

Oracle® Solaris Cluster Data Service for SWIFTAlliance Access Guide

SPARC Platform Edition

Copyright © 2000, 2011, Oracle and/or its affiliates. All rights reserved.

License Restrictions Warranty/Consequential Damages Disclaimer

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

Warranty Disclaimer

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

Restricted Rights Notice

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Hazardous Applications Notice

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Trademark Notice

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group in the United States and other countries.

Third Party Content, Products, and Services Disclaimer

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédé sous licence par X/Open Company, Ltd.

Contents

Preface	5
1 Installing and Configuring Oracle Solaris Cluster HA for Alliance Access	11
Overview of Installing and Configuring HA for Alliance Access	12
Oracle Solaris Cluster HA for Alliance Access Overview	12
Planning the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration	13
Configuration Restrictions	13
Configuration Requirements	14
Installing and Configuring Alliance Access	14
▼ How to Install and Configure Alliance Access	15
Verifying the Installation and Configuration of Alliance Access	19
▼ How to Verify the Installation and Configuration of Alliance Access	19
Installing the HA for Alliance Access Packages	20
▼ How to Install the HA for Alliance Access Packages	20
Registering and Configuring Oracle Solaris Cluster HA for Alliance Access	22
▼ How to Register and Configure Oracle Solaris Cluster HA for Alliance Access as a Failover Service	22
Verifying the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration ...	23
▼ How to Verify the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration	24
Understanding the Oracle Solaris Cluster HA for Alliance Access Fault Monitor	24
Resource Properties	24
Probing Algorithm and Functionality	25
Debugging Oracle Solaris Cluster HA for Alliance Access	25
▼ How to turn on debugging for Oracle Solaris Cluster HA for Alliance Access	25
Index	27

Preface

Oracle Solaris Cluster Data Service for SWIFTAlliance Access Guide explains how to install and configure Oracle Solaris Cluster data services.

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, AMD64, and Intel 64. In this document, x86 refers to the larger family of 64-bit x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

This document is intended for system administrators with extensive knowledge of Oracle software and hardware. Do not use this document as a planning or presales guide. Before reading this document, you should have already determined your system requirements and purchased the appropriate equipment and software.

The instructions in this book assume knowledge of the Oracle Solaris Operating System and expertise with the volume-manager software that is used with Oracle Solaris Cluster software.

Using UNIX Commands

This document contains information about commands that are specific to installing and configuring Oracle Solaris Cluster data services. The document does *not* contain comprehensive information about basic UNIX commands and procedures, such as shutting down the system, booting the system, and configuring devices. Information about basic UNIX commands and procedures is available from the following sources:

- Online documentation for the Oracle Solaris Operating System
- Oracle Solaris Operating System man pages
- Other software documentation that you received with your system

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-2 Shell Prompts

Shell	Prompt
Bash shell, Korn shell, and Bourne shell	\$
Bash shell, Korn shell, and Bourne shell for superuser	#
C shell	machine_name%
C shell for superuser	machine_name#

Related Documentation

Information about related Oracle Solaris Cluster topics is available in the documentation that is listed in the following table. All Oracle Solaris Cluster documentation is available at <http://www.oracle.com/technetwork/indexes/documentation/index.html>.

Topic	Documentation
Data service administration	<i>Oracle Solaris Cluster Data Services Planning and Administration Guide</i> Individual data service guides
Concepts	<i>Oracle Solaris Cluster Concepts Guide</i>
Software installation	<i>Oracle Solaris Cluster Software Installation Guide</i>
System administration	<i>Oracle Solaris Cluster System Administration Guide</i>
Hardware administration	<i>Oracle Solaris Cluster 3.3 Hardware Administration Manual</i> Individual hardware administration guides
Data service development	<i>Oracle Solaris Cluster Data Services Developer's Guide</i>
Error messages	<i>Oracle Solaris Cluster Error Messages Guide</i>
Command and function reference	<i>Oracle Solaris Cluster Reference Manual</i>

For a complete list of Oracle Solaris Cluster documentation, see the release notes for your release of Oracle Solaris Cluster at <http://www.oracle.com/technetwork/indexes/documentation/index.html>.

