

Oracle® Solaris Cluster Data Service for SAP Guide

Copyright © 2000, 2010, Oracle and/or its affiliates. All rights reserved.

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.

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.

If this is software or related software 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.

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 which 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.

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

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd.

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.

Contents

Preface	5
1 Installing and Configuring HA for SAP	11
HA for SAP Overview	11
Installing and Configuring HA for SAP	12
Planning the HA for SAP Installation and Configuration	14
Configuration Restrictions	14
Configuration Requirements	14
Standard Data Service Configurations	15
Configuration Considerations	16
Configuration Planning Questions	17
Packages and Support	18
Upgrading HA for SAP	19
▼ How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource	19
Preparing the Nodes and Disks	20
▼ How to Prepare the Nodes	20
Installing and Configuring SAP and Database	20
▼ How to Install SAP and the Database	20
▼ How to Install and Enable an SAP Scalable Application Server	21
▼ How to Enable Failover SAP Instances to Run in a Cluster	24
▼ How to Configure an SAP J2EE Engine Cluster With Your Oracle Solaris Cluster HA for SAP With an Application Server	27
▼ How to Configure an SAP J2EE Engine With Your Oracle Solaris Cluster HA for SAP With Central Instance	28
▼ How to Configure an SAP Web Dispatcher With Your Oracle Solaris Cluster HA for SAP	28
Configuring Your Highly Available Database	29
Where to Go From Here	29

Verifying the SAP Installation	30
▼ How to Verify SAP and the Database Installation With Central Instance	30
▼ How to Verify an SAP Failover Application Server	31
Verifying an SAP Scalable Application Server	32
Installing the HA for SAP Packages	33
▼ How to Install the HA for SAP Packages	33
Setting Up SAP on Non-Global Zones for HAStoragePlus Configuration	35
▼ How to Set Up SAP on Non-Global Zones for HAStoragePlus Configuration	35
Registering and Configuring HA for SAP	37
HA for SAP Extension Properties	37
▼ How to Register and Configure HA for SAP With Central Instance	41
▼ How to Register and Configure HA for SAP as a Failover Data Service	42
▼ How to Register and Configure HA for SAP as a Scalable Data Service	43
Setting Up a Lock File	45
▼ How to Set Up a Lock File for Central Instance or the Failover Application Server	45
▼ How to Set Up a Lock File for Scalable Application Server	46
Verifying the HA for SAP Installation and Configuration	47
▼ How to Verify HA for SAP Installation and Configuration and Central Instance	47
▼ How to Verify the Installation and Configuration of HA for SAP as a Failover Data Service	47
▼ How to Verify the Installation and Configuration of HA for SAP as a Scalable Data Service	48
Understanding HA for SAP Fault Monitor	49
HA for SAP Fault Probes for Central Instance	49
HA for SAP Fault Probes for Application Server	50
Index	53

Preface

Oracle Solaris Cluster Data Service for SAP Guide explains how to install and configure HA for SAP.

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://docs.sun.com>.

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>
Overview	<i>Oracle Solaris Cluster Overview</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://docs.sun.com>.

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, Support, and Training

See the following web sites for additional resources:

- Documentation (<http://docs.sun.com>)
- Support (<http://www.oracle.com/us/support/systems/index.html>)
- Training (<http://education.oracle.com>) – Click the Sun link in the left navigation bar.

Oracle Welcomes Your Comments

Oracle welcomes your comments and suggestions on the quality and usefulness of its documentation. If you find any errors or have any other suggestions for improvement, go to <http://docs.sun.com> and click Feedback. Indicate the title and part number of the documentation along with the chapter, section, and page number, if available. Please let us know if you want a reply.

Oracle Technology Network (<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/technology/obe/start/index.html) (<http://www.oracle.com/technology/obe/start/index.html>).
- Download [Sample Code](http://www.oracle.com/technology/sample_code/index.html) (http://www.oracle.com/technology/sample_code/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 HA for SAP

This chapter explains how to install and configure HA for SAP.

Note – If you are using the Solaris 10 OS, you can install and configure this data service to run in the non-global zone. HA for SAP is supported in non-global zones.

This chapter contains the following sections.

- “HA for SAP Overview” on page 11
- “Installing and Configuring HA for SAP” on page 12
- “Planning the HA for SAP Installation and Configuration” on page 14
- “Upgrading HA for SAP” on page 19
- “Preparing the Nodes and Disks” on page 20
- “Installing and Configuring SAP and Database” on page 20
- “Configuring Your Highly Available Database” on page 29
- “Verifying the SAP Installation” on page 30
- “Installing the HA for SAP Packages” on page 33
- “Registering and Configuring HA for SAP” on page 37
- “Setting Up a Lock File” on page 45
- “Verifying the HA for SAP Installation and Configuration” on page 47
- “Understanding HA for SAP Fault Monitor” on page 49

HA for SAP Overview

Use the information in this section to understand how HA for SAP makes SAP highly available.

For conceptual information about failover and scalable services, see the *Oracle Solaris Cluster Concepts Guide*.

HA for SAP provides fault monitoring and automatic failover for the SAP application to eliminate single points of failure in an SAP system. The following table lists the data services

that best protect SAP components in a Oracle Solaris Cluster configuration. You can configure HA for SAP as a failover application or a scalable application.

TABLE 1-1 Protection of SAP Components

SAP Component	Protected by
SAP database	The data service for the database that you are using, for example: <ul style="list-style-type: none"> ■ For the MaxDB database, the data service is Solaris Cluster HA for MaxDB. ■ For the Oracle database, the data service is HA for Oracle.
SAP central instance	HA for SAP The resource type is <code>SUNW.sap_ci</code> or <code>SUNW.sap_ci_v2</code> .
SAP application server	HA for SAP The resource type is <code>SUNW.sap_as</code> or <code>SUNW.sap_as_v2</code> .
NFS file system	HA for NFS

HA for SAP requires a functioning cluster with the initial cluster framework already installed. See the *Oracle Solaris Cluster Software Installation Guide* for details on initial installation of clusters and data service software. Register HA for SAP after you successfully install the basic components of the Oracle Solaris Cluster and SAP software.

