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. From the navigation panel, select Administration ➤ Users & Groups.
2. On the Group tab, select a user group from the Name column.
3. On the Users tab, in the Unassigned Members box, select the user name.
4. Click the right arrow.
5. To unassign a user, click the name in the Assigned Members box and click the left arrow.
6. Click **Save**.

Setting Permissions on User Groups

To perform the tasks in this section, you must log in to the SAS Web Client as a super or customer administrator. (The default super administrator is `admin`.)

If you change permissions while a user is logged on to the SAS Web Client or SA Client, the user must log out and log in again for the changes to take effect.

Setting the Customer Permissions

In SA, you can associate a customer with a number of resources, including servers, folders, application configurations, and OS installation profiles. By setting the customer permission, you control the access that the users of a group have to the resources associated with the customer. 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, as well as permission for the Customers feature.

To control the access to the resources associated with a customer, perform the following steps:

1. From the navigation panel, select Administration ➤ Users & Groups.
2. On the Groups tab, select an entry in the Name column. Another set of tabs appears, including the Customers tab.
3. On the Customers tab, for each customer listed, select Read, Read & Write, or None.
4. Click **Save**.

Setting the Facility Permissions

In SA, a facility can be associated with resources such as servers and IP ranges. To modify a server of a particular facility, a user must belong to a group that has Read & Write permission for the facility.

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, as well as permission for the Facilities feature.

To control the access to the resources associated with a facility, perform the following steps:

1. From the navigation panel, select Administration ➤ Users & Groups.
2. On the Groups tab, select an entry in the Name column. Another set of tabs appears, including the Facilities tab.
3. On the Facilities tab, select Read, Read & Write, or None.
4. Click **Save**.

Setting the Device Group Permissions

To control access to the servers in a public device group, select a permission on the Device Groups tab. (You cannot control access to a private device group, which is visible only to the user who created it.)

If the Device Groups tab lists no device groups, then access to servers is not controlled by membership in device groups; however, access to servers is still controlled by their association with customers and facilities. If the Device Groups tab lists at least one device group, then access is denied to unlisted device groups (the equivalent of a None permission).

Access control based on device groups is optional. By default, membership in a device group does not restrict access. In contrast, for servers associated with customers or facilities, the default permission is None, which prohibits access.

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.

A device group can contain other device groups. However, permissions are not inherited by the contained (children) device groups.

The permissions on the Device Groups tab 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 check boxes selected on the Other tab. Also, the Managed Servers and Groups check box must be selected on the Features tab. To add devices to a device group being used as an Access Control Group, the user must be a member of the Super Administrators group.

To control access to servers that belong to a device group, perform the following steps:

1 From the navigation panel, select Administration ➤ Users & Groups.

2 On the Groups tab, select an entry in the Name column. Another set of tabs appears, including the Device Groups tab.

3 On the Device Groups tab, note the check box below **Assign**. If this check box is selected, then access to managed servers is not based on device groups.

4 Deselect the check box below **Assign**.

5 Click **Assign**.

The Select Groups page appears. (See Figure 4.)

Figure 4 Select Groups Page

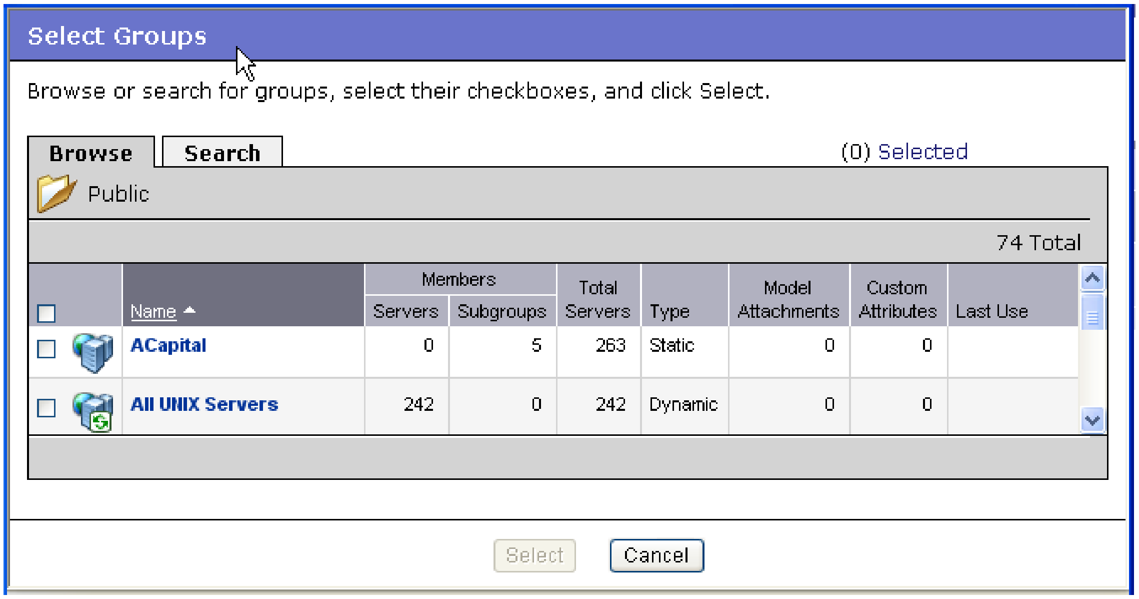

Select Groups

Browse or search for groups, select their checkboxes, and click Select.

		Members		Total		Model	Custom	
□	Name ▲	Servers	Subgroups	Servers	Type	Attachments	Attributes	Last Use
□	ACapital	0	5	263	Static	0	0	
□	All UNIX Servers	242	0	242	Dynamic	0	0	

6 On the Select Groups page, use the Browse or Search tab to locate the device groups.

7 On the Browser or Search tab, click on the device group name and then click **Select**.

8 On the Device Groups tab, for each device group listed, select the check box and click the button for the appropriate access.

To allow viewing (but not modification) of the servers in a device group, select the Read permission. To allow both viewing and modification, select the Read & Write permission.

9 Click **Save**.

Setting the General Feature Permissions

The Features tab of the SAS Web Client includes many tasks, including managing the servers and running the wizards. If the check box for a feature is unselected, then the SAS Web Client does not display the related links in the navigation panel.

To allow the users in a group the ability to view and execute a task on the Features tab, perform the following:

1. From the navigation panel, select Administration ➤ Users & Groups.
2. On the Group tab, select a user group from the Name column.
3. Another set of tabs appears, including the Features tab. (See Figure 5.)

Figure 5 Features Tab

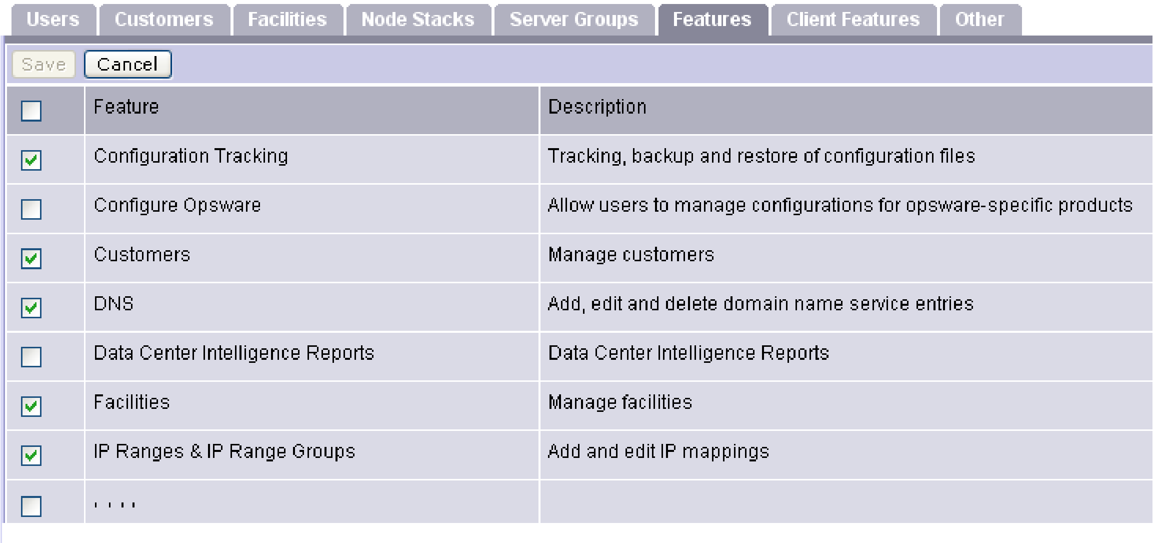

□	Feature	Description
☑	Configuration Tracking	Tracking, backup and restore of configuration files
□	Configure Opsware	Allow users to manage configurations for opsware-specific products
☑	Customers	Manage customers
☑	DNS	Add, edit and delete domain name service entries
□	Data Center Intelligence Reports	Data Center Intelligence Reports
☑	Facilities	Manage facilities
☑	IP Ranges & IP Range Groups	Add and edit IP mappings
□	. . . .	

4. On the Features tab, select the check box for each feature that should be enabled for the user group. To prevent (and hide) a feature, deselect the check box.
5. Click **Save**.

Setting the SA Client Features Permissions

The Client Features tab of the SAS Web Client lists permissions for the actions performed with the SA Client. These actions are for features such as Application Configuration and Software Policy Management.

To set these permissions for the SA Client perform the following steps:

1. From the navigation panel, select Administration ➤ Users & Groups.
2. On the Group tab, select a user group from the Name column. Another set of tabs appears, including the Client Features tab.
3. On the Client Features tab, select the appropriate permission buttons.
4. Click **Save**.

Setting the Other Features Permissions

The Other tab of the SAS Web Client contains the following permissions:

- **General Permissions**: Allows users in a user group to edit shared scripts or run “my scripts” as root. The Features tab also has script-related permissions: Scripts, and Wizard: Run Scripts.
- **Server and Device Group Permissions**: Enables users in a user group to perform particular tasks on managed servers. The Allow Run Refresh Jobs permission lets users specify a job to update the servers list. The Manage Public Servers Group permission enables users to create device groups, modify the group properties, and change the group membership (through rule changes, or adding and deleting servers). All users can view all public device groups.

  The Model Public Servers Group permission lets users add custom attributes. (These permissions apply to public, not private device groups. Only the user who creates a private device group can view or modify it.) The Features tab also has a permission related to managing servers: Managed Servers and Group.
- **Job Permissions**: Allows users in a user group to view and schedule jobs, which include operations such as Audit Servers, Snapshots, Push Configurations, and Audit Configurations. The View All Jobs permission lets users view the details and schedules of jobs created by all users. The Edit All Jobs permission enables users to view or modify the schedules of jobs created by all users and to view the job details of all users. Without these permissions, users can view and schedule only their own jobs.

To set the permissions on the Other tab, perform the following steps:

1. From the navigation panel, select Administration ➤ Users & Groups.
2. On the Group tab, select a user group from the Name column. Another set of tabs appears, including the Other tab.
3. On the Other tab, select the check boxes to assign permissions to this user group.
4. Click **Save**.

Setting Folder Permissions

To perform this task, your user or user group must have the Edit Folder Permission on the target folder. When you create a folder, it has the same permissions and customer as its parent folder. If you are changing the permissions of a folder that has children, you are prompted to apply the changes to the children.

To set the permissions of a folder, perform the following steps:

1. In the SA Client, navigate to the folder.
2. From the Actions menu, select **Folder Properties.**
3. In the Folder Properties window, select the Permissions tab.
4. On the Permissions tab, click **Add** to allow certain user groups to access the folder.
5. For each user group and user displayed on the Permissions tab, select a check box such as Write Objects Within Folder. To delegate the setting of permissions for this folder, select Edit Folder Permissions.

Adding OGFS Permissions

You can add OGFS permissions with the SAS Web Client or with the aaa command-line utility. For syntax and examples of the aaa utility, see the *SA User's Guide: Server Automation*.

To add an OGFS permission in the SAS Web Client, perform the following steps:

1 From the navigation panel, select Administration ➤ Users & Groups.

2 On the Group tab, select a user group from the Name column. Another set of tabs appears, including the OGFS Permissions tab.

3 On the OGFS Permissions tab click **Add Permission**. The Add OGFS Permissions window appears. (See Figure 6.)

Figure 6 Add OGFS Permissions Window

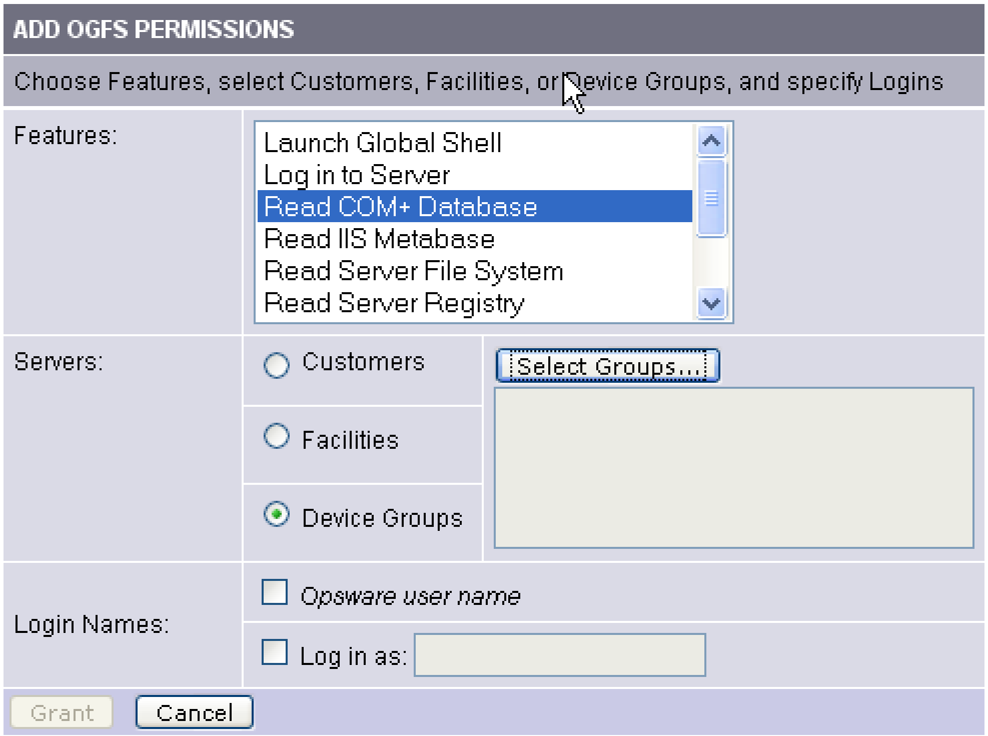

ADD OGFS PERMISSIONS

Choose Features, select Customers, Facilities, or Device Groups, and specify Logins

Features:
Launch Global Shell
Log in to Server
Read COM+ Database
Read IIS Metabase
Read Server File System
Read Server Registry

Servers:
○ Customers
○ Facilities
● Device Groups

Login Names:
□ *Opsware user name*
□ Log in as:

4 In the Add OGFS Permissions window, select a feature.

For descriptions of these features, see Table 3 on page 17.

5 If you selected a feature other than Launch Global Shell, select the managed servers this permission applies to.

You can select servers associated with a customer, facility, or device group. If you want to select servers associated with multiple resources (for example, two device groups) then you must add a separate OGFS permission for each resource.

6 For Login Name, select the user account (login) on the managed servers.

The operation indicated by the Feature field will run on the managed server as the user indicated by Login Name.

7 Click **Grant**.

Setting Private User Group Permissions

The SA Administrator can set the feature, resource, or folder permissions on a private user group.

Perform the following steps to set permissions on a private user group:

1. From the navigation panel, select Administration ➤ Users & Groups. The Users & Groups: View Users window appears.
2. Select the user from the User Name column. The Users & Groups: Edit User window appears.
3. From the View As drop-down menu, select the user name and then click **Edit**.
4. In the Users & Groups: Edit Group - `<user name>` window select Features, Customers, or Client features tab to assign the permissions.
5. Refer to the following sections to assign the permissions:

   Setting the Customer Permissions on page 27

   Setting the Facility Permissions on page 28

   Setting the Device Group Permissions on page 28

   Setting the General Feature Permissions on page 30

   Setting the SA Client Features Permissions on page 30

   Setting Folder Permissions on page 31

Managing Super and Customer Administrators

These users are the security administrators who assign permissions to user groups. To manage super and customer 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.

Viewing Super and Customer Administrators

To see which users are super or customer administrators, perform the following steps:

1. From the navigation panel, select Administration ➤ Users & Groups.
2. Select the Administrators tab. (See Figure 7.)

Figure 7 Administrators Tab

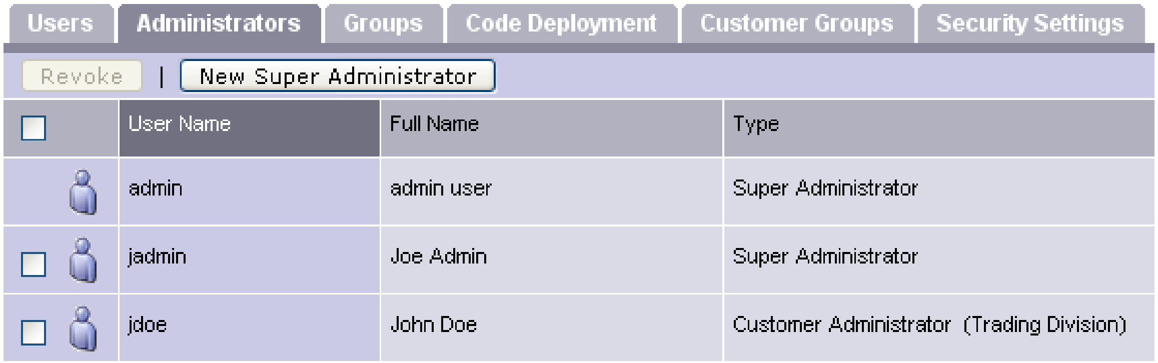

□	User Name	Full Name	Type
	admin	admin user	Super Administrator
□	jadmin	Joe Admin	Super Administrator
□	jdoe	John Doe	Customer Administrator (Trading Division)

3 On the Administrators tab, note the Type field, which identifies the user as either a super or customer administrator. The name of the customer group is in parentheses.

Creating a Super Administrator

To create a super administrator, perform the following steps:

1 Create a new user who will be the super administrator. For instructions, see Creating a User on page 24.

2 From the navigation panel, select Administration ➤ Users & Groups.

3 Select the Administrators tab.

4 Click **New Super Administrator**.

5 On the Add Super Administrators page, select one or more user names.

6 Click **Save**.

Deleting a Super Administrator

To delete a super administrator, perform the following steps:

1 From the navigation panel, select Administration ➤ Users & Groups.

2 Select the Administrators tab.

3 Select the check box for the user.

4 Click **Revoke**. This action revokes super administrator privileges from the user, but does not delete the user from SA.

5 To delete the user from SA, follow the instructions in Deleting a User on page 25.

Delegating User Group Management to a Customer Administrator

A customer administrator is a user who can manage a subset of user groups. The subset is determined by customer permissions and customer groups. For a full explanation of the relationships between these objects, see Customer Administrators on page 21.

To delegate the management of a user group, perform the following steps:

1 Identify the user who will be responsible for user group management.

This user will be a customer administrator. If the user does not exist, follow the instructions in Creating a User on page 24.

2 Decide which user group will be managed by the user identified in the preceding step.

For instructions on viewing the user group, see Creating a User Group on page 26.

3 Note the customer permissions of the user group.

For instructions on viewing these permissions, see Setting the Customer Permissions on page 27.

4 From the navigation panel, select Administration ➤ Users & Groups.

5 Select the Customer Groups tab.

6 Click **New Group**.

7 Enter the customer group name. (See Figure 8.)

Figure 8 New Customer Group Window

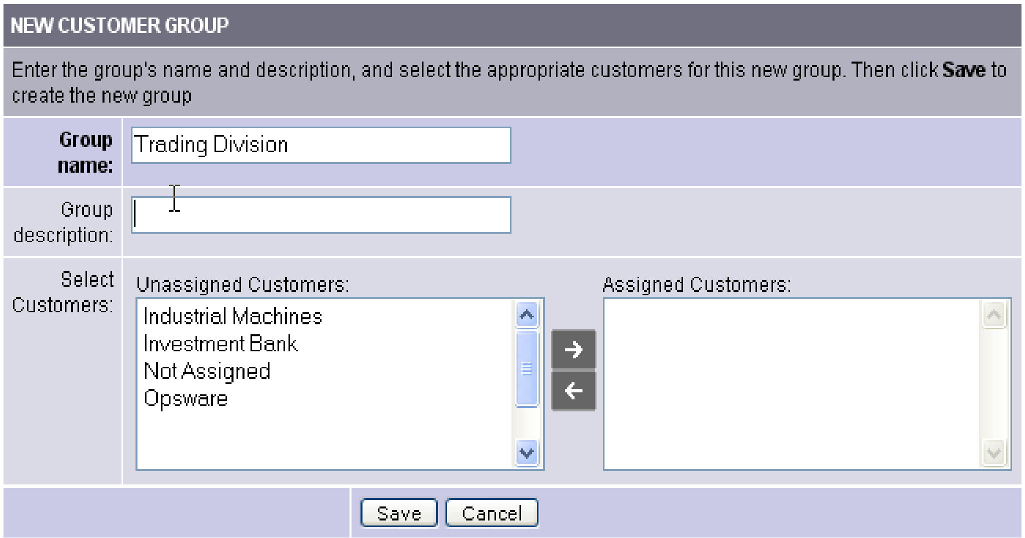

NEW CUSTOMER GROUP

Enter the group's name and description, and select the appropriate customers for this new group. Then click **Save** to create the new group

Group name: Trading Division

Group description:

Select Customers:

Unassigned Customers:

Industrial Machines
Investment Bank
Not Assigned
Opsware

Assigned Customers:

8 Click **Save**.

9 Add the customers you noted in step 3 to the customer group.

10 Add the user you identified in step 1 to the customer group.

The user of step 1 is now the customer administrator who can manage the user group of step 2.

11 (Optional) Verify that the user is listed as a customer administrator by following the instructions in Viewing Super and Customer Administrators on page 33.

Managing Passwords and Login Settings

An SA user can change his or her own password on the Profile page of the SAS Web Client. A super administrator can change the password of other users, as well as perform other password management tasks described in the following sections.

Changing Passwords

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 SAS Web Client.

To change the password of an SA user in the SAS Web Client, perform the following steps:

1 From the navigation panel, select Administration ➤ Users & Groups.

2 On the User tab, select a user name.

3 On the User Identification tab, click Change Password.

4 Enter the new password, confirm it, and click **Save**.

Specifying Password Character Requirements

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

1 From the navigation panel, select Administration➤ System Configuration. The Select a Product page appears.

2 Under Select a Product, click SAS Web Client. The Modify Configuration Parameters page appears.

3 On the Modify Configuration Parameters page, set the owm.features.MiniPasswordPolicy.allow parameter 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.MiniPasswordPolicy.allow to false.

4 Set the values for the password parameters listed inTable 8.

5 Click **Save**.

6 To apply these parameter changes to other cores in a multimaster mesh, you must restart the other cores.

Table 8 Password Requirements on the Modify Configuration Parameters Page

password requirement	parameter	allowed values	default value
maximum number of repeating, consecutive characters	owm.pwpolicy.maxRepeats	must be greater than 0	2
minimum number of characters	owm.pwpolicy.minChars	positive integer	6
minimum number of non-alphabetic characters	owm.pwpolicy.minNonAlphaChars	must be less than value of owm.pwpolicy.minChars	0

Resetting Initial Passwords

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

1 From the navigation panel, select Administration➤ Users & Groups.

2 Select the Security Settings tab.

3 Select Reset.

Setting Password Expiration

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

1 From the navigation panel, select Administration➤ Users & Groups.

2 Select the Security Settings tab.

3 In the check boxes next to the Expiration label, select the number of days for the password expiration and the number of grace logins.

A grace login allows the user to log in with the old password. Typically, the grace login is set to 1, enabling the user to log in to the SAS Web Client, access the My Profile page, and change the password.

4 To specify the number of previous passwords allowed by users, select Retention and enter a value.

This setting prohibits users from re-using the same set of passwords. For example, if the value is 10, the users are not be allowed to re-use their previous 10 passwords.

For information on Login Failure and Account Inactivity, see Suspending a User on page 26.

Specifying Session Timeout

You can specify the timeout interval (in minutes) of inactive SA Client sessions. When a session times out, the user must re-enter the password or log out.

To specify the timeout for SA Client sessions, perform the following steps:

1 From the navigation panel of the SAS Web Client, select Administration ➤ Users & Groups.

2 Select the Security Settings tab.

3 Select Session Inactivity and specify the number of minutes.

The Session Inactivity parameter does not affect SAS Web Client sessions. The default session timeout for the SAS Web Client is 60 minutes. To change the default, you edit a configuration file and restart the OCC core component. For instructions on editing the configuration file, contact your HP support representative.

Setting the User Agreement

If you enable the user agreement, when users log in with the SA Client or the SAS Web Client, a dialog appears with a specified message. To continue the log in procedure, the users must click **Agree**. (In some products, a user agreement dialog is called a login approval screen.)

To set the user agreement, perform the following steps:

1 From the navigation panel, select Administration➤ Users & Groups.

2 Select the Security Settings tab.

3 In the User Agreement section, select Display and enter text in the Message field.

Setting the Banner

If you enable the banner, after the users log in, the specified text appears in a banner at the top of the SA Client and the SAS Web Client. To set the banner, perform the following steps:

1 From the navigation panel, select Administration➤ Users & Groups.

2 Select the Security Settings tab.

3 In the Banner Settings section, select Display

4 In the Message field, enter the text to be displayed in the banner.

5 To set the background color of the banner, select an item from Color Code or enter a hex value in the adjacent field.

External LDAP Directory Service with SA

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 SAS Web Client, they enter their LDAP user names and passwords.

Imported Users

With the SAS Web Client, you search for users in the external LDAP and then you import selected users into SA. You can limit the search results by specifying a filter. The import process fetches the following user attributes from the LDAP:

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

After the import process, you may edit the preceding list of attributes within the SAS Web 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 SAS Web Client are not propagated back to the external LDAP directory server, and vice versa.

Imported users are managed in the same way as users created by the SAS Web Client. For example, you use the SAS Web Client to assign imported users to user groups and to delete imported users from SA. If you delete an imported user with the SAS Web Client, the user is not deleted from the external LDAP directory.

If you use external authentication, you can still create separate users with the SAS Web Client. However, this practice is not recommended.

To see which users have been imported, view the Users tab of the SAS Web Client and note the users with External in the Credential Store column.

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 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 or 2008)
- Novell eDirectory 8.7
- SunDS 5.2

Modify the nsswitch.conf File (Linux)

In order to use LDAP authentication on Linux, the `nsswitch.conf` file must be modified as follows:

```
passwd: files ldap
group:  files ldap
```

The typical value for these entries is `compat`. However, this value can cause the SA gateway to fail and/or interfere with OGFS functionality/access.

Alternatively, these values may also work:

```
passwd:       compat
group:        compat
passwd: files ldap
group:  files ldap

passwd_compat:  ldap
group_compat:   ldap
```

Using an LDAP Directory Server with SA

To use an LDAP directory server with SA, perform the following basic steps:

1. Add the `aaa.ldap` entries to the `twistOverrides.conf` file with a text editor. See “Modifying the Web Services Data Access Engine Configuration File” on page 40.
2. Get the SSL server certificate from the LDAP directory server. See “Importing a Server Certificate from the LDAP into SA” on page 42. (Use of SSL is not required, but strongly recommended.)
3. Edit the loginModule.conf file with a text editor. See “Configuring the JAAS Login Module (loginModule.conf)” on page 43.
4. Restart the Web Services Data Access Engine:

   Linux: `/etc/init.d/opsware-sas restart twist`

   HP-UX: `/sbin/init.d/opsware-sas restart twist`

   AIX: `/etc/rc.d/init.d/opsware-sas restart twist`
5. Use the SAS Web Client to import users from the LDAP directory server into SA. See “Importing External LDAP Users and User Groups” on page 44.

In a mulitmaster mesh, you must perform steps 1 - 4 on each Web Services Data Access Engine.

Modifying the Web Services Data Access Engine Configuration File

To modify `twistOverrides.conf`, perform the following steps:

1. Log in as root to the system running the Web Services Data Access Engine, an SA Core Component.
2. In a text editor, open this file:

   `/etc/opt/opsware/twist/twistOverrides.conf`
3. In the text editor, add the necessary properties (listed in Table 9) to the `twistOverrides.conf` file. Although not required, the SSL properties are recommended. For examples of the lines required for the `twistOverrides.conf` file see, the sections that follow Table 9.
4. Save the `twistOverrides.conf` file and exit the text editor.
5. Make sure that the Unix `twist` user has write access to the `twistOverrides.conf` file.

Table 9 Properties in twistOverrides.conf for an External LDAP

Property	description
`aaa.ldap.hostname`	The host name of the system running the LDAP directory server.
`aaa.ldap.port`	The port number of the LDAP directory server.
`aaa.ldap.search.binddn`	The BIND DN (Distinguished Name) for LDAP is required by the search of the import user operation. A blank value denotes an anonymous BIND.
`aaa.ldap.search.pw`	The BIND password for LDAP is required by the search for the import user operation. This value is encrypted when the Web Services Data Access Engine is restarted. A blank value denotes an anonymous BIND.
`aaa.ldap.search.filter.template`	The search filter template is used, with optional filter substitution, as the filter in the LDAP search for the user import. Any dollar sign ($) character in the template will be replaced by the filter string specified in the Import Users page of the SAS Web Client. (The default value is an asterisk (*) which matches all entries.)
`aaa.ldap.search.base.template`	The configurable template allows support for a range of DIT configurations and schema in the LDAP service. The search base template string is used for the “search base” in the LDAP search operations for the user import.
`aaa.ldap.search.naming.attribute`	The naming attribute allows support for a range of schema in the LDAP services. Some use `uid`, others use `cn`, and so on. The value of this attribute is used for the internal user ID in SA.

Table 9 Properties in twistOverrides.conf for an External LDAP (cont'd)

Property	description
aaa.ldap.search.naming.display.name	The naming attribute allows support for a range of schema in the LDAP services. Some use cn, others use displayName, and so on. The value of this attribute is used for the Full Name of SA user.
aaa.ldap.ssl	SSL: A value of true enables SSL.
aaa.ldap.secureport	SSL: The secure port of the LDAP directory server.
aaa.ldap.usestarttls	SSL: A value of true enables Start TLS.
aaa.ldap.servercert.ca.fname	SSL: The fully qualified file name of the server CA certificate.
aaa.ldap.clientcert	SSL: A value of true enables client certificate use.
aaa.ldap.clientcert.fname	SSL: The fully qualified file name of the client certificate.
aaa.ldap.clientcert.ca.fname	SSL: The fully qualified file name of the client CA certificate.