Related Third-Party Web Site References

Third-party URLs that are referenced in this document provide additional related information.

Note – Oracle is not responsible for the availability of third-party web sites mentioned in this document. Oracle does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Oracle will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Documentation and Support

See the following web sites for additional resources:

- [Documentation](http://www.oracle.com/technetwork/indexes/documentation/index.html) (<http://www.oracle.com/technetwork/indexes/documentation/index.html>)
- [Support](http://www.oracle.com/us/support/systems/index.html) (<http://www.oracle.com/us/support/systems/index.html>)

Oracle Software Resources

[Oracle Technology Network](http://www.oracle.com/technetwork/index.html) (<http://www.oracle.com/technetwork/index.html>) offers a range of resources related to Oracle software:

- Discuss technical problems and solutions on the [Discussion Forums](http://forums.oracle.com) (<http://forums.oracle.com>).
- Get hands-on step-by-step tutorials with [Oracle By Example](http://www.oracle.com/technetwork/tutorials/index.html) (<http://www.oracle.com/technetwork/tutorials/index.html>).

Getting Help

If you have problems installing or using Oracle Solaris Cluster, contact your service provider and provide the following information:

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the Oracle Solaris Operating System (for example, Oracle Solaris 10)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 3.3)

Use the following commands to gather information about each node on your system for your service provider.

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/clnode show-rev</code>	Displays Oracle Solaris Cluster release and package version information

Also have available the contents of the `/var/adm/messages` file.

Installing and Configuring Oracle Solaris Cluster HA for Alliance Access

This chapter explains how to install and configure HA for Alliance Access.

Note – In the current version, SWIFTAlliance Access is renamed as Alliance Access. You can use the Oracle Solaris Cluster HA for Alliance Access with both the product names. HA for Alliance Access 7.0 is supported for Oracle Solaris Cluster 3.3 only, and not for Oracle Solaris Cluster releases prior to 3.3.

This Chapter contains the following sections.

- “Overview of Installing and Configuring HA for Alliance Access” on page 12
- “Oracle Solaris Cluster HA for Alliance Access Overview” on page 12
- “Planning the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration” on page 13
- “Installing and Configuring Alliance Access” on page 14
- “Verifying the Installation and Configuration of Alliance Access” on page 19
- “Installing the HA for Alliance Access Packages” on page 20
- “Registering and Configuring Oracle Solaris Cluster HA for Alliance Access” on page 22
- “Verifying the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration” on page 23
- “Understanding the Oracle Solaris Cluster HA for Alliance Access Fault Monitor” on page 24
- “Debugging Oracle Solaris Cluster HA for Alliance Access” on page 25

Overview of Installing and Configuring HA for Alliance Access

Table 1–1 lists the tasks for installing and configuring Oracle Solaris Cluster HA for Alliance Access. Perform these tasks in the order that they are listed.

TABLE 1–1 Task Map: Installing and Configuring Oracle Solaris Cluster HA for Alliance Access

Task	For Instructions, Go To
1. Plan the installation.	“Planning the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration” on page 13
2. Install Oracle Solaris Cluster HA for Alliance Access Packages.	“How to Install and Configure Alliance Access” on page 15
3. Verify installation and configuration.	“Verifying the Installation and Configuration of Alliance Access” on page 19
4. Register and Configure Oracle Solaris Cluster HA for Alliance Access.	“Registering and Configuring Oracle Solaris Cluster HA for Alliance Access” on page 22
5. Verify Oracle Solaris Cluster HA for Alliance Access Installation and Configuration.	“Verifying the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration” on page 23
6. Understand Oracle Solaris Cluster HA for Alliance Access fault monitor.	“Understanding the Oracle Solaris Cluster HA for Alliance Access Fault Monitor” on page 24
7. Debug Oracle Solaris Cluster HA for Alliance Access.	“Debugging Oracle Solaris Cluster HA for Alliance Access” on page 25

Oracle Solaris Cluster HA for Alliance Access Overview

The HA agent is written to work with Alliance Access versions 6.0, 6.2, 6.3, and 7.0. Alliance Access is a trademark of SWIFT.