Installing and Configuring HA for SAP

Table 1-2 lists the tasks for installing and configuring HA for SAP. Perform these tasks in the order that they are listed.

TABLE 1-2 Task Map: Installing and Configuring HA for SAP

Task	Instructions
Plan the SAP installation	“Planning the HA for SAP Installation and Configuration” on page 14
Upgrade HA for SAP	“How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource” on page 19
Prepare the nodes and disks	“How to Prepare the Nodes” on page 20
Install SAP and the database	“How to Install SAP and the Database” on page 20

TABLE 1-2 Task Map: Installing and Configuring HA for SAP (Continued)

Task	Instructions
Install the SAP application server	<p>If you plan to install the SAP application server as a failover service, go to “How to Enable Failover SAP Instances to Run in a Cluster” on page 24</p> <p>If you plan to install the SAP application server as a scalable service, go to “How to Install and Enable an SAP Scalable Application Server” on page 21</p>
Configure your highly available database	<p>“Configuring Your Highly Available Database” on page 29</p>
Verify the SAP installation	<p>If you plan to install a failover SAP application server, go to the following sections:</p> <ul style="list-style-type: none"> ■ “How to Verify SAP and the Database Installation With Central Instance” on page 30 ■ “How to Verify an SAP Failover Application Server” on page 31 <p>If you plan to install a scalable SAP application server, go to “Verifying an SAP Scalable Application Server” on page 32.</p>
Install HA for SAP packages	<p>“Installing the HA for SAP Packages” on page 33</p>
Register and configure HA for SAP	<p>If you plan to configure HA for SAP as a failover data service, go to:</p> <ul style="list-style-type: none"> ■ “How to Register and Configure HA for SAP With Central Instance” on page 41 ■ “How to Register and Configure HA for SAP as a Failover Data Service” on page 42 <p>If you plan to configure HA for SAP as a scalable data service, go to:</p> <ul style="list-style-type: none"> ■ “How to Register and Configure HA for SAP With Central Instance” on page 41 ■ “How to Register and Configure HA for SAP as a Scalable Data Service” on page 43
Set up a lock file	<p>“Setting Up a Lock File” on page 45</p>
Verify HA for SAP installation and configuration	<p>If you plan to configure HA for SAP as a failover data service, go to:</p> <ul style="list-style-type: none"> ■ “How to Verify HA for SAP Installation and Configuration and Central Instance” on page 47 ■ “How to Verify the Installation and Configuration of HA for SAP as a Failover Data Service” on page 47 <p>If you plan to configure HA for SAP as a scalable data service, go to “How to Verify the Installation and Configuration of HA for SAP as a Scalable Data Service” on page 48.</p>
Understand HA for SAP fault monitor	<p>“Understanding HA for SAP Fault Monitor” on page 49</p>

Planning the HA for SAP Installation and Configuration

This section contains the information you need to plan your HA for SAP installation and configuration.

Note – HA for SAP can be configured to run in a whole root or a sparse root non-global zone, if required.

Configuration Restrictions



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

Use the restrictions in this section to plan the installation and configuration of HA for SAP. This section provides a list of software and hardware configuration restrictions that apply to HA for SAP.

For restrictions that apply to all data services, see the release notes for your release of Oracle Solaris Cluster.

- **Limit node names as outlined in the SAP installation guide.** This limitation is an SAP software restriction.

Configuration Requirements



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

Use the requirements in this section to plan the installation and configuration of HA for SAP. These requirements apply to HA for SAP only. You must meet these requirements before you proceed with your HA for SAP installation and configuration. For information about updates to SAP profiles, see the SAP documentation at <http://service.sap.com/ha>.

For requirements that apply to all data services, see [Chapter 1, “Planning for Oracle Solaris Cluster Data Services,” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*](#).


- **After you create all the file systems for the database and for SAP software, create the mount points, and put the mount points in the `/etc/vfstab` file on all the cluster nodes.** See the SAP installation guides, *Installation of the SAP R/3 on UNIX* and *R/3 Installation on UNIX-OS Dependencies*, for details about how to set up the database and SAP file systems.

- **Create the required groups and users on all the cluster nodes.** See the SAP installation guides, *Installation of the SAP R/3 on UNIX* and *R/3 Installation on UNIX-OS Dependencies*, for details about how to create SAP groups and users.
- **Configure HA for NFS on the cluster that hosts the central instance if you plan to install some external SAP application servers.** See *Oracle Solaris Cluster Data Service for Network File System (NFS) Guide* for details about how to configure HA for NFS.
- **Install application servers on either the same cluster that hosts the central instance or on a separate cluster.** If you install and configure any application server outside of the cluster environment, HA for SAP does not perform fault monitoring and does not automatically restart or fail over those application servers. You must manually start and shut down application servers that you install and configure outside of the cluster environment.
- **Use an SAP software version with automatic enqueue reconnect mechanism capability.** HA for SAP relies on this capability. SAP 4.0 software with patch information and later releases should have automatic enqueue reconnect mechanism capability.

Standard Data Service Configurations


Use the standard configurations in this section to plan the installation and configuration of HA for SAP. HA for SAP supports the standard configurations in this section. HA for SAP might support additional configurations. However, you must contact your Enterprise Services representative for information about additional configurations.