Example: twistOverrides.conf for Microsoft Active Directory Without SSL

```
aaa.ldap.search.binddn=cn=Administrator,cn=users,dc=example,dc=com
aaa.ldap.search.pw=secret
aaa.ldap.hostname=myservername.internal.example.com
aaa.ldap.port=389
aaa.ldap.search.filter.template=(&(objectclass=user)(cn=$))
aaa.ldap.search.base.template=cn=users,dc=example,dc=com
aaa.ldap.search.naming.attribute=samaccountname
aaa.ldap.search.naming.display.name=cn
```

Example: twistOverrides.conf for Microsoft Active Directory With SSL

```
aaa.ldap.search.binddn=cn=Administrator,cn=users,dc=example,dc=com
aaa.ldap.search.pw=secret
aaa.ldap.hostname=myservername.internal.example.com
aaa.ldap.secureport=636
aaa.ldap.ssl=true
aaa.ldap.servercert.ca.fname=/var/opt/opsware/crypto/twist/cert.pem
aaa.ldap.search.filter.template=(&(objectclass=user)(cn=$))
aaa.ldap.search.base.template=cn=users,dc=example,dc=com
aaa.ldap.search.naming.attribute=samaccountname
aaa.ldap.search.naming.display.name=cn
```

Example: twistOverrides.conf for Novell eDirectory Without SSL

```
aaa.ldap.search.binddn=cn=admin,o=example
aaa.ldap.search.pw=secret
aaa.ldap.hostname=myservername.internal.example.com
aaa.ldap.port=389
aaa.ldap.search.filter.template=(&(objectclass=inetorgperson)(uid=$))
aaa.ldap.search.base.template=o=example
aaa.ldap.search.naming.attribute=uid
aaa.ldap.search.naming.display.name=cn
```

Example: twistOverrides.conf for Novell eDirectory With SSL

```
aaa.ldap.search.binddn=cn=admin,o=example
aaa.ldap.search.pw=secret
aaa.ldap.hostname=myservername.internal.example.com
aaa.ldap.secureport=636
aaa.ldap.ssl=true
aaa.ldap.servercert.ca.fname=/var/opt/opsware/crypto/twist/ldapcert.pem
aaa.ldap.search.filter.template=(&(objectclass=inetorgperson)(uid=$))
aaa.ldap.search.base.template=o=example
aaa.ldap.search.naming.attribute=uid
aaa.ldap.search.naming.display.name=cn
```

Example: twistOverrides.conf for SunDS Without SSL

```
aaa.ldap.search.binddn=cn=Directory Manager
aaa.ldap.search.pw=secret
aaa.ldap.hostname=myservername.internal.example.com
aaa.ldap.port=389
aaa.ldap.search.filter.template=(&(objectclass=inetorgperson)(uid=$))
aaa.ldap.search.base.template=ou=people,dc=example,dc=com
aaa.ldap.search.naming.attribute=uid
aaa.ldap.search.naming.display.name=cn
```

Example: twistOverrides.conf for SunDS With SSL

```
aaa.ldap.search.binddn=cn=Directory Manager
aaa.ldap.search.pw=secret
aaa.ldap.hostname=myservername.internal.example.com
aaa.ldap.secureport=636
aaa.ldap.ssl=true
aaa.ldap.servercert.ca.fname=/var/opt/opsware/crypto/twist/ldapcert.pem
aaa.ldap.search.filter.template=(&(objectclass=inetorgperson)(uid=$))
aaa.ldap.search.base.template=ou=people,dc=example,dc=com
aaa.ldap.search.naming.attribute=uid
aaa.ldap.search.naming.display.name=cn
```

Importing a Server Certificate from the LDAP into SA

For SSL, the necessary certificates must be extracted from the LDAP and copied over to SA.

To import a server certificate from the LDAP into SA, perform the following steps:

1. Extract the server certificate from the external LDAP. 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 Web Services Data Access Engine configuration file (`twistOverrides.conf`). For example, the `twistOverrides.conf` file could have the following line:

```
aaa.ldap.servercert.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=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.

Configuring the JAAS Login Module (loginModule.conf)

To configure the JAAS login module, perform the following steps:

1 Log in as root to the system running the Web Services Data Access Engine, an SA Core Component.

2 In a text editor, open this file:

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

3 In the text editor, modify the `loginModule.conf` file so that it contains the following lines:

```
/** Login configuration for JAAS modules **/
```

```
TruthLoginModule {
  com.opsware.login.TruthLoginModule sufficient debug=true;
  com.opsware.login.LdapLoginModule sufficient debug=true;
};
```

4 Save the `loginModule.conf` file and exit the text editor.

Importing External LDAP Users and User Groups

Before importing external LDAP users, you must configure SA for use with your LDAP server. See Using an LDAP Directory Server with SA on page 39 in this chapter for more information.

Importing LDAP Users Using the SAS Web Client

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

➤ This method does not import LDAP user groups. If you want to import users and user groups, see Importing LDAP Users and Groups Using the LDAP Authentication Configuration Tool on page 45.

To import external users, perform the following steps:

1 In the SAS Web Client, from the navigation panel, select Administration ➤ Users & Groups.

2 Select the Users tab. The page lists the existing SA users.

3 On the Users tab, click **Import External Users**.

The page displays the users in the LDAP that match the search filter. The default filter is an asterisk (*), indicating that all users are selected. If a check box does not appear to the left of the user name, then the user already exists in SA and cannot be imported.

If SA cannot connect to the LDAP, check for error messages in the following file:

```
/var/log/opsware/twist/stdout.log
```

4 To change the search filter, enter a value in the field to the left of **Change Filter**. For example, to fetch only those user names beginning with the letter A, you enter A* in the field.

➤ In order to avoid retrieving large lists of user names, you should be very strict in your use of the search filter so that you can retrieve a manageable list of users.

5 If you modified the search filter in the preceding step, click **Change Filter**. The page displays the users in the LDAP that match the search filter.

6 You can assign users to the user groups listed at the bottom of the page or you can assign them later.

7 Select the check boxes for the users you want to import. To import all users displayed, select the top check box.

8 On the Import Users page, click **Import**.

Importing LDAP Users and Groups Using the LDAP Authentication Configuration Tool

The LDAP Authentication Configuration tool allows you to import both LDAP users and user groups into the SA Model Repository. It is, however, a more complex process that requires some preparation. This tool can be run from the command line or by selecting and running the LDAP Authentication Configuration tool from the SA Client APX Library.

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 10 LDAP Authentication Configuration Tool Prerequisites

Prerequisite	Description
Hostname	The fully-qualified host name (FQHN) or IP address of the LDAP directory server that SA is to use.
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 Certification Authority (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 Certification Authority (CA) certificates, in Privacy Enhanced Mail (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: 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	1 The complete path to the file containing the trusted Certification Authority (CA) certificates, in Privacy Enhanced Mail (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 Privacy Enhanced Mail (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 Directory Information Tree (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 Distinguished Name (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.

Table 10 LDAP Authentication Configuration Tool Prerequisites

Prerequisite	Description
Bind Distinguished Name (DN)	Required only if anonymous search is disabled. The bind DN for the user who has read access to the DIT.
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 Username	The attribute for the unique username. For Active Directory the default is `SAMAccountName`. For Novell eDirectory the default is `cn`. For all other vendors, the default is `uid`.
Attribute For User Display Name	The attribute for the user display name. For Active Directory the default is `displayName`. For Novell eDirectory the default is `fullName`. For all other vendors, the default is `cn`.
Base DN	The base distinguished name (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 sub tree search, therefore, the search filter is only applicable to users at or below the base DN.
Search Filter Template	The Search Filter Template is used, with optional filter substitution, as the filter in the LDAP search for the user import. Any dollar sign (`$`) character in the template is replaced by the filter string specified in the Import Users page of the SAS Web Client. (The default value is an asterisk (`*`) which matches all entries.) For Active Directory the default is `(&(sAMAccountName=$)(objectCategory=person)(objectClass=user)(sAMAccountType=805306368))`. For Novell eDirectory the default is `(&(cn=$)(objectClass=person))`. For all other vendors, the default is `uid=$`.

The LDAP Authentication Configuration Tool Process

When you run the LDAP Authentication Configuration tool, you will be prompted depending on whether your LDAP Directory server requires SSL authentication or not and/or whether anonymous search is allowed or not.

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 the LDAP Authentication Configuration tool:

```
./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: **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 the LDAP Authentication Configuration tool:

```
./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 the LDAP Authentication Configuration tool:

```
./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 the LDAP Authentication Configuration tool:

```
./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 the LDAP Authentication Configuration tool:

```
./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 the LDAP Authentication Configuration tool:

```
./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.

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

Example LDAP Authentication Configuration Tool 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 anonymous 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=Administrator,cn=users,dc=hyrule,dc=local [N] :

You have entered the following information:
LDAP Directory Server FQHN/IP                     : sample-centos.example.com
LDAP Directory Server Port                        : 8389
SSL Enabled?                                      : false
Bind DN                                           : cn=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=$)(objectCategory=person)(objectClass=user)
(sAMAccountType=805306368))
LDAP Search Base Distinguished Name (DN)             :
cn=users,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)(sAMAcco
untType=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.

LDAP Users & Groups Synchronization Configuration
Do you want to configure users & groups synchronization? [Y] :
LDAP User Group Synchronization Configuration
Enter the LDAP search base distinguished name (DN) for the user groups
[cn=users,dc=hyrule,dc=local]
 :
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 :
(&(sAMAccountName=*)(objectCategory=person)(object
Class=user)(sAMAccountType=805306368))
Found 4 LDAP users

Parsing search results ...
Searching LDAP directory server for all LDAP user gruops ...

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=$)(objectCategory=person)(o
bjectClass=user)(sAMAcc
ountType=805306368))
aaa.ldap.search.base.template=cn=users,dc=hyrule,dc=local
aaa.ldap.enable.users.groups.sync=true
aaa.ldap.search.usergroup.naming.attribute=SamAccountName
aaa.ldap.search.usergroup.membership.naming.attribute=member
aaa.ldap.search.usergroup.base.template=cn=users,dc=hyrule,dc=local
aaa.ldap.search.usergroup.filter.template=(&(cn=$)(objectCategory=group))

Are you sure? [Y] :
Saving LDAP configuration ...
LDAP configuration successfully saved.
Do you want to schedule a recurring job for LDAP users & user groups
synchronization? [Y] :
Select one of the following recurring schedule for LDAP users & user groups
synchronization job:

   1) Daily
   2) Weekly
   3) Monthly

Enter 1, 2, or 3  [3] : 1
Scheduling users & user groups synchronization job ...
LDAP users & user groups synchronization job has been successfully schedule.
Job ID=110001
```

Viewing Imported LDAP User Groups in the SAS Web Client

After you have imported LDAP users and user groups using the LDAP Authentication Configuration tool, you can view the user groups in the SAS Web Client. Log in to the SAS Web Client and select Users and Groups in the navigation panel, then the Groups tab. Figure 9 shows a screen similar to what you will see:

Figure 9 LDAP User Groups in the SAS Web Client

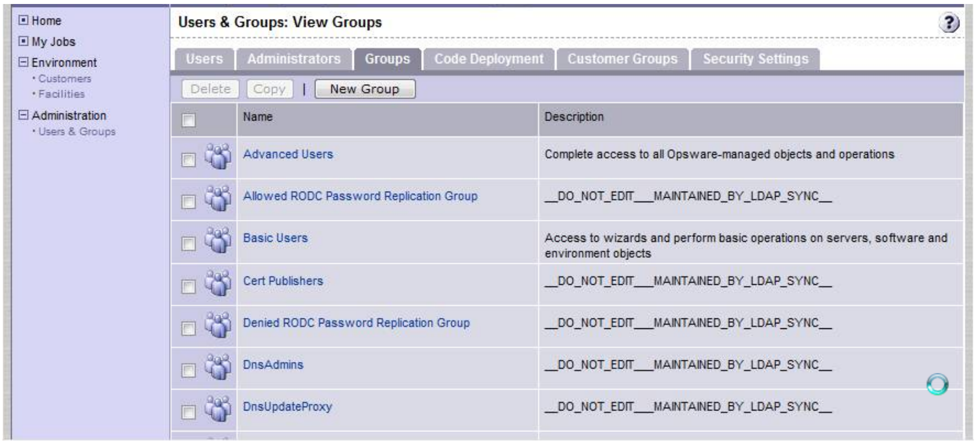

□	Name	Description
□	Advanced Users	Complete access to all Opsware-managed objects and operations
□	Allowed RODC Password Replication Group	__DO_NOT_EDIT___MAINTAINED_BY_LDAP_SYNC__
□	Basic Users	Access to wizards and perform basic operations on servers, software and environment objects
□	Cert Publishers	__DO_NOT_EDIT___MAINTAINED_BY_LDAP_SYNC__
□	Denied RODC Password Replication Group	__DO_NOT_EDIT___MAINTAINED_BY_LDAP_SYNC__
□	DnsAdmins	__DO_NOT_EDIT___MAINTAINED_BY_LDAP_SYNC__
□	DnsUpdateProxy	__DO_NOT_EDIT___MAINTAINED_BY_LDAP_SYNC__

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

Synchronizing LDAP Users

After you have imported users from the LDAP directory server, you can use the LDAP Authentication Configuration tool to synchronize LDAP users.

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 the LDAP Authentication Configuration tool:

5 `./ldap_config.sh`

6 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 since 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.

RSA SecurID®/SA 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 document describes how to take advantage of SecurID authentication in your SA system, it does not attempt to 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 assimilated into your existing SecurID environment. As far as the RSA authentication server is concerned, the Web Services Data Access Engine (twist) 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 take advantage of the feature:

➤ The first two tasks must be performed on every Web Services Data Access Engine host in your Multimaster Mesh or in installations that have multiple installed 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(s) and restarting after editing the `loginModule.conf` file to enable SecurID authentication in SA.
- Creating/modifying users in the SAS Java Client or SAS Web 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 much stronger authentication than a user password.

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 known as a *passcode*.

Table 11 shows the most typical authentication methods and which are supported by SA/ SecurID integration.

Table 11 SecurID Authentication Methods

Authentication method	**description**
Normal Authentication	The most used method. The user's PIN is assigned (*provisioned*). 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 due to the fact 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 know the IP address of the SecurID authentication server and be able to communicate with it in a secure manner.

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 You must 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 or both 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 the component:

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

HP-UX: `/sbin/init.d/opsware-sas stop twist`

AIX: `/etc/rc.d/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(s) on all servers:

Linux: `/etc/init.d/opsware-sas start twist`

HP-UX: `/sbin/init.d/opsware-sas start twist`

AIX: `/etc/rc.d/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 RAS SecurID authentication server (ACE server) and then must either be created or modified in the SAS Client to use SecurID authentication.

In either the SAS Client or the SAS Web 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 See "Managing Users and User Groups" on page 24.

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.

Code Deployment Permissions

Permissions to perform CDR operations are based on user membership in user groups predefined specifically for CDR. Users must also have the necessary permissions for the customer associated with the servers. Except for the Super User group, CDR operations are customer specific. A member of the Super User group can perform CDR operations on the servers of any customer.

➤ The SAS Web Client might still show the legacy term CDS. However, all documentation references use SA Code Deployment & Rollback term CDR.

The SAS Web Client includes predefined user groups that have specific permissions to perform CDR operations. SA administrators create and add users to these user groups to grant them permissions to perform specific CDR operations, based on their role in an organization. When logged on to the SAS Web Client, users see only the services, synchronizations, and sequences that they have authorization to perform because of their user group membership. Users are assigned to these groups as part of the Create User process.

See "Code Deployment User Groups" on page 278 in Appendix for more information.

See the *SA User's Guide: Server Automation* for information about the process to deploy code and content to managed servers.

➤ When a user requests a service operation, synchronization, or sequence, an e-mail notification is sent to the individuals assigned to actually perform the requested service operation or synchronization.

Adding Members to a Code Deployment User Group

Permissions to perform specific Code Deployment operations are granted based on a user's membership in specific Code Deployment user groups.

1. From the navigation panel, select Administration ➤ Users & Groups. The Manage Users: View Users page appears.
2. Select the Code Deployment tab.
3. Select the code deployment user group that you want to modify by clicking the hyperlinked user group name. The Users and Groups: Edit Code Deployment Group - [group name] page appears.
4. From the drop-down list, choose the customer whose group membership to modify.

➤ Code Deployment permission is assigned based on the associated SA customer. You cannot select Customer Independent, Not assigned, and SA customers and modify their group membership.

5. To add a user to the group, select the name in the left box, and then click the right arrow.
6. Click **Save** when you finish moving the user names to the box on the right.

   A confirmation page appears.
7. Click **Continue**.

   The Users & Groups: View Code Deployment Group page appears. You can continue modifying Code Deployment Groups, or you can select another function.

User and Security Reports

The Reporting feature in 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.

SA allows you to run the 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

2 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, 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 rollback, 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. In other words, the cure must not be worse than the disease. 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 that stores comprehensive logs of who did what on which server when in a central, secure repository.
- **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 users with 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 strictly 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.

In addition, 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 Lightweight Directory Access Protocol (LDAP) system. For example, if a company has an existing Microsoft Active Directory implementation, they can synchronize with the directory server to re-use 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 pre-defined 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. *Feature 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.

Another dimension that security administrators can control in SA is 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-like hierarchy, and setup fine-grained permissions for creation, viewing, modification, and execution. This model provides for 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.

➤ See User and User Group Setup and Security on page 13 in Chapter 1 for information about 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 including 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 6 months), and they can easily create a data warehouse that stores the audit trail (and other SA data) for longer periods of time.

Signed MD5 Checksums for Packages in the Software Repository

When SA users upload software to the Software Repository, SA automatically computes an MD5 checksum for the package. SA uses a combination of the software package contents and a secret key that only the Software Repository knows to generate the MD5 checksum. This prevents users from tampering with the software in the Software Repository. Before SA installs software on a managed server, it re-computes the MD5 checksum for the software before permitting the download. This helps ensure that the software installed by SA is the exact 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 reside in a single physical location. 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 an actual paying customer, a cost center, or simply 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, though 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 which 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**: The software policy modeling feature in SA 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 preventing any undetectable modification to audit logs. Since 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.

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 their local desktop or server. The SA browser communicates securely using HTTPS to the SA Command Center. Users provide username and passwords to login to SA; the credentials are authenticated either within SA, or optionally within 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 using 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 backend 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 which is generated when SA is installed. The component uses its public-key certificate when authenticating itself to another component. In this fashion, most inter-process communication is strongly authenticated, encrypted using the strongest ciphers available, and integrity checked.

Figure 10 Component Communication.

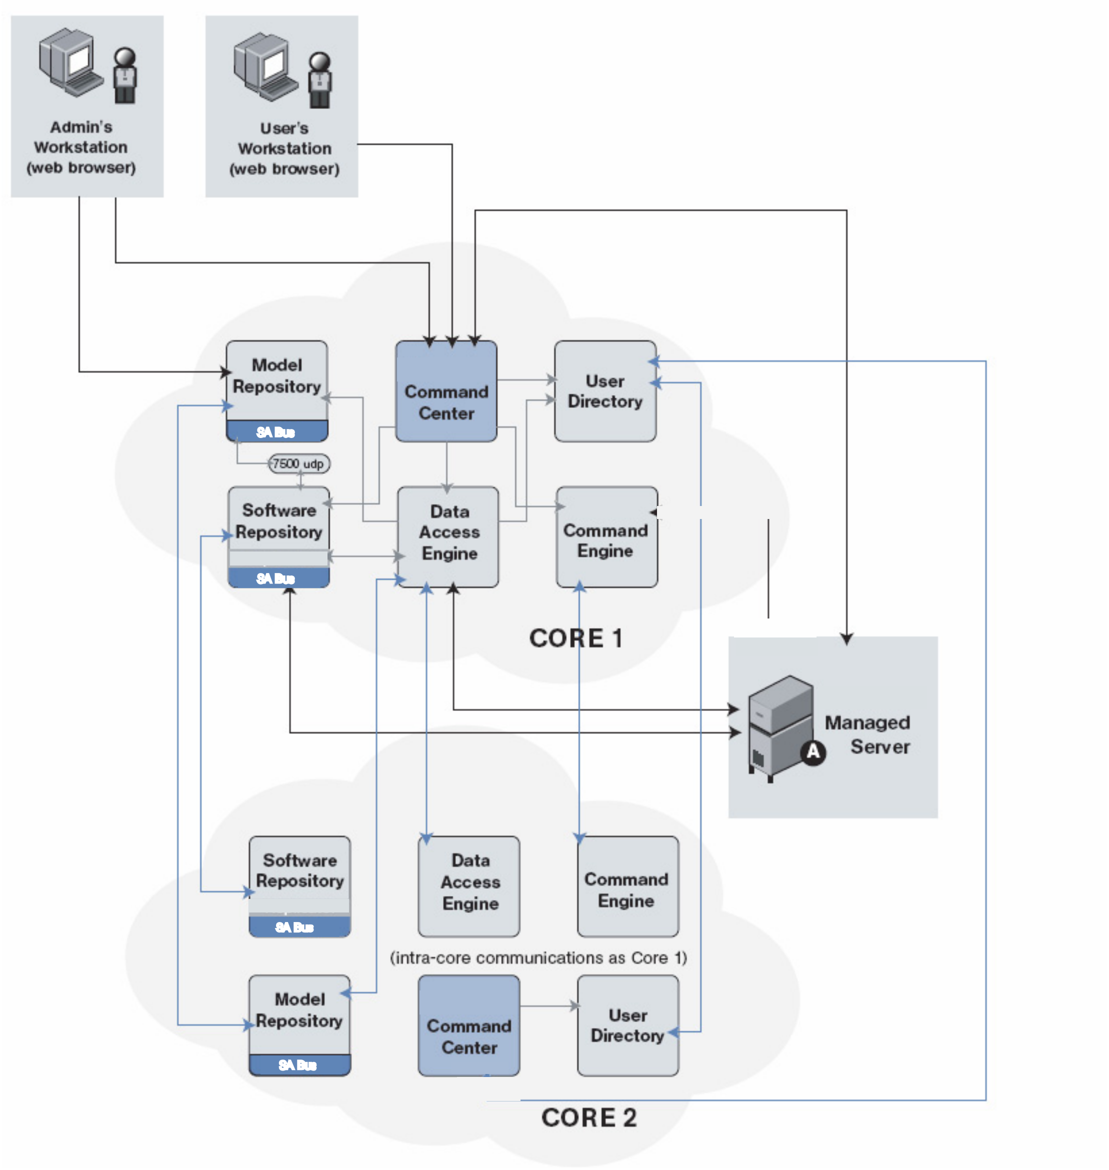

Communication between Agents and SA Core Components

The SA Server Agent also 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) helps 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 successfully 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 re-computes an MD5 checksum for the package along with a secret key it knows. Only if the MD5 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, since 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 expiry.
- 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 theSA instance of Apache.

Communication Between SA Cores

If you are running SA across multiple data centers SA automatically synchronize 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 doesn't 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

A Satellite, rather than a full SA core, installed at secondary locations to enable management of remote servers is as seamless as management of data center servers. The Satellite consists of an SA Gateway and a Software Repository Cache. A Satellite Gateway provides 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 in 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 either 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**: Enables 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 quickly and easily identify vulnerabilities as soon as you learn about them, and deploy the appropriate fixes without consuming extra resources.

Recognizing that no single standards 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:

1. **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/*)
2. **Microsoft (MS) Security Guide**: A standard established by Microsoft that details the configuration settings to harden Windows servers. (*http://www.microsoft.com/*)
3. **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 today's 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 7.80 provides a *Core Recertification Tool* that allows you to recertify SA Cores and Agents. The Core Recertification Tool automates and speeds up the process of issuing new security certificates as compared to previous releases. This tool is separate from and compatible with the existing Agent Recertification, also known as Server Recert Custom Extension utility.

Major advantages of the Core Recertification Tool are:

- The ability to regenerate all of SA certificates before their expiration, which effectively shorten their life span
- The ability to mitigate certificate compromises.

SA is a closed Public Key Infrastructure (PKI) system which utilizes X.509 v3 certificates to facilitate authentication, authorization, and secure network communications. An X.509 certificate is a form of identification which 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 certification authorities (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 change 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.

➤ 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 in a different hosts or under a different directories, will not be supported by this tool.

Agent vs. Core Recertification

There is an important distinction between Agent and Core Recertification. SA provides a Custom Extension*, Recert_Server*, to handle Agent Recertification, however, Recert_Server cannot regenerate an Agent CA's identity.

Core Recertification regenerates all of SA's certification authorities' identities and their respective hierarchies. Therefore, Core Recertification implies Agent Recertification. This document focuses on Core Recertification.

Upgrading after Core Recertification

Core recertification does not update the crypto database (CADB) on all cores. Only the First Core has the latest CADB. You can determine the First Core by running the command:

```
./corerecert --status
```

in `/opt/opsware/oi_util/OpswareCertTool/oi_utils/` of the core in which you performed the recertification.

Before upgrading to a newer SA release or patch, you must do the following:

1 Copy the CADB (`/var/opt/opsware/crypto/cadb/realm/*`) from the First Core to the same directory on the core server being upgraded.

2 On the core server being upgraded, issue the following commands:

```
rm -rf /var/opt/opsware/crypto/oi
rm -rf /var/opt/opsware/crypto/gateway
rm -rf /var/opt/opsware/crypto/dhcp
rm -rf /var/opt/opsware/crypto/word_upload
```

Core Recertification Phases

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

Table 12 describes the Core Recertification phases:

Table 12 Core Recertification Phases

Phas e		Description
1-3		Backup 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 utility. All the existing crypto materials are backup into the `crypto.`*<session number>* directories. Each Core component has its own backup directory.
4		[*Only required when using Code Deployment & Rollback (CDR)*] 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.
6	a	**Mesh Restart**: Restart the Mesh so that it trusts both the new and old CA hierarchies.
	b.	**Start Scheduled Mesh Restart**: Using the configuration file parameter, you can schedule a delayed start for the Multimaster Mesh Core restart that is appropriate for your maintenance window(s).
7		Recertify the Gateways.
8		Recertify the Agents.
9	a.	Recertify the Core components; issue the command `touch /var/opt/opsware/crypto/twist/upgradeInProgress` on First Core; Mesh restart; Regenerate Signatures.
	b.	Check Mesh Restart status. If the Mesh has successfully restarted, all the Core components are now using the new crypto material while still trust the old crypto material.
12		[*Optional*] Remove old CAs in the Multimaster Mesh. Required only when CAs have been compromised or no longer wish to trust the old CAs.
13	a.	[*Optional*] Remove the old CA hierarchies. Required only when CAs have been compromised or no longer wish to trust the old CA hierarchies.
	b.	[*Optional*] Create jobs. Required only 13a is also required.

Figure 11 shows the flow and phases of the recertification process:

Figure 11 Core Recertification Phases and Flow

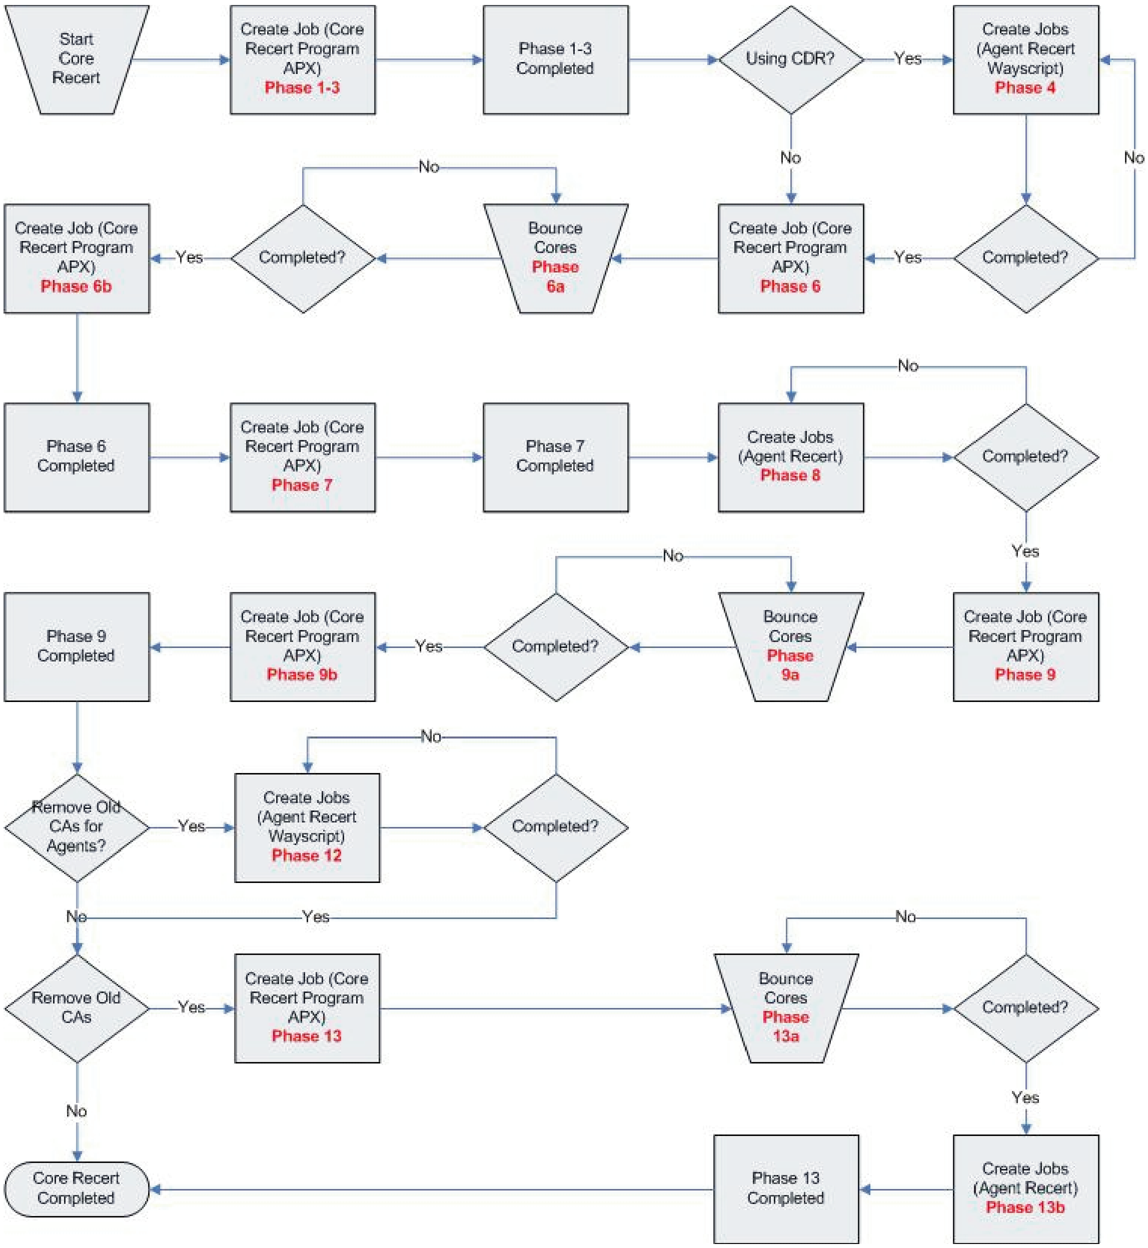
Start Core Recert
Create Job (Core Recert Program APX) Phase 1-3
Phase 1-3 Completed
Using CDR?
Yes
Create Jobs (Agent Recert Wayscript) Phase 4
No
No
Create Job (Core Recert Program APX) Phase 6b
Yes
Completed?
No
Bounce Cores Phase 6a
Create Job (Core Recert Program APX) Phase 6
Yes
Completed?
Phase 6 Completed
Create Job (Core Recert Program APX) Phase 7
Phase 7 Completed
Create Jobs (Agent Recert) Phase 8
No
Completed?
Yes
Phase 9 Completed
Create Job (Core Recert Program APX) Phase 9b
Yes
Completed?
No
Bounce Cores Phase 9a
Create Job (Core Recert Program APX) Phase 9
Remove Old CAs for Agents?
Yes
Create Jobs (Agent Recert Wayscript) Phase 12
No
Completed?
No
Yes
Remove Old CAs
Yes
Create Job (Core Recert Program APX) Phase 13
Bounce Cores Phase 13a
No
Completed?
Yes
No
Core Recert Completed
Phase 13 Completed
Create Jobs (Agent Recert) Phase 13b

Agent Recertification Phases

Three of the phases depicted in Figure 11 are *Agent Recertification phases*:

- **Phase 4**: Distributing new Agent CA. This phase is required only if you use *Code Deployment & Rollback (CDR)*. If you do not use CDR, you should skip this phase. 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 following properties, which you must specify in the Core Recertification configuration file:

Table 13 Core Recertification Configuration File: Agent Recertification Properties

Property Name	Req ?	Description	Example
`agent_recert.all.facilities.start_time=<HH:mm>`	Yes	The start time for the Agent Recertification phase. Users 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, `HH:mm`, where `00 <= HH < 24` and `00 <= mm < 60`. 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.	`agent_recert.all.facilities.start_time=18:30`
`agent_recert.facility.<facility_name>.start_time=<HH:mm>`	No	If present, the start time of the given facility will be used instead of `agent_recert.all.facilities.start_time`.	`agent_recert.facility.sacramento.start_time=08:00`

Table 13 **Core Recertification Configuration File: Agent Recertification Properties (cont'd)**

Property Name	Req ?	Description	Example
`agent_recert.all.facilities.duration=<hours>`	Yes	The duration, in hours, for the Agent Recertification job. Duration dictates how long the Agent Recertification job runs before stopping. Once the duration has elapsed and the success rate has not been reached, the Agent Recertification job will continue at the next start time. Users can overwrite this value for a given facility by specifying the `agent_recert.facility.<facility_name>.duration property.` Duration must be an integer value between 1 and 24.	`agent.recert.all.facilities.duration=8`
`agent_recert.facility.<facility_name>.duration=<hours>`	No	If present, the duration of the given facility will be used instead of `agent_recert.all.facilities.duration`	`agent_recert.facility.sacramento.duration=10`