The Oracle Solaris Cluster HA for Alliance Access data service provides a mechanism for orderly startup, shutdown, fault monitoring, and automatic failover of the Oracle Solaris Cluster service. The Oracle Solaris Cluster components protected by the Oracle Solaris Cluster HA for Alliance Access data service are the following.

TABLE 1–2 Protection of Components

Component	Protected by
Alliance Access	Oracle Solaris Cluster HA for Alliance Access

Note – The fault monitoring for Alliance Access is switched off by default. If the Alliance Access application fails, the agent will not restart the Alliance Access application automatically. This behavior was explicitly requested by SWIFT. It will enable you to operate the application in a way that the probe does not interfere with the normal behavior of some of the following Alliance Access features:

- Operator manually triggering the Alliance Access restart function , for example, to run Alliance Access in housekeeping mode.
- Automatic or scheduled Alliance Access restart, for example, to run database backup and other maintenance or end-of-day processes.
- Any graceful Alliance Access restart or recovery, in case of a Alliance Access transient local error.

The HA agent provides the start, stop, takeover, and switchover functionality. This means that when a node fails, the other node will automatically start the Alliance Access application. The HA agent also provides an option to turn on fault monitoring for Alliance Access at registration time. However, this option is not recommended by SWIFT.

Planning the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration

This section contains the information you need to plan your Oracle Solaris Cluster HA for Alliance Access installation and configuration.

Configuration Restrictions

This section provides a list of software and hardware configuration restrictions that apply to Oracle Solaris Cluster HA for Alliance Access only.



Caution – Your data service configuration might not be supported if you do not observe these restrictions.

- You can configure the Oracle Solaris Cluster HA for Alliance Access *only as a HA agent and not as a scalable agent*.
- You can install the Alliance Access software on a global file system. Best practice is to use a failover file system.
- Only one Alliance Access instance is supported by this agent.
- You can install Alliance Access 7.0 only with the embedded database option.
- You can upgrade to Alliance Access 7.0 only with the non-root installation procedure.

For restrictions that apply to all data services, see the *Oracle Solaris Cluster Release Notes*.

Configuration Requirements

These requirements apply to Oracle Solaris Cluster HA for Alliance Access only. You must meet these requirements before you proceed with your Oracle Solaris Cluster HA for Alliance Access installation and configuration. Follow the Alliance Access installation guide for the installation of the mandatory patch levels and the installation of the software.



Caution – Your data service configuration might not be supported if you do not adhere to these requirements.

Oracle Solaris Cluster components and their dependencies

Configure the Oracle Solaris Cluster HA for Alliance Access data service to protect a Oracle Solaris Cluster instance and its respective components. These components, and their dependencies, are briefly described next.

Component	Description
Alliance Access	<p>→ <i>SUNW.LogicalHost</i> resource</p> <p>→ <i>SUNW.HAStoragePlus</i> resource</p> <p>The <i>SUNW.HAStoragePlus</i> resource manages the Alliance Access System Mount points and ensures that Oracle Solaris Cluster is not started until these are mounted.</p>

The Oracle Solaris Cluster component has two configuration and registration files in the `/opt/SUNWscsaa/util` directory. These files enable you to register the Oracle Solaris Cluster component with Solaris Cluster.

Within these files, the appropriate dependencies have already been defined. You must update the `saa_config` file before you run the `saa_register` script.

Installing and Configuring Alliance Access

This section describes the procedure to install and configure Alliance Access.

References will be made to some user-accessible directories for Alliance Access throughout the following sections.

Note – HA for Alliance Access can be configured to run in a whole root or a sparse root non-global zone for Alliance Access version 6.0, 6.2, 6.3, and 7.0 if required.

▼ How to Install and Configure Alliance Access

Use this procedure to install and configure Alliance Access.

1 Create the resources for Alliance Access.

■ Create a resource group for Alliance Access.

```
# clresourcegroup create [-n node-zone-list] swift-rg
```

```
-n node-zone-list
```

Specifies a comma-separated, ordered list of zones that can master the resource group.

The format of each entry in the list is node. In this format, node specifies the node name and zone specifies the name of a non-global Solaris zone. To specify the global zone, or to specify a node without non-global zones, specify only node. This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

■ Create a logical host.