FIGURE 1-1 Four-Node Cluster With Central Instance, Application Servers, and Database




CLUSTER 1

FIGURE 1-2 Two-Node Cluster With Central Instance, NFS, Database, and Non-HA External Application



Note – The configuration in [Figure 1-2](#) was a common configuration under previous Oracle Solaris Cluster releases. To use the Oracle Solaris Cluster software to the fullest extent, configure SAP as shown in [Figure 1-1](#) or [Figure 1-3](#).

FIGURE 1-3 Two-Node Cluster With Central Instance, Database, and Development Node



Configuration Considerations

Use the information in this section to plan the installation and configuration of HA for SAP. The information in this section encourages you to think about the impact your decisions have on the installation and configuration of HA for SAP.

Failover and Scalable Applications

- **Retrieve the latest patch for the sapstart executable.** This patch enables HA for SAP users to configure a lock file. For details about the benefits of this patch in your cluster environment, see [“Setting Up a Lock File” on page 45](#).

- **Read all the related SAP online service-system notes for the SAP software release and database that you are installing on your Oracle Solaris Cluster configuration.** Identify any known installation problems and fixes.
- **Consult SAP software documentation for memory and swap recommendations.** SAP software uses a large amount of memory and swap space.
- **Generously estimate the total possible load on nodes that might host the central instance, the database instance, and the application server, if you have an internal application server.** This consideration is especially important if you configure the cluster to ensure that the central instance, database instance, and application server will all exist on one node if failover occurs.

Scalable Applications

- **Ensure that the *SAPSIDadm* home directory resides on a cluster file system.** This consideration enables you to maintain only one set of scripts for all application server instances that run on all nodes. However, if you have some application servers that need to be configured differently (for example, application servers with different profiles), install those application servers with different instance numbers, and then configure them in a separate resource group.
- **Install the application server's directory locally on each node instead of on a cluster file system.** This consideration ensures that another application server does not overwrite the `log/data/work/sec` directory for the application server.
- **Use the same instance number when you create all application server instances on multiple nodes.** This consideration ensures ease of maintenance and ease of administration because you will only need to use one set of commands to maintain all application servers on multiple nodes.
- **Create separate scalable application server instances for each SAP logon group.**
- **Create an SAP lock file on the local instance directory.** This consideration prevents a system administrator from manually starting an application instance that is already running.

Configuration Planning Questions

Use the questions in this section to plan the installation and configuration of HA for SAP. Insert the answers to these questions into the data service worksheets in [Appendix C, “Data Service Configuration Worksheets and Examples,”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide*. See “Configuration Considerations” on page 16 for information that might apply to these questions.

- What resource groups will you use for network addresses and application resources and the dependencies between them?

- What is the logical hostname (for failover services) for clients that will access the data service?
- Where will the system configuration files reside?
 See “[Configuration Guidelines for Oracle Solaris Cluster Data Services](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for the advantages and disadvantages of placing the SAP binaries on the local file system as opposed to the cluster file system.

Packages and Support

Table 1–3 and Table 1–4 list the packages that HA for SAP supports.

TABLE 1–3 HA for SAP Packages From Oracle Solaris Cluster 3.0 7/01

Resource Type	Description
SUNW.sap_ci	Added support for failover central instance.
SUNW.sap_as	Added support for failover application servers.

The *_v2 resource types are the latest version of the resource types (RT) for HA for SAP. The *_v2 resource types are a superset of the original RTs. Whenever possible, use the latest RTs provided.

TABLE 1–4 HA for SAP Package From Oracle Solaris Cluster 3.0 12/01

Resource Type	Description
SUNW.sap_ci	Same as Oracle Solaris Cluster 3.0 7/01. See Table 1–3.
SUNW.sap_as	Same as Oracle Solaris Cluster 3.0 7/01. See Table 1–3.
SUNW.sap_ci_v2	Added the Network_resources_used resource property to the Resource Type Registration (RTR) file.
	Retained support for failover central instance.
SUNW.sap_as_v2	Added the Network_resources_used resource property to RTR file. Added support for scalable application servers.
	Retained support for failover application servers.

Upgrading HA for SAP

As [Table 1–3](#) and [Table 1–4](#) illustrate, the HA for SAP package from Oracle Solaris Cluster 3.0 7/01 does not support a scalable application server and the `Network_resources_used` resource property. Therefore, you have the following upgrade options.

- Retain (do not upgrade) the existing `SUNW.sap_ci` and `SUNW.sap_as` resource types. Choose this option if any of the following statements apply to you.
 - You cannot schedule downtime.
 - You do not want the `Network_resources_used` resource property.
 - You do not want to configure a scalable application server.
- Upgrade a resource type.

See “[How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource](#)” on page 19 for the procedure on how to upgrade a resource type.
- Convert a failover application resource to a scalable application resource.

See “[How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource](#)” on page 19 for the procedure on how to convert a failover application resource to a scalable application resource.

▼ How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource

Use this procedure to upgrade a resource type or to convert a failover application server resource to a scalable application server resource. This procedure requires that you schedule downtime.

- 1 **Disable the existing resource.**
- 2 **Delete the existing resource from the resource group.**
- 3 **Delete the existing resource type if no other resource uses it.**
- 4 **Register the new resource type.**
- 5 **Choose the next step depending on the task you are performing.**
 - If you are upgrading the resource type for the central instance, skip to [Step 7](#).
 - If you are converting a failover application server resource to a scalable application server resource, proceed to [Step 6](#).

- 6 Create the new application server resource group.
- 7 Add the scalable application resource to the resource group.

Next Steps Go to [“How to Prepare the Nodes”](#) on page 20.

Preparing the Nodes and Disks

This section contains the procedures you need to prepare the nodes and disks.

▼ How to Prepare the Nodes

Use this procedure to prepare for the installation and configuration of SAP.