Table 13 Core Recertification Configuration File: Agent Recertification Properties (cont'd)

Property Name	Req ?	Description	Example
agent_recert.all.facilities.success_rate=<whole percentage>	Yes	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. You can overwrite this value for a given facility by specifying the agent_recert.facility.<facility_name>.success_rate property. Success rate must be an integer value between 1 and 100.	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
agent_recert.all.facilities.job_notification=<email addresses>	No	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.	agent_recert.all.facilities.job_notification=admin@example.com
agent_recert.facility.<facility_name>.job_notification=<email addresses>	No	If present, the job notification for the given facility will be used instead of agent_recert.all.facilities.job_notification.	agent_recert.facility.sacramento.job_notification=admin3@example.com

Agent Recertification Job Flow

Figure 12 shows the Agent Recertification job flow:

Figure 12 Agent Recertification Job Flow

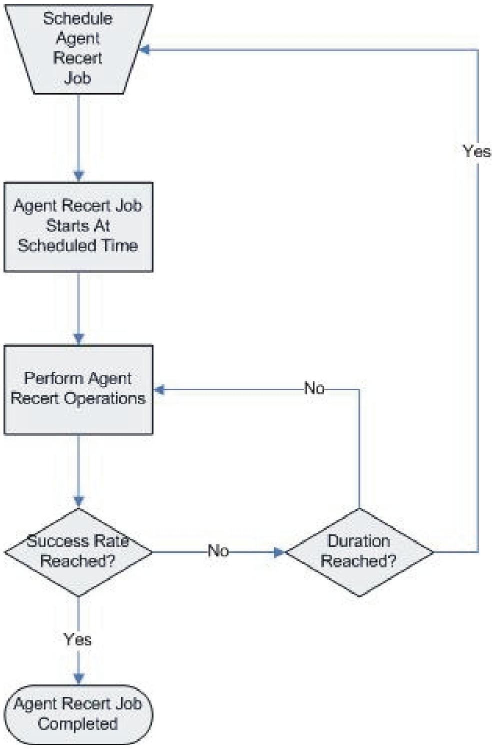
Schedule Agent Recert Job
Yes
Agent Recert Job Starts At Scheduled Time
Perform Agent Recert Operations
No
Success Rate Reached?
No
Duration Reached?
Yes
Agent Recert Job Completed

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/OpswareCertTool/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>]
```

Arguments

Table 14 describes the valid arguments for the Core Recertification utility:

Table 14 Core Recertification Utility 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/OpswareCertTool/recert_utils/corerecert.conf`.
`--doit`	Reruns or forces a rerun of a given Core Recertification phase. This is useful when certain newly added components has missed by 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 `corerecert` executable.
`-s, --status`	Displays the current status of the recertification process.
`-d, --debug`	Sets Core Recertification to debug mode, debug logs are available in `/tmp/recerttool.log`.
`--summary`	Prints out the current status summary, shorter version of `--status`.
`--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.

‼ 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 support before adding new Core components while a Core Recertification is in progress.*

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 support before performing Core Recertification.

Security Considerations

The following security issues should be considered:

Crypto Database Files

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

The SA crypto database consists of the file:

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

This files 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 backup 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. Therefore, 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 use has root access to the Core host.

Creating the Core Recertification User

In order to use the Core Recertification utility, 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

   – - Core Recertification (**Client ➤ Core Recert**)

   – -Agent Recert (**Client ➤ Agent Recert**)

3 Add the Core Recertification user to the SA System Administrators user group.

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 Recert Prerequisites

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

- Select a new password to protect the crypto materials and decide on 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

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 is comprised of multiple phases and most of them require the crypto database password. It is very crucial to protect the crypto database password.

!! Some of the Core Recertification tasks are accomplished by Application 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 SAS Web 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.

Since 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 allow 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 support before using running the Core Recertification Tool*.

Configuring Core Recert

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. Since 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.:

Table 15 Core Recertification Configuration File: Properties

Property Name	Req ?	Description	Example
		Global Properties	
`username=<username>`	Yes	User name of the user who has privilege to perform Core Recertification operations	`username=jdoe`
`password=<password>`	Yes	Password of the user who has privilege to perform Core Recert operations.	`password=dontask`
		Agent Recertification Properties	

Table 15 Core Recertification Configuration File: Properties (cont'd)

Property Name	Req ?	Description	Example
`agent_recert.using_cdr=<0 | 1>`	No	Indicates Code Deployment & Rollback (CDR) feature is being used. If CDR feature is not being used, the Agent CA push will be skipped. Also, the old Agent CA cleanup phase will be skipped unless `agent_recert.cleanup_old_agent_ca property` is set to 0. Valid values are 1 (true) or 0 (false). Any other value will result in an invalid property error. Default: 0.	`agent_recert.using_cdr=0`
`agent_recert.cleanup_old_agent_ca=<0 | 1>`	No	Indicates whether to clean up the old Agent CA after Core Recertification. If the CDR feature is not used, cleanup of old Agent CA phase is not necessary and should 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.	`agent_recert.cleanup_old_agent_ca=0`

Table 15 Core Recertification Configuration File: Properties (cont'd)

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

Table 15 Core Recertification Configuration File: Properties (cont'd)

Property Name	Req ?	Description	Example
`agent_recert.all.facilities.duration=<HH>`	Yes	The default duration, in hours, for the Agent Recertification operation in all facilities. Duration must be an integer value between 1 and 24. You can override the duration for a given facility by specifying the `agent_recert.facility.<facility name>.duration` property	`agent_recert.all.facilities.duration=2`
`agent_recert.facility.<facility name>.duration=<HH>`	No	Overrides the default duration for a specific facility.	`agent_recert.facility.yellow.duration=10`
`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`

Table 15 Core Recertification Configuration File: Properties (cont'd)

Property Name	Req ?	Description	Example
`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.yellow.job_notification=saadmin@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 in `/tmp/core_recert.out` in the Global Shell. Default: 0.	`debug =1`
`base_core_server_ref=<n>`	No	Server reference of the host from which you launch Core Recertification.	`base_core_server_ref=10010`

Table 15 Core Recertification Configuration File: Properties (cont'd)

Property Name	Req ?	Description	Example
`job_schedule= <YYYY:MM:DD:HH:mm>`	No	Job schedule for the current Core Recertification phase jobs. 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`. If this property is not specified, the job starts immediately.	`job_schedule= 2009:02:12:23:05`
`job_schedule.gateway_recert. <facility name>= <YYYY:MM:DD:HH:mm>`	No	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`. If this property is not specified, the `job_schedule` property for the gateway recertification phase is used.	`job_schedule.gateway_recert.<facility name>= 2009:02:12:23:05`
`job_notification= <email_address>`	No	Job notification for all Core Recertification phase jobs. You can override this value for a given phase by specifying the `job_notification. <phase_number>` property	`job_notification= admin@example.com>`

`job_notification. <phase_number>= <email_address>`	No	Job notification for a specified Core Recertification phase.	`job_notification.7=`
			`saadmin@example.com`

Table 15 Core Recertification Configuration File: Properties (cont'd)

Property Name	Req ?	Description	Example
`job_notification.gateway_recert.<facility name>=<email_address>`	No	Job notification for the Gateway Recert phase for a given facility.	`job_notification.gateway_recert.yellow=admin@acme.com`
`cleanup_old_opsware_ca=<0 | 1>`	No	Specifies whether to clean old SA CA after Core Recert. 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. The valid values are 1 (true) or 0 (false). Any other value will result in an invalid property error. Default: 0 (false)	`cleanup_old_opsware_ca=1`

Ensure that All Cores are Running/Resolve Conflicts

Before performing Core Recertification, HP strongly recommends that you run **System Diagnosis** from the SA Command Center, on all Cores to be recertified to ensure that they are running correctly. You should also use the SA Command Center **Multimaster Tools** to detect and resolve any transaction conflicts.

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/OpswareCertTool/recert_utils/discover_mesh -p
```

3 Check the output to make sure it reflects your current Mesh setup. If not, contact HP support before proceeding with Core Recertification.

Recertifying SA Cores

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 Edit:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert.conf
```

to ensure the information is correct.

4 Run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

to ensure Core Recertification is not currently in progress.

5 Run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/discover_mesh -p
```

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

6 Run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 1
```

from the command line to initialize Core Recertification.

7 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

until it has indicated Phase 4 is in progress.

8 If you are not using CDR, go to step 10 on page 86. Otherwise, run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 4
```

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

9 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

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

➤ 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 8 on page 86. You only need to reschedule the facilities that had errors. You do not need to reschedule the Agent Recert job for the successful facilities.

10 Run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 6
--doit
```

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

11 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

until it has indicated `mesh_restart_pending`.

At this point, you must work with the SA System Administrator to restart the mesh.

➤ 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 10 on page 86.

12 After the mesh has successfully restarted, run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 6
```

from the command line to continue phase 6.

13 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

until Phase 7 should be started. If you encounter any errors during this stage, make sure you resolve the errors and go back to step 12 on page 87.

14 Run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 7
```

from the command line to start phase 7.

15 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

until Phase 8 should be started. If you encounter any errors during this stage, make sure you resolve the errors and go back to step 14 on page 87.

16 Run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 8
```

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

17 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

until all Agents have successfully been recertified.

➤ This step could take days depending on 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 during this stage, resolve the errors and go back to step 16 on page 87. You only need to reschedule the facilities that had errors. You do not need to reschedule the Agent Recertification job for the successful facilities.

18 Run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/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.

19 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/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 18 on page 87.

this point, you must work with your SA System Administrator to restart the mesh.

➤ 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 18 on page 87. You only need to reschedule the facilities that had errors. You do not need to reschedule the Agent Recertification job for the successful facilities.

20 After the mesh has successfully restarted, the Recertification User must run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 9
```

from the command line to continue phase 9.

21 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

until Phase 12 should be started. If you encounter any errors during this stage, make sure you resolve the errors and go back to step 20 on page 88.

22 If you do not intend to remove the Agent CA, skip to step 24 on page 88. Otherwise, run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 12
```

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

23 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

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

➤ 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 22 on page 88. You only need to reschedule the facilities that had errors. You do not need to reschedule the Agent Recertification job for the successful facilities.

24 Run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 13
--doit
```

from the command line to start phase 13. If you do not want to remove the old CAs, a Mesh restart is not required in this phase.

25 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

until it indicates `mesh_restart_pend`.

At this point, you must work with the SA System Administrator to restart the mesh.

➤ 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 tostep 24 on page 88.

26 After the mesh has successfully restarted, run:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --phase 13
```

from the command line to continue phase 13.

27 Monitor the progress on screen by running:

```
/opt/opsware/oi_utils/OpswareCertTool/recert_utils/corerecert --status
```

until it indicates that Core Recertification has completed successfully.

3 Multimaster Mesh Administration

This section documents how to administer and maintain an existing 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 Planning and Installation Guide* or consult your SA Support Representative.

Multimaster Facilities Administration

In the SAS Web Client, the term *Facility* refers to the collection of servers that a single SA Core or Satellite manages. For more information about Facilities and how they fit into Multimaster Mesh architecture, see the *SA Simple/Advanced Installation Guide*.

The SAS Web Client provides administrative capabilities for Facilities. You can modify information such as the Facility name, Facility ID, and custom attributes used to provide parameters to SA (for example, to customize displays or provide settings to use during installation or configuration of packaged software).

➤ Facilities must be manually associated with Customers in a Multimaster Mesh. If you have not yet associated Customers with your facility, see the the *SA Policy Setter Guide*.

Modifying a Facility’s Name

Perform the following steps to modify a Facility’s name:

1 Log in to the SAS Web Client.

2 From the Navigation panel, click **Environment ➤ Facilities**. The **Facilities** page appears and displays the names of the Facilities in your mesh.

3 Click the name of the Facility that you want to modify. The **Facilities: Edit Facility** page appears with the **Properties** tab automatically selected, as Figure 13 shows.

Figure 13 Properties Tab of the Edit Facility Page

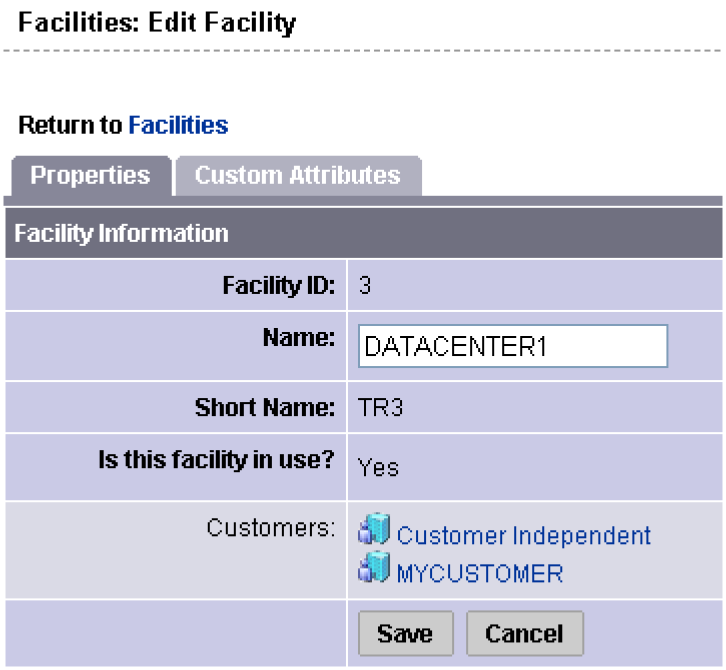

Facilities: Edit Facility

Facility Information

Facility ID:	3
Name:	DATACENTER1
Short Name:	TR3
Is this facility in use?	Yes
Customers:	Customer Independent MYCUSTOMER

4 You can change the SAS Web Client display name of the Facility by editing the **Name** field or click the **Return to Facilities** link to exit without making any changes.

5 To save the new Facility name, click **Save**. The SAS Web Client displays a message that confirms that the properties for that Facility were updated.

Adding or Modifying a Facility's Custom Attributes

Perform the following steps to modify a Facility's custom attributes:

1 Log in to the SAS Web Client.

2 From the Navigation panel, click **Environment ➤ Facilities**. The **Facilities** page appears and displays the names of the Facilities in your mesh.

3 Click the name of the Facility that you want to modify custom attributes for. The **Facilities: Edit Facility** page appears with the **Properties** tab automatically selected, as Figure 13 shows.

4 Select the **Custom Attributes** tab.

The **Custom Attributes** page appears and displays any *name-value pairs* associated with the Facility.

5 Click the name of an attribute to display the **Facilities: Edit Attribute for [facility name]** dialogue from which you can:

- make changes to the Facility's custom attribute values
- click **New** to add a new custom attribute: specify a name and a value for the attribute

‼ Be careful when you update or remove existing custom attribute settings as it can affect or disrupt operation of the operational environment. Contact your SA Support Representative to help you determine the appropriate changes to make when you update the information or settings for a specific Facility.

6 When you finish making updates to the Facility properties or custom attributes, click the **Return to Facilities** link.

➤ Contact your SA Support Representative if you need to make other changes not described in this section to the Facility properties.

Multimaster Mesh Conflict Administration

Multimaster Mesh administration tasks typically involve detecting, preventing, and resolving the various types of conflicts that can occur when you have users in multiple Facilities updating multiple, replicated Model Repositories. This section provides information on Multimaster Mesh administration within SA and contains the following topics:

- Types of Conflicts
- Model Repository Multimaster Component Conflicts
- Causes of Conflicts
- User Overlap
- User Duplication of Actions
- Out of Order Transactions

Types of Conflicts

A Multimaster Mesh installation can experience two types of conflicts:

- Transaction Conflicts
- Model Repository Multimaster Component Conflicts

Transaction Conflicts

A transaction conflict can occur when updates are made to the same record in different Model Repositories.

When transaction conflicts occur, the SA Multimaster components detect them and can be configured to email alerts to responsible parties.

The SA Core Components themselves cannot resolve the conflicts. SA administrators must use the Multimaster Tools in the SAS Web Client to resolve the conflicts at the target databases when they occur to ensure that the transaction updates are not lost. The administrator can also use the Multimaster Status Monitor described in the *SA Simple/ Advanced Installation Guide* to detect where there are conflicts.

Model Repository Multimaster Component Conflicts

Replication conflicts occur in an environment in which concurrent updates to the same data at multiple sites occurs. For example, when two transactions originating from users accessing different Model Repositories update the same row at nearly the same time, a conflict can occur. If you need basic information about database conflicts, see your *Oracle Database* documentation.

The probability of Multimaster conflicts occurring 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 Facilities in the mesh
- The number of installed SAS Web Clients — more users making updates
- The propensity for users to make changes in more than one Facility by using different SAS Web Clients

SA Transaction Conflict Handling

When SA flags a conflict, the following occurs:

1. The transaction is canceled.
2. All database rows affected by the transaction are locked, thereby preventing further changes to those rows.
3. The Outbound Model Repository Multimaster Component propagates the transaction lock to all remote databases in the mesh, thereby locking the rows in all Facilities.
4. An alert message with the conflict information is emailed to a user-configured mailing list.
5. The Inbound Model Repository Multimaster Component continues to the next transaction.

If the Inbound or Outbound Model Repository Multimaster Component encounters an exception that prevents it from going to the next transaction, it sends an email to the user-configured mailing list discribing the problem and shuts down.

➤ An SA administrator must manually resolve conflicts using the SAS Web Client. Resolving the conflict unlocks the rows. See Best Practices for Resolving Database Conflicts on page 96 in this chapter for more information.

Causes of Conflicts

Conflicts are typically caused by:

- User Overlap
- User Duplication of Actions
- Out of Order Transactions

User Overlap

Conflicts occur when a user concurrently makes a change using the SAS Web 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.

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 (OCLI) to upload the package `carol.conf`.
2. Carol immediately logs in to the SAS Web 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 SAS Web Client in Phoenix.
4. When the data eventually propagates from Seattle, SA generates a conflict because the data already exists in Phoenix.

Out of Order Transactions

Out of order transactions occur only in a Multimaster Mesh with three or more facilities.

Transactions between two Facilities should 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 command-line interface (OCLI) in one Facility, and immediately uses the SA Client to add the package to a Software Policy in different Facility.

The occurrence of out of order transactions can be aggravated by concurrent updates in different Facilities and/or problems with inter-Facility network connections.

For example:

1. Henry uses the OCLI 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.

Best Practices for Preventing Multimaster Conflicts

There are measures you can take to keep the number of Multimaster conflicts to a minimum.

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.
- 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 hasn't 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 Multimaster Mesh on page 106 in this chapter for more information.

  See the *SA Planning and 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 SAS Web Client facilitates this by allowing you to set permissions by Customer, Facility, and User Group types.

  See Permissions Reference on page 205 of this guide for more information about User Groups and SA Permissions.

Examining the State of the Multimaster Mesh

You can examine the state of a Multimaster Mesh by launching the SAS Web Client and selecting the **Multimaster Tools** option

When you select the **Multimaster Tools** option, the **Multimaster Tools: State View** page appears. In addition to a color-coded legend that shows possible transaction states:

Table 16 Multimaster Transaction State Color Code

Color	State
Red	Conflict
Orange	Not Sent
Yellow	Not Received
Gray	Unable to Connect
green	Good

The **Multimaster Tools: State View** page also provides:

- An overview of the health of the Multimaster Mesh achieved by an automatic check of all Facilities in the mesh.
- A display of the state of the last five transactions in the mesh and a list of all conflicting and unpublished transactions.
- The time that the SAS Web Client generated and cached the data. Click **Refresh** to refresh the time.

System Diagnosis Tools

Administrators can also use the System Diagnosis tools in the SAS Web Client to view information about the health of the multimaster components.

See SA Maintenance on page 131 in Chapter 5 for more information.

Multimaster State Monitoring Utility

If you prefer a command line utility, you can use the Multimaster State Monitoring Utility to get information about conflicts in your mesh. For more information, see the *SA Simple/ Advanced Installation Guide*.

Best Practices for Resolving 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. You should see your *Oracle Database Administration* documentation for more detailed and up-to-date information about identifying and resolving data and transaction conflicts.

This section contains the following topics:

- Types of Conflicts
- Guidelines for Resolving Each Type of Conflict

Types of Conflicts

The following are the most typically seen conflicts:

Table 17 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).

Table 17 Types of Conflicts (cont'd)

Conflict	description
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 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 SAS Web Client UI (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 SA 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 depend 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 SAS Web Client in the London Facility, John creates Node A1 as a subordinate node of Node A.
2. From the v 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.

Model Repository Multimaster Component Conflicts

This section provides information on resolving model repository, multimaster component conflicts and contains the following topics:

- Overview of Resolving Model Repository Multimaster Component Conflicts
- Resolving a Conflict by Object
- Resolving a Conflict by Transaction

Overview of Resolving Model Repository Multimaster Component Conflicts

SA administrators can view and resolve multimaster conflicts in any SAS Web Client by using the Multimaster Tools. The Multimaster Tools are available in all SAS Web Clients.

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 SAS Web Client of a single Facility. Do not attempt to resolve the same conflict multiple times from the SAS Web Client of different facilities.

If you see a large volume of conflicts that you cannot resolve by using the Multimaster Tools, contact your SA Support Representative for assistance synchronizing databases.

Resolving a Conflict by Object

Perform the following steps to *resolve conflicting transactions individually*:

1 Log on to the SAS Web Client.

2 From the Navigation panel, click **Administration ➤ Multimaster Tools**. The **Multimaster Tools: State View** page appears, showing a summary of all transactions and, if they exist, all conflicts. See Figure 14.

Figure 14 Transaction Table Showing Conflicts

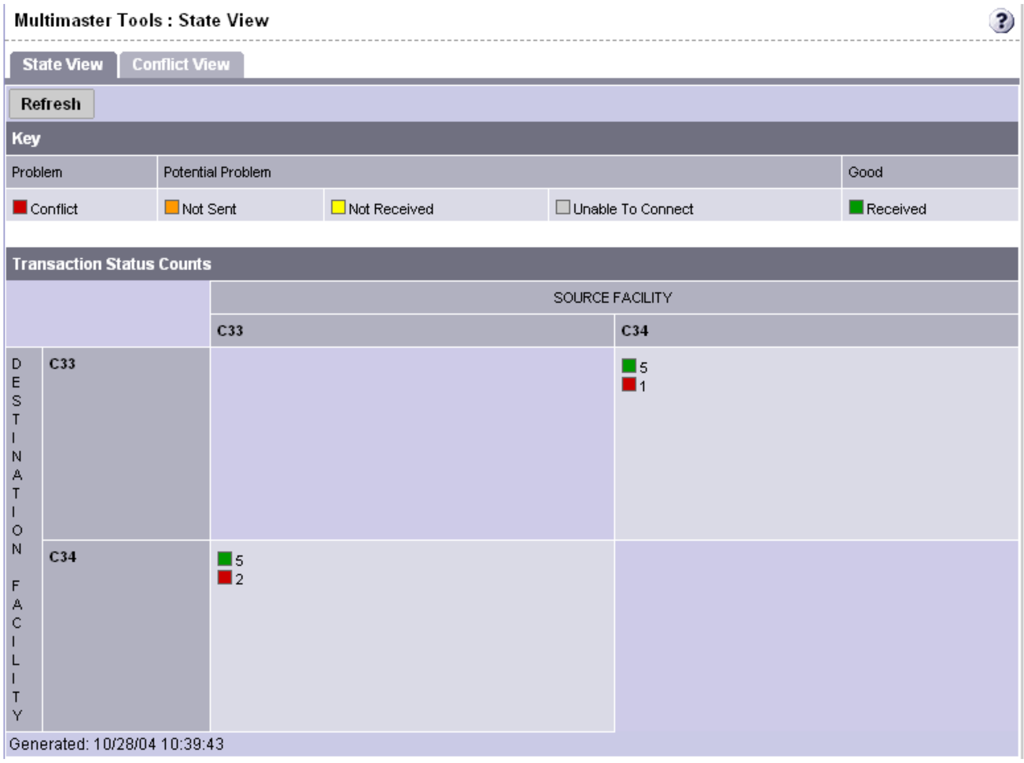

Multimaster Tools : State View

Key

Problem	Potential Problem			Good
🟥 Conflict	🟧 Not Sent	🟨 Not Received	□ Unable To Connect	🟩 Received

Transaction Status Counts

		SOURCE FACILITY	
		C33	C34
DESTINATION FACILITY	C33		🟩 5 🟥 1
	C34	🟩 5 🟥 2	

Generated: 10/28/04 10:39:43

Different types of transaction statuses are indicated by color-coded boxes:

— **Green**: The last five transactions that were successfully sent.

— **Orange**: All transactions that have not been published (sent to other facilities).

— **Red**: All conflicts.

Each box is displayed in a color scheme to indicate the status and success of the transaction. A key that explains the significance of the colors, like the one shown in Figure 15, is listed at the top of the page.

Figure 15 Conflict Color Key

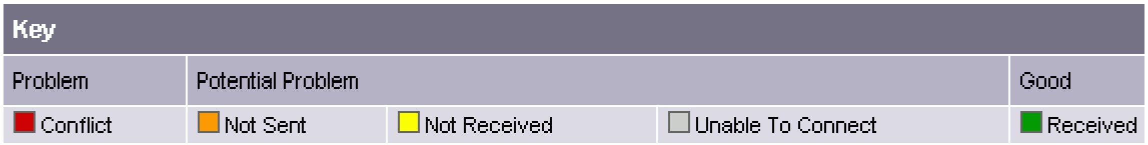

Key

Problem	Potential Problem			Good
🟥 Conflict	🟧 Not Sent	🟨 Not Received	□ Unable To Connect	🟩 Received

Red boxes indicate that one or more transactions between facilities are in conflict and need to be resolved.

3 To resolve a conflict, select the **Conflict View** tab. The **Multimaster Tools: Conflict View** page appears, as shown in Figure 16.

Figure 16 Multimaster Tools: Conflict Page showing all Conflicts in a Multimaster Mesh

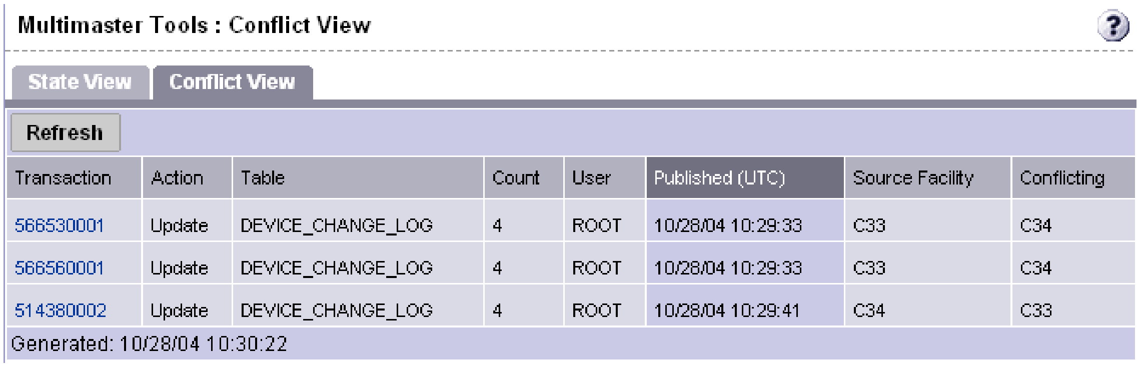

Transaction	Action	Table	Count	User	Published (UTC)	Source Facility	Conflicting
566530001	Update	DEVICE_CHANGE_LOG	4	ROOT	10/28/04 10:29:33	C33	C34
566560001	Update	DEVICE_CHANGE_LOG	4	ROOT	10/28/04 10:29:33	C33	C34
514380002	Update	DEVICE_CHANGE_LOG	4	ROOT	10/28/04 10:29:41	C34	C33

Generated: 10/28/04 10:30:22

The page lists each transaction by ID number (clickable link), the actions that caused the conflict, the database objects affected by the conflict, the user responsible for the conflict (listed by the IP of the SAS Web Client where the user made the change), when the offending action occurred, the source Facility that originated the transaction, and the facilities where the transaction conflicted.

The page might show a conflict where the data is the same in both facilities but a conflict exists, because the same change was made in both facilities. Even though the data is correct, the conflict still exists and must be resolved. See Best Practices for Resolving Database Conflicts on page 96 in this chapter for more information.

4 To resolve a conflict, click the transaction ID number link. You see the Multimaster Tools: Transaction Differences page, which shows a comparison of the objects between facilities, with any differences shown in red, as illustrated in Figure 17.

Figure 17 Transaction Differences Page Showing Conflicts Between Facilities

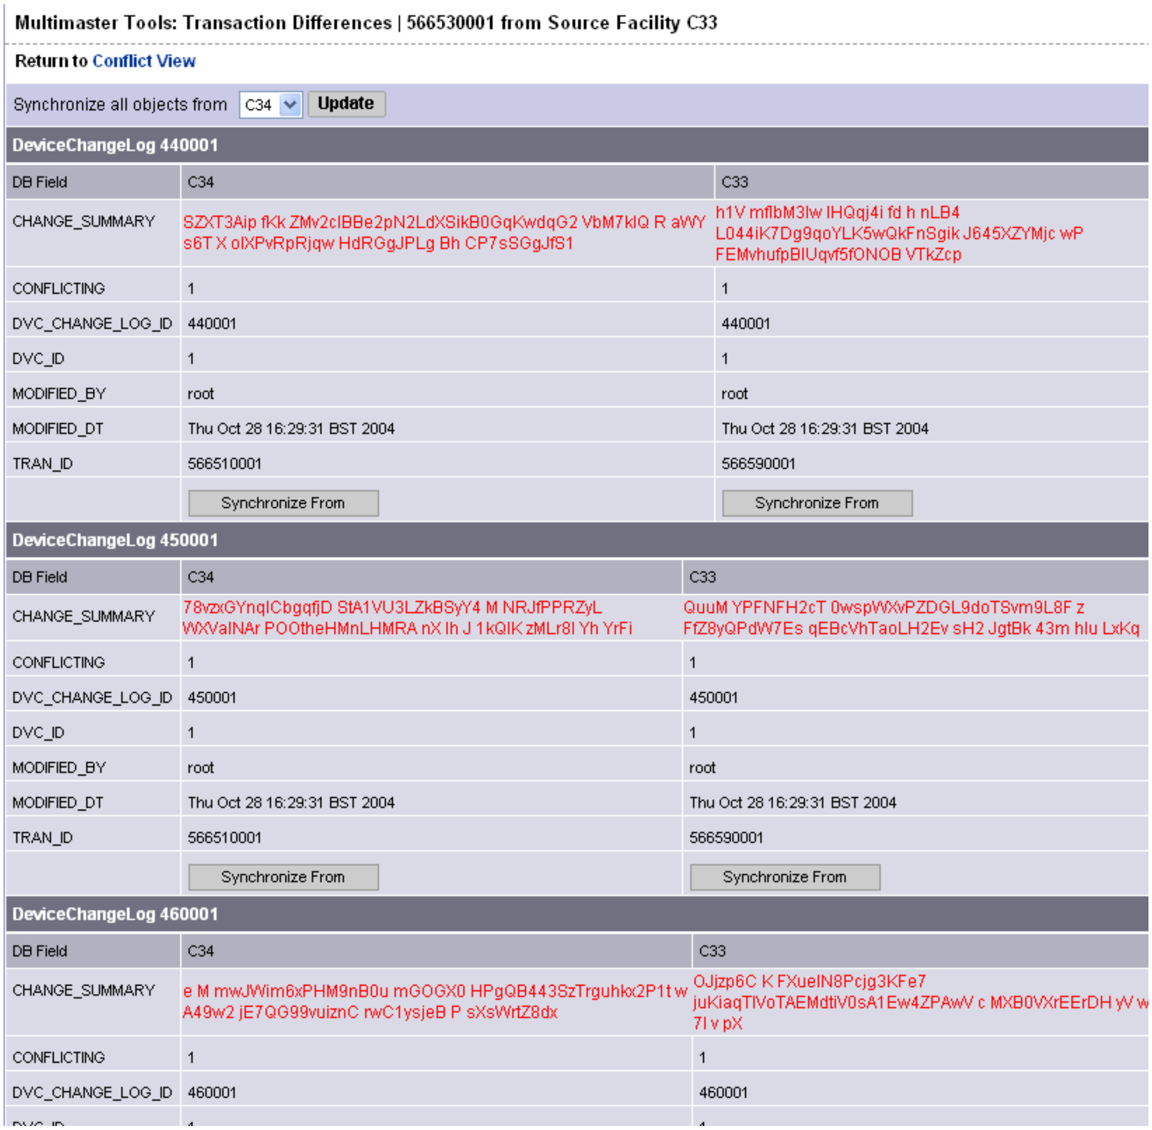

Multimaster Tools: Transaction Differences | 566530001 from Source Facility C33

Return to Conflict View

Synchronize all objects from C34 Update

DeviceChangeLog 440001

DB Field	C34	C33
CHANGE_SUMMARY	SZXT3Aip fKk ZMv2clBBe2pN2LdXSikB0GqKwdqG2 VbM7klQ R aWY s6T X olXPvRpRjqw HdRGgJPLg Bh CP7sSGgJfS1	h1V mflbM3lw IHQqj4i fd h nLB4 L044iK7Dg9qoYLK5wQkFnSgik J645XZYMjc wP FEMvhufpBIUqvf5fONOB VTkZcp
CONFLICTING	1	1
DVC_CHANGE_LOG_ID	440001	440001
DVC_ID	1	1
MODIFIED_BY	root	root
MODIFIED_DT	Thu Oct 28 16:29:31 BST 2004	Thu Oct 28 16:29:31 BST 2004
TRAN_ID	566510001	566590001
	Synchronize From	Synchronize From

DeviceChangeLog 450001

DB Field	C34	C33
CHANGE_SUMMARY	78vzxGYnqICbgqfjD StA1VU3LZkBSyY4 M NRJfPPRZyL WXValNAr POOtheHMnLHMRA nX Ih J 1kQIK zMLr8I Yh YrFi	QuuM YPFNFH2cT 0wspWXvPZDGL9doTSvm9L8F z FfZ8yQPdW7Es qEBcVhTaoLH2Ev sH2 JgtBk 43m hlu LxKq
CONFLICTING	1	1
DVC_CHANGE_LOG_ID	450001	450001
DVC_ID	1	1
MODIFIED_BY	root	root
MODIFIED_DT	Thu Oct 28 16:29:31 BST 2004	Thu Oct 28 16:29:31 BST 2004
TRAN_ID	566510001	566590001
	Synchronize From	Synchronize From

DeviceChangeLog 460001

DB Field	C34	C33
CHANGE_SUMMARY	e M mwJWim6xPHM9nB0u mGOGX0 HPgQB443SzTrguhkx2P1t w A49w2 jE7QG99vuiznC rwC1ysjeB P sXsWrtZ8dx	OJjzp6C K FXueIN8Pcjg3KFe7 juKiaqTIVoTAEMdtiV0sA1Ew4ZPAwV c MXB0VXrEErDH yW w 7I v pX
CONFLICTING	1	1
DVC_CHANGE_LOG_ID	460001	460001
DVC_ID	1	1

5 Click **Synchronize From** at the bottom of the conflict listing to resolve.

The Multimaster Tools insert or delete conflict objects in the transaction where necessary, and then propagate the change to every Facility in the Multimaster Mesh.