Add the hostname and IP address in the `/etc/inet/hosts` file on all cluster nodes or zones that can master the resource group. Register the logical host and add it to the resource group.

```
# clreslogicalhostname create -g swift-rg -l swift-lh swift-saa-lh-rs
```

■ Create the device group and file system.

See “[Planning the Global Devices, Device Groups, and Cluster File Systems](#)” in *Oracle Solaris Cluster Software Installation Guide* for instructions on how to create global file systems.

■ Create an HAStoragePlus resource.

It is recommended to create a HAStoragePlus failover resource to contain the Alliance Access application and configuration data instead of using the global file system.

```
# clresource create -g swift-rg \  
-t SUNW.HAStoragePlus \  
-x FilesystemMountPoints=/global/saadg/alliance swift-ds
```

■ Bring the resource group online.

```
# clresouregroup online -M swift-rg
```

■ Create the configuration directory.

This directory contains Alliance Access information and creates a link from the `/usr`

```
# cd /global/saadg/alliance
```

```
# mkdir swa
```

```
# ln -s /global/saadg/alliance/swa /usr/swa
```

Note – If you install Alliance Access in a sparse root zone, that is if the `/usr` directory is inherited in read-only mode through a loopback mount, the link needs to be created within the global zone.

- For Alliance Access 7.0:

```
# mkdir -p /global/saadg/alliance/var/opt/swift
```

On all cluster nodes or zones that can master the resource group:

```
# ln -s /global/saadg/alliance/var/opt/swift /var/opt/swift
```

2 Install Alliance Access software.

Perform the following steps on all cluster nodes or zones that can master the resource group. The steps vary between different versions of Alliance Access. You must perform the steps corresponding to the version of Alliance Access you are using.

- For Alliance Access 6.0 only: Create the users `all_admin`, `all_usr` and the group `alliance` on all cluster nodes or zones that can master the resource group with the same user ID and group ID.

```
# groupadd -g groupid alliance
```

```
# useradd -m -g alliance -d /export/home/all_admin -s /usr/bin/ksh all_admin
```

```
# useradd -m -g alliance -d /export/home/all_usr -s /usr/bin/ksh all_usr
```

Create a project called `swift` and assign the users `all_admin` and `all_usr` to it.

```
# projadd -U all_admin,all_usr swift
```

- For Alliance Access 6.2, 6.3, and 7.0: Create the user `all_admin` and the groups `alliance` and `sagsnlg` on all cluster nodes or zones that can master the resource group with the same user ID and group ID. Also, create a project called `swift` and assign the users `all_admin` to it.

```
# groupadd -g groupid alliance
```

```
# groupadd -g groupid sagsnlg
```

```
# useradd -m -g alliance -G sagsnlg -d /export/home/all_admin -s \
/usr/bin/ksh all_admin
```

```
# projadd -U all_admin swift
```

- Set the values of the resource controls for the project `swift`.

For Alliance Access 6.0 only:

```
# projmod -s -K "project.max-sem-ids=(privileged,128,deny)" swift
```

```
# projmod -s -K "process.max-sem-nsems=(privileged,512,deny)" swift
```

```
# projmod -s -K "process.max-sem-ops=(privileged,512,deny)" swift
```

```
# projmod -s -K "project.max-shm-memory=(privileged,4294967295,deny)" swift
```



```
# projmod -s -K "project.max-shm-ids=(privileged,128,deny)" swift
# projmod -s -K "process.max-msg-qbytes=(privileged,4194304,deny)" swift
# projmod -s -K "project.max-msg-ids=(privileged,500,deny)" swift
# projmod -s -K "process.max-sem-messages=(privileged,8192,deny)" swift
```

For Alliance Access 6.2 and 6.3 only:

```
# projmod -s -K "project.max-sem-ids=(privileged,1320,deny)" swift
# projmod -s -K "project.max-shm-ids=(privileged,1500,deny)" swift
# projmod -s -K "project.max-shm-memory=(privileged,4294967295,deny)" swift
# projmod -s -K "project.max-msg-ids=(privileged,800,deny)" swift
# projmod -s -K "process.max-sem-nsems=(privileged,512,deny)" swift
# projmod -s -K "process.max-sem-ops=(privileged,512,deny)" swift
# projmod -s -K "process.max-msg-qbytess=(privileged,10485760,deny)" swift
# projmod -s -K "process.max-msg-messages=(privileged,8192,deny)" swift
# projmod -s -K "process.max-stack-size=(basic,33554432,deny)" swift
# projmod -s -K "process.max-data-size=(basic,8.0EB,deny)" swift
# projmod -s -K "process.max-file-descriptor=(basic,1000,deny)" swift
```