- 1 Become superuser on all the nodes.
- 2 Configure the `/etc/nsswitch.conf` file so that HA for SAP starts and stops correctly in the event of a switchover or a failover.

On each node or zone that can master the logical host that runs HA for SAP, include one of the following entries for group in the `/etc/nsswitch.conf` file.

```
group:  
group: files [NOTFOUND=return] nis  
group: file [NOTFOUND=return] nisplus
```

HA for SAP uses the `su user` command to start and probe SAP. The network information name service might become unavailable when a cluster node's public network fails. Adding one of the entries for group in the `/etc/nsswitch.conf` file ensure that the `su(1M)` command does not refer to the NIS or NIS+ name services if the network information name service is unavailable.

Next Steps Go to [“How to Register and Configure HA for SAP as a Scalable Data Service”](#) on page 43.

Installing and Configuring SAP and Database

This section contains the procedures that you need to install and configure SAP and the database.

▼ How to Install SAP and the Database

Use this procedure to install SAP and the database.

- 1 **Become superuser on one of the nodes in the cluster where you are installing the central instance.**
- 2 **Install SAP binaries on a cluster file system.**

If you are installing SAP 7.0 and later versions or NW2004 SR1, use the logical host while installing the software. During configuration the SAP installer prompts you to skip all the steps that ask you to replace the physical-hostname references with logical-hostname references. These updates to file names and file contents with such references are not required, because when you install SAP 7.0 and later versions or NW2004 SR1 using the logical host, the SAP installer creates files and file contents using the logical hostname. Thus the updates are already taken care of.

Note – Before you install SAP software on a cluster file system, use the `scstat(1M)` command to verify that the Oracle Solaris Cluster software is fully operational.

- a. **For all the SAP-required kernel parameter changes, edit the `/etc/system` file on all the cluster nodes that will run the SAP application.**

After you edit the `/etc/system` file, reboot each node. See the SAP document *R/3 Installation on UNIX-OS Dependencies* for details about kernel parameter changes.

- b. **See the SAP document *Installation of the SAP R/3 on UNIX* for details about how to install the central instance, the database, and the application server instances.**

See “[How to Install and Enable an SAP Scalable Application Server](#)” on page 21 for the procedure on how to install a scalable application server in a Oracle Solaris Cluster environment.

Next Steps Go to “[How to Enable Failover SAP Instances to Run in a Cluster](#)” on page 24 or “[How to Install and Enable an SAP Scalable Application Server](#)” on page 21.

▼ **How to Install and Enable an SAP Scalable Application Server**

Use this procedure to install scalable application server instances. This procedure assumes that you installed the central instance and the database. This procedure includes additional steps for SAP 6.10, SAP 6.20, SAP 6.40, SAP 7.0 and later users to ensure that HA for SAP can manage and bring online SAP 6.10, SAP 6.20, SAP 6.40, SAP 7.0 and later services. SAP 6.10, SAP 6.20, SAP 6.40, SAP 7.0 and later versions create one `startsap` script and one `stopsap` script. Other SAP versions create one `startsap` script and one `stopsap` script for each service you create. This difference accounts for the additional steps for SAP 6.10, SAP 6.20, SAP 6.40, SAP 7.0 and later users.

To configure a J2EE engine cluster with your Oracle Solaris Cluster HA for SAP, first complete this procedure. Then complete the additional steps in [“How to Configure an SAP J2EE Engine Cluster With Your Oracle Solaris Cluster HA for SAP With an Application Server”](#) on page 27.

To configure an SAP Web dispatcher with your Oracle Solaris Cluster HA for SAP agent, first complete this procedure. Then complete the additional steps in [“How to Configure an SAP Web Dispatcher With Your Oracle Solaris Cluster HA for SAP”](#) on page 28.

Tip – The following file system layout ensures ease of use and prevents data from being overwritten.

- Cluster file systems
 - `/sapmnt/SAPSID`
 - `/usr/sap/SAPSID` -> all subdirectories except the *app-instance* subdirectory
 - `/usr/sap/SAPSID/home` -> the *SAPSIDadm* home directory
 - `/usr/sap/trans`
- Local file systems
 - `/usr/sap/local/SAPSID/app-instance`

1 Create all SAP directories on cluster file systems.

- Ensure that the central instance and the database can fail over.
- Set up the lock file on the cluster file system for the central instance to prevent a multiple startup from a different node.
 - For the procedure on how to set up a lock file on the central instance, see [“How to Set Up a Lock File for Central Instance or the Failover Application Server”](#) on page 45.
- Ensure that all application servers can use the SAP binaries on a cluster file system.

2 Install the central instance and the database on a cluster file system.

See the SAP document *Installation of the SAP R/3 on UNIX* for details about how to install the central instance and the database.

3 On all nodes that host the scalable application server, create a local directory for the `data/log/sec/work` directories and the log files for starting and stopping the application server.

Create a local directory for each new application server.

Example:

```
# mkdir -p /usr/sap/local/SUN/D03

# chown sapsidadm:sapsys /usr/sap/local/SUN/D03
```



Caution – You must perform this step. If you do not perform this step, you will inadvertently install a different application server instance on a cluster file system. The two application servers will then overwrite each other.

- 4 Set up a link to point to the local application server directory from a cluster file system, so that the application server, the start up log file, and the stop log file is installed on the local file system.

Example:

```
# ln -s /usr/sap/local/SAPSID/D03 /usr/sap/SAPSID/D03
```

- 5 Install the application server on all nodes.
- 6 If you are using SAP 6.10, SAP 6.20, or SAP 6.40, perform the following sub steps. Otherwise, go to [Step 7](#).

- a. Become user *sapsidadm*.