The Multimaster Tools: Object Synchronization Results page appears, displaying the results of the transaction synchronization, as shown in Figure 18.

Figure 18 Object Synchronization Results Page

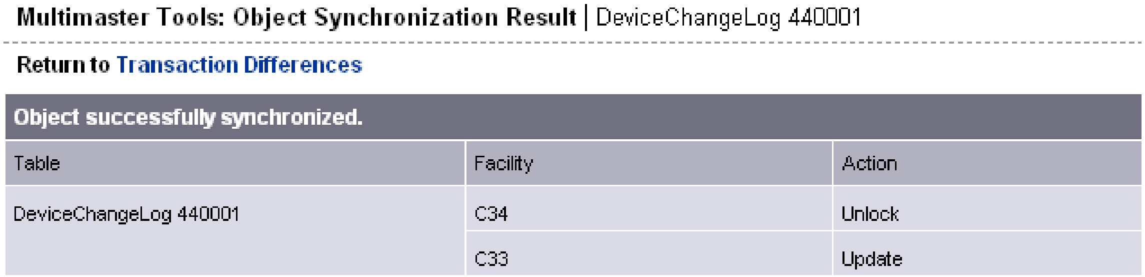

Multimaster Tools: Object Synchronization Result | DeviceChangeLog 440001

Return to Transaction Differences

Object successfully synchronized.

Table	Facility	Action
DeviceChangeLog 440001	C34	Unlock
	C33	Update

6 Click **Return to Transaction Differences**. The **Multimaster Tools: Transaction Difference** page displays again. Notice that the conflict object you synchronized shows on the page as being identical between the facilities, as shown in Figure 19.

Figure 19 Single Object Resolved

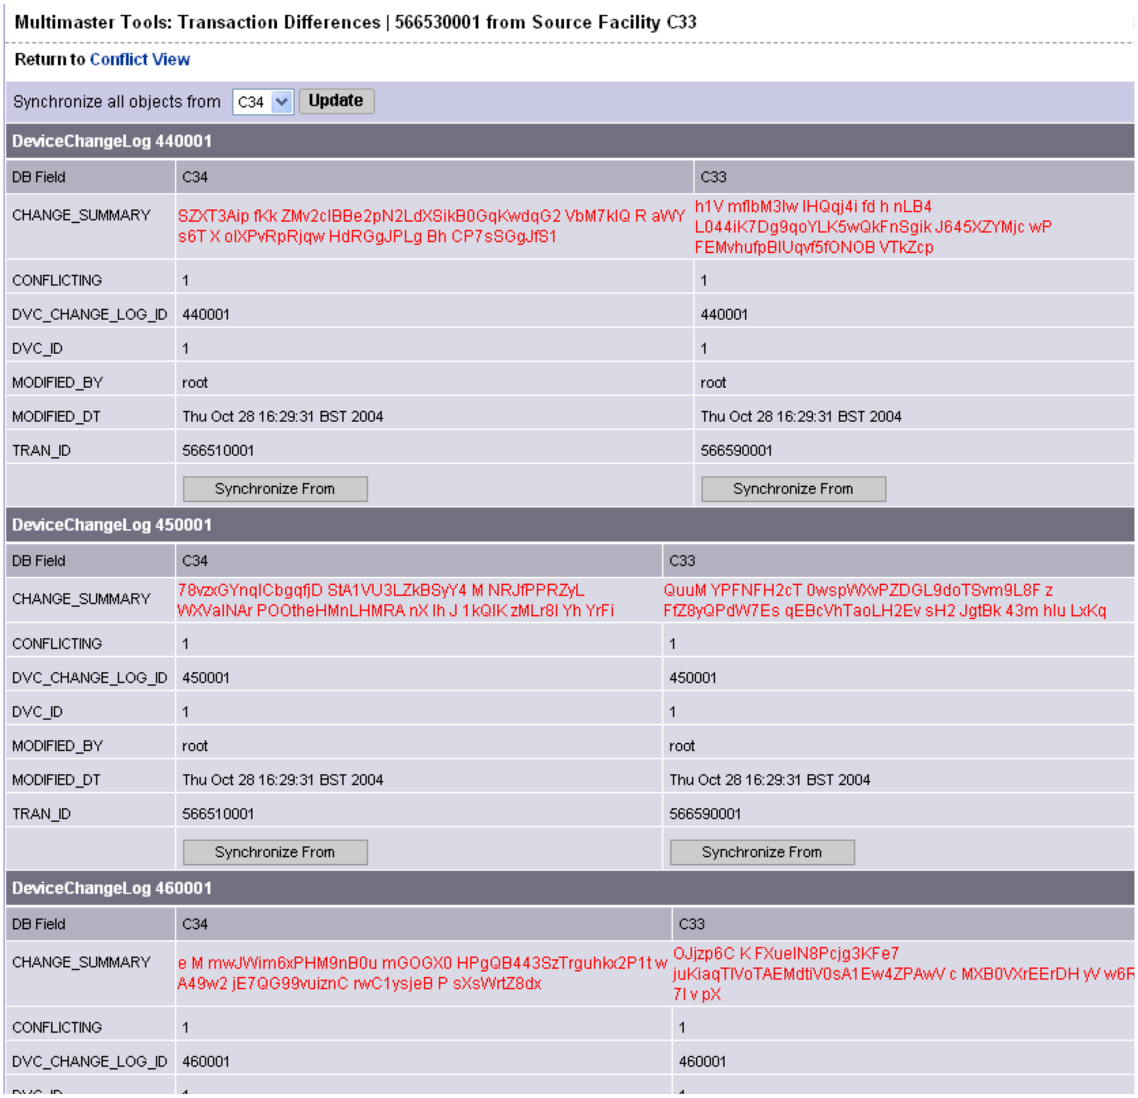

Multimaster Tools: Transaction Differences | 566530001 from Source Facility C33

Return to Conflict View

Synchronize all objects from C34 Update

DeviceChangeLog 440001

DB Field	C34	C33
CHANGE_SUMMARY	SZXT3Aip fKk ZMv2clBBe2pN2LdXSikB0GqKwdqG2 VbM7kIQ R aWY s6T X olXPvRpRjqw HdRGgJPLg Bh CP7sSGgJfS1	h1V mflbM3lw IHQqj4i fd h nLB4 L044iK7Dg9qoYLK5wQkFnSgik J645XZYMjc wP FEMvhufpBIUqvf5fONOB VTkZcp
CONFLICTING	1	1
DVC_CHANGE_LOG_ID	440001	440001
DVC_ID	1	1
MODIFIED_BY	root	root
MODIFIED_DT	Thu Oct 28 16:29:31 BST 2004	Thu Oct 28 16:29:31 BST 2004
TRAN_ID	566510001	566590001
	Synchronize From	Synchronize From

DeviceChangeLog 450001

DB Field	C34	C33
CHANGE_SUMMARY	78vzxGYnqlCbgqfjD StA1VU3LZkBSyY4 M NRJfPPRZyL WXValNAr POOtheHMnLHMRA nX Ih J 1kQIK zMLr8I Yh YrFi	QuuM YPFNFH2cT 0wspWXvPZDGL9doTSvm9L8F z FfZ8yQPdW7Es qEBcVhTaoLH2Ev sH2 JgtBk 43m hlu LxKq
CONFLICTING	1	1
DVC_CHANGE_LOG_ID	450001	450001
DVC_ID	1	1
MODIFIED_BY	root	root
MODIFIED_DT	Thu Oct 28 16:29:31 BST 2004	Thu Oct 28 16:29:31 BST 2004
TRAN_ID	566510001	566590001
	Synchronize From	Synchronize From

DeviceChangeLog 460001

DB Field	C34	C33
CHANGE_SUMMARY	e M mwJWim6xPHM9nB0u mGOGX0 HPgQB443SzTrguhkx2P1t w A49w2 jE7QG99vuiznC rwC1ysjeB P sXsWrtZ8dx	OJjzp6C K FXueIN8Pcjg3KFe7 juKiaqTIVoTAEMdtiV0sA1Ew4ZPAwV c MXB0VXrEErDH yV w6R 7I v pX
CONFLICTING	1	1
DVC_CHANGE_LOG_ID	460001	460001
DVC_ID	1	1

7 Continue synchronizing the objects in the transaction until all objects in the transaction are synchronized. (Repeat steps 3 and 4.) When all objects in the transaction are synchronized, **Mark Resolved** appears at the bottom of the page, as Figure 20 shows.

Figure 20 When All Conflicts Are Resolved, the Mark Resolved Button Appears

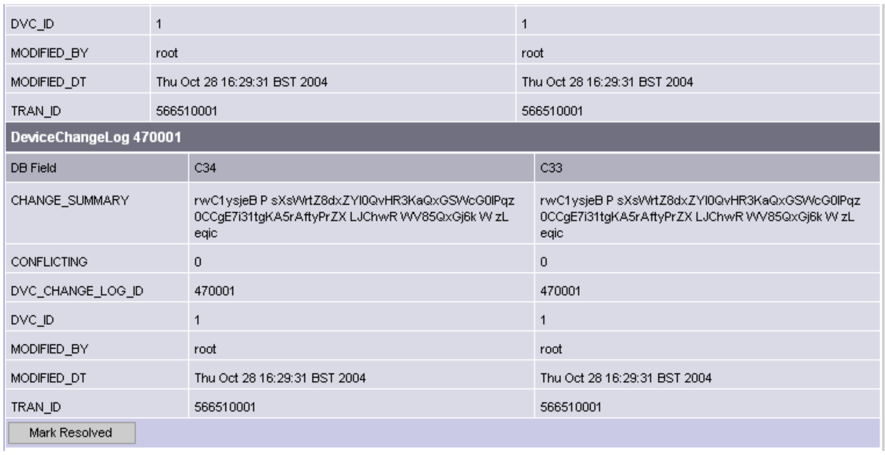

DVC_ID	1	1
MODIFIED_BY	root	root
MODIFIED_DT	Thu Oct 28 16:29:31 BST 2004	Thu Oct 28 16:29:31 BST 2004
TRAN_ID	566510001	566510001

DeviceChangeLog 470001

DB Field	C34	C33
CHANGE_SUMMARY	rwC1ysjeB P sXsWrtZ8dxZYI0QvHR3KaQxGSWcG0lPqz 0CCgE7i31tgKA5rAftyPrZX LJChwR WV85QxGj6k W zL eqic	rwC1ysjeB P sXsWrtZ8dxZYI0QvHR3KaQxGSWcG0lPqz 0CCgE7i31tgKA5rAftyPrZX LJChwR WV85QxGj6k W zL eqic
CONFLICTING	0	0
DVC_CHANGE_LOG_ID	470001	470001
DVC_ID	1	1
MODIFIED_BY	root	root
MODIFIED_DT	Thu Oct 28 16:29:31 BST 2004	Thu Oct 28 16:29:31 BST 2004
TRAN_ID	566510001	566510001

8 Click **Mark Resolved**. The **Multimaster Tools: Mark Conflict Resolved** page appears, as Figure 21 shows. The page displays the results of marking a transaction resolved.

Figure 21 Multimaster Tools Mark Conflict Resolved Page

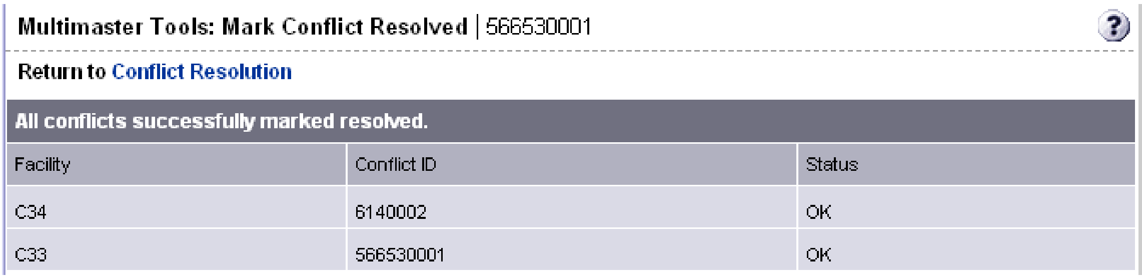

Multimaster Tools: Mark Conflict Resolved | 566530001

Return to Conflict Resolution

All conflicts successfully marked resolved.

Facility	Conflict ID	Status
C34	6140002	OK
C33	566530001	OK

After it is marked resolved, the transaction disappears from the State and Conflicts views after SA refreshes the data in the Multimaster Tools.

9 Click **Conflict Resolution** to return to the Conflict view.

Resolving a Conflict by Transaction

Perform the following steps If you know that *synchronizing all objects from one Facility* will resolve the conflict:

1 Log on to the SAS Web Client.

2 From the Navigation panel, click **Administration ➤ Multimaster Tools**. The **Multimaster Tools: State View** page appears, showing a summary of all transactions and, if they exist, all conflicts.

3 To resolve a conflict, select the **Conflict View** tab. The Multimaster Tools: Conflict View page appears, as shown in Figure 22.

Figure 22 Multimaster Tools: Conflict Page showing all Conflicts in a Multimaster Mesh

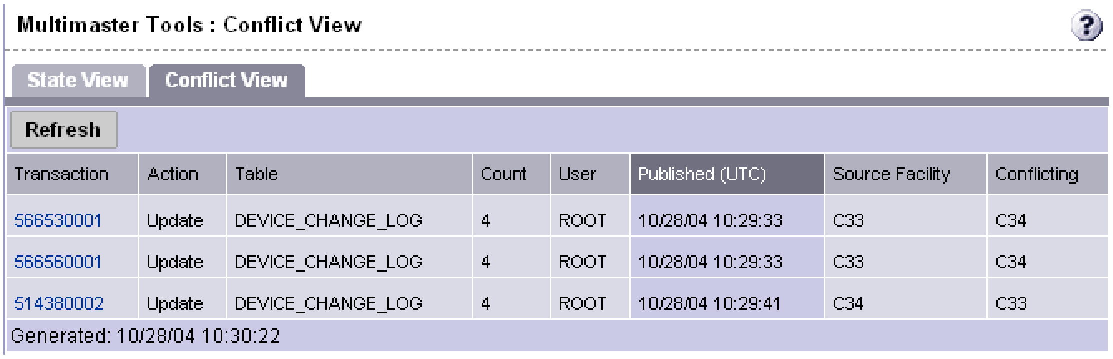

Multimaster Tools : Conflict View

Transaction	Action	Table	Count	User	Published (UTC)	Source Facility	Conflicting
566530001	Update	DEVICE_CHANGE_LOG	4	ROOT	10/28/04 10:29:33	C33	C34
566560001	Update	DEVICE_CHANGE_LOG	4	ROOT	10/28/04 10:29:33	C33	C34
514380002	Update	DEVICE_CHANGE_LOG	4	ROOT	10/28/04 10:29:41	C34	C33

Generated: 10/28/04 10:30:22

The page lists each transaction by ID number (clickable link), the actions that caused the conflict, the database objects affected by the conflict, the user responsible for the conflict (listed by the IP of the SAS Web Client where the user made the change), when the offending action occurred, the source Facility that originated the transaction, and the facilities where the transaction conflicted.

4 Click the link of the transaction you want to resolve. You now see the Multimaster Tools: Transaction Differences page, as shown in Figure 23.

Figure 23 Transaction Differences Page Showing Conflicts Between Facilities

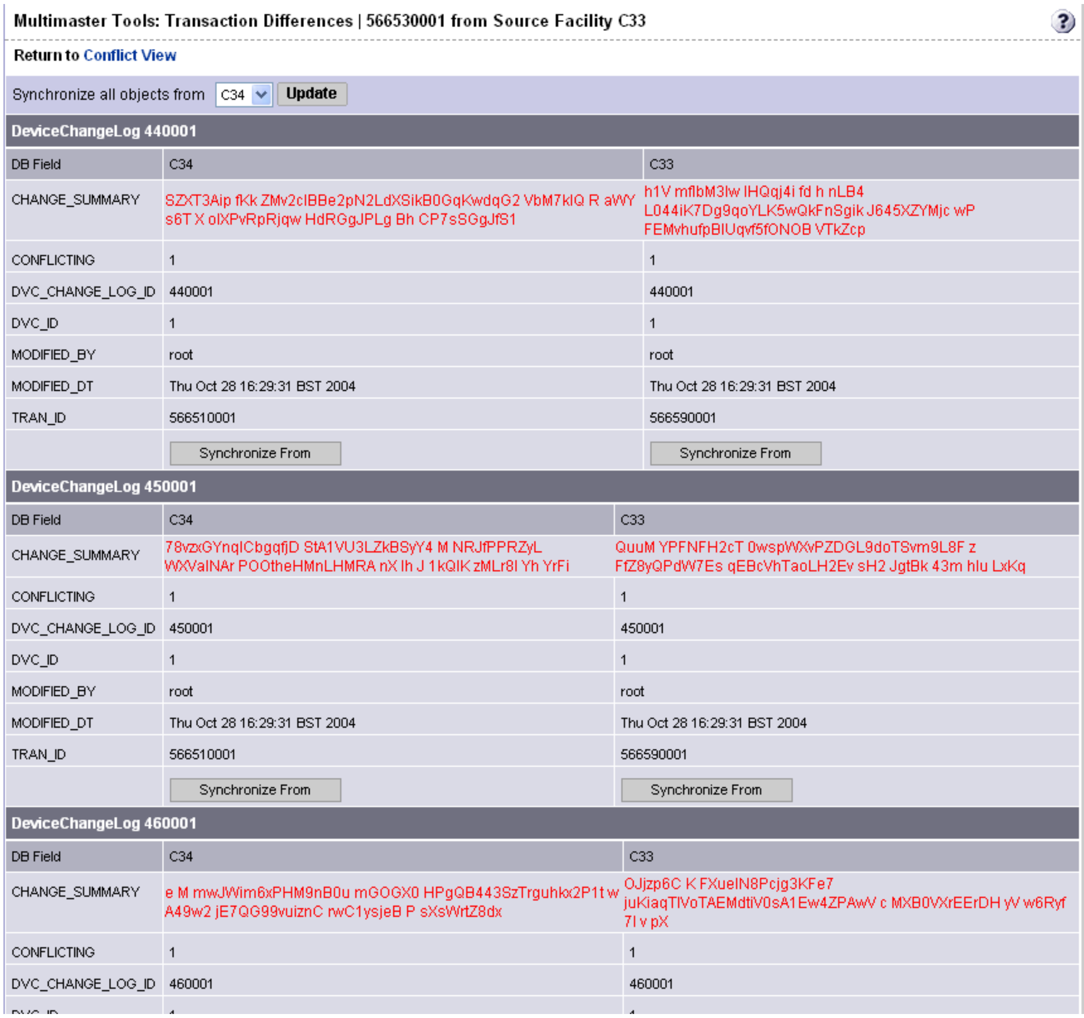

Multimaster Tools: Transaction Differences | 566530001 from Source Facility C33

Return to Conflict View

Synchronize all objects from C34

DeviceChangeLog 440001

DB Field	C34	C33
CHANGE_SUMMARY	SZXT3Aip fKk ZMv2clBBe2pN2LdXSikB0GqKwdqG2 VbM7klQ R aWY s6T X olXPvRpRjqw HdRGgJPLg Bh CP7sSGgJfS1	h1V mflbM3lw lHQqj4i fd h nLB4 L044iK7Dg9qoYLK5wQkFnSgik J645XZYMjc wP FEMvhufpBlUqvf5fONOB VTkZcp
CONFLICTING	1	1
DVC_CHANGE_LOG_ID	440001	440001
DVC_ID	1	1
MODIFIED_BY	root	root
MODIFIED_DT	Thu Oct 28 16:29:31 BST 2004	Thu Oct 28 16:29:31 BST 2004
TRAN_ID	566510001	566590001
	Synchronize From	Synchronize From

DeviceChangeLog 450001

DB Field	C34	C33
CHANGE_SUMMARY	78vzxGYnqlCbgqfjD StA1VU3LZkBSyY4 M NRJfPPRZyL WXValNAr POOtheHMnLHMRA nX Ih J 1kQIK zMLr8l Yh YrFi	QuuM YPFNFH2cT 0wspWXvPZDGL9doTSvm9L8F z FfZ8yQPdW7Es qEBcVhTaoLH2Ev sH2 JgtBk 43m hlu LxKq
CONFLICTING	1	1
DVC_CHANGE_LOG_ID	450001	450001
DVC_ID	1	1
MODIFIED_BY	root	root
MODIFIED_DT	Thu Oct 28 16:29:31 BST 2004	Thu Oct 28 16:29:31 BST 2004
TRAN_ID	566510001	566590001
	Synchronize From	Synchronize From

DeviceChangeLog 460001

DB Field	C34	C33
CHANGE_SUMMARY	e M mwJWim6xPHM9nB0u mGOGX0 HPgQB443SzTrguhkx2P1t w A49w2 jE7QG99vuiznC rwC1ysjeB P sXsWrtZ8dx	OJjzp6C K FXuelN8Pcjg3KFe7 juKiaqTlVoTAEMdtiV0sA1Ew4ZPAwV c MXB0VXrEErDH yV w6Ryf 7l v pX
CONFLICTING	1	1
DVC_CHANGE_LOG_ID	460001	460001
DVC_ID	1	1

5 From the **Synchronize all objects from** drop-down list at the top of the page, select the Facility to use as the correct source of data, as Figure 24 shows.

Figure 24 By Transaction

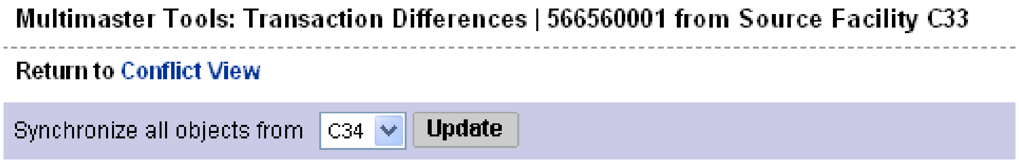

Multimaster Tools: Transaction Differences | 566560001 from Source Facility C33

Return to Conflict View

Synchronize all objects from C34

See Best Practices for Resolving Database Conflicts on page 96 in this chapter for more information

6 Click **Update** beside the drop-down list. The Multimaster Tools: Transaction Synchronization Results page appears, as shown in Figure 25.

Figure 25 Transaction Synchronization Results for All Objects in Transaction

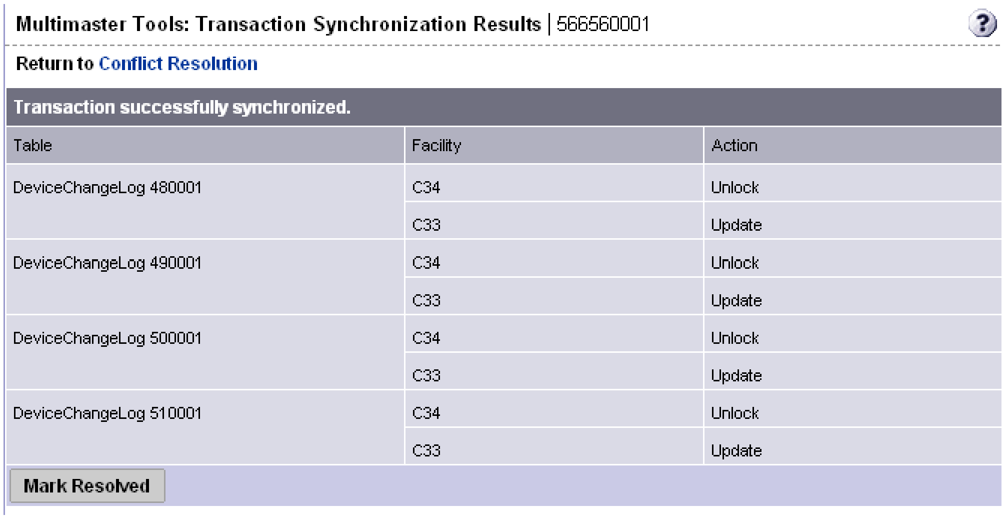

Multimaster Tools: Transaction Synchronization Results | 566560001

Return to Conflict Resolution

Transaction successfully synchronized.

Table	Facility	Action
DeviceChangeLog 480001	C34	Unlock
	C33	Update
DeviceChangeLog 490001	C34	Unlock
	C33	Update
DeviceChangeLog 500001	C34	Unlock
	C33	Update
DeviceChangeLog 510001	C34	Unlock
	C33	Update

This page shows the results of the synchronization and prompts you to mark the conflicts resolved.

7 Click **Mark Resolved**. The **Multimaster Tools: Mark Conflict Resolved** page appears. The page displays the results of marking a transaction resolved.

8 Click the link to return to the Conflict view. After it is marked resolved, the transaction disappears form the State and Conflicts views after SA refreshes the data in the Multimaster Tools.

Network Administration for 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
- SAS Web Clients cannot contact the multimaster central Data Access Engine

Production experience for multimaster configurations supports the performance data that Table 18 shows.

Table 18 Performance Data for Multimaster Configurations

# Facilities	Duration Network Outage	# multimaster conflicts *
8 facilities (SA core installed in each Facility)	12 hour outage (1 Facility looses 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 SAS Web Clients in other Facilities increases the number of conflicts.		

Network connectivity issues include SA Bus or multicast routing problems.

Multimaster Alert Emails

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 SA Support Representative or See SA Notification Configuration on page 193 in Chapter 7 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 SA Support Representative for assistance troubleshooting and resolving SA problems that affect the multimaster operation.

See Table 19 for error messages.

Table 19 Multimaster Error Messages

Subject Line	Type of Error	Details
`vault.ApplyTransactionError`	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.
`vault.configValueMissing`	SA Problem	No value was specified for a given configuration parameter. Log into the SAS Web Client and provide the value for this configuration parameter. Contact your SA Support Representative for assistance setting SA configuration values.

Table 19 Multimaster Error Messages (cont'd)

Subject Line	Type of Error	Details
`vault.DatabaseError`	Multimaster Transaction Conflict	An error occurred while querying the database for updates to send to other databases or while applying updates from other databases. Restart the Model Repository Multimaster Component.
`vault.InitializationError`	SA Problem	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. Resolve the error condition. Restart the Model Repository Multimaster Component.
`vault.ParserError`	Multimaster Transaction Conflict	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. 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.
`vault.SOAPError`	Multimaster Transaction Conflict	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. Run the SA Admin Multimaster Tools and verify that the transaction data does not contain special characters SOAP might be unable to interpret.

Table 19 Multimaster Error Messages (cont'd)

Subject Line	Type of Error	Details
`vault.UnknownError`	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.

4 Satellite Administration

This section describes basic SA Satellite topologies and concepts and the following adminstrative tasks:

- Viewing Satellite Facilities
- Enabling the Display of Realm Information
- Viewing the Realm of a Satellite Managed Server
- Viewing and Managing Satellite Gateway Information
- Satellite Software Repository Cache Management
- Updating Software in the Satellite Software Repository Cache
- Satellite Software Repository Cache Management

Overview of the SA Satellite

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.

➤ 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 Planning and Installation Guide*.

Satellites can be installed using various topologies. For detailed information about Satellite topologies, see the *SA Planning and Installation Guide*.

Figure 26, shows a Satellite linked to a single SA Core communicating through the First Core's Management Gateway.

Figure 27, shows two Satellites linked to an SA Core each communication directly with the First Core's Management Gateway. Communication between the Satellites, when required, travels from one Satellite to the First Core Management Gateway, then to the other Satellite.