For Alliance Access 7.0 only:

```
# chown -R all_adm:alliance /global/saadg/alliance
# projmod -s -K "project.max-sem-ids=(privileged,1320,deny)" swift
# projmod -s -K "project.max-shm-ids=(privileged,1500,deny)" swift
# projmod -s -K "project.max-shm-memory=(privileged,4294967295,deny)" swift
# projmod -s -K "project.max-msg-ids=(privileged,800,deny)" swift
# projmod -s -K "process.max-sem-nsems=(privileged,512,deny)" swift
# projmod -s -K "process.max-sem-ops=(privileged,512,deny)" swift
# projmod -s -K "process.max-msg-qbytess=(privileged,10485760,deny)" swift
# projmod -s -K "process.max-msg-messages=(privileged,8192,deny)" swift
# projmod -s -K "process.max-stack-size=(basic,33554432,deny)" swift
# projmod -s -K "process.max-data-size=(basic,8.0EB,deny)" swift
# projmod -s -K "process.max-file-descriptor=(basic,4096,deny)" swift
```

The previous values are examples. For more accurate values, refer to the latest SWIFT documentation release notes of the corresponding version.

- For Alliance Access 6.0 only:

Assign the project swift as the default project for all_adm and all_usr by editing the file /etc/user_attr and typing the following two lines at the end of the file.

```
all_adm:::project=swift
```

```
all_usr:::project=swift
```

For Alliance Access 6.2, 6.3, and 7.0 only:

Assign the project `swift` as the default project for `all_adm` by editing the file `/etc/user_attr` and typing the following line at the end of the file.

```
all_adm:::project=swift
```

- For Alliance Access 6.0 only:

Repeat the installation of the software on the other node or zone that can master the resource group, but you must end the installation before the Alliance Access software licensing step.

3 For Alliance Access 6.0 only: Continue configuring Alliance Access application.

To enable clients to connect to the failover IP address, create a file named `interfaces.rpc` on the data subdirectory of the Alliance Access software.

If you are using the same file system as shown in the examples, this directory will be `/global/saadg/alliance/data`. This file must contain the IP address of the logical host as configured within the Alliance Access resource group.

```
# cd /global/saadg/alliance/data
```

```
# chown all_adm:alliance interfaces.rpc
```

If Alliance Messenger is licensed, create a file called `interfaces.mas` and add the cluster logical IP address used to communicate with SAM.

```
# cd /global/saadg/alliance/data
```

```
# chown all_adm:alliance interfaces.mas
```

4 Add a symbolic link and entries.

- Add the symbolic link `/usr/swa` on all cluster nodes or zones that can master the resource group, see the last bullet in [Step 1](#).
- For Alliance Access 7.0 only: add the symbolic link `/var/opt/swift` on all cluster nodes or zones that can master the resource group. See the last bullet in [Step 1](#).
- For Alliance Access 6.0, 6.2, and 6.3: entries in `/etc/services` have to be added on all cluster nodes or zones that can master the resource group. This can be done as root by running the `/usr/swa/apply_alliance_ports` script.

For Alliance Access 7.0: as the root user, run the `saa_rootpostinstall.ksh` script on all cluster nodes or zones that can master the resource group:

```
# /global/saadg/alliance/install/saa_rootpostinstall.ksh
```

- For Alliance Access 6.0 only:

The `rc.alliance` and `rc.swa_boot` scripts in `/etc/init.d` must remain in place. Any references to these files in `/etc/rc?.d` need to be removed and the access rights must be as follows:

```
# cd /etc/init.d
# chmod 750 rc.alliance rc.swa_boot
# chown root:sys rc.alliance rc.swa_boot
```

If the Alliance Access Installer displays “Start this SWIFTAlliance at Boottime”, choose **No**.