- b. Make a copy of the `startsap` script and the `stopsap` script, and save these files in the *SAPSIDadm* home directory. The file names that you choose specify this instance.

```
# cp /usr/sap/SAPSID/SYS/exe/run/startsap \  
$SAPSID_HOME/startsap_instance-number
```

```
# cp /usr/sap/SAPSID/SYS/exe/run/stopsap \  
$SAPSID_HOME/stopsap_instance-number
```

- 7 Edit the start script and the stop script so that the start up log file and the stop log file will be node specific under the home directories of users *sapsidadm* and *oraspsid* or *oradbsid*.

Example:

```
# vi startsap_03
```

Before:

```
LOGFILE=$R3S_LOGDIR/'basename $0.log'
```

After:

```
LOGFILE=$R3S_LOGDIR/'basename $0'_uname -n'.log
```

- 8 If you are using the MaxDB database, perform the following sub steps. Otherwise, skip to [Step 9](#).
 - a. Create a `.XUSER.62` file in the *SAPSIDadm* home directory as described in “[How to Install and Configure MaxDB](#)” in *Oracle Solaris Cluster Data Service for MaxDB Guide*.

Note – This user whose home directory contains the `.XUSER.62` file is an internal database user who has permissions to start, stop, and query the database.

- b. Verify that the `.XUSER.62` file is correctly configured by using the `-d` option of the `R3trans` command. You must be an `SAPSIDadm` user to perform this verification.**
- 9 Copy the application server (with the same `SAPSID` and the same instance number) on all nodes that run the scalable application server.**

The nodes that run the scalable application server are in the resource group `node-zone-list`, which is located in the scalable application server .
- 10 Ensure that you can start and stop the application server from each node, and verify that the log files are in the correct location. For further information about updates on SAP profiles, see the SAP documentation at <http://service.sap.com/ha>.**

Note – Update the files under the `SAPSIDadm` home directory so that they refer to the corresponding logical hosts. This ensures that the necessary environment variables are sourced.

- 11 Create the SAP logon group if you use a logon group.**

Next Steps To configure a J2EE engine cluster with your Oracle Solaris Cluster HA for SAP, go to “[How to Configure an SAP J2EE Engine Cluster With Your Oracle Solaris Cluster HA for SAP With an Application Server](#)” on page 27.

To configure an SAP Web dispatcher with your Oracle Solaris ClusterHA for SAP agent, go to “[How to Configure an SAP Web Dispatcher With Your Oracle Solaris Cluster HA for SAP](#)” on page 28.

If you are not using a J2EE engine cluster or an SAP Web dispatcher with your Oracle Solaris Cluster HA for SAP, go to “[Configuring Your Highly Available Database](#)” on page 29.

▼ **How to Enable Failover SAP Instances to Run in a Cluster**

During SAP installation, the SAP software creates files and shell scripts on the server on which you installed the SAP instance. These files and scripts use physical server names. To run the SAP software with Oracle Solaris Cluster software, replace references to a physical server with references to a network resource (logical hostname). Use this procedure to enable SAP to run in a cluster. The previous statement does not apply to SAP 7.0 and later versions and NW2004 SR1.

To configure a J2EE engine cluster with your Oracle Solaris Cluster HA for SAP with a central instance, after you complete this procedure, complete the additional steps in “[How to Configure an SAP J2EE Engine With Your Oracle Solaris Cluster HA for SAP With Central Instance](#)” on page 28.

To configure a J2EE engine cluster with your Oracle Solaris Cluster HA for SAP with an SAP Application Server, go to “[How to Configure an SAP J2EE Engine Cluster With Your Oracle Solaris Cluster HA for SAP With an Application Server](#)” on page 27.

To configure SAP Web dispatcher with your Oracle Solaris Cluster HA for SAP agent, first complete this procedure. Then complete the additional steps in “[How to Configure an SAP Web Dispatcher With Your Oracle Solaris Cluster HA for SAP](#)” on page 28.

- 1 **Make backup copies of the files you will modify in [Step 5](#) through [Step 11](#).**
- 2 **Log in to the node on which you installed the SAP software.**
- 3 **Shut down the SAP instances (central instance and application server instances) and the database.**
- 4 **If you are using SAP 6.10, SAP 6.20, or SAP 6.40, proceed to [Step 5](#). Otherwise, go to [Step 6](#).**
- 5 **Make a copy of the `startsap` script and the `stopsap` script, and save these files in the `SAPSIDadm` home directory. The file names that you choose must specify this instance.**

```
# cp /usr/sap/SAPSID/SYS/exe/run/startsap \
$SAPSID_HOME/startsap_logical-hostname_instance-number
# cp /usr/sap/SAPSID/SYS/exe/run/stopsap \
$SAPSID_HOME/stopsap_logical-hostname_instance-number
```

- 6 **If you are using SAP 7.0 and later versions or NW2004 SR1 proceed to [Step 7](#). Otherwise, go to [Step 8](#).**
- 7 **Create a wrapper script that calls the `startsap` and `stopsap` scripts and proceed to [Step 9](#).**

```
# echo "/usr/sap/SAPSID/SYS/exe/run/startsap instance-name logical-host" > \
$SAPSID_HOME/startsap_instance-number

# chmod 755 $SAPSID_HOME/startsap_instance-number

# echo "/usr/sap/SAPSID/SYS/exe/run/stopsap instance-name logical-host" > \
$SAPSID_HOME/stopsap_instance-number

# chmod 755 $SAPSID_HOME/stopsap_instance-number
```

- 8 **Become user `sapsidadm`, and then perform the following tasks to change references of physical server name to logical hostname.**

- In the *SAPSIDadm* home directory, modify all the file names that reference a physical server name.

```
# su - sapsidadm
# mkdir backup
# cp *.* backup
# for i in .*physical-hostname.*> do> mv $i `echo $i | sed \
"s/_physical-hostname/"`> done
# for i in *physical-hostname*> do> mv $i `echo $i | \
sed "s/physical-hostname/logical host/"`> done
```

- In the *SAPSIDadm* home directory, modify all the file contents, except log file contents, that reference a physical server name.
- In the SAP profile directory, modify all the file names that reference a physical server name.