Figure 26 Single SA Core with a Single Satellite

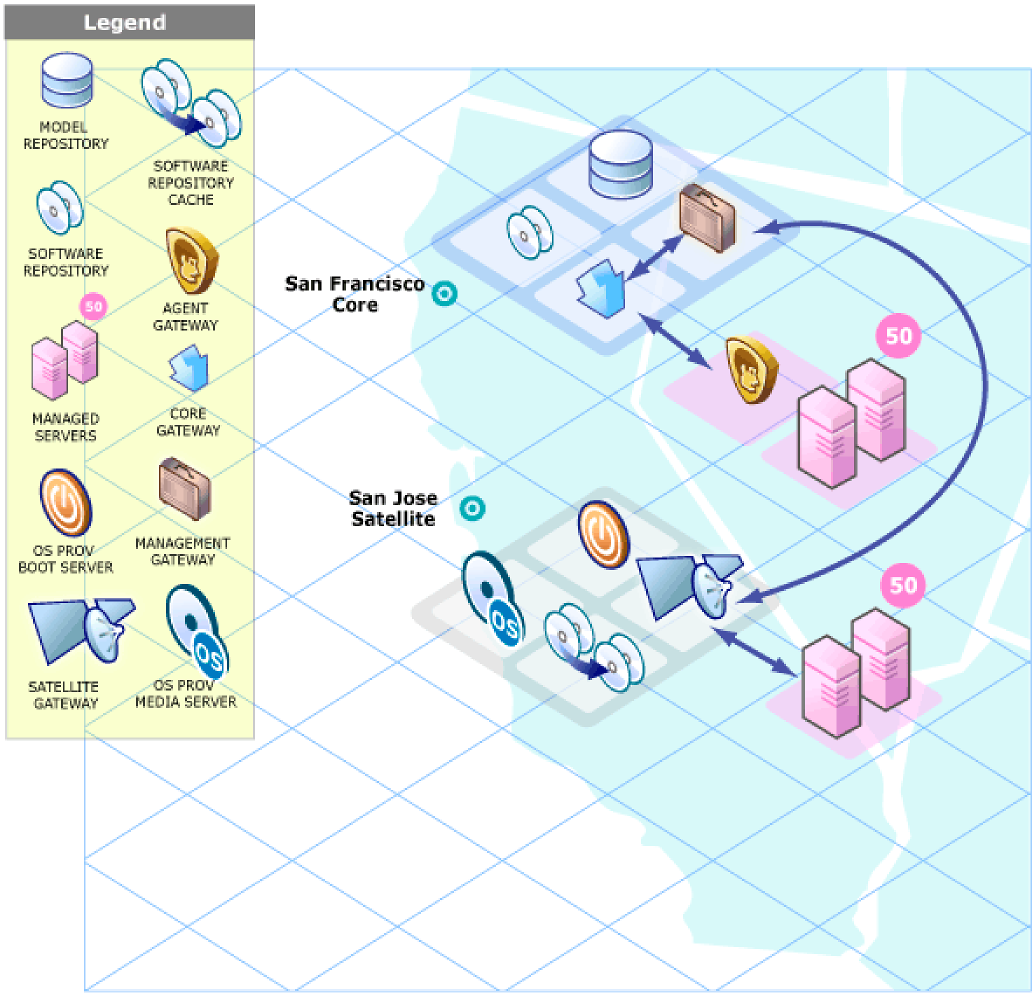

Figure 27 Single SA Core with Multiple Satellites

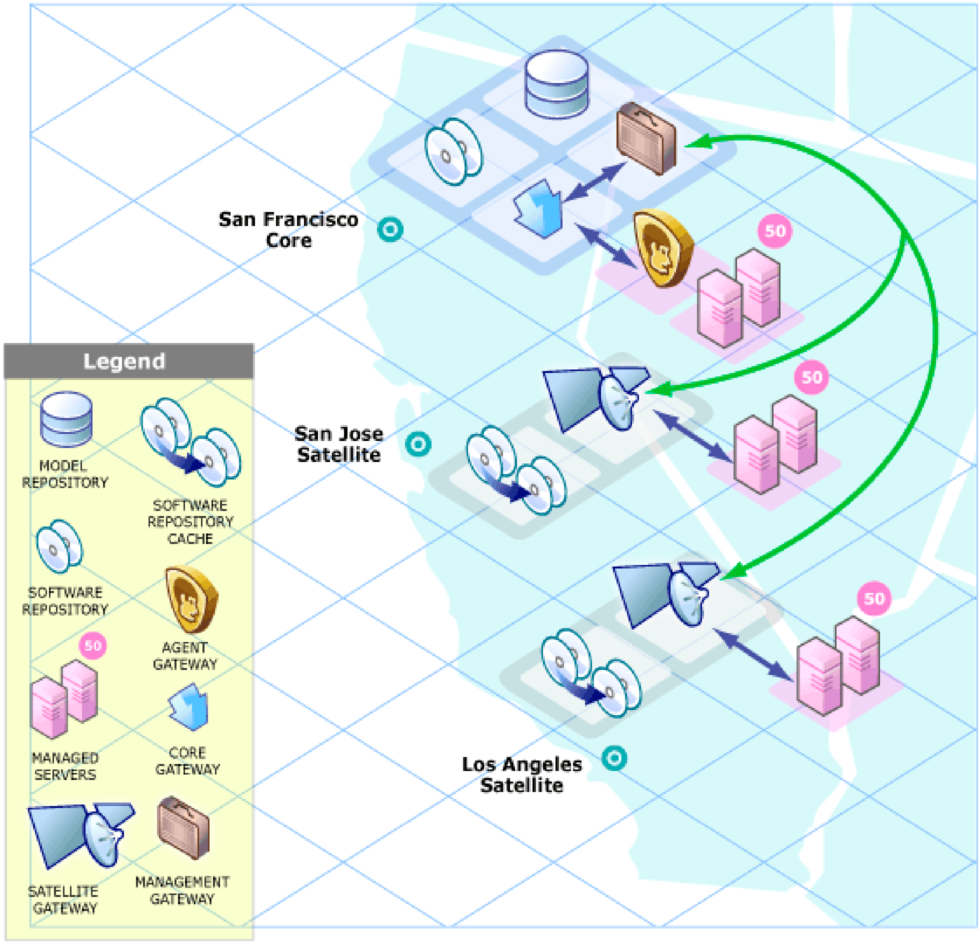

Management Gateway

Satellite communication with an SA First Core is achieved, either directly or through a network of Gateways, through a *Management Gateway* that resides on the First Core and in the same IP address space as the servers that the First Core manages. The Management Gateway communicates with the First Core through the Core Gateway.

Facilities and Realms

A *Facility* encompasses the Managed Servers that reside in a single physical location. A Facility can be all or part of a data center, server room, or computer lab.

To deal with the potential problem of Facilities in the same SA Core with overlapping IP address spaces, SA uses the concept of *Realms*. An SA Realm is a *routable IP address space*, which is serviced by one or more Gateways. A Facility can contain multiple *Realms*. Each IP address space requires a separate Realm. Typically, each physical building is modeled as a Facility that has as many Realms as needed.

The Realm allows each Managed Server in a Facility to be identified by its *Realm and IP address combination*. Since separate Facilities can contain duplicate IP addresses, this Realm/IP address combination allows SA to differentiate between Managed Servers in different Facilities but with the same IP address and route traffic accordingly.

For more information about Facilities and Realms, see the *SA Simple/Advanced Installation Guide*.

Satellite Information and Access

This section discusses the following topics:

- Permissions Required for Managing Satellites
- Viewing Satellite Facilities
- Enabling the Display of Realm Information
- Viewing the Realm of a Satellite Managed Server
- Viewing and Managing Satellite Gateway Information

Permissions Required for Managing Satellites

To access the *Manage Gateway* feature, 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 Chapter 1, User and User Group Setup and Security, on page 13 of this guide.

Viewing Satellite Facilities

The **Facilities** page in the SAS Web Client lists the Core and Satellite facilities. In particular, the Facilities page displays **Unreachable Facilities**, as shown in Figure 28.

Figure 28 Facilities Channel

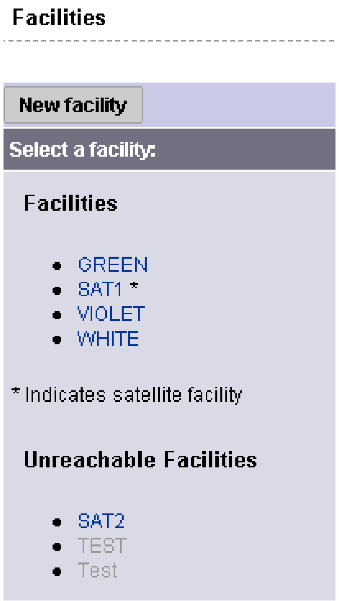

Facilities

Select a facility:

Facilities

- GREEN
- SAT1 *
- VIOLET
- WHITE

* Indicates satellite facility

Unreachable Facilities

- SAT2
- TEST
- Test

Clicking the link for a Facility, and then selecting the **Realms** tab displays the configured bandwidth of the connections between the Realms in that Facility, as shown in Figure 29.

Figure 29 Realms and Connection Bandwidth in Facilities

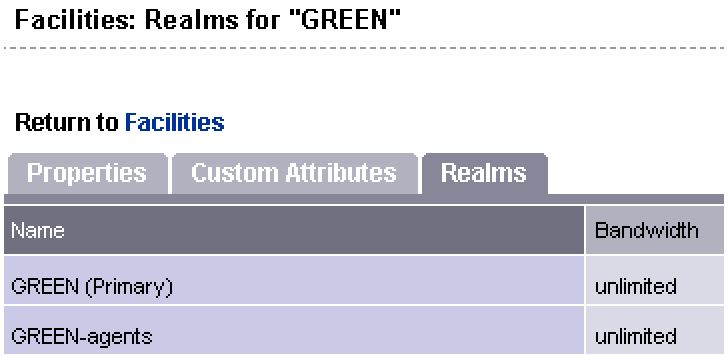

Facilities: Realms for "GREEN"

Return to Facilities

Name	Bandwidth
GREEN (Primary)	unlimited
GREEN-agents	unlimited

Additionally, you can view the Facilities that contain Realms by clicking **Administration ➤ System Configuration** as shown in Figure 30.

Figure 30 Satellite Configuration

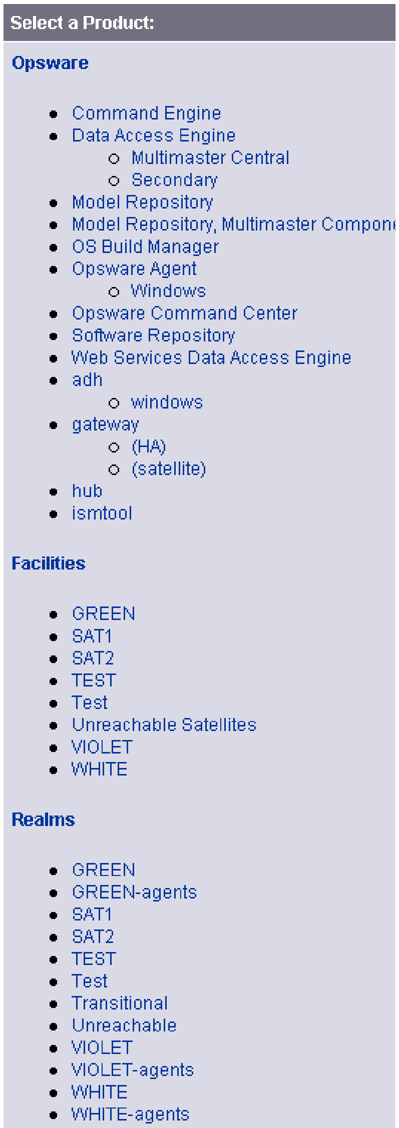

Select a Product:

Opsware

- Command Engine
- Data Access Engine
  - Multimaster Central
  - Secondary
- Model Repository
- Model Repository, Multimaster Compone
- OS Build Manager
- Opsware Agent
  - Windows
- Opsware Command Center
- Software Repository
- Web Services Data Access Engine
- adh
  - windows
- gateway
  - (HA)
  - (satellite)
- hub
- ismtool

Facilities

- GREEN
- SAT1
- SAT2
- TEST
- Test
- Unreachable Satellites
- VIOLET
- WHITE

Realms

- GREEN
- GREEN-agents
- SAT1
- SAT2
- TEST
- Test
- Transitional
- Unreachable
- VIOLET
- VIOLET-agents
- WHITE
- WHITE-agents

Enabling the Display of Realm Information

By default, the SAS Web Client does not display Realm information, which is needed by users who manage Gateways and Software Repository Caches.

To enable access to the Realm information, perform the following tasks:

1 Log on to the SAS Web Client as a user who is a member of the SA *Administrators group* and to a group that has the *Configure Opsware permission*.

2 From the Navigation panel, click **Administration ➤ System Configuration**.

3 Select the **SA System Web Client** link.

4 On the **System Configuration** page, for the parameter `owm.features.Realms.allow`, set the value to `true`.

5 Click **Save**.

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 SAS Web Client displays additional information that identifies the server corresponding to the IP address, as shown in Figure 31.

Figure 31 Server Properties Page Showing the Realm of a Managed Server

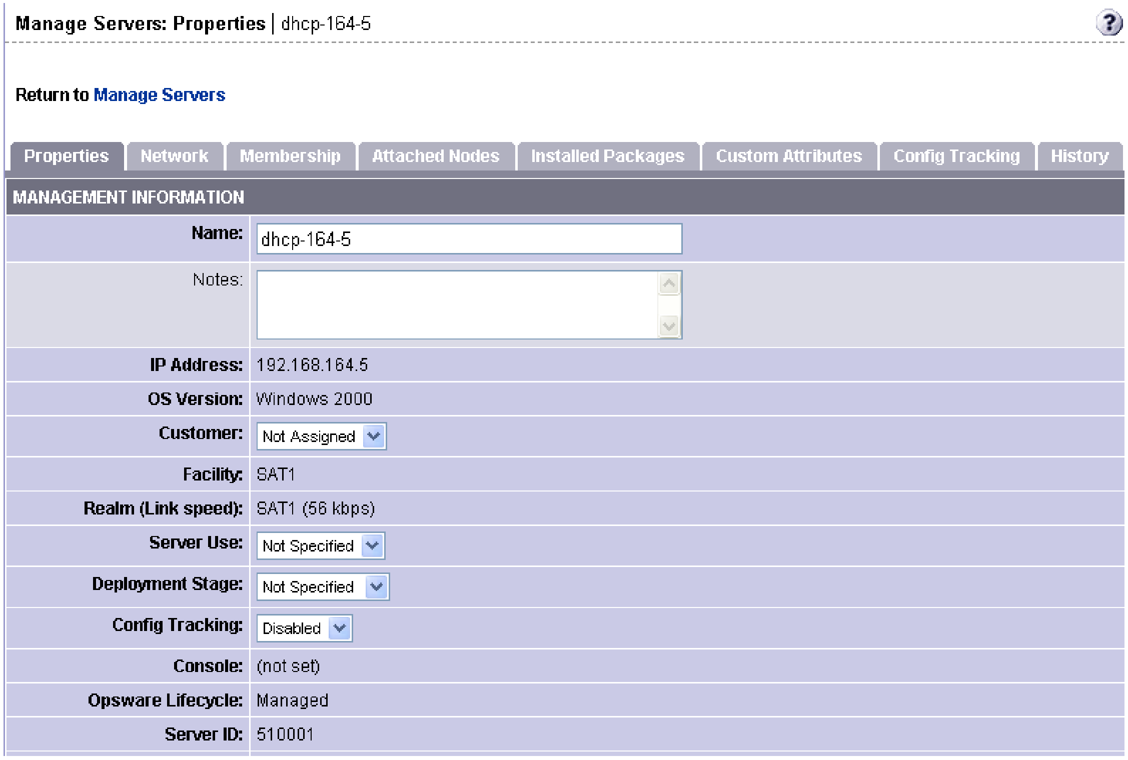

Manage Servers: Properties | dhcp-164-5

Return to Manage Servers

MANAGEMENT INFORMATION

Name:	dhcp-164-5
Notes:	
IP Address:	192.168.164.5
OS Version:	Windows 2000
Customer:	Not Assigned
Facility:	SAT1
Realm (Link speed):	SAT1 (56 kbps)
Server Use:	Not Specified
Deployment Stage:	Not Specified
Config Tracking:	Disabled
Console:	(not set)
Opsware Lifecycle:	Managed
Server ID:	510001

Viewing and Managing Satellite Gateway Information

To access Satellite Gateway information, in the SAS Web Client Navigation panel, click **Administration ➤ Gateway**. The **Manage Gateway** page appears, as shown in Figure 32. From the list of Gateways on the left, select the Gateway you want to view information for, and then click the Page Selection link for the page you want to view.

Figure 32 Status Page of the Manage Gateway Feature

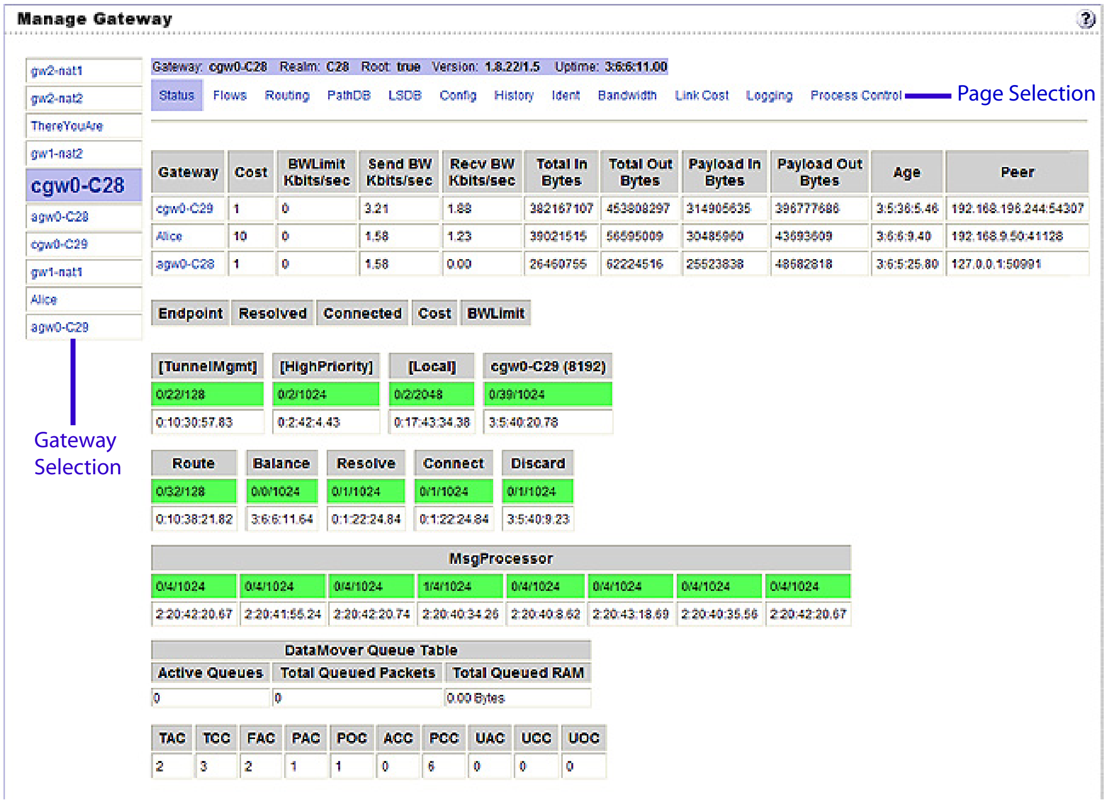

Page Selection

Gateway Selection

Gateway: cgw0-C28 Realm: C28 Root: true Version: 1.8.22/1.5 Uptime: 3:6:6:11.00

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	382167107	453808297	314905635	396777686	3:5:36:5.46	192.168.196.244:54307
Alice	10	0	1.58	1.23	39021515	56595009	30485960	43693609	3:6:6:9.40	192.168.9.50:41128
agw0-C28	1	0	1.58	0.00	26460755	62224516	25523838	48682818	3:6:5:25.80	127.0.0.1:50991

Endpoint	Resolved	Connected	Cost	BWLimit

[TunnelMgmt]	[HighPriority]	[Local]	cgw0-C29 (8192)
0/22/128	0/2/1024	0/2/2048	0/39/1024
0:10:30:57.83	0:2:42:4.43	0:17:43:34.38	3:5:40:20.78

Route	Balance	Resolve	Connect	Discard
0/32/128	0/0/1024	0/1/1024	0/1/1024	0/1/1024
0:10:38:21.82	3:6:6:11.64	0:1:22:24.84	0:1:22:24.84	3:5:40:9.23

MsgProcessor							
0/4/1024	0/4/1024	0/4/1024	1/4/1024	0/4/1024	0/4/1024	0/4/1024	0/4/1024
2:20:42:20.67	2:20:41:55.24	2:20:42:20.74	2:20:40:34.26	2:20:40:8.62	2:20:43:18.69	2:20:40:35.56	2:20:42:20.67

DataMover Queue Table		
Active Queues	Total Queued Packets	Total Queued RAM
0	0	0.00 Bytes

TAC	TCC	FAC	PAC	POC	ACC	PCC	UAC	UCC	UOC
2	3	2	1	1	0	6	0	0	0

Use the Manage Gateway page for the following tasks:

- Obtain debugging and status information about Gateways and the tunnels between Gateways.
- Perform specific tasks on Gateways, such as changing the bandwidth limits or tunnel cost between Gateway instances, restarting Gateway processes, or changing the logging levels for Gateway processes.

Viewing Diagnostic and Debugging Information

1 From the SAS Web Client Navigation panel, click **Administration ➤ Gateway**. The **Manage Gateway** page appears.

2 From the list of Gateways on the left, select the Gateway that you want to view information for. The **Status** page for that Gateway appears.

The Status page displays the following information for the Gateway:

- A table of Active Tunnels. This table includes:
  - — 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 timestamp indicating when the message high-water mark was last reached to troubleshoot Gateway issues. The timestamp is displayed in the format `DD:HH:mm:ss`.

3 To view the details and statistics for a tunnel between Gateways, click the link for the Gateway that *terminates* the tunnel, as Figure 33 shows.

Figure 33 Manage Gateway — Status Page

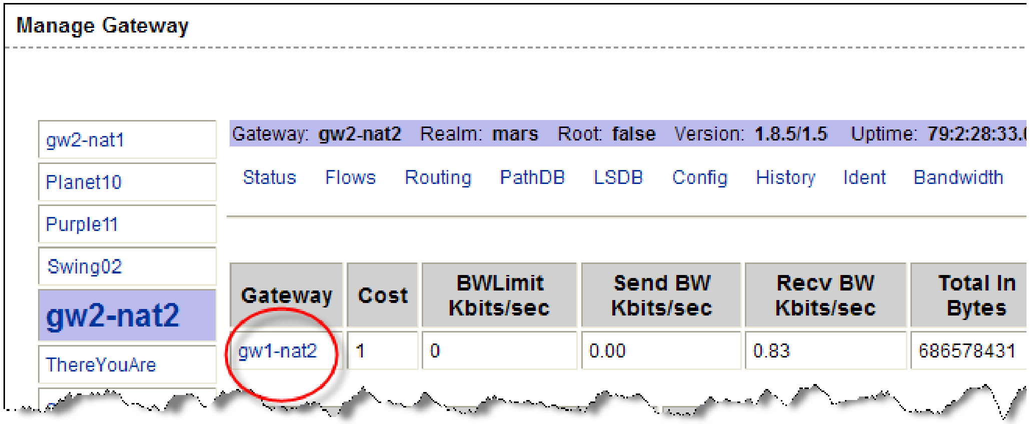

Manage Gateway

Gateway: **gw2-nat2** Realm: **mars** Root: **false** Version: **1.8.5/1.5** Uptime: **79:2:28:33.**

Gateway	Cost	BWLimit Kbits/sec	Send BW Kbits/sec	Recv BW Kbits/sec	Total In Bytes
gw1-nat2	1	0	0.00	0.83	686578431

The page refreshes and displays the tunnel details and statistics.

4 To view the following pages containing diagnostic information, click the link for the page in the menu bar.

- **Flows page**: Displays information about all open connections for the selected Gateway.
- **Routing page**: 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.

When a tunnel collapses, by default, routing information is retained in the routing table for two minutes to provide continuity for the mesh.

- **Path database (PathDB) page**: 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*.
- **Link State database (LSDB) page**: 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.
- **Configuration (Config) page**: Displays the *Gateway Properties file* for the Gateway you have selected. This page includes 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.

Above the properties values, the **Properties Cache** field appears. 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** page 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. From the SAS Web Client Navigation panel, click **Administration ➤ Gateway**. The **Manage Gateway** page appears.
2. From the Page Selector, click **Ident**. The page refreshes with an interface to the real-time connection identification database.
3. In the text field, enter the *protocol* and *source port* for an active connection (for example, TCP:25679).
4. Click **Lookup**.

   The page refreshes with the client Realm and client IP address — where the connection came from.

Changing the Bandwidth Usage or Link Cost Between Gateways

1. From the SAS Web Client Navigation panel, click **Administration ➤ Gateway**. The **Manage Gateway** page appears.
2. To specify a bandwidth limit for a connection:
   a. From the Page Selector, click `Bandwidth`. The page refreshes with fields in which you can specify the bandwidth for the connection between Gateway instances.
   b. Specify two Gateway instance names that are connected by a tunnel.
   c. Specify the bandwidth limit you want in kilobits per second (Kbps). Specify zero (0) to remove bandwidth constraints for the connection.
   d. Click **Apply**.
3. To set a link cost for a connection:
   a. From the Page Selector, click **Link Cost**. The page refreshes with fields in which you can specify the link cost for the connection between Gateway instances.
   b. Specify two Gateway instance names that are connected by a tunnel.
   c. Specify the cost you want in the **Cost** field.
   d. Click **Apply**.

Viewing the Gateway Log or Change the Log Level

➤ 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. From the SAS Web Client Navigation panel, click **Administration ➤ Gateway**. The **Manage Gateway** page appears.
2. From the Page Selector, click **Logging**. The page refreshes with a `tail` of the Gateway log file.
3. To change the logging level, select an option: `LOG_INFO`, `LOG_DEBUG`, or `LOG_TRACE`.
4. Click **Submit**.

Restarting or Stopping a Gateway Process

1. From the SAS Web Client Navigation panel, click **Administration ➤ Gateway**. The **Manage Gateway** page appears.
2. From the Page Selector, click **Process Control**. The page refreshes.
3. To restart the Gateway process, click **Restart**.
4. To stop the Gateway watchdog and the Gateway, click **Shutdown**.

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. Additionally, the Manage Gateway UI is unavailable after stopping the process.

To restart the Gateway after stopping it from the Manage Gateway page, you must log onto the server running the Gateway component and manually restart the process.

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.

Since 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 Repostory Cache and, optionally, your OS Provisioning Media and Boot servers.

Availability of Satellite Software Repository Cache Content

Software Repository content is not automatically replicated to the Satellite Software Repository cache, therefore, not all content is available locally for Satellites in a mesh. You must manually update the Satellite's Software Repository Cache with the software you want to install locally. On-demand updates are available only when the caching policy for the Realm of the Software Repository Cache is `on-demand`.

SA can only warn you that the requested software is not available locally and that you must update content from the First Core Software Repository or another Satellite Software Repository Cache. SA keeps track of whether a package is available locally.

Instead, when SA is attempting to remediate requested software that is not available locally onto a managed server, the SAS Web 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

➤ The SAS Web 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 on page 128 in this chapter for more information.

The flowchart in Figure 34 illustrates the logic that the Software Repository Cache uses to update packages in a Satellite.

Figure 34 Software Repository Cache Update Logic

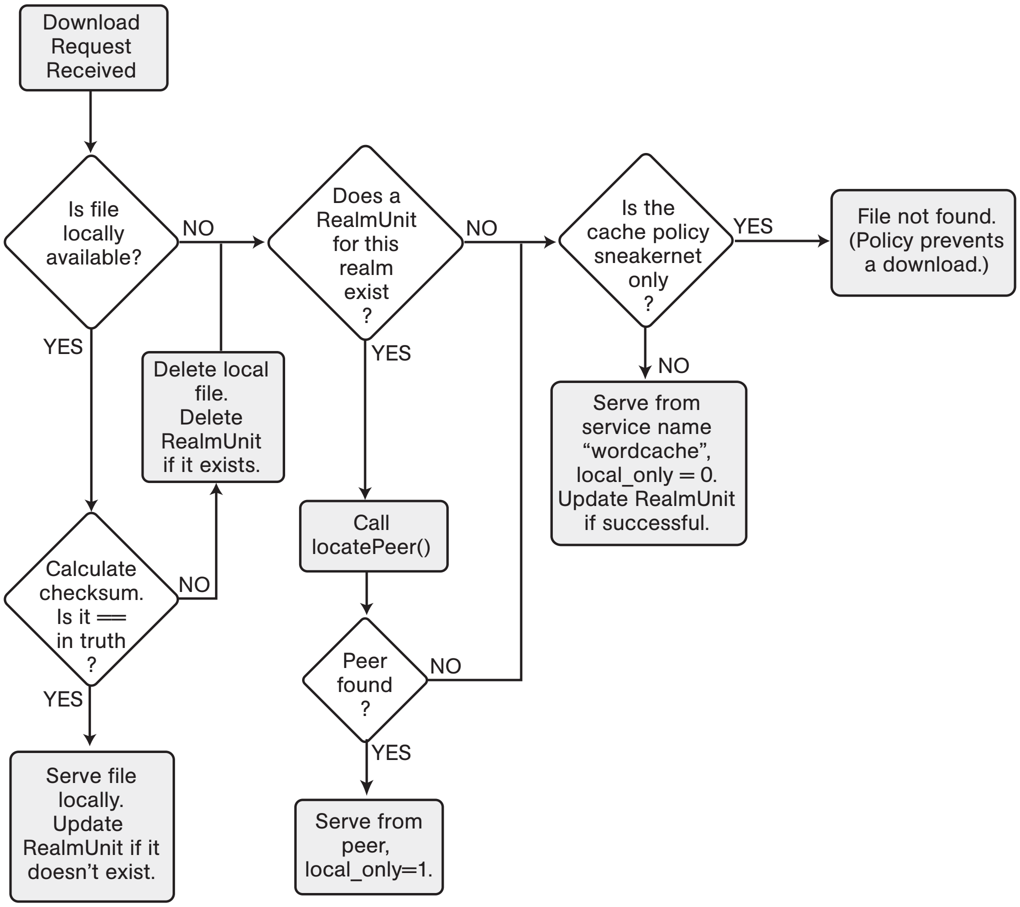
Download Request Received
Is file locally available?
NO
YES
Delete local file. Delete RealmUnit if it exists.
Calculate checksum. Is it == in truth ?
NO
YES
Serve file locally. Update RealmUnit if it doesn't exist.
Does a RealmUnit for this realm exist ?
NO
YES
Call locatePeer()
Peer found ?
NO
YES
Serve from peer, local_only=1.
Is the cache policy sneakernet only ?
YES
NO
File not found. (Policy prevents a download.)
Serve from service name "wordcache", local_only = 0. Update RealmUnit if successful.

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 From the SAS Web Client Navigation panel, click **Administration ➤ System Configuration**. The **Select a Product** page appears.

2 Click the link of the Realm for which you want to set the Software Repository Cache update policy. The configuration values for that Facility appear.

3 For the parameter `word.caching_policy`, set the caching policy value by selecting the **Use default value** option or the **Use value** option and enter JIT (On-Demand Update) or `SNEAKERNET` (Manual Update), See Figure 35.

Figure 35 Software Repository Cache Configuration Parameters

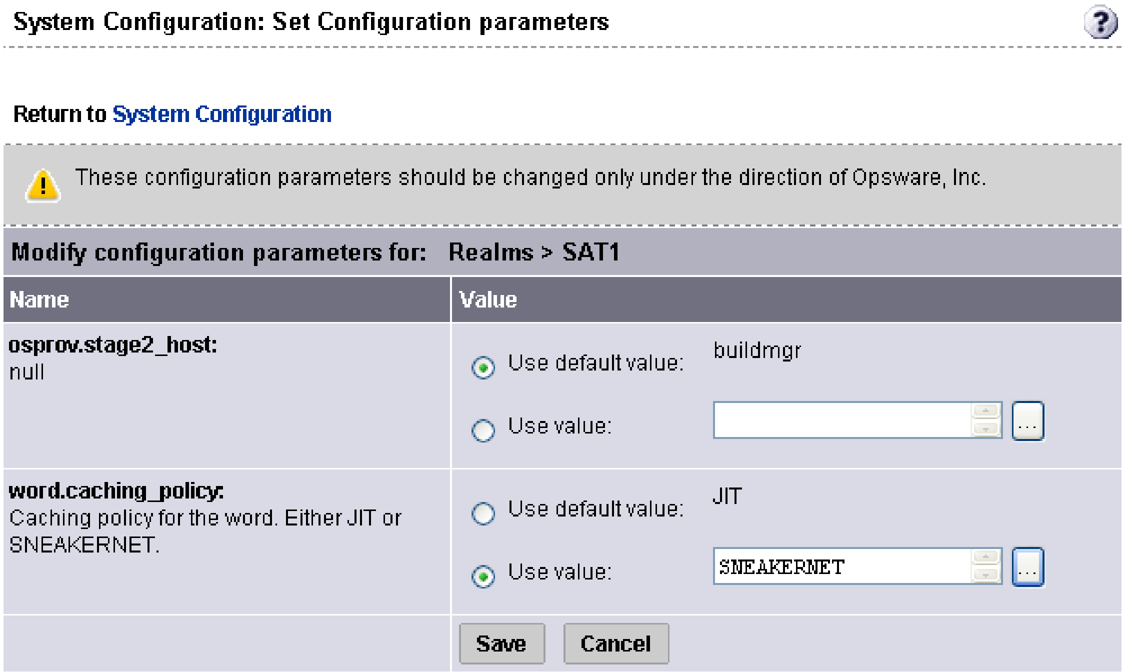

System Configuration: Set Configuration parameters

Return to System Configuration

These configuration parameters should be changed only under the direction of Opsware, Inc.