You must copy the `rc.alliance` and `rc.swa_boot` scripts to all other cluster nodes or zones that can master the resource group:

```
# scp rc.alliance rc.swa_boot node2:/etc/init.d
```

Note – You must not configure to automatically start at boot time through the `saa_configbootstrap` command for Alliance Access 6.2 and 6.3.

5 Install Alliance Access Remote API (RA).

- Install RA after Alliance Access on shared storage using the following options:
Instance RA1 (default), user `all_admin`
- Alliance Access 6.0 only:
Copy the files in the home directory of the `all_admin` and `all_usr` user to all cluster nodes or zones that can master the resource group.
- Alliance Access 6.2 and 6.3 only:
Copy the files in the home directory of the `all_admin` user to all cluster nodes or zones that can master the resource group. Copy the `root/InstallShield` directory to all cluster nodes or zones that can master the resource group.

Verifying the Installation and Configuration of Alliance Access

This section contains the procedure you need to verify the installation and configuration.

▼ How to Verify the Installation and Configuration of Alliance Access

This procedure does not verify that your application is highly available because you have not yet installed your data service.

- 1 **Log in as user `all_admin` into the node or zone where the `swift-rg` resource group is online. Ensure that you can display X11 clients. The profile of user `all_admin` starts the GUI automatically.**

- 2 **Start the Alliance Access application.**
 - a. **Within the GUI, select ALLiance and Start ALLiance Servers.**
 - b. **Verify that the Alliance Access instance started.**
- 3 **Stop the Alliance Access application.**
 - a. **Within the GUI, select ALLiance and Stop ALLiance Servers.**
 - b. **Verify that the Alliance Access instance stopped.**

Installing the HA for Alliance Access Packages

If you did not install the HA for Alliance Access packages during your initial Oracle Solaris Cluster installation, perform this procedure to install the packages. To install the packages, use the `installer` program.

Note – You need to install the HA for Alliance Access packages in the global cluster and not in the zone cluster.

▼ How to Install the HA for Alliance Access Packages

Perform this procedure on each cluster node where you want the HA for Alliance Access software to run.

You can run the `installer` program with a command-line interface (CLI) or with a graphical user interface (GUI). The content and sequence of instructions in the CLI and the GUI are similar.

Note – Even if you plan to configure this data service to run in non-global zones, install the packages for this data service in the global zone. The packages are propagated to any existing non-global zones and to any non-global zones that are created after you install the packages.

Before You Begin Ensure that you have the Oracle Solaris Cluster installation media.

If you intend to run the `installer` program with a GUI, ensure that your `DISPLAY` environment variable is set.

- 1 **On the cluster node where you are installing the data service packages, become superuser.**

2 Load the Oracle Solaris Cluster installation media into the DVD-ROM drive.

If the Volume Management daemon `vol(1M)` is running and configured to manage DVD-ROM devices, the daemon automatically mounts the DVD-ROM on the `/cdrom` directory.

3 Change to the installation wizard directory of the DVD-ROM.

- **If you are installing the data service packages on the SPARC platform, type the following command:**

```
# cd /cdrom/cdrom0/Solaris_sparc
```

- **If you are installing the data service packages on the x86 platform, type the following command:**

```
# cd /cdrom/cdrom0/Solaris_x86
```

4 Start the installation wizard.

```
# ./installer
```

5 When you are prompted, accept the license agreement.**6 From the list of Oracle Solaris Cluster agents under Availability Services, select the data service for Alliance Access.****7 If you require support for languages other than English, select the option to install multilingual packages.**

English language support is always installed.

8 When prompted whether to configure the data service now or later, choose Configure Later.

Choose Configure Later to perform the configuration after the installation.

9 Follow the instructions on the screen to install the data service packages on the node.

The installation wizard displays the status of the installation. When the installation is complete, the wizard displays an installation summary and the installation logs.

10 (GUI only) If you do not want to register the product and receive product updates, deselect the Product Registration option.

The Product Registration option is not available with the CLI. If you are running the installation wizard with the CLI, omit this step.

11 Exit the installation wizard.

- 12 Unload the installation media from the DVD-ROM drive.
 - a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.
 - b. Eject the DVD-ROM.

```
# eject cdrom
```

Registering and Configuring Oracle Solaris Cluster HA for Alliance Access

This section contains the procedures you need to configure Oracle Solaris Cluster HA for Alliance Access.