9 As user *sapsidadm*, add entries for the parameter *SAPLOCALHOST*.

Add this entry to the *SAPSID_Service-StringSystem-Number_logical-hostname* profile file under the */sapmnt/SAPSID/profile* directory.

For central instance:

```
SAPLOCALHOST=ci-logical-hostname
```

This entry enables the external application server to locate the central instance by using the network resource (logical hostname).

For application server:

```
SAPLOCALHOST=as-logical-hostname
```

10 Become user *sapsidadm*, add the following entries for *enq* in the *DEFAULT.PFL* profile file under the */sapmnt/SAPSID/profile* directory.

```
rdisp/enqname=ci-logical-hostname_COK_00
```

11 Become user *orasapsid* or *oradbsid*, and then perform the following tasks.

Note – This step is not applicable for SAP 7.0 and later versions or NW2004 SR1.

- In the *oraSAPSID* or *oradbsid* home directory, modify all the file names that reference a physical server name.
 - In the *oraSAPSID* or *oradbsid* home directory, modify all the file contents, except log file contents, that reference a physical server name.
- 12 Ensure that the */usr/sap/tmp* directory owned by user *sapsidadm* and group *sapsys* exists on all nodes or zones that can master the failover SAP instance. For further information about updates to SAP profiles, see the SAP documentation at <http://service.sap.com/ha>.

Next Steps To configure a J2EE engine cluster with your Oracle Solaris Cluster HA for SAP with a Central Instance, go to [“How to Configure an SAP J2EE Engine With Your Oracle Solaris Cluster HA for SAP With Central Instance”](#) on page 28.

To configure a J2EE engine cluster with your Oracle Solaris Cluster HA for SAP with an SAP Application Server, go to [“How to Configure an SAP J2EE Engine Cluster With Your Oracle Solaris Cluster HA for SAP With an Application Server”](#) on page 27.

To configure an SAP Web dispatcher with your Oracle Solaris Cluster HA for SAP agent, go to [“How to Configure an SAP Web Dispatcher With Your Oracle Solaris Cluster HA for SAP”](#) on page 28.

If you are not using a J2EE engine cluster or an SAP Web dispatcher with your Oracle Solaris Cluster HA for SAP agent, go to [“Configuring Your Highly Available Database”](#) on page 29.

▼ How to Configure an SAP J2EE Engine Cluster With Your Oracle Solaris Cluster HA for SAP With an Application Server

After you have completed [“How to Install and Enable an SAP Scalable Application Server”](#) on page 21, perform the following steps.

Note – The SAP J2EE engine is started by the SAP dispatcher, which is under the protection of the Oracle Solaris ClusterHA for SAP. If the SAP J2EE engine fails, the SAP dispatcher will restart it.

- 1 **Using the SAP J2EE Admintool GUI, change the ClusterHosts parameter to list the logical host for the central instance and port pair under the dispatcher/Manager/ClusterManager.**

```
logical-host-ci:port
```

- 2 **Configure the `com.sap.instanceId` property in the file `j2ee-install-dir/additionalproperties`.**

```
com.sap.instanceId = logical-host-as_SAPSID_SYSNR
```

- 3 **Configure the `sapbasis.ashost` property in the file `j2ee-install-dir/server/services/security/work/R3Security.properties`.**

```
sapbasis.ashost = logical-host-as
```

Next Steps Go to [“Configuring Your Highly Available Database”](#) on page 29

▼ How to Configure an SAP J2EE Engine With Your Oracle Solaris Cluster HA for SAP With Central Instance

After you have completed “[How to Enable Failover SAP Instances to Run in a Cluster](#)” on page 24 or “[How to Enable Failover SAP Instances to Run in a Cluster](#)” on page 24, perform the following steps.

Note – The SAP J2EE engine is started by the SAP dispatcher which is under the protection of the Oracle Solaris Cluster HA for SAP. If the SAP J2EE engine fails, the SAP dispatcher will restart it.

- 1 **Using the SAP J2EE Admintool GUI, change the ClusterHosts parameter to list all logical hosts for the application server and port pair under dispatcher/Manager/ClusterManager.**

as1-1h:port;as2-1h:port ...

- 2 **Configure the `com.sap.instanceId` property in the file `j2ee-install-dir/additionalproperties`.**

`com.sap.instanceId = logical-host-ci_SAPSID_SYSNR`

- 3 **Configure the `sapbasis.ashost` property in the file `j2ee-install-dir/server/services/security/work/R3Security.properties`.**

`sapbasis.ashost = logical-host-ci`

- 4 **Configure the `host` property in the file `SDM-dir/program/config/flow.xml`.**

`host = logical-host-ci`

Next Steps Go to “[Configuring Your Highly Available Database](#)” on page 29.

▼ How to Configure an SAP Web Dispatcher With Your Oracle Solaris Cluster HA for SAP

After you have configured the SAP Web dispatcher with your Oracle Solaris Cluster HA for SAP, perform the following steps.

Note – The SAP Web dispatcher has the capability of auto restart. If the SAP Web dispatcher fails, the SAP Web dispatcher watch dog process will restart. Currently, there is no Oracle Solaris Cluster agent available for the SAP Web dispatcher.

- 1 **Ensure that SAP Web dispatcher has an instance number different than the Central Instance and the application server instances.**

For example, SAPSYSTEM = 66 is used in the profile for the SAP Web dispatcher.