Modify configuration parameters for: Realms > SAT1

Name	Value
osprov.stage2_host: null	● Use default value: buildmgr ○ Use value:
word.caching_policy: Caching policy for the word. Either JIT or SNEAKERNET.	○ Use default value: JIT ● Use value: SNEAKERNET
	Save Cancel

4 Click **Save** to apply your configuration change. Since, by default, the Software Repository Cache polls for configuration changes every five minutes, it make take up to five minutes for your change to take affect.

On-demand Updates

Enabling On-Demand Updates allows software to be downloaded to the Satellite Software Rpository Cache when it is not yet locally available as soon as that software is requested. 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 on page 125.

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 Repositorye 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.

➤ 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, you 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 on page 126 and Staging Files to a Software Repository Cache on page 128.

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.

➤ 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.

☑ 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 on page 128.

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 DCML Exchange Tool (DET). Using the DET, you 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 SAS Web 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 SAS Web Client. If you want to export packages from some other category of software in the SAS Web 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, you 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 Cach 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 only 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.

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 utilities, including the Microsoft Patch Database (`mssecure.cab`), the Microsoft Baseline Security Analyzer (`mbsacli.exe`), or the Windows `chain.exe` utility to the Software Repository, 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 on page 128 in this chapter for information about how to stage files.

5 SA Maintenance

While maintaining SA, you might encounter the following types of problems:

- **Operational problems**: processes failing or becoming unresponsive (Data Access Engine, Command Engine, 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 SAS Web 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 there are *unreachable servers* existing in the managed environment communication is disrupted.

➤ Many problems with the Code Deployment & Rollback (CDR) feature are caused by errors with the CDR configuration and setup. See the *SA User's Guide: Server Automation* for information about CDR configuration.

☑ When using the System Diagnosis function in an environment with multiple Facilities, System Diagnosis can only be run on one Facility at a time.

SA Diagnosis

The *System Diagnosis tool* allows you to check the functionality of the SA Core Components and the ability of servers running in the managed environment to interact with the SA Core.

You can troubleshoot most of the errors that occur within the SA Core with the SA Diagnosis tool.

This section provides information about how to diagnose SA problems and contains the following topics:

- SA Component Troubleshooting
- System Diagnosis Testing Process
- Data Access Engine Tests
- Software Repository Tests
- Web Services Data Access Tests
- Command Engine Tests
- Model Repository Multimaster Component Tests
- Running an SA Core Component System Diagnosis

SA Component Troubleshooting

The following mechanisms for troubleshooting SA are available:

- Running the SA Diagnosis tool (a tool for debugging common problems with SA Core Components). Running an SA Core Component System Diagnosis on page 136.
- Reviewing error logs for components.See SA Components Logs on page 147.

System Diagnosis Testing Process

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 Component.

  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:

- Server Agent Standalone Tests
- Server Agent Comprehensive Tests
- 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:

- Data Access Engine
- Software Repository (and Word Store)
- Web Services Data Access Engine
- Command Engine
- Server Agents on SA Core servers
- Model Repository Multimaster Component
- OS Build Manager

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 Planning and 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.

Running an SA Core Component System Diagnosis

To access the System Diagnosis tool, your user must belong to the Administrators group. The SAS Web Client has access to all the Server Agents running on the SA Core Component servers.

Perform the following taskss to run a system diagnosis of the SA Core Components:

1 Log on to the SAS Web Client.

2 From the Navigation panel, click **Administration ➤ System Diagnosis**. The **System Diagnosis: Begin Diagnosis** page appears.

3 Select the components that you want to test. By default, all components are selected (the Data Access Engine, the Software Repository, Command Engine, and Web Services Data Access Engine; in multiple core environments, there is also a selection for the Model Repository Multimaster Component). See Figure 36.

Figure 36 System Diagnosis Page Showing SA Components Selected for Testing

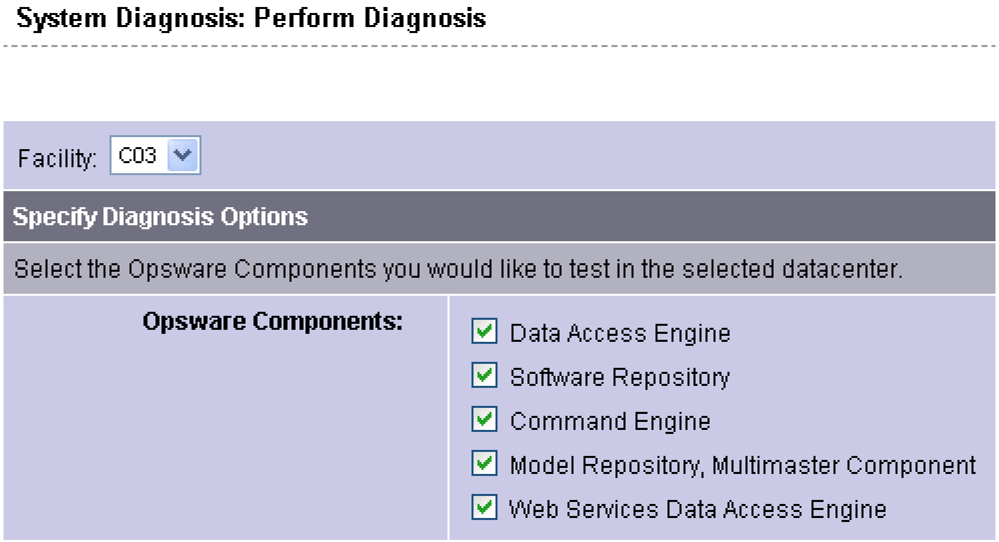

System Diagnosis: Perform Diagnosis

Facility: C03

Specify Diagnosis Options

Select the Opsware Components you would like to test in the selected datacenter.

Opsware Components:

☑ Data Access Engine
☑ Software Repository
☑ Command Engine
☑ Model Repository, Multimaster Component
☑ Web Services Data Access Engine

4 Click **Run Diagnosis**.

The **System Diagnosis: Performing Diagnosis** window appears, which displays a progress bar while the tests are running, as Figure 37 shows.

Figure 37 System Diagnosis Progress Bar

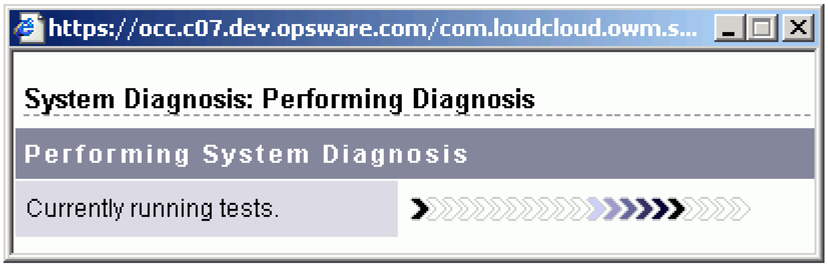

https://occ.c07.dev.opsware.com/com.loudcloud.owm.s...

System Diagnosis: Performing Diagnosis

Performing System Diagnosis

Currently running tests.

When all the tests are complete, the window closes and the **System Diagnosis: Failed Tests** page appears in the main SAS Web Client window. If all tests passed, the **System Diagnosis: Successful Tests** page appears.

5 To review the results of a test, click the Test Name link in the **Test** column. The **System Diagnosis: Test Information** page appears. If the test contained an error, error information appears at the bottom of the page.

The SA Core Health Check Monitor (HCM)

The Health Check Monitor (HCM) includes a suite of tests to check the status of an SA Core. HCM requires SA version 6.0.2 or later. The scripts that comprise the HCM are installed by the SA Installer. There is some functional overlap between HCM and the System Diagnosis Tool described in SA Diagnosis on page 132.

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. You 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`.

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>]...]
```

Validating Core Components by 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 20 describes the HCM command-line arguments. For a description of the `opsware-sas` options for starting and stopping a core, see Table 25.

Table 20 Options for the HCM Local Test Script

option	description
`mode`	The set of tests to run. The `mode` can be one of the following strings: • `status`: 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. • `verify_post`: Same as `status`. • `verify_pre`: Runs tests that validate the conditions necessary for the specified components to operate. • `verify_functionality`: Runs tests that are similar to the tests run by the `status` mode; however, they might take longer to run. Therefore, you might choose to skip these tests to save time. • `health`: Runs the tests of the `status`, `verify_pre`, and `verify_functionality` modes and provides an overview of the overall state of the specified components.
`component`	The internal name of the core component (see note below). If this option is not specified, then all components are validated. To view the internal names of the components installed on the local server, enter the following command: `/etc/init.d/opsware-sas list`
`name=value`	Options that control how the tests are run. Allowed values: • `terse=[true|false]`: If `true`, summarizes the results of all successful tests for each component in a single SUCCESS message; however, the results of failed tests are displayed individually. By default, this option is set to `false`. (This option is passed to the individual tests.) • `parsable=[true|false]`: If `true`, summarizes the results from all tests for each component with a single SUCCESS or FAILURE message. By default, this option is set to `false`. (This option is passed to the individual tests.) • `verify_filter=<regex>`: Runs only the tests whose file names match the regular expression you enter. For example, specifying `verify_filter="OPSW"` runs only tests with file names that contain the string `OPSW`, such as `100_OPSWcheck_host_spin.sh`. By default, this option is not defined. (This option is not passed to the individual tests.) 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, `verify_filter` uses the file name of the file it is pointing to for comparisons.

➤ You can find a list of the internal name used for certain Core Components and their standard names in Internal and External Component Names on page 159.

Overview of HCM Global Tests

A *global* HCM test checks an entire SA Core. You run these tests by executing the `run_all_probes.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 multi-server core, the global tests connect to the other core servers using SSH. All connections are made as `root`. Authentication is performed by specifying the `root` password or the key file on the command line. If both are specified, then the `root` password is used. One of these authentication methods must be specified unless the server is the local host.

Validating a Core By 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 tablespace 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="<system>[:<password>] [<system>[:<password>]]..."
[keyfile=<keyfiletype>:<keyfile>[:<passphrase>]]
```

Table 21 describes the options for this syntax.

Table 21 Options for the HCM Global Test Script

option	description
`list`	Lists the available tests.
`run`	Runs the specified tests.

Table 21 Options for the HCM Global Test Script (cont'd)

option	description
test	The name of the test to run. If no tests are specified, all tests are run. When shipped, the script includes the following tests: • check_opsware_services: Runs the local tests on all specified servers by running the following command remotely on each core server: /etc/init.d/opsware-sas health • check_MM_state: For a multimaster source core, checks the multimaster state of the core. • check_time: In a multi-server core, verifies that the system clocks are in sync across core servers. • check_opsware_version: Validates that the versions of all the components in the core are the same version. • check_database_tables: Validates that the Model Repository tablespace usage is within acceptable limits. For more information on tablespaces, see “Oracle Setup for Model Repository” in the *SA Planning and Installation Guide*. • check_OS_resources: Validates whether the virtual memory and disk space on SA partitions is within acceptable thresholds.
system:password	Specifies a remote core server (host name or IP address) and optional root password for the server.
keyfiletype	Specifies the type of key file to use. Allowed values: • rsa_key_file • dsa_key_file.
keyfile	Specifies the file containing the current server's SSH private key.
passphrase	Specifies the passphrase that was used to encrypt the SSH private key.

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 root passwords or to use SSH public/private keys.

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>
```

Extensibility of the Health Check Monitor

This section is intended for advanced system administrators with experience in Unix shell programming and SA administration.

The Health Check Monitor (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 pre-defined 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 Validating Core Components by Running HCM Local Tests on page 137.) 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_OPSWportping.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 on page 143.

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-sas with the verify_pre option. If you specify the health option of opsware-sas, the scripts in all three directories are executed. The following table describes the mapping between the directory names and the mode options.

Table 22 Modes of opsware-sas 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 non-zero 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 on page 144. 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 1.5.2 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 pre-defined names, see the name=value option in Table 20.

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; therefore, 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 Validating a Core By Running HCM Global Tests on page 139.) 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 on page 139, 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 mulitmaster 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 handing 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 on page 146.

- **Exit Code**: The script returns an exit code of zero to indicate success or non-zero 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 on page 145. 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 `"sys1:pw1 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 on page 140.

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 SAS 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.

SA Components Logs

SA components record events in log files that are useful for troubleshooting. 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`, `grep`, or `vi`.

The log file for a component resides on the server where the component is installed.

By default, the logging debug levels are configured for the highest value (indicating higher priority). The default for the maximum log file size is 10 MB. When the specified maximum file size is reached, additional logs are created. To change the log levels or file sizes, contact your support representative for assistance.

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*
```

Data Access Engine Logs

These logs are in the following files:

```
/var/log/opsware/spin/spin.err*
/var/log/opsware/spin/spin.log*
```

In a core with multiple Data Access Engines, each server running an engine has a set of these log files.

HP Live Network (HPLN) Logs

These logs are in the following location:

```
/var/log/opsware/hpln
```

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 NFSD. 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 Planning and 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
/var/log/opsware/rvrd/rvrdlog*
```

To configure the log file name, log file size, or logging level, in the SAS Web Client, go to Administration ➤ System Configuration ➤ Model Repository Multimaster Component.

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*
```

SAS Web Client Logs

The SAS Web Client does not generate its own logs. The SAS Web Client uses JBoss server, which writes to the following log files:

```
/var/log/opsware/occ/server.log*
```

```
/var/log/opsware/httpsProxy/*log*
```

Software Repository Logs

These logs 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*
```

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 1og.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. 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 1og.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.

Gateway Logs

These logs are in the following files:

/var/log/opsware/*gateway-name*/opswgw.log*

Global File System Logs

These logs are in the following files:

```
/var/log/opsware/hub/OPSWhub.log*
/var/log/opsware/ogfs/ogsh.err*
/var/log/opsware/adapter/adapter.err*
/var/log/opsware/agentcache/agentcache.log
/var/log/opsware/spoke/spoke-*.log
/var/log/opsware/spoke/stdout.log
```

HTTPS Server Proxy Logs

These logs are found in:

```
/cust/apache/servers/https-Proxy/logs
```

The log file ssl_request_log can grow quite large and should be inspected if you are concerned about disk space availability.

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)

The Global Shell audit logs are on the server where the Global File System (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 on page 152.

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:

```
jdoe@m185:051202182224813:13  jdoe@m185:051202182224790:12
2006/01/28-12:40:19.622 User.Id=2610003 User.Name=jdoe
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:

```
jdoe@m185:051202182224813:13
```

This ID field has the following syntax:

```
opsware-user@ogfs-host:YYMMDDHHmmssSSS: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:YYMMDDHHmmssSSS: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:

```
23f1d546cc657137fa012f78d0adfdd56095c3b5
```

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 steps that follow 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 jdoe opened a Remote Terminal to the host (Device.Name) toro.example.com. The event ID is jdoe@m235:060413184452579:59.

   ```
   jdoe@m235:060413184452595:60 jdoe@m235:060413184452579:59 2006/04/
   13-18:44:52.728 User.Id=6220044 User.Name=jdoe 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 `jdoe@m235:060413184452579:59`, then you would enter the following command:

```
tail -f jdoe*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=jdoe@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 tampered with, `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 parameter in the System Configuration page of the SAS Web Client. For instructions on accessing the System Configuration page, see Configuring the Global Shell Audit Logs on page 155.

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. The System Configuration page of the SAS Web Client contains 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 23. For instructions on accessing the System Configuration page of the SAS Web Client, see Configuring the Global Shell Audit Logs on page 155.

Table 23 Parameters for Global Shell Audit Log Configuration

Parameter	description	default value
`audit.root.dir`	The root directory for audit streams and scripts.	`/var/opt/opsware/ogfs/mnt/audit/`
`audit.script.archive_days`	Audit script files older than this value (in days) are deleted. 0 means files are never deleted.	100
`audit.script.archive_size`	Maximum amount of disk space (in MB) used by all audit script files. Older files are removed first. 0 means no maximum.	100
`audit.signature.algorithm`	Signature algorithm to use when signing audit streams.	`RSA-SHA1`
`audit.signature.key_path`	Location of the private key used when signing audit streams.	`/var/opt/opsware/crypto/waybot/waybot.srv`
`audit.stream.archive_days`	Audit stream files older than this value (in days) are deleted. 0 means files are never deleted.	10
`audit.stream.archive_size`	Maximum amount of disk space (in MB) used by all audit stream files. Older files are removed first. 0 means no maximum.	1000
`audit.stream.file_keep`	Maximum number of rotated audit stream files.	50
`audit.stream.file_size`	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 parameters such as the maximum log file size. For a list of the parameters, see Table 23 on page 154. To configure the parameters, perform the following steps:

1. In the SAS Web Client, under Administration click the System Configuration link.
2. On the “System Configuration: Select Product” page, click the hub link.
3. On the “System Configuration: Set Configuration Parameters” page, you can change parameters such as audit.root.dir.
4. Click **Save**.

Start Script for SA

SA provides a multipurpose Start script.

```
/etc/init.d/opsware-oracle
```

You can use the Start 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.

When running the script on a Core Server, the Start script performs the necessary prerequisite checks for each component installed on the local system.

The Start script runs in the background when a server running a component reboots; thus, ensuring that the multiuser boot process will not hang until SA has fully started

➤ When an SA core’s components are distributed across multiple servers, the Start script cannot interact directly with remote servers to start or stop the remote components. However, the Start 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 Start script uses timeout values to allow for different boot times and speed differences among servers. If any of the prerequisite checks fail, the Start script terminates with an error.

Dependency Checking by the Start Script

The Start script has knowledge of SA component dependencies and starts SA components in the correct order. The prerequisite checks verify that dependencies are met before the Start 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 Start 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

Starting the Oracle Database (Model Repository)

The SA Start 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
```

Logging by the Start Script

The Start script writes to the following logs:

Table 24 Start Script Logging

log	notes
/var/log/opsware/startup	When the server boots, the Start script logs the full text (all text sent to stdout) of the start process for all SA components installed on the local system.
stdout	When invoked from the command line, the Start script displays the full text of the start process for the components.
syslog	When the server boots, the Start script runs as a background process and sends status messages to the system event logger.

Syntax of the Start Script

The SA Start 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 25 lists the options for the SA Start script. To see the options of the Health Check Monitor (HCM) also invoked with opsware-sas, see Table 20.

Table 25 Options for the SA Start Script

option	description
list	Displays all components that are installed on the local system and managed by the Start script. The Start script displays the components in the order that they are started.

Table 25 Options for the SA Start Script (cont’d)

option	description
`start`	Starts all components installed on the local system in the correct order. When you use the `start` option to start a specific component, the Start script performs the necessary prerequisitite checks, then starts the component. The `start` option does not start the Oracle database (Model Repository), which must be up and running before the SA components can be started. Some SA components, such as the Web Services Data Access Engine (twist), can take longer to start. For these components, you can run the Start script with the `start` option so that the Start script runs on the local system as a background process and logs errors and failed checks to the component’s log file. NOTE: When you use the `start` option to start multiple components installed on a server, the Start script will always run the `/etc/init.d/opsware-sas` command with the `startsync` option.
`startsync`	The `startsync` option starts all components installed on the local system in a synchronous mode. When you use the `startsync` option, the Start script runs in the foreground and displays summary messages of its progress to `stdout`.
`restart`	Stops and starts all components installed on the local system in a synchronous mode. First, the Start script stops all local components in reverse older; then, executes the `startsync` option to restart the components in the correct order.
`stop`	Stops all components installed on the local system in the correct order. This option does not stop the Oracle database.

Starting an 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

To start a core that has been installed on multiple servers, perform the following steps:

1. Find out which servers contain which SA core components. To list the components installed on a particular server, log in to the server as `root` and enter the following command:

   ```
   /etc/init.d/opsware-sas list
   ```

2. Log in as `root` to the server with the Model Repository and start the Oracle listener and database:

   ```
   /etc/init.d/opsware-oracle start
   ```

3. In the order listed in Details: Start Order for SA Components on page 158, log in as `root` to each core server and enter the following command:

   ```
   /etc/init.d/opsware-sas start
   ```

Starting an SA Core Component

You can specify one or more components to start as long as 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-sas` 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`, you would enter the following command to start the OS Provisioning Build Manager:

   ```
   /etc/init.d/opsware-sas start buildmgr
   ```

Alternatively, you can enter the `startsync` option when starting a component on a server. See Table 25 on page 156 in this chapter for a description of the `startsync` option.

Details: Start Order for SA Components

The Start script starts SA components in the following order. When stopping an SA core, the components are stopped in the reverse order.

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. `dhcpd`: A component of the OS Provisioning feature

6 pxe: The PXE boot environment

7 spin: The Data Access Engine

8 mm_wordbot: A component of the Software Repository

9 waybot: The Command Engine

10 smb: A component of the OS Provisioning feature

11 twist: The Web Services Data Access Engine

12 da: The Deployment Automation component

13 buildmgr: The OS Provisioning Build Manager

14 opswgw-agw0-<facility>: The agent-side Gateway for the facility in which the core is running

15 opswgw-agws: The Agent Gateway

16 opswgw-lb: A Gatway component

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 SAS Web Client

23 httpsProxy: A component of the SAS Web Client

24 opsware-agent: The Server Agent

SA Core Component Internal Names

For legacy reasons, certain SA Core Components are referred to in this documentation using internal naming. Table 26 shows the internal and external names of SA components.

Table 26 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 (SAS Web Client)
opswgw-agw0	Agent Gateway
opswgw-mgws0	Master Gateway
spin	Data Access Engine
spoke	A component of the Global File System

Table 26 Internal and External Component Names (cont'd)

Internal Name	External Name
truth	Model Repository
twist	Web Services Data Access Engine
vault/vaultdaemon	Model Repository Multimaster Component
way/waybot	Command Engine
word	Software Repository

Mass Deletion of Backup Files

SA includes a script that you can run as a cron job for performing mass deletions of backup files. Backup files are created by configuration tracking. They can accumulate quickly and take up disk space. Consequently, performance when viewing backup history in the SAS Web Client can be sluggish, and the information that displays might be cluttered with out-of-date configuration tracking data.

When the backup deletion script is run, it deletes all backed up files with the exception that it always keeps one copy of the latest version of every file ever backed up. If you want to delete those files, use the process for deleting backups individually or a few at a time that is covered in the *SA User's Guide: Server Automation*.

The script is called backup_delete.pyc. It is located on any server hosting a Data Access Engine, in the following directory:

```
/opt/opsware/spin/util
```

The script is run using a configuration file that contains the script arguments such as host name, port number, whether you want full or incremental backups, the backup retention period, the name of the log file to use, email addresses for notifications, and the email server to use. See Table 27, Configuration File Options, for the arguments, their values, and their descriptions.

Syntax of Backup Deletion Script

```
backup_delete.pyc [options]
```

Usage: `backup_delete.py [-c <conf_filename>]`

Deleting Backup Files with the Mass Deletion Script

Perform the following steps to use the mass deletion script to delete backup files:

1. Log in as root to the server where the Data Access Engine is installed.
2. Make sure that /opt/opsware/pylibs is in your PYTHONPATH environment variable.
3. Create a file that contains the arguments and values that you want SA to use with the mass deletion script. See Table 27 on page 161, Configuration File Options, for the available arguments.

For example, the following file specifies that a host called spin.yourcore.example.com, on port 1004 will have incremental backups that are three months old deleted. In addition, a log file called run.log, located in /tmp will be used to capture events, and email will be sent to user@example.com from user1@example.com reporting that the mass deletion was performed successfully.

```
host: spin.yourcore.example.com
port: 1004
inc: 1
time: 3m
logfile: /tmp/run.log
emailto: user@example.com
emailserver: smtp.example.com
emailfrom: user1@example.com
emailsuccess: 1
```

Table 27 Configuration File Options

arguments	values	description
host	host: [hostname], for example host: spin.yourcore.example.com	Host name of the Data Access Engine
port	port: [port number], for example port: 1004	Port of the Data Access Engine (defaults to 1004)
full	Set value to 1 to enable, for example full:1	Delete full backups. You must specify Either full or inc.
inc	Set value to 1 to enable, for example inc:1	Delete incremental backups. You must specify either full or inc.
time	time: [digits][dmy], for example, 6d equals six days. 3m equals three months. 1y equals one year.	Retention period beyond which backups should be deleted.
hostsfile	hostsfile: [filename] The hostsfile should contain the name of each host on a line by itself, for example <hostname> <hostname>	The script deletes backups on every managed server in your system, unless you provide a hostsfile that contains a specific list of servers on which to perform the mass backup deletion.
logfile	logfile: [filename], for example logfile: /tmp/run.log	File to use for log events.
emailto	emailto: [email address], for example emailto: user@example.com	Optional email notification recipient.

Table 27 Configuration File Options (cont'd)

arguments	values	description
emailserver	emailserver: [server name], for example emailserver: smtp.example.com	The SMTP server to send email through. Optional if emailto not specified, otherwise required.
emailfrom	emailfrom: [email address], for example emailfrom: user1@example.com	Email address to appear in the From: line. Optional if emailto not specified, otherwise required.
emailsuccess	Set value to 1 to enable, for example emailsuccess: 1	Send email even if no errors occurred deleting backups and more than one backup was deleted.

4 Optionally, if you want to run the script as a cron job, create a crontab entry.

For example, to run the job at 3:00 AM daily, create the following entry:

```
0 3 * * * env PYTHONPATH=/opt/opsware/spin/util/
backup_delete.pyc -c <path>/<your_backup_filename.conf>
```

The crontab entry must be all on one line.

5 If you do not plan to run the script as a cron job, enter the following command at the prompt:

```
# python /opt/opsware/spin/util/backup_delete.pyc\-c /[conf_filename]
```

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:

- **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.

- **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.
- **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 mesh can be the central data access engine.

Reassigning the Data Access Engine to a Secondary Role

If you 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 SAS Web Client as a user that belongs to SA Administrators group.The SAS Web Client home page appears.
2. From the Navigation panel, click **Administration ➤ Opsware Software**. The **Software** page appears.
3. Click the **spin** link. The **Opsware Software | spin** page appears.
4. Select the **Members** tab. The list of Managed Servers that are hostin a Data Access Engine appears.
5. Select the check box for the **additional Data Access Engine server**.
6. From the **Tasks** menu, select **Re-Assign Node**.
7. Select the option for the **Service Levels | Opsware | spin node**.
8. Click **Select**.
9. Navigate the node hierarchy by clicking the following nodes:
   - Opsware
   - spin
   - Secondary
10. Click **Re-Assign**.
11. 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.

In most case, 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 your support representative.

Perform the following steps to designate the multimaster central data access engine:

1. Log into the SAS Web Client as a user that belongs to the SA System Administrators group.
2. From the Navigation panel, click **Opsware Software** under Administration. The **Opsware Software** page appears.
3. Click the **spin** link.
4. Select the **Servers** tab.
5. Select the check box for the Data Access Engine server for the new core.
6. From the **Server** menu, select **Re-Assign Node**.
7. Select the option for the **Service Levels | Opsware | spin | node**.
8. Click **Select**.
9. Navigate the node hierarchy by clicking each node: **Opsware | Spin | Multimaster Central**.
10. Click **Re-Assign**.
11. Restart the Multimaster Central Data Access Engine.

    ```
    /etc/init.d/opsware-sas restart spin
    ```

Scheduling Audit Result and Snapshot Removal

Because Audit Results (results of an audit) 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 snapshots removal, perform the following steps:

1. Log into the SAS Web Client as the admin user with the password you supplied during the interview. Log in by opening a browser and entering the IP address of the server running the SAS Web Client.

   The SAS Web Client should be installed and listening. The SAS Web Client home page appears.
2. From the navigation panel, click System Configuration under Administration. The Select a Product page appears.
3. Click the link for the SAS Web Client. The configuration page appears.
4. Under Select a Product, click the Data Access Engine link. The configuration page for the Data Access Engine appears.

5 To set the number of days to elapse before an audit results or snapshot are deleted, modify the following parameters:

- Scroll down to the spin.cronbot.delete_snapshots.cleanup_day parameter, and in the Use value field enter the number of days that must elapse before all non-archived snapshots will be deleted. If you select the Use default value setting, no snapshots will deleted.
- Scroll down to the spin.cronbot.delete_audits.cleanup_days, and in the Use value field enter the number of days that must elapse before all non-archived Audit Results will be deleted. If you select the Use default value setting, no snapshots will deleted.

  When you are finished, at the bottom of the page, click **Save**.

Web Services Data Access Engine Configuration Parameters

This section discusses how to change Web Services Data Access Engine configuration parameters using the SAS Web Client or by editing the configuration file.

➤ The Web Services Data Access Engine must be restarted for any parameter modifications you have made to take affect.

Changing a Configuration Parameter

This section describes how to change the parameters displayed by the SAS Web Client. However, the SAS Web Client does not list all of the Web Services Data Access Engine parameters. If you want to change an unlisted parameter, follow the instructions in the next section.

To change a parameter in the SAS Web Client, perform the following steps:

1 Log into the SAS Web Client as a member of the Administrators group (admin) and from the navigation panel, click System Configuration under Administration. The Select a Product page appears.

2 Under Select a Product, click Web Services Data Access Engine.

3 Update the parameters you want to change.

4 Click **Save**.

5 Restart the Web Services Data Access Engine.

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 SAS Web Client.) The fully-qualified name of the configuration file follows:

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

➤ During an upgrade of SA, the `twist.conf` file is replaced, but the `twistOverrides.conf` file is preserved. When you upgrade to a new version of SA, to retain the configuration settings, you must edit the `twistOverrides.conf` file. The properties in `twistOverrides.conf` override those specified in `twist.conf`. T Unix `twist` user must have write access to the `twistOverrides.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.

➤ You must belong to the Administrators group (admin) in order 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`.

Table 28 Configuration File for SA Web Services API 2.2

Property	Default	Description
`twist.webservices.debug.level`	`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.
`twist.webservices.locale.country`	`US`	The country Internationalization parameter for the Localizer utility. Currently only the `US` code is supported.
`twist.webservices.locale.language`	`en`	Sets the language Internationalization parameter for the Localizer utility. Currently only the `en` code is supported.
`twist.webservices.caching.windowsize`	`120`	In minutes, the size of the sliding window maintaining the server event cache.
`twist.webservices.caching.windowslide`	`15`	In minutes, the sliding scope for the window maintaining the server event cache.
`twist.webservices.caching.safetybuffer`	`5`	In minutes, the safety buffer for the sliding window maintaining the server event cache.
`twist.webservices.caching.minwindowsize`	`30`	In minutes, the minimum size of the sliding window that maintains the server event cache.

Table 28 Configuration File for SA Web Services API 2.2 (cont'd)

Property	Default	Description
`twist.webservices.caching.maxwindowsize`	`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 defaut value is 1280Mb. To do so, perform the following tasks:

1. Using a text editor, open the file:

   `/etc/opt/opsware/twist/twistOverrides.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 `twistOverrides.conf` parameter blank, the default value (1280m) specified in `twist.sh` is used.

Changing Software Repository Mirroring Parameters

This section discusses how to change Word Mirroring configuration parameters using the SAS Web Client.

Changing a Configuration Parameter

This section describes how to change the configuration parameters displayed by the SAS Web Client.

To change a parameter in the SAS Web Client, perform the following steps:

1. Log into the SAS Web Client as a member of the Administrators group (`admin`). From the navigation panel, click System Configuration under Administration. The Select a Product page appears.
2. Under Select a Product, click Software Repository.
3. Update the parameters you want to change.
4. Click **Save**.

5 Restart the all instances of the Software Repository for the SA Core. If the change is global, restart all instances of the Software Repository for all cores in the Multimaster Mesh.

Software Repository Mirroring Configuration Parameters

You can modify the following configuration parameters:

Table 29 Software Repository Mirroring Parameters

Parameter	Type	Allowed Values	Default	Description
`word.enable_content_mirroring`	Flag	1, 0	0	Enables (1)/ Disables (0) Software Repository mirroring.
`word.mirror_job_period`	Minutes	A positive integer	60	The frequency (in minutes) between . Software Repsository mirroring job runs.

6 Monitoring SA

Overview of SA Monitoring

SA has a built-in system diagnosis function in the SAS Web Client, which allows you to diagnosis the functionality 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 chapter provides information for performing basic monitoring of the components listed above and for the following additional SA components:

- Server Agent
- Agent Cache
- SAS Web Client
- Model Repository
- Spoke
- Gateways
- OS Build Manager
- OS Boot Server
- OS Media Server

The commands and other information shown in this document are identical to those in the SAS Web Client.