▼ How to Register and Configure Oracle Solaris Cluster HA for Alliance Access as a Failover Service

This procedure assumes that you installed the data service packages during your initial Oracle Solaris Cluster installation.

Steps 1 to 6 will normally already be done in order to prepare for the installation of the Alliance Access software. See [“How to Install and Configure Alliance Access” on page 15](#). Typically, you should go directly to step 7.

- 1 Become superuser on one of the nodes in the cluster that will host Oracle Solaris Cluster.
- 2 Register the SUNW.gds resource type.

```
# clresourcetype register SUNW.gds
```
- 3 Register the SUNW.HAStoragePlus resource type.

```
# clresourcetype register SUNW.HAStoragePlus
```
- 4 Create a failover resource group.

```
# clresourcegroup create [-n node-zone-list] swift-rg
```

`-n node-zone-list` Specifies a comma-separated, ordered list of zones that can master the resource group. The format of each entry in the list is `node`. In this format, `node` specifies the node name and `zone` specifies the name of a non-global Solaris zone. To specify the global zone, or to specify a node without non-global zones, specify only `node`. This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

5 Create a resource for the Oracle Solaris Cluster Disk Storage.

```
# clresource create -g swift-rg \
-t SUNW.HAStoragePlus \
-x FilesystemMountPoints=/global/saadg/alliance swift-ds
```

6 Create a resource for the Oracle Solaris Cluster Logical Hostname.

```
# clreslogicalhostname create -g swift-rg\
-h swift-lh swift-lh-rs
```

7 Enable the failover resource group that now includes the Oracle Solaris Cluster Disk Storage and Logical Hostname resources.

```
# clresourcegroup online -M swift-rg
```

8 Create a resource for Alliance Access.

a. Before running this script, check that the names of the resources match what is configured in `/opt/SUNWscsaa/util/saa_config`.

```
# /opt/SUNWscsaa/util/saa_register
```

b. Run the registration script provided as part of the Alliance Access HA agent.

9 Start the Alliance Access instance manually.

```
su - all_adm
The GUI will open up. From within the GUI, select the menu
Alliance - Start Alliance Servers
```

10 Stop the Alliance Access manually.

```
su - all_adm
The GUI will come up. Stop the application from within the GUI.
```

11 Enable each Oracle Solaris Cluster resource.

```
# clresource status -g swift-rg
# clresource enable swift-saa-rs
```

Verifying the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration

This section describes the procedure to verify that you have installed and configured your data service correctly.

▼ How to Verify the Oracle Solaris Cluster HA for Alliance Access Installation and Configuration

1 Become superuser on one of the nodes in the cluster that will host Oracle Solaris Cluster.

2 Ensure all the Oracle Solaris Cluster resources are online with `cluster status`.

```
# cluster status
```

For each Oracle Solaris Cluster resource that is not online, use the `clresource` command as follows.

```
# clresource enable resource-name
```

3 Run the `clresourcegroup` command to switch the Oracle Solaris Cluster resource group to another cluster node, such as `node2`.

```
# clresourcegroup switch swift-rg -h node2
```

4 Check that Alliance Access is stopped on the first node and that the application is restarted on the second node.

When using a failover file system, this should disappear on the first node and will be mounted on the second node.

Understanding the Oracle Solaris Cluster HA for Alliance Access Fault Monitor

This section describes the Oracle Solaris Cluster HA for Alliance Access fault monitor's probing algorithm or functionality, and states the conditions, messages, and recovery actions associated with unsuccessful probing.

For conceptual information on fault monitors, see the *Oracle Solaris Cluster Concepts Guide*.

Resource Properties

The Oracle Solaris Cluster HA for Alliance Access fault monitor uses the same resource properties as resource type `SUNW.gds`, refer to the `SUNW.gds(5)` man page for a complete list of resource properties used.