- 2 **Activate the Internet Communication Frame Services manually after you install the SAP Web Application Server.**

For more information, see SAP OSS note 517484.

Next Steps Go to “[Configuring Your Highly Available Database](#)” on page 29.

Configuring Your Highly Available Database

SAP supports various databases. See the appropriate Oracle Solaris Cluster data services guide for details about how to configure the resource type, resource group, and resource for your highly available database. For example, see the following documentation for more information:

- If you plan to use MaxDB with SAP, see *[Oracle Solaris Cluster Data Service for MaxDB Guide](#)*.
- If you plan to use Oracle with SAP, see *[Oracle Solaris Cluster Data Service for Oracle Guide](#)*.

Additionally, see the appropriate Oracle Solaris Cluster data services guide and the appropriate chapter of your database installation book for details about other resource types to configure with your database. The Oracle Solaris Cluster data services guides include details about how to configure other resource types for databases. For example, if you plan to use the MaxDB database or the Oracle database, set up the SUNW.HAStoragePlus resource type. For more information, see the following sections:

- If you plan to use the MaxDB database, see “[Synchronizing the Startups Between Resource Groups and Device Groups](#)” in *[Oracle Solaris Cluster Data Services Planning and Administration Guide](#)* and “[Configuring the HAStoragePlus Resource Type to Work With HA for MaxDB](#)” in *[Oracle Solaris Cluster Data Service for MaxDB Guide](#)*.
- If you plan to use the Oracle database, see “[Synchronizing the Startups Between Resource Groups and Device Groups](#)” in *[Oracle Solaris Cluster Data Services Planning and Administration Guide](#)*.

Where to Go From Here

Go to “[How to Verify SAP and the Database Installation With Central Instance](#)” on page 30 or “[Verifying an SAP Scalable Application Server](#)” on page 32.

Verifying the SAP Installation

This section contains the procedures you need to verify the SAP installation.

▼ How to Verify SAP and the Database Installation With Central Instance

Use this procedure to verify the SAP central instance. Perform the following steps on all the potential nodes on which the central instance can run.

- 1 **Create the failover resource group to hold the network and central instance resources.**

```
# clresourcegroup create sap-ci-resource-group [-n node-zone-list]
```

Note – Use the `-n` option to the `clresourcegroup(1CL)` command to select the set of nodes on which the SAP central instance can run.

- 2 **Verify that you have added to your name service database all the network resources that you use.**
- 3 **Add a network resource (logical hostname) to the failover resource group.**

```
# clreslogicalhostname create -g sap-ci-resource-group \  
-d sap-ci-resource
```

Note – If you require a fully qualified hostname, you must specify the fully qualified name with the `-h` option and you cannot use the fully qualified form in the resource name.

- 4 **Enable the resource group.**

Run the `clresourcegroup(1CL)` command to move the resource group into a managed state and bring the resource group online.

```
# clresourcegroup online -emM sap-ci-resource-group
```

- 5 **Log in to the cluster member that hosts the resource group for the central instance.**
- 6 **Ensure that the database is running.**
- 7 **If you are you using the MaxDB database, create a `.XUSER.62` file in the `SAPSIDadm` home directory as described in “How to Install and Configure MaxDB” in *Oracle Solaris Cluster Data Service for MaxDB Guide*. Verify that the `.XUSER.62` file is correctly configured by using the `-d` option of the `R3trans` command. You must be an `SAPSIDadm` user to perform this verification. If you are not using the MaxDB database, proceed to the next step.**

- 8 Manually start the central instance.
- 9 Start the SAP GUI by using the logical hostname, and verify that SAP initializes correctly.
The default dispatcher port is 3200.
- 10 Manually stop the central instance.
- 11 Switch this resource group to another cluster member that can host the central instance.
`# clresourcegroup switch -n node sap-ci-resource-group`
- 12 Repeat the steps [Step 5](#) to [Step 10](#) until you verify startup and shutdown of the central instance on each cluster node that can host the central instance.

Next Steps Go to [“How to Verify an SAP Failover Application Server”](#) on page 31.

▼ How to Verify an SAP Failover Application Server

Use this procedure to verify SAP and the database installation for the failover application server. Perform the following steps on all the potential nodes on which the failover application server can run.

- 1 Create the failover resource group to hold the network and application server resources.

```
# clresourcegroup create sap-as-fo-resource-group
```

Note – Use the `-p` option to the `clresourcegroup` command to select the set of nodes on which the SAP application server can run.

```
# clresourcegroup create sap-as-fo-resource-group [-p node-zone-list]
```

- 2 Verify that you added to your name service database all the network resources that you use.
- 3 Add a network resource (logical hostname) to the failover resource group.

```
# clreslogicalhostname create -g sap-as-fo-resource-group \  
-d sap-as-fo-logical-hostname
```

Note – If you require a fully qualified hostname, you must specify the fully qualified name with the `-h` option and you cannot use the fully qualified form in the resource name.

4 Enable the resource group.

Run the `clresourcegroup(1CL)` command to move the resource group into a managed state and bring the resource group online.

```
# clresourcegroup online -emM sap-as-of-resource-group
```

5 Log in to the cluster member that hosts the application server resource group.**6 If you are you using the MaxDB database, proceed to [Step 7](#). Otherwise go to [Step 9](#).****7 Create a `.XUSER.62` file in the `SAPSIDadm` home directory as described in [“How to Install and Configure MaxDB”](#) in *Oracle Solaris Cluster Data Service for MaxDB Guide*.****8 Verify that the `.XUSER.62` file is correctly configured by using the `-d` option of the `R3trans` command.**

Note – You must be an `SAPSIDadm` user to perform this verification.