The information contained in this document should be used when the System Diagnosis feature cannot be used because the SAS Web Client cannot be reached 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.

The type of monitoring information described in this document 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

➤ 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 an intelligent agent 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 SAS Web 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 output:

```
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
                      /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/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/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/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/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/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/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/opsware/agent/pylibs/shadowbot/daemonbot.pyc
                     --conf /etc/opt/opsware/agent/agent.args
```

Agent URL

```
https://<hostname>:1002
```

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 “OpswareError”

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 Agent that is available. The Discovery and Agent Deployment (ODAD) feature obtains the agent installation binaries from the Agent Cache component during agent deployment.

Agent Cache Ports

The Agent Cache uses port 8081.

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/opsware/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. You use the SAS Web 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 **Solaris** or **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.
```

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 URL

```
https://occ.<data_center>
```

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`

Load Balancing Gateway Monitoring

The Load Balancing Gateway provides High Availability and horizontal scaling in an SA core.

When you run the HP BSA Installer, it installs a Load Balancing Gateway with the Command Center component.

Load Balancing Gateway Ports

By default, the Load Balancing Gateway uses the port 8080.

Monitoring Processes for the Load Balancing Gateway

In all configurations, the Load Balancing 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 Command Center component:

```
# ps -eaf | grep -v grep | grep opswgw | grep lb
```

Running this command should produce output similar to the following output:

```
root  32149  1  0  Sep27  ?  00:00:00  [opswgw-watchdog-2.1.1: lb]
                     --PropertiesFile /etc/opt/opsware/opswgw-lb/opswgw.properties --BinPath /opt/
                     opsware/opswgw/bin/opswgw
root  32156  32149  0  Sep27  ?  00:24:31  [opswgw-gateway-2.1.1: lb]
                     --PropertiesFile /etc/opt/opsware/opswgw-lb/opswgw.properties --BinPath /opt/
                     opsware/opswgw/bin/opswgw --Child true
```

Load Balancing Gateway Logs

The Load Balancing 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 terminate)

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 on page 162.

For information about the Gateway mesh topology, see the *SA Overview and Architecture Guide*.

Monitoring Processes for the Data Access Engine

On Solaris, execute the command on the server running the Data Access Engine component:

```
# /usr/ucb/ps auxwww | grep -v grep | grep spin | grep -v java
```

Running this command should produce output similar to the following output:

```
root  8010  0.5  0.84541631552  ?  S  19:36:42  4:56  /opt/opsware/bin/
                     python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/opt/opsware/
                     spin/spin.args
root  8008  0.0 0.1  4040  2080  ?  S  19:36:42  0:00  /opt/opsware/bin/
                     python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/opt/opsware/
                     spin/spin.args
root  8026  0.0  0.53224018224  ?  S  19:36:57  0:01  /opt/opsware/bin/
                     python /opt/opsware/spin/certgenmain.pyc --start --conf
                     /etc/opt/opsware/spin/spin.args
```

On **Solaris**, you see multiple process that look like the first line of the output above; however, there should be only one process that contains `certgenmain` in the output.

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/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  /opt/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  /opt/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`.

- `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 **Solaris**, execute the command on the server running the Command Center component and on the server running the Slice Component bundle:

```
# /usr/ucb/ps auxwww | grep -v grep | grep \/opt\/opsware\/twist
```

Running this command should produce output similar to the following output:

```
twist  9274  0.0  1.416748054040  ?  S  Aug 08 410:33 /opt/opsware/
                   j2sdk1.4.2_10/bin/java -server -Xms16m -Xmx128m -Dtwist.port=1026 ......
                   -classpath opt/opsware/j2sdk1.4.2_10/jre ......
twist  9238  0.0  0.1  1088  744  ?  S  Aug 08  0:00 /bin/sh /opt/
                   opsware/twist/watchdog.sh start 60
```

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 **Solaris**, execute the command on the server running the Command Engine component:

```
# /usr/ucb/ps auxwww | egrep '(COMMAND$|waybot)' | grep -v grep
```

Running this command should produce output similar to the following output:

```
USER  PID  %CPU  %MEM  SZ  RSS  TT      S  START  TIME  COMMAND
root 1246   0.0   0.1 4040 2064  ?      S Sep 24  0:00 /opt/opsware/bin/
                    python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/opt/opsware/
                    waybot/waybot.args
root 1248   0.0   0.41596814592 ?       S Sep 24  2:19 /opt/opsware/bin/
                    python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/opt/opsware/
                    waybot/waybot.args
```

On Solaris, the Command Engine has two processes — one process for the daemon monitor and one process for the server.

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 URL

```
https://way.<data_center>:1018
```

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 is where all software managed by SA is stored.

Software Repository Ports

The Software Repository uses the following ports:

- 1003 (Encrypted)
- 1006 (Cleartext)
- 1005 (Replicator administrative user interface)
- 5679 (Multimaster Software Repository)

Monitoring Processes for the Software Repository

On **Solaris**, execute the command on the server running the Software Repository component:

```
# /usr/ucb/ps auxwwww | grep -v grep | grep mm_wordbot
```

Running this command should produce output similar to the following output:

```
root  8625  0.0  0.1 4048 1912 ?  S  Aug 08  0:00 /opt/opsware/bin/
                     python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/opt/opsware/
                     mm_wordbot/mm_wordbot.args
root  8627  0.0  0.52034418600  ?  S  Aug 08  7:38  /opt/opsware/bin/
                     python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/opt/opsware/
                     mm_wordbot/mm_wordbot.args
root  8675  0.0  0.1  4032  1904  ?  S  Aug 08  0:00  /opt/opsware/bin/
                     python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/opt/opsware/
                     mm_wordbot/mm_wordbot-clear.args
root  8677  0.0  0.210104  8096  ?  S  Aug 08  0:01  /opt/opsware/bin/
                     python /opt/opsware/pylibs/shadowbot/daemonbot.pyc --conf /etc/opt/opsware/
                     mm_wordbot/mm_wordbot-clear.args
```

On Solaris, the Software Repository has four running processes — two processes for the encrypted Software Repository and two for the cleartext Software Repository.

On **Linux**, execute the command on the server running the Software Repository component:

```
# ps auxwwww | grep -v grep | grep mm_wordbot
```

Running this command should produce output similar to the following output:

```
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 /opt/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  /opt/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/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 cleartext Software Repository.

Software Repository URL

```
https://theword.<data_center>:1003
```

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

If you have Software Repository mirroring enabled, you can montor the status by:

1. Logging in to the SAS Web Client as a member of the SA Administrator group or with Multimaster Tools permissions.
2. Select **Administration ➤ Software Repository Mirroring** from the navigation panel.
3. You will see a screen similar to that shown in Figure 38:

Figure 38 Software Repository Mirroring Status Scteen

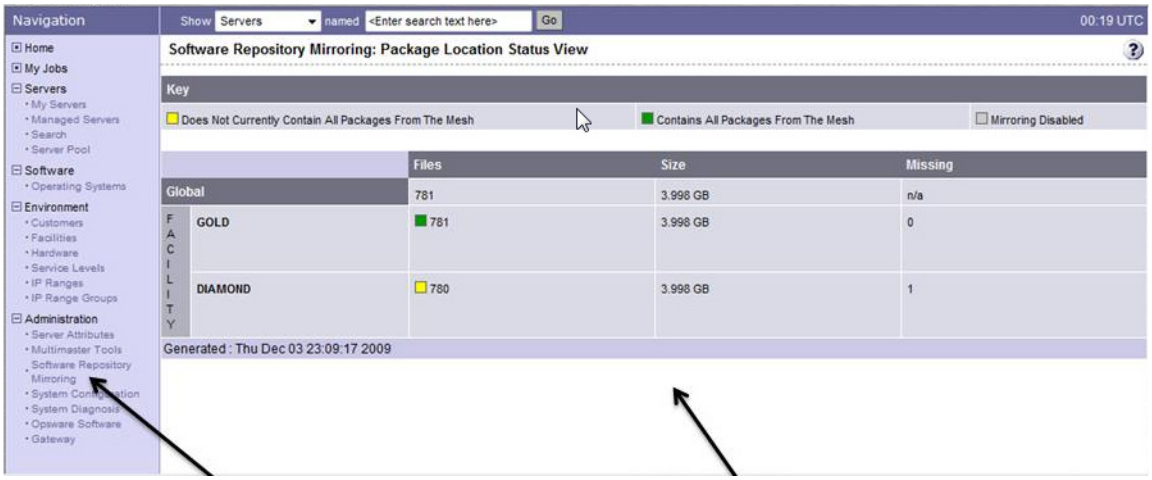

Software Repository Mirroring: Package Location Status View

Key

🟨 Does Not Currently Contain All Packages From The Mesh   🟩 Contains All Packages From The Mesh   □ Mirroring Disabled

		Files	Size	Missing
Global		781	3.998 GB	n/a
FACILITY	GOLD	🟩 781	3.998 GB	0
	DIAMOND	🟨 780	3.998 GB	1

Generated : Thu Dec 03 23:09:17 2009

The status screen displays the Software Repository hosts in your Multimaster Mesh by facility name. Data displayed includes:

- **Files**: the number of files in the host’s Software Repository
- **Size**: the approximate total disk space used by the Software Repository files
- **Missing**: the number of files that should be mirrored by the Facility's Software Repository but that have not yet been replicated. This state is also indicated by the color coded boxes in the Files column:
  - — **Green**: Contains all packages from the Mesh
  - — **Yellow**: Does not contains all packages from the Mesh
  - — **Grey**: Mirroring Disabled

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 Simple/Advanced 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 **Solaris** or **Linux**, execute the command on the server running Oracle:

```
# ps -fu oracle | grep pmon
```

Running this command should produce output similar to the following output:

```
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 **Solaris** or **Linux**, use this command to monitor the Oracle Listener process:

```
# ps -fu oracle | grep tnslsnr
```

Running this command should produce output similar to the following output:

```
oracle    2021     1  0 21:22 ?    00:00:01 /u01/app/oracle/product/10.2.0/
db_1/bin/tnslsnr 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 [lxmcpen()+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 Simple/Advanced 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 **Solaris**, execute the command on the server where you installed the Infrastructure Component bundle:

```
# /usr/ucb/ps auxwww | grep -v grep | grep vault | grep -v twist
```

Running this command should produce output similar to the following output:

```
root  3884  0.0  0.1  2792  1568  ?  S  Jul 26  0:00  /opt/opsware//bin/
                     python /opt/opsware//pylibs/shadowbot/etc/daemonizer.pyc
                     --runpath /var/log/opsware/vault --cmd /opt/opsware/j2sdk1.4.2_10/bin/java
                     -classpath /opt/opsware/vault ...... -ms120m -mx1024m -DCONF=/etc/opt/opsware/
                     vault/
                     -DHOSTNAME= com.Opsware.vault.Vault
root  3885  0.0  0.1 1096  848  ?  S  Jul 26  0:00  /bin/sh -c /opt/
                     opsware/j2sdk1.4.2_10/bin/java -classpath /opt/opsware/vault/cl
root  3887  0.0  3.9194192155784  ?  S  Jul 26  2:34  /opt/opsware/
                     j2sdk1.4.2_10/bin/java -classpath /opt/opsware/vault ...... -ms120m -mx1024m
                     -DCONF=/etc/opt/opsware/vault/
                     -DHOSTNAME= com.loudcloud.vault.Vault
```

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 should produce output similar to the following output:

```
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 /opt/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  /opt/opsware/
                    j2sdk1.4.2_10/bin/java -classpath /opt/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`

Condition to monitor in the logs: The string “Traceback”

To configure the log file name, log file size, or logging level, in the SAS Web Client, go to Administration ➤ System Configuration ➤ Model Repository Multimaster Component.

Global File System Monitoring

The Global Shell feature is installed as part of any Slice Component bundle, and dynamically constructs a virtual file system — the Global File System (OGFS).

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 (though 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. On Solaris, a Python program that communicates with a custom 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 should produce output similar to the following output:

```
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 /opt/opsware/adapter/SunOS/bin/rotator.py
                     /opt/opsware/......
    7617 /opt/opsware/ogfsutils/bin/python2.4 /opt/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 /opt/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 should produce 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 /opt/opsware/bin/python /opt/
                     opsware/hub/bin/rotator.py /opt/opsware/j2sdk1.4.2_10/b......
root 8868 0.1 1.8 1256536 76672 ? S Sep29 35:51 /opt/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 /opt/opsware/bin/python /opt/
                     opsware/adapter/bin/adapter......
root 8908 0.0 0.0 13088 684 ? S Sep29 0:10 /opt/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 /opt/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 /opt/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 should produce 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 should produce output similar to the following output:

```
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 should produce output similar to the following:

```
137 1322cd8 43a9 272 1 ogdrv (Opsware GFS driver v1.13)
138 13ac227 338df 18 1 ogfs (Opsware 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 Global File System (OGFS) and provides access to run commands inside an OGFS session.

Spoke Ports

The Spoke uses port 8020.

Monitoring Processes for the Spoke

On **Solaris**, execute the command on the server running the Slice Component bundle:

```
# /usr/ucb/ps auxwww | grep -v grep | grep Spoke
```

Running this command should produce output similar to the following:

```
root  4831  0.1  1.316426451168 pts/1  S  Jul 26  167:58  /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/logging.bootstrap
                    -Dweblogic.security.SSL.ignoreHostnameVerification=true ...... -classpath /
                    opt/opsware/spoke/lib/HTTPClient-hacked.jar: ...... com.opsware.spoke.Spoke
```

On **Linux**, execute the command on the server running the Slice Component bundle:

```
# ps -ef | grep -v grep | grep spoke
```

Running this command should produce 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 on page 111.

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 should produce output similar to the following output:

```
root  17092  1  0  Sep21  ?  00:00:00  [opswgw-watchdog-2.1.1: cgw0-C43]
                    --PropertiesFile /etc/opt/opsware/opswgw-cgw0-C43/opswgw.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/opswgw.properties
                    --BinPath /opt/opsware/opswgw/bin/opswgw --Child true
```

```
# ps -eaf | grep -v grep | grep opswgw | grep agw
```

Running this command should produce output similar to the following output:

```
root  17207  1  0  Sep21  ?  00:00:00  [opswgw-watchdog-2.1.1: agw0-C43]
                    --PropertiesFile /etc/opt/opsware/opswgw-agw0-C43/opswgw.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/opswgw.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 should produce output similar to the following output:

```
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 SAS Web 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 **Solaris** or **Linux**, execute the command on the server running the OS Build Manager component:

```
# ps -eaf | grep -v grep | grep buildmgr
```

Running this command should produce 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.security -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 processes used to provide this support include the Internet Software Consortium DHCP server, and Sun Solaris TFTP and NFS.

These applications are installed by the HP BSA 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.

7 SA Notification Configuration

This section describes user-definable configuration parameters that allow you to modify contact information in the SAS Web 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 Planning and Installation Guide*.

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.

Server Agents read system configuration values at installation time only. If you change any configuration values, all Agents' configurations must be updated manually. Contact Technical Support 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 HP 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 SAS Web Client Help:

   ```
   pref.occ.support.href
   ```

   ```
   pref.occ.support.text
   ```

5. Save the file and exit the editor.
6. Restart the Command Center 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 parameter opsware.mailserver to determine the address of the mail server to use for email notification. 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 tasks:

1 Log into the SAS Web Client as the admin user. The SAS Web Client home page appears.

2 From the Navigation panel, click **Administration ➤System Configuration**. The **Select a Product** page appears.

3 Under **Select a Product**, click the *link for the Facility name*. The **Configuration** page for the Facility appears.

4 In the value field for opsware.mailserver, enter the host name of your mail server.

5 Click **Save** to apply the changes. The configuration page refreshes and a message should appear indicating that the update was successful.

Configuring the Command Engine Notification Email

1 From the Navigation panel, click **Administration ➤System Configuration**. The **Select a Product** page appears.

2 Under Select a Product, click **Command Engine**.

3 In the value field for way.notification.email.fromAddr, enter the From: email address for the emails that will be sent by the Command Engine to notify users about scheduled jobs.

4 Click **Save** to apply the changes.

5 Exit the SAS Web Client and restart the Command Engine.

```
/etc/init.d/opsware-sas restart occ.server
```

6 If HP Server Automation 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 Chapter 5, “Starting an SA Core” on page 157 of this guide.

Setting Email Alert Addresses for an SA Core

Server Agents read system configuration values at installation time only. If you change any configuration values, all Agents’ configurations must be updated manually. Contact Technical Support for help making these changes, or in making any other changes in SA System Configuration.

Perform the following tasks to configure email alert addresses. SA Core installation uses the default value (EMAIL_ADDR) for these parameters.

1 Log on to the SAS Web Client as admin user. The SAS Web Client home page appears.

2 From the Navigation panel, click **Administration ➤System Configuration**. The **Select a Product** page appears.

3 Under **Select a Product**, click the **Server Agent** link. The configuration page for the Agent appears.

4 Configure the following required email alert addresses:

- In the field, acsbar.ErrorEmailAddr, enter the address that HP Server Automation should use to send email warnings when configuration tracking limit is exceeded (for example, when the configuration tracking feature stopped backing up configuration files and databases).

- In the field, acsbar.emailFromAddr, enter the address that the Agent should use as the email From: address when sending emails. For example, agent@yourdomain.com.

- In the field, CronbotAlertAddress, enter the email address that the Server Agent should use to alert the recipient about failed scheduled jobs.

- In the field, CronbotAlertFrom, enter the email address that the Agent should use as the email From: address in the emails about failed scheduled jobs. For example, agent@yourdomain.com.

5 Click **Save** to apply the changes. The configuration page refreshes and a message should appear indicating that the update was successful.

Configuring Email Alert Addresses for a Multimaster Mesh

Perform the following tasks to configure email alert addresses for Multimaster alerts. SA Core installation uses the default value EMAIL_ADDR for these parameters.

1 Log on to the SAS Web Client as admin user. The SAS Web Client home page appears.

2 From the Navigation panel, click **Administration ➤System Configuration**. The **Select a Product** page appears.

3 Click the **Model Repository, Multimaster Component** link. The **Configuration** page for the Model Repository Multimaster Component appears.

4 Configure the following email parameters:

- In the field, sendMMErrorsTo, enter the email address to which multimaster conflicts will be sent.

- In the field, sendMMErrorsFrom, enter the address that HP Server Automation will use as the email From: address for Multimaster conflicts alert emails.

5 Click **Save** to apply the changes. The configuration page refreshes and a message should appear indicating that the update was successful.

6 Restart the Model Repository Multimaster Component in all SA Cores in the Multimaster Mesh. See Chapter 5, “Starting an SA Core Component” on page 158 of this guide.

Configuring Email Notification Addresses for Code Deployment and Rollback (CDR)

Code Deployment and Rollback (CDR) is deprecated but still supported in SA 7.80. It will not be supported in a future major release.

You can configure email notification addresses for SA Code Deployment & Rollback (CDR). When users request that a service operation or synchronization be performed on their behalf, an email notification is sent to the individuals assigned to perform the requested service operation or synchronization.

Perform the following tasks to configure email notification addresses for CDR. SA Core installation uses the default value `EMAIL_ADDR` for these parameters.

1 Log on to the SAS Web Client as `admin` user. The SAS Web Client home page appears.

2 From the Navigation panel, click **Administration ➤System Configuration**. The **Select a Product** page appears.

3 Click the link for the **SAS Web Client**. The **Configuration** page appears, as Figure 39 shows.

Figure 39 CDR Email Notification Configuration Parameters

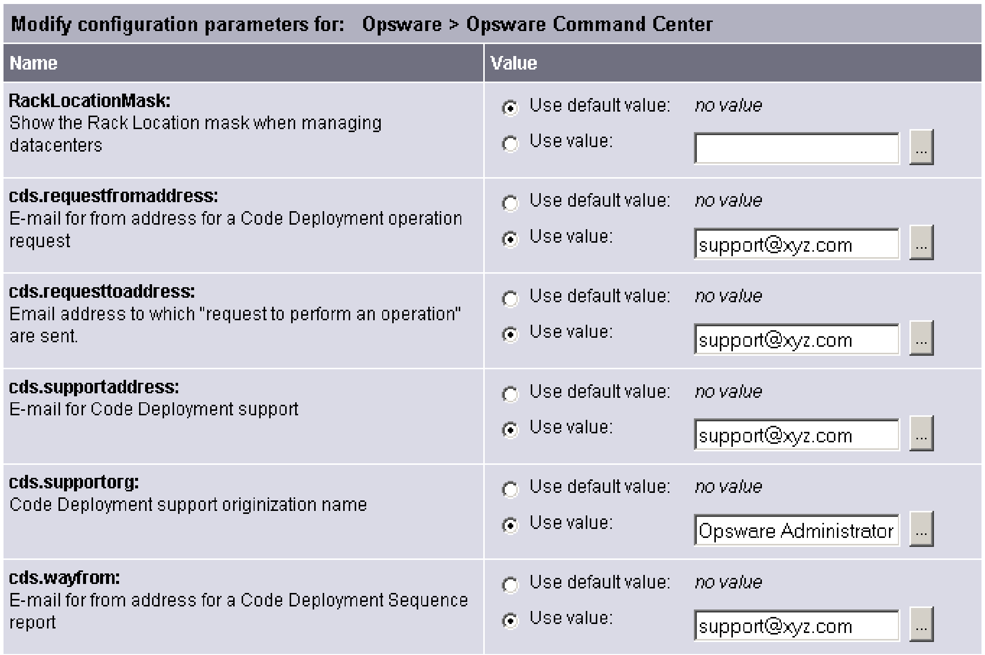

Modify configuration parameters for: Opsware > Opsware Command Center

Name	Value
RackLocationMask: Show the Rack Location mask when managing datacenters	● Use default value: *no value* ○ Use value:
cds.requestfromaddress: E-mail for from address for a Code Deployment operation request	○ Use default value: *no value* ● Use value: support@xyz.com
cds.requesttoaddress: Email address to which "request to perform an operation" are sent.	○ Use default value: *no value* ● Use value: support@xyz.com
cds.supportaddress: E-mail for Code Deployment support	○ Use default value: *no value* ● Use value: support@xyz.com
cds.supportorg: Code Deployment support originization name	○ Use default value: *no value* ● Use value: Opsware Administrator
cds.wayfrom: E-mail for from address for a Code Deployment Sequence report	○ Use default value: *no value* ● Use value: support@xyz.com

4 Customize the following parameters to include the following email notification information:

- In the field, `cds.requesttoaddress`, enter the email address to use in the `To:` field of a request notification email.
- In the field, `cds.requestfromaddress`, enter the email address to use in the `From:` field of a request notification email.
- In the field, `cds.wayfrom`, enter the email address to use in the `From:` field of emails sent following completion of a sequence.

- In the field, `cds.supportaddress`, enter the email address to use for a Facility’s' support organization or contact person.
- In the field, `cds.supportorg`, enter the display name (how it appears in the SAS Web Client) of a Facility’s support organization.

5. Click **Save** to apply the changes. The configuration page refreshes and a message should appear indicating that the update was successful.
6. Restart the Command Engine and the Model Repository Multimaster Component. See Chapter 5, “Starting an SA Core Component” on page 158 of this guide.

8 Global Shell: Windows Subauthentication Package

It is a well-known attribute of Microsoft® Windows that 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 rationale applies to the SA Agent, as it does to all other services. 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 NT4, Windows 2000, Windows XP, and Windows Server 2003 and Windows Server 2008 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 username, 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 Configuration Guide* 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 (LSASS) and takes the form of a process running the lsass.exe application on a Windows server.

The design of LSASS 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, LSASS 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 and Windows Server 2008.

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 NT4, Windows 2000, Windows Server 2003, and Windows Server 2008) and is used in an identical way on each operating system.

➤ All Windows NT4 operating systems must have the Microsoft patch Q828035 installed to support the ogshcap.dll.

Figure 1 illustrates the subauthentication process in SA.

Figure 1 SA Subauthentication Process Flow

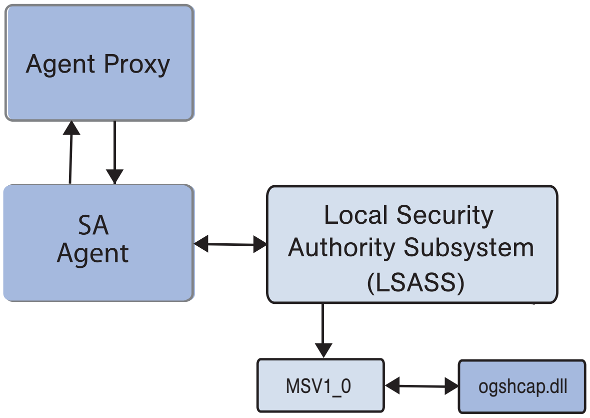

In the case of the SA Agent, the Agent passes a NULL password along with the user name when it calls a special Windows API to request subauthentication by the SA subauthentication package (ogshcap.dll). The Windows API then calls the MSV1_0 authentication package which, in turn, passes the credentials, 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 LSASS. In turn, LSASS 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 `CreateProcessAsUser()` 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.