Probing Algorithm and Functionality

The fault monitoring for Alliance Access is switched off by default. If the Alliance Access application fails, the agent will not restart the Alliance Access application automatically. This behavior was explicitly requested by SWIFT. It will enable you to operate the application in a way that the probe does not interfere with the normal behavior of some Alliance Access features like:

- operator manually triggering the Alliance Access restart function , for example, to run Alliance Access in housekeeping mode.
- automatic or scheduled Alliance Access restart, for example, to run database backup and other maintenance or end-of-day processes.
- any graceful Alliance Access restart or recovery, in case of a Alliance Access transient local error.

The HA agent will update the resource status message to output Degraded - SAA Instance offline.

The Alliance Access application will cause a failover only when it does not start on the current node.

The HA agent provides an option to turn on fault monitoring for Alliance Access at registration time. However, this option is not recommended by SWIFT. The optional probing checks for the existence of the Alliance Access instance by calling the `alliance` command that is part of the application and by evaluating its return code. If the Alliance Access instance is not running, return code 100 is sent to `SUNW.gds`, which in turn will perform an automatic restart depending on the configuration of the resource properties.

Debugging Oracle Solaris Cluster HA for Alliance Access

▼ How to turn on debugging for Oracle Solaris Cluster HA for Alliance Access

Each Oracle Solaris Cluster component has a `DEBUG` file in `/opt/SUNWscsaa/etc` directory, where `saa` is a three-character abbreviation for the respective Oracle Solaris Cluster component.

These files enable you to turn on debugging for all Oracle Solaris Cluster instances or for a specific Oracle Solaris Cluster instance on a particular node with Oracle Solaris Cluster. If you

require debugging to be turned on for Oracle Solaris Cluster HA for Alliance Access across the entire Oracle Solaris Cluster installation, repeat this step on all nodes or zones that can master the resource group.

1 Edit the `/etc/syslog.conf` file.

Change `daemon.notice` to `daemon.debug`.

```
# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.notice;mail.crit      /var/adm/messages
*.alert;kern.err;daemon.err                   operator
#
```

Change the `daemon.notice` to `daemon.debug` and restart `syslogd`. The following output, from the command `grep daemon /etc/syslog.conf`, shows that `daemon.debug` has been set.

```
# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.debug;mail.crit      /var/adm/messages
*.alert;kern.err;daemon.err                   operator
#
# pkill -1 syslogd
#
```

2 Edit the `/opt/SUNWscsaa/etc/config` file.

Change `DEBUG=` to `DEBUG=ALL` or `DEBUG=resource`

```
# cat /opt/SUNWscsaa/etc/config
#
# Copyright 2003 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# Usage:
#     DEBUG=<RESOURCE_NAME> or ALL
#
DEBUG=ALL
#
```

Note – To turn off debugging, reverse the previous steps.

Index

A

- Alliance Access application
 - configuring, 14–19
 - installing, 14–19
 - versions supported, 12

C

- clnode command, 8
- commands
 - clresource, 24
 - clresourcegroup, 23
 - clresourcetype, 22
 - cluster status, 24
 - node information, 8
- configuration
 - directory for SAA, 15
 - registering Oracle Solaris Cluster HA for Alliance Access, 22–23
 - requirements, 14
 - verifying, 19–20
- configuring, 13–14
 - Alliance Access, 14–19

D

- debugging, how to turn on, 25–26

F

- fault monitor
 - probing algorithm and functionality in, 25
- Fault Monitor, Resource Properties, 24
- fault monitor
 - understanding, 24–25

G

- global zone, 20

H

- HA for Alliance Access
 - installing, 20–22
 - software packages, installing, 20–22
- help, 8–9

I

- installing
 - Alliance Access, 14–19
 - HA for Alliance Access, 20–22

L

- local zones, *See* non-global zones

M

messages file, 9

N

non-global zones, 20

P

packages, 20–22

prtconf -v command, 8

prtdiag -v command, 8

psrinfo -v command, 8

R

resource

- Alliance Access, 23

- logical host, 23

- SUNW.HAStoragePlus, 23

resource group, 15, 22

- logical host, 15

resource types

- SUNW.gds, 22

- SUNW.HAStoragePlus, 22

restrictions, 13–14

- zones, 20

S

show - rev subcommand, 8

showrev -p command, 8

software packages, 20–22

T

technical support, 8–9

V

/var/adm/messages file, 9

versions supported, 12

Z

zones, 20