➤ 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 a new file (ogshcap.dll) in the SA 7.80 release. During an Agent installation, the ogshcap.dll file is copied to the following source location:

```
%SystemDrive%:\Program Files\Opsware\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.

➤ When a post-install reboot is required to get the latest version of the DLL, the reboot performs a move operation where 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 checksum 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\Opsware\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\Opsware\cogbot\hmac.key size: 36 bytes
[08/Jun/2005 20:59:18] [TRACE] Successfully called CloseHandle(C:\Program
Files\Common Files\Opsware\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\Opsware\cogbot\hmac.key size: 36 bytes
```

```
[08/Jun/2005 20:59:18] [TRACE] Successfully called CreateFileMapping() for
C:\WINDOWS\system32\ogshcap.dll
[08/Jun/2005 20:59:18] [TRACE] Successfully called CreateFileMapping() for
C:\Program Files\Common Files\Opsware\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\system32\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\Opsware\bin\ogshcap.dll
[08/Jun/2005 20:59:18] [TRACE] C:\Program Files\Opsware\agent\bin\ogshcap.dll
size:
40960 bytes
[08/Jun/2005 20:59:18] [TRACE] C:\Program Files\Common
Files\Opsware\cogbot\hmac.key size: 36 bytes
[08/Jun/2005 20:59:18] [TRACE] Successfully called CreateFileMapping() for
C:\Program Files\Opsware\agent\bin\ogshcap.dll
[08/Jun/2005 20:59:18] [TRACE] Successfully called CreateFileMapping() for
C:\Program Files\Common Files\Opsware\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\Opsware\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 successful
[08/Jun/2005 20:59:18] [TRACE] Win32InstallN() = 1
[08/Jun/2005 20:59:18] [INFO] Installation completed successfully.
[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 on page 204 for more information.

A Permissions Reference

Permissions Required for the SAS Web Client

The following table lists the feature permissions according to tasks that can be performed with the SAS Web Client.

Table 1 Permissions Required for SAS Web Client Tasks

Task	Feature Permission
OS PROVISIONING	
Prepare OS	Wizard: Prepare OS
Edit OS nodes	Operating Systems
View servers in the server pool	Server Pool
CONFIGURATION TRACKING	
Create or edit tracking policy	Configuration Tracking Managed Servers and Groups
Reconcile tracking policy	Configuration Tracking Managed Servers and Groups
Perform configuration backup	Configuration Tracking Managed Servers and Groups
View backup history, restore queue	Configuration Tracking Managed Servers and Groups
Enable or disable tracking	Configuration Tracking Managed Servers and Groups
SERVER MANAGEMENT	
Edit server properties	Managed Servers and Groups
Edit server network properties	Managed Servers and Groups
Edit server custom attributes	Managed Servers and Groups
Deactivate server	Deactivate
Delete server	Managed Servers and Groups
Re-assign customer	Managed Servers and Groups
View servers (read-only access)	Managed Servers and Groups
Run server communications test	Managed Servers and Groups

Table 1 Permissions Required for SAS Web Client Tasks (cont'd)

Task	Feature Permission
Lock servers	Managed Servers and Groups
Set scheduled job to refresh server list	Allow Run Refresh Jobs
REPORTS	
Create or view reports	Data Center Intelligence Reports
MANAGE ENVIRONMENT	
Create or edit customer	Customers
Create or edit facility	Facilities
IP RANGES AND RANGE GROUPS	
IP Ranges	IP Ranges and Range Groups Model: Hardware Model: SA
IP Range Groups	IP Ranges and Range Groups Model: Hardware Model: SA
SYSTEM CONFIGURATION	
Manage users and groups	(Administrators group only)
Define server attributes	Server Attributes
Run system diagnosis tools	System Diagnosis
Manage SA System configuration	Configure SA
Run SA multimaster tools	Multimaster
Gateway management	Manage Gateway
OTHER TASKS	
Run custom extension	Wizard: Custom Extension
Deploy code	See "Code Deployment User Groups" on page 278.

Storage Visibility and Automation Permissions

You must have certain permissions to perform actions in the Storage Visibility and Automation feature. See the *Storage Visibility and Automation Installation & Administration Guide* for a description of these permissions.

Permissions Required for the SA Client

The following tables in this section summarize the permissions required for the SA Client features.

- Application Configuration Management Permissions
- Device Group Permissions
- SA Discovery and Agent Deployment Permissions
- Job Permissions
- Patch Management for Windows - Permissions
- Patch Management for Solaris - Permissions
- Solaris Patch Policy Management - Permissions
- Patch Management for Other Unix - Permissions
- Software Management Permissions
- Script Execution Permissions
- Audit and Remediation Permissions
- Service Automation Visualizer Permissions
- Virtual Server Permissions
- OS Provisioning Permissions
- Compliance View Permissions
- Server Property and Reboot Permissions
- Server Objects Permission

More Information for Security Administrators

In some organizations, security administrators work with many applications and do not specialize in SA. To learn about SA quickly, security administrators can refer to Process Overview for Security Administration on page 22 - This short section lists the overall tasks for setting up security in SA.

Application Configuration Management Permissions

Table 2 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?

In addition to the feature permissions listed in Table 2, every user action also requires the Managed Servers and Groups feature permission.

In Table 2, 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 SAS Web Client. In Table 2, 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

Table 2 Application Configuration Management Permissions Required for User Actions

User Action	Feature 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 Configuration	Manage Application Configurations: Read & Write and Manage Configuration Templates: Read	None	Read & Write

Table 2 Application Configuration Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
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 and Manage Configuration Templates: Read and Manage Installed Configuration and Backups on Servers: Read & Write	Read & Write	Read

Table 2 Application Configuration Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
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 Configuration Audit	Allow Configuration Compliance Scan: Yes and Manage Application Configurations: Read and Manage Configuration Templates: Read	Read	Read

Table 2 Application Configuration Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
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: Read & Write	None	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

Table 2 Application Configuration Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
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

Table 3 lists the actions that users can perform with application configurations for each permission. Table 3 has the same data as Table 2, but is sorted by permission. Although not indicated in Table 3, the Managed Servers and Groups permission is required for all OS provisioning actions.

For security administrators, Table 3 answers this question: If a user is granted a particular permission, what actions can the user perform?

Table 3 User Actions Allowed by Application Configuration Management Permissions

Feature 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

Table 3 User Actions Allowed by Application Configuration Management Permissions (cont'd)

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permission (App Config, App Config Template)
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
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

Table 3 User Actions Allowed by Application Configuration Management Permissions (cont'd)

Feature 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 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 Configuration Template	None	Read & Write
Manage Configuration Templates: Read & Write (cont.)	Set Application Configuration Template to Run as Script	None	Read & Write
	View Application Configuration Template	None	Read

Device Group Permissions

To use the Device Groups feature in the SA Client, you must have the permissions described in the Table 4. For a list of tasks that require the Model Public Device Group permission, see Table 15.

Table 4 Device Groups Feature Permissions

User Action	Feature Permission
Creating a public static device group	Manage Public Device Group

Table 4 Device Groups Feature Permissions (cont'd)

User Action	Feature Permission
Creating a public dynamic device group	Manage Public Device Group
Adding a server to a public static device group	Manage Public Device Group
Adding a server to a public dynamic device Group	Manage Public Device Group
Removing a server from a public static device group	Manage Public Device Group
Removing servers from a public dynamic device group	Manage Public Device Group
Moving a public device group	Manage Public Device Group
Duplicating a public device group	Manage Public Device Group
Deleting a public device group	Manage Public Device Group
Adding devices to a device group being used as an Access Control Group	Manage Public Device Group and Super Administrator

SA Discovery and Agent Deployment Permissions

To use Discovery and Deployment (ODAD) in the SA Client, you must have the permissions described in the Table 5.

Table 5 ODAD Feature Permissions

User Action	Feature Permission
Deploy (Install) Agent with ODAD	Allow Deploy Agent: Yes
Scan Network with ODAD	Allow Scan Network: Yes
View Servers Running Agents	Managed Servers and Groups
Add servers to a facility	Facilities - Manage Facilities

In addition to the feature permissions listed in the preceding table, the following is also required:

- Read access to facilities where you will scan for servers and manage servers
- **Features ➤ Managed Servers and Groups** must be enabled
- **Client Features ➤ Unmanaged Servers ➤ Allow Manage Server** set to Yes
- **Client Features ➤ Unmanaged Servers ➤ Allow Scan Network** set to Yes
- Read access must be set to customer `Opsware`

SA-OO Integration Permissions

The following permissions are required to perform the SA-OO integration actions:

Table 6 OO/SA Integration Permissions

User Action	Permission
Configure SA-OO integration	Administer Flow Integrations
Run flows in the SA Client as an SA user	Run Flow

Job Permissions

To manage jobs in the SA Client, you must have the permissions described in the Table 7. When you select the Edit All Jobs 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 7 Job Management Permissions

User Action	Feature Permission
Enable Approval Integration	Manage Approval Integration
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 All Jobs View All Jobs
End (Cancel) Job	Edit All Jobs View All Jobs
Delete Schedule	Edit All Jobs View All Jobs

Patch Management for Windows - Permissions

Table 8 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?

➤ In addition to the feature permissions listed in Table 8, every user action also requires the Managed Servers and Groups feature permission.

In Table 8, most of the entries in the User Action column correspond to menu items in the SA Client. In addition to feature permissions, server permissions are required on the managed servers affected by the patching operation.

➤ If the Allow Install Patch permission is set to Yes, then the Manage Patch and the Manage Windows Patch Policies permissions are automatically set to Read.

Table 8 Windows Patch Management Permissions Required for User Actions

User Action	Feature 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 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

Table 8 Windows Patch Management Permissions Required for User Actions

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)
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
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 9 lists the actions that users can perform for each Patch Management permission. Table 9 has the same data as Table 8, but is sorted by feature permission. Although it is not indicated in Table 9, the Managed Servers and Groups permission is required for all Patch Management actions.

For security administrators, Table 9 answers this question: If a user is granted a particular feature permission, what actions can the user perform?

Table 9 User Actions Allowed by Windows Patch Management Permissions

Feature 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 Device Group)	Read & Write
	Open Patch Policy (View)	N/A

Table 9 User Actions Allowed by Windows Patch Management Permissions

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)
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) Scan Patch Compliance	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 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 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 on page 225. For permissions on Solaris patch policies, see Solaris Patch Policy Management - Permissions on page 222.

Table 10 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?

➤ In addition to the feature permissions listed in Table 10, every user action also requires the Managed Servers and Groups feature permission.

In Table 10, most of the entries in the User Action column correspond to menu items in the SA Client. In addition to feature permissions, server permissions are required on the managed servers affected by the patching operation.

➤ If the Allow Install Patch permission is set to Yes, then the Manage Patch permission is automatically set to Read. If you plan to use Solaris patch policies, you should also set Manage Software Policy to Read or Read & Write. For more information, see Solaris Patch Policy Management - Permissions on page 222.

Table 10 Solaris Patch Management Permissions Required for User Actions

User Action	Feature 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 11 lists the actions that users can perform for each Solaris Patch Management permission. Table 11 has the same data as Table 10, but is sorted by feature permission. Although it is not indicated in Table 11, the Managed Servers and Groups permission is required for all Patch Management actions.

For security administrators, Table 11 answers this question: If a user is granted a particular feature permission, what actions can the user perform?

Table 11 User Actions Allowed by Solaris Patch Management Permissions

Feature 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 12 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 Customer Constraints, Folders, and Software Policies on page 16.

Table 12 Solaris Patch Policy Management Permissions Required for User Actions

User Action	Feature 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

Table 12 Solaris Patch Policy Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Detach 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
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 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 on page 220. You can use software policies with Unix patches. For more information, see Software Management Permissions on page 227.

Table 13 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?

➤ In addition to the feature permissions listed in Table 13, every user action also requires the Managed Servers and Groups feature permission.

In Table 13, most of the entries in the User Action column correspond to menu items in the SA Client. In addition to feature permissions, server permissions are required on the managed servers affected by the patching operation.

➤ If the Allow Install Patch permission is set to Yes, then the Manage Patch permission is automatically set to Read. If you plan to use policies, you should also set Manage Software Policy to Read or Read & Write.

Table 13 Unix Patch Management Permissions Required for User Actions

User Action	Feature 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 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

Table 13 Unix Patch Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)
Export Patch	or Manage Policy: Read and Package	N/A
Delete Patch	Manage Patch: Read & Write	N/A

Table 14 lists the actions that users can perform for each Patch Management permission. Table 14 has the same data as Table 13, but is sorted by feature permission. Although it is not indicated in Table 14, the Managed Servers and Groups permission is required for all Patch Management actions.

For security administrators, Table 14 answers this question: If a user is granted a particular feature permission, what actions can the user perform?

Table 14 User Actions Allowed by Unix Patch Management Permissions

Feature 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

Table 14 User Actions Allowed by Unix Patch Management Permissions (cont'd)

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)
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

Software Management Permissions

Table 15 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 Customer Constraints, Folders, and Software Policies on page 16.

To install software, you must belong to a user group that has the install software feature permissions. This user group must also have folder permissions for the software you want to install.

Table 15 Software Management Permissions Required for User Actions

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Software Policy			
Create Software Policy	Manage Software Policy: Read & Write	N/A	Write
Delete Software Policy	Manage Software Policy: Read & Write	N/A	Write

Table 15 Software Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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
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

Table 15 Software Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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 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

Table 15 Software Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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: Read	N/A	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

Table 15 Software Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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 Permissions	N/A	N/A	Edit Folder Permissions
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) 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

Table 15 Software Management Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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 Destination Folder: Write

Table 16 lists the actions that users can perform for each Software Management permission. Table 16 has the same data as Table 15, but is sorted by feature permission. For security administrators, Table 16 answers this question: If a user is granted a particular feature permission, what actions can the user perform?

Table 16 User Actions Allowed by Software Management Permissions

Feature 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
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

Table 16 User Actions Allowed by Software Management Permissions (cont'd)

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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
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

Table 16 User Actions Allowed by Software Management Permissions (cont'd)

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Manage Software Policy: Read And Allow Attach/Detach Software Policy: Yes And Model Public Device Groups: Yes (Required if you are attaching the software policy to a public device group)	Attach Software Policy	Read & Write	Read
	Detach Software Policy	Read & Write	Read
Manage Software Policy: Read And Allow Remediate Servers: Yes And Model Public Device Groups: Yes (Required if you remediate a public device group)	Remediate	Read & Write	Read
Manage Software Policy: Read And Allow Attach/Detach Software Policy: Yes And Allow Install/Uninstall Software: Yes And Model Public Device Groups: Yes (Required if you remediate a public device group)	Install/ Uninstall Software	Read & Write	Read

Table 16 User Actions Allowed by Software Management Permissions (cont'd)

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
Manage Software Policy: Read And Allow Run ISM Control: Yes And Model Public Device Groups: Yes (Required if you run ISM Control on a public device group)	Run ISM Control	Read & Write	Read
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

Script Execution Permissions

Table 17 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 Customer Constraints, Folders, and Software Policies on page 16.

Table 17 Script Execution Permissions Required for User Actions

User Action	Feature 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 except script contents) a Non Super User Server Script	Manage Server Script: Read	N/A	Execute
Opening (Viewing all script properties including script contents) a Non Super User Server Script	Manage Server Script: Read	N/A	Read
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	Execute
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	Read
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

Table 17   Script Execution Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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 Allow Control of Super User Server Scripts: Yes	N/A	Write
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

Table 17 Script Execution Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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
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 18 lists the actions that users can perform for each Script Execution permission. Table 18 has the same data as Table 17, but is sorted by feature permission. For security administrators, Table 18 answers this question: If a user is granted a particular feature permission, what actions can the user perform?

Table 18 User Actions Allowed by Script Execution Permissions

Feature 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 User Server Scripts: Yes	Creating a Super User Server Script	N/A	Write

Table 18 User Actions Allowed by Script Execution Permissions (cont'd)

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
	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
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	Read and Write	N/A
N/A	Running Non Super User Server Script	Read and Write	Execute

Table 18 User Actions Allowed by Script Execution Permissions (cont'd)

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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 19 Script Execution Permissions Required for Software Management

User Action	Feature 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 (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 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

Table 19 Script Execution Permissions Required for Software Management

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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
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

Audit and Remediation Permissions

Table 20 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?

➤ In addition to the feature permissions listed in Table 20, every user action also requires the Managed Servers and Groups feature permission.

Server Permissions for Audit and Remediation

Audit and Remediation actions require both feature and server feature permissions. For example, the Create Audit action requires the feature permission “Manage Audit: Read & Write” and the Managed Servers and Groups feature permission. This action also needs Read permission on the server referenced by the Audit. In Table 20, 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 SAS Web Client.

If an Audit and Remediation object (such as a Snapshot Specification) references multiple servers, at least 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. but customer and facility permissions do control access to servers which are referenced by Audit and Remediation objects, such as Snapshot Specifications and Audits.

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 20 Audit and Remediation Permissions Required for User Actions

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Snapshot Specification			
View contents of Snapshot Specification	Manage Snapshot Specification: Read	N/A	Read
Schedule and run a Snapshot Specification	Manage Snapshot Specification: Read	N/A	Read
Create Snapshot Specification	Manage Snapshot Specification: Read & Write	N/A	Read & Write
Create Application Configuration Rule	Manage Snapshot Specification: Read & Write Allow Create Task Specific Policy: Yes	Write Server File System	Read & Write

Table 20 Audit and Remediation Permissions Required for User Actions (cont'd)

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Create COM+ Rule	Manage Snapshot Specification: Read & Write Allow Create Task Specific Policy: Yes	Read COM+ Database	Read & Write
Create Custom Script Rule	Manage Snapshot Specification: Read & Write Allow Create Task Specific Policy: Yes Allow Create Custom Script Policy Rules: Yes.	Write Server File System	Read & Write
Create Files	Manage Snapshot Specification: Read & Write Allow Create Task Specific Policy: Yes	Write Server File System	Read & Write
Create IIS Metabase Rule	Manage Snapshot Specification: Read & Write Allow Create Task Specific Policy: Yes	Read IIS Metabase	Read & Write
Create Registry Rule	Manage Snapshot Specification: Read & Write Allow Create Task Specific Policy: Yes	Read Server Registry	Read & Write
Link Audit Policy into Snapshot Specification	Manage Snapshot Specification: Read & Write Manage Audit Policy: Read LIbrary Folder: Read	N/A	Read & Write
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

Table 20 Audit and Remediation Permissions Required for User Actions (cont'd)

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
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
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

Table 20 Audit and Remediation Permissions Required for User Actions (cont'd)

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
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
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
Schedule and run an Audit	Manage Audit: Read Manage Audit Result: Read & Write	N/A	Read
Create an Audit	Manage Audit: Read & Write	N/A	Read & Write
Create Application Configuration Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes	Write Server File System	Read & Write
Create COM+ Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes	Read COM+ Database	Read & Write

Table 20 Audit and Remediation Permissions Required for User Actions (cont'd)

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Create Custom Script Rule	Manage Audit: Read & Write Allow Create Custom Script Policy Rules: Yes Allow Create Task Specific Policy: Yes	Write Server File System	Read & Write
Create Discovered Software Rule	Manage Audit: Read & Write Manage Server Modules: Read Allow Create Task Specific Policy: Yes	N/A	Read & Write
Create Files Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes	Write Server File System	Read & Write
Create Hardware Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes	N/A	Read & Write
Create IIS Metabase Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes	Read IIS Metabase	Read & Write
Create Internet Information Server Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes	N/A	Read & Write
Create Registered Software Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes Manage Server Modules: Read	N/A	Read & Write
Create Runtime State Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes Manage Server Modules: Read	N/A	Read & Write

Table 20 Audit and Remediation Permissions Required for User Actions (cont'd)

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Create Software Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes	N/A	Read & Write
Create Storage Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes Manage Server Modules: Read	N/A	Read & Write
Create Weblogic Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes Manage Server Modules: Read	N/A	Read & Write
Create .NET Framework Configurations Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes Manage Server Modules: Read	N/A	Read & Write
Create Windows Registry Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes	Read Server Registry	Read & Write
Create Windows Services Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes	N/A	Read & Write
Create Windows/Unix Users and Groups Rule	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes Manage Server Modules: Read	N/A	Read & Write
Link an Audit Policy into an Audit	Manage Audit: Read & Write Allow Create Task Specific Policy: Yes Manage Audit Policy: Read SA Client Library Folder: Read	N/A	Read & Write

Table 20 Audit and Remediation Permissions Required for User Actions (cont'd)

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
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 Library Folder: Read & Write	N/A	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

Table 20 Audit and Remediation Permissions Required for User Actions (cont'd)

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
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: 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
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

Table 20 Audit and Remediation Permissions Required for User Actions (cont'd)

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
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
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

Table 20 Audit and Remediation Permissions Required for User Actions (cont'd)

User Action	Feature Permission	OGFS Permission	Server Permission (Customer, Facility, Device Group)
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 Allow Remediate Audit/ Snapshot Results: Yes	Read Server Registry	Read & Write
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 21 lists the actions that users can perform for each Audit and Remediation permission. Table 21 has the same data as Table 20, but is sorted by feature permission. Although it is not indicated in Table 21, the Managed Servers and Groups permission is required for all Audit and Remediation actions.

For security administrators, Table 21 answers this question: If a user is granted a particular feature Audit and Remediation permission, what actions can the user perform?

Table 21 User Actions Allowed by Audit and Remediation Permissions

Feature Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Allow Create Custom Script Rule Policy: No and Manage Audit: Read	View Custom Script Rule: Audit	N/A	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 and Allow Remediate Audit/ Snapshot Results: No and Manage Audit or Manage Snapshot: Read	View Audit or Snapshot, No Remediation	N/A	Read

Table 21 User Actions Allowed by Audit and Remediation Permissions (cont'd)

Feature Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Manage Snapshot Specification: Read and Allow Remediate Audit/ Snapshot Results: Yes and Manage Audit or Manage Snapshot: Read & Write	Remediate Audit/Snapshot Results	N/A	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
Manage Audit: Read & Write and Allow Create Task Specific Policy: Yes	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

Table 21 User Actions Allowed by Audit and Remediation Permissions (cont'd)

Feature Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Manage Audit: Read & Write and Allow Create Custom Script Policy Rules: Yes and Allow Create Task Specific Policy: Yes	Create Custom Scripts Rule	Write Server File System	Read & Write
Manage Audit: Read & Write and Manage Server Module: Read and Allow Create Task Specific Policy	Create the following Audit Rules: • Discovered Software • Registered Software • Runtime State • Storage • Weblogic • Windows .NET Framework Configurations • Windows Users and Groups	N/A	Read & Write
Manage Audit Results: Read	View Audit Results	N/A	Read
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

Table 21 User Actions Allowed by Audit and Remediation Permissions (cont'd)

Feature Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Manage Snapshot Specification: Read & Write and Allow Create Task Specific Policy	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
	Create Windows Registry Rule	Read Server Registry	Read & Write
Manage Snapshot Specification: Read & Write and Manage Server Module: Read and Allow Create Task Specific Policy	Create the following Snapshot Rules: • Discovered Software • Registered Software • Runtime State • Storage • Weblogic • Windows .NET Framework Configurations • Windows Users and Groups	N/A	Read & Write

Table 21 User Actions Allowed by Audit and Remediation Permissions (cont'd)

Feature Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Manage Snapshot Specification: Read & Write and Create Custom Script Policy Rule and Allow Create Task Specific Policy	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
	Create Windows Registry Rule	Read Server Registry	Read & Write

Table 21 User Actions Allowed by Audit and Remediation Permissions (cont'd)

Feature Permission	User Action	OGFS Permission	Server Permission (Customer, Facility, Device Group)
Manage Audit Policy: Read & Write Manage Server Module: Read	Create the following Snapshot Rules: • Discovered Software • Registered Software • Runtime State • 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

Service Automation Visualizer Permissions

Table 22 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 22, most of the entries in the User Action column correspond to menu items in the SA Client. In addition to feature 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.

➤ SAV permissions required to scan a server are the same for both physical servers and virtual servers.

For details on this feature, see the “Service Automation Visualizer” chapter in the *SA User’s Guide: Application Automation*.

Table 22 SAV Permissions Required for User Actions

User Action	Feature 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
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

Table 22 SAV Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Source Server Permission (Customer, Facility)	Folder Permission
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

➤ 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.

Virtual Server Permissions

Table 23 specifies the virtual server permission 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 23, most of the entries in the User Action column correspond to menu items in the SA Client.

In addition to feature permissions, server read permissions are required on hypervisor servers. After you create a new virtual server, its permissions are treated just like a physical server, including OS Provisioning.

To clone a virtual machine, you must have read-write permission to at least one customer, so the clone can be assigned to that customer. The following types of permissions are required for cloning:

- Read (or Read-Write) permissions to the source hypervisor customer

- Read (or Read-Write) permissions to the source virtual machine customer
- Read (or Read-Write) permissions to the target hypervisor customer
- Read (or Read-Write) permissions to the source hypervisor and virtual machine, and target hypervisor facility

For details on this feature, see the “Virtual Server Management” in the *SA User’s Guide: Server Automation*.

Table 23 Virtual Server Permissions Required for User Actions

User Action	Feature Permission	Hypervisor Server Permission (Customer, Facility, Device Group)
View the Virtual Servers feature in the SA Client navigation panel View virtualization object in a virtual server’s device explorer	View Virtual Servers: Yes (Note: If this permission is set to No, you will still be able to see virtual servers in the All Managed Servers list.)	Read on the hypervisor server
Refresh a hypervisor server	View Virtual Servers: Yes	Read on the hypervisor server
Create, modify, or remove a virtual server	View Virtual Servers: Yes Manage Virtual Servers: Yes	Read on the hypervisor server
Clone a virtual server	View Virtual Servers: Yes Manage Virtual Servers: Yes Customer: Yes Manage Servers and Groups: Yes	Read on the hypervisor server Read on the virtual server
Start and stop a Solaris zone	View Virtual Servers: Yes Manage Virtual Servers: Yes Administer Virtual Servers: Yes	Read on the hypervisor server
Power on, power off, suspend, or reset a VMware virtual machine	View Virtual Servers: Yes Manage Virtual Servers: Yes Administer Virtual Servers: Yes	Read on the hypervisor server
Use the ESX VM Creation wizard	Allow Configuration of Network Booting Server	Read/Write to the Facility and Customer Folder
View the user login name to VMware ESXi servers.	Manage Credentials: Read	Read on the hypervisor server

Table 23 Virtual Server Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Hypervisor Server Permission (Customer, Facility, Device Group)
Modify the user credentials, authenticate with the user credentials and validate the connection status on VMware ESXi servers.	Manage Credentials: Read & Write	Read/Write on the hypervisor server

Table 24 lists the actions that users can perform for each virtual server permission. Table 24 has the same data as Table 23, but is sorted by feature permission. For security administrators, Table 24 answers this question: If a user is granted a particular feature permission, what actions can the user perform?

Table 24 User Actions Allowed by Virtual Server Permissions

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)
View Virtual Servers: Yes Manage Virtual Servers: No	View the Virtual Servers feature in the SA Client navigation panel and virtual servers in the Contents pane. View virtualization object in a virtual server's device explorer Refresh a hypervisor server	Read on the hypervisor server
View Virtual Servers: Yes Manage Virtual Servers: Yes	Create, modify, or remove a virtual server Clone a virtual server	Read on the hypervisor server
View Virtual Servers: Yes Manage Virtual Servers: Yes Administer Virtual Servers: Yes	Start and stop a Solaris zone Power on, power off, suspend, or reset a VMware virtual machine	Read on the hypervisor server
Manage Credentials: Read	View the user login name to VMware ESXi hypervisor servers.	Read on the hypervisor server
Manage Credentials: Read & Write	Modify the user credentials, authenticate with the user credentials and validate the connection status on VMware ESXi servers.	Read/Write on the hypervisor server

OS Provisioning Permissions

The following section describes the OS Provisioning permissions required by users to perform specific actions in the SA. For security administrators, the following table answers this question: To perform a particular action, what permissions does a user need?

In Table 25, 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 SAS Web Client.

With the OS Provisioning feature in the SAS Web Client, in order to create and save an OS sequence you must save it in a folder, so you will need write permissions to the folder.

See "Customer Permissions and Folders" on page 73 in this chapter for more information.

Table 25 OS Provisioning Permissions Required for User Actions

User Action	Feature 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

Table 25 OS Provisioning Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
Add Device Group to OS Build Plan	Any of the feature permission combination 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 section) + Manage OS Build Plan: Read & Write, or 3) Manage Public Device Groups (SAS Web Client) (from Others tab, Servers and Device Group Permission section) + Manage OS Build Plan: Read & Write	None	Folder containing the OS Build Plan: 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

Table 25 OS Provisioning Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
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 + Manage OS Build Plan: Read & Write	None	Folder containing the Software Policy: Read + Folder containing the OS Build Plan: 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)	Manage Servers and Groups + Manage OS Build Plan: Read + Allow Execute OS Build Plan: Yes	Read & Write	Folder (`/Opsware /Tools/OS Provisioning`) contains the Run OS Build Plan web extension: Execute + Folder containing the OS Build Plan: Read
OS Sequence			
Create OS Sequence	Manage OS Sequence: Read & Write	None	Write
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

Table 25 OS Provisioning Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
Run OS Sequence (From server or from OS sequences)	Manage OS Sequence: Read and Allow Execute OS Sequence: Yes	Read & Write	Read
View unprovisioned servers	SA Web 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

Table 25 OS Provisioning Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permission
Attach Device Group	Any of the following combinations are valid: 1) Manage Servers and Groups + Manage OS Sequence: Read & Write 2) Manage Public Device Group (NOTE: under Client Features, Servers section) + + Manage OS Sequence: Read & Write 3) Manage Public Device Group (SAS Web Client) (NOTE: under Others tab, Server and Device Group Permissions section)	NA	Folder containing the OS Sequence: Write
OS Installation Profile			
Create, edit, delete OS installation profile	Wizard: Prepare OS	Read	N/A
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 `/Opsware /Tools/OS Provisioning/ Manage Boot Clients` folder

Table 26 lists the actions that users can perform for each OS Provisioning permission. Table 26 has the same data as Table 25, but is sorted by feature permission.

For security administrators, Table 26 answers this question: If a user is granted a particular feature permission, what actions can the user perform?

Table 26 User Actions Allowed in the SA Client by OS Provisioning Permissions

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder
Manage OS Sequence: Read	View OS sequence	Read	Read
Manage OS Sequence: Read & Write Manage OS Sequence: Read & Write	Run OS sequence	Write	Write
	Create OS sequence	Read	Write
Allow Execute OS Sequence: Yes	Run OS sequence	Write	Read
Allow Execute OS Sequence: No	View OS sequence	N/A	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
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 27 Manage Boot Client (MNC) Utility Permissions

Feature Permission	User Action	Server Permission (Customer, Facility, Device Group)	Folder
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

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 28 Compliance View Permissions Required for User Actions

User Action	Feature 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 on page 244, Table 21.	Read & Write
Software		
Remediate	Manage Software Policy: Read Allow Remediate Servers: Yes	Read & Write

Table 28 Compliance View Permissions Required for User Actions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)
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 Or Allow Install Patch: Yes Or Allow Uninstall Patch: Yes Or Allow Install/Uninstall Software Or Allow Remediate Servers	Read & Write
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 on page 207 for permissions required for remediating application configurations.	Read & Write

Server Property and Reboot Permissions

Table 29 specifies the permissions required by users to modify server properties, reboot servers, and deactivate servers. For security administrators, the table answers this question: To perform a particular action, what permissions does a user need?

Table 29 Server Property and Reboot Permissions Required for User Actions

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)
Deactivate Server	Deactivate	Read & Write
Modify Property: Server Name or Description	N/A	Read & Write
Reboot Server	Reboot Server: Yes	Read & Write

Server Objects Permission

Table 30 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 30 Server Object Permissions

User Action	Feature 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

Table 30 Server Object Permissions (cont'd)

User Action	Feature Permission	Server Permission (Customer, Facility, Device Group)	Folder Permissions
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

Predefined User Group Permissions

The following table lists the permissions of the predefined user groups for the features in the SAS Web Client. An X in a table cell indicates that the group has permission to use the feature. The headings in the table columns abbreviate the names of the user groups as follows:

- **Basic**: Basic Users
- **Inter**: Intermediate Users
- **Adv**: Advanced Users
- **OSA**: SA System Administrators
- **Admin**: Administrators

Table 31 SAS Web Client Permissions of the Predefined User Groups

Feature Name	basic	inter	adv	osa	admin
FEATURE TAB					
Configuration Tracking	X	X	X		
Configure SA				X	
Customers	X	X	X		X
DNS	X	X	X		
Data Center Intelligence Reports			X	X	
Facilities	X	X	X		X
IP Ranges and Range Groups	X	X	X		
ISM Controls	X	X	X		
Manage Gateway				X	

Table 31 SAS Web Client Permissions of the Predefined User Groups (cont'd)

Feature Name	basic	inter	adv	osa	admin
Managed Servers and Groups	X	X	X	X	
Model: Hardware	X	X	X		
Model: Opsware			X		
Model: Service Levels	X	X	X		
Multimaster				X	
Operating Systems		X	X		
Server Attributes			X	X	
Server Pool		X	X		
System Diagnosis			X	X	
Wizard: Custom Extension			X	X	
Wizard: Prepare OS	X	X	X		
OTHER TAB					
Deactivate		X	X		
Allow Run Refresh Jobs					
Manage Public Device Groups				X	
Model Public Device Groups					
View All Jobs					
Edit All Jobs					

Only the Administrator group also has permission to manage SA users and user groups, a feature not listed on the SAS Web Client tabs.

The following table lists the permissions of the predefined user groups for the SA Client features.

The table cells contain the following abbreviations:

- **R**: Read (only)
- **RW**: Read & Write
- **Y**: Yes
- **N**: No or None

Table 32 SA Client Permissions of the Predefined User Groups

Feature Name	basic	inter	adv	osa	admin
APPLICATION CONFIGURATION					
Configuration	N	R	RW	N	N
Configuration Files	N	R	RW	N	N

Table 32 SA Client Permissions of the Predefined User Groups (cont'd)

Feature Name	basic	inter	adv	osa	admin
Configuration on Servers	N	R	RW	N	N
Allow Check Consistency on Servers	N	N	Y	N	N
COMPLIANCE					
Audit Templates	N	R	RW	N	N
Audit Results	N	R	RW	N	N
Snapshot Templates	N	R	RW	N	N
Snapshots (specific to servers)	N	R	RW	N	N
Selection Criteria	N	R	RW	N	N
Allow General Snapshot Management	N	Y	Y	N	N
AGENT DEPLOYMENT					
Allow Deploy Agent	N	N	Y	N	N
Allow Scan Network	N	N	Y	N	N
PATCH MANAGEMENT					
Manage Patch	N	N	RW	N	N
Manage Patch Policy	N	N	RW	N	N
Allow Install Patch	N	N	Y	N	N
Allow Uninstall Patch	N	N	Y	N	N
Manage Patch Compliance Rules	N	N	N	N	N
SCRIPT MANAGEMENT (FRESH INSTALL - 7.0/7.5)					
Manage Server Scripts	RW	RW	RW	RW	N
Allow Control of Super User Server Scripts	N	N	Y	Y	N
Run Ad Hoc Scripts	Y	Y	Y	Y	N
Run Ad Hoc & Source Visible Server Scripts As Super User	N	Y	Y	Y	N
Manage OGFS Script	RW	RW	RW	RW	N

When HP Server Automation is first installed, default permissions are assigned to the top-level folders of the SAS Web Client. The following table lists these default permissions. The table uses the following abbreviations for permissions:

- **L**: List Contents of Folder
- **R**: Read Objects Within Folder

- **W**: Write Objects Within Folder
- **P**: Edit Folder Permissions

Table 33 Default Top-Level Folder Permissions of the Predefined User Groups

Folder	basic	inter	adv	osa	admin
/	L	L	W	L	P
/Opsware		L	L	L	P
/Opsware/Tools		L	L	L	P
/Opsware/Tools/ISMTOOL		R	W		P
/Package Repository		R	W		P
/Package Repository/All AIX		R	W		P
/Package Repository/All AIX/AIX <version>		R	W		P
/Package Repository/All HP-UX		R	W		P
/Package Repository/All HP-UX/ HP-UX <version>		R	W		P
/Package Repository/All Red Hat Linux		R	W		P
Package Repository/All Red Hat Linux/Red Hat Linux <version>		R	W		P
/Package Repository/All SunOS		R	W		P
/Package Repository/All SunOS/ SunOS <version>		R	W		P
/Package Repository/All SuSE Linux		R	W		P
/Package Repository/All SuSE Linux/SuSE Linux <version>		R	W		P
/Package Repository/All Windows		R	W		P
/Package Repository/All Windows/ Windows <version>		R	W		P

Private User Groups (PUG) Permissions

Private user groups are new as of SA 7.0. Private user groups are created for each pre-existing user during a pre-SA 7.0 core upgrade to SA 7.50 or later.

The following table lists the script management and folder permissions of the private user groups who were a member of the following pre-defined group(s) before the core was upgraded from a pre-SA 7.0 release to SA 7.50. The headings in the table columns abbreviate the names of the user groups as follows:

- **Basic**: Basic Users
- **Inter**: Intermediate Users
- **Adv**: Advanced Users
- **OSA**: SA System Administrators
- **Admin**: Administrators

Table 34 ***Script Management and folder permissions of*** **Private User Groups**

Feature Name	basic	inter	adv	osa	admin
SCRIPT MANAGEMENT (Upgrade from pre-7.0 to 7.5)					
Manage Server Scripts	RW	RW	RW	RW	N
Allow Control of Super User Server Scripts	N	N	Y	Y	N
Run Ad-hoc Scripts	Y	Y	Y	Y	N
Run Ad Hoc & Source Visible Server Scripts As Super User	N	Y	Y	Y	N
Manage Server Scripts	RW	RW	RW	RW	N
Manage OGFS Script	N	N	N	N	N
Folder permission on `/Migrated/Scripts/Shared Scripts` folder	LRX	LRX	LRXW	LRXW	LE

Code Deployment User Groups

The following tables describe the capabilities of the Code Deployment user groups. For more information, see the Accessing Code Deployment & Rollback section of the *SA User's Guide: Server Automation*.

Table 35

Code Deployment User Group	**Description**
Super User	Can define, request, or perform any code deployment operation on hosts designated for either staging or production. Because a Super User can perform operations on hosts associated with any customer, only a few users should belong to this group.
History Viewer	Can view a log of operations (service operations, synchronizations and sequences) that have been previously executed from the Code Deployment feature. Viewing this information can help you determine the status of particular deployment operations, and whether they completed successfully.

Table 36

Code Deployment User Group	**Description**
Service Editor	Can define a service, and modify or delete service definitions.
Production Service Performer	Can directly perform or request performance of service operations on hosts designated for use in production.
Staging Service Performer	Can directly perform or request performance of service operations on hosts designated for use in staging.
Production Service Requester	Can request performance of service operations on hosts designated for use in production.
Staging Service Requester	Can request performance of service operations on hosts designated for use in staging.

Table 37

Code Deployment User Group	**Description**
Synchronization Editor	Can define a synchronization, and modify or delete the synchronization definition.
Synchronization Performer	Can directly perform or request performance of a synchronization action.
Synchronization Requester	Can request performance of a synchronization action.

Table 38

Code Deployment User Group	**Description**
Sequence Editor	Can define a sequence, and modify or delete the sequence definition.
Production Sequence Performer	Can directly perform or request performance of a sequence of actions on hosts designated for use in production.
Staging Sequence Performer	Can directly perform or request performance of a sequence of actions on hosts designated for use in staging.
Production Sequence Requester	Can request performance of a sequence of actions on hosts designated for use in production.
Staging Sequence Requester	Can request performance of a sequence of actions on hosts designated for use in staging.

Index