9 Manually start the application server.**10 Start the SAP GUI by using the logical hostname, and verify that SAP initializes correctly.****11 Manually stop the application server.****12 Switch this resource group to another cluster member that can host the application server.**

```
# clresourcegroup switch -n nodesap-as-fo-resource-group
```

13 Repeat the steps [Step 5](#) to [Step 10](#) until you verify startup and shutdown of the application server on each cluster node that can host the application server.

Next Steps Go to [“Installing the HA for SAP Packages”](#) on page 33.

Verifying an SAP Scalable Application Server

If you installed scalable application server instances in [“How to Install and Enable an SAP Scalable Application Server”](#) on page 21, you verified the installation of an SAP scalable application server in [Step 10](#) of [“How to Install and Enable an SAP Scalable Application Server”](#) on page 21.

Where to Go From Here

Go to [“Installing the HA for SAP Packages”](#) on page 33.

Installing the HA for SAP Packages

If you did not install the HA for SAP 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 SAP packages in the global cluster and not in the zone cluster.

▼ How to Install the HA for SAP Packages

Perform this procedure on each cluster node where you are installing the HA for SAP packages.

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 `vold(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 SAP.

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
```

Next Steps Go to [“Registering and Configuring HA for SAP”](#) on page 37.

Setting Up SAP on Non-Global Zones for HAStoragePlus Configuration

Use the following procedure to configure a HAStoragePlus resource for non-global zones.

Note –

- The entries in the `/etc/vfstab` file for cluster file systems should contain the global keyword in the mount options.
 - The SAP binaries that will be made highly available using the HAStoragePlus resource should be accessible from the non-global zones.
 - In non-global zones, file systems that are used by different resources in different resource groups must reside in a single HAStoragePlus resource that resides in a scalable resource group. The nodelist of the scalable HAStoragePlus resource group must be a superset of the nodelists of the application resource groups that have resources which depend on the file systems. These application resources that depend on the file systems must have a strong resource dependency set to the HAStoragePlus resource. In addition, the dependent application resource group must have a strong positive resource group affinity set to the scalable HAStoragePlus resource group.
-

▼ How to Set Up SAP on Non-Global Zones for HAStoragePlus Configuration

- 1 On any node in the cluster, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Create the scalable resource group with non-global zones that contain the HAStoragePlus resource.

```
# clresourcegroup create \  
-p Maximum primaries=m \  
-p Desired primaries=n \  
[-n node-zone-list] hasp-resource-group
```

- | | |
|------------------------|---|
| -p Maximum primaries=m | Specifies the maximum number of active primaries for the resource group. |
| -p Desired primaries=n | Specifies the number of active primaries on which the resource group should attempt to start. |
| -n node-zone-list | In the node list of a HAStoragePlus resource group, specifies the list of nodename: zonename pairs as the node list of the HAStoragePlus resource group, where the SAP instances can come online. |

hasp-resource-group Specifies the name of the scalable resource group to be added. This name must begin with an ASCII character.

3 Register the resource type for the HAStoragePlus resource.

```
# clresourcetype register HAStoragePlus
```

4 Create the HAStoragePlus resource *hasp-resource* and define the SAP filesystem mount points and global device paths.

```
# clresource create -g hasp-resource-group -t SUNW.HAStoragePlus \
-p GlobalDevicePaths=/dev/global/dsk/d5s2,dsk/d6 -p affinityon=false \
-p FilesystemMountPoints=/sapmnt/JSC,/usr/sap/trans,/usr/sap/JSC hasp-resource
```

-g hasp-resource-group Specifies the resource group name.

GlobalDevicePaths Contains the following values:

- Global device group names, such as *sap-dg*, *dsk/d5*
- Paths to global devices, such as */dev/global/dsk/d5s2*, */dev/md/sap-dg/dsk/d6*

FilesystemMountPoints Contains the following values:

- Mount points of local or cluster file systems, such as */local/mirrlogA*, */local/mirrlogB*, */sapmnt/JSC*, */usr/sap/JSC*

The HAStoragePlus resource is created in the enabled state.

5 Register the resource type for the SAP application.

```
# clresourcetype register resource-type
```

resource-type Specifies the name of the resource type to be added. For more information, see the release notes for your release of Oracle Solaris Cluster.

6 Create a SAP resource group.

```
# clresourcegroup create [-n node-zone-list] -p
RG_affinities==+hastorageplus-rg resource-group-1
```

resource-group-1 Specifies the SAP services resource group.

7 Add the SAP application resource to *resource-group-1* and set the dependency to *hastorageplus-1*.

```
# clresource create -g resource-group-1 -t SUNW.application \
[-p "extension-property[{node-specifier}]"=value, ?] \
-p Resource_dependencies=hastorageplus-1 resource
```

8 Bring the failover resource group online.

```
# clresourcegroup online resource-group-1
```

Registering and Configuring HA for SAP

This section contains the procedures that you need to configure HA for SAP.

HA for SAP Extension Properties

Use the extension properties in “[HA for SAP Extension Properties for the Central Instance](#)” on page 37 and “[HA for SAP Extension Properties for the Application Servers](#)” on page 40 to create your resources. Use the command line `clresource create -g resource-group -t resource-type -p property=value resource` to configure extension properties when you create your resource. Use the procedure in Chapter 2, “[Administering Data Service Resources](#),” in *Oracle Solaris Cluster Data Services Planning and Administration Guide* to configure the extension properties if you have already created your resources. You can update some extension properties dynamically. You can update others, however, only when you create or disable a resource. The Tunable entries indicate when you can update each property. See [Appendix A, “Standard Properties,”](#) in *Oracle Solaris Cluster Data Services Planning and Administration Guide* for details about all Oracle Solaris Cluster properties.

HA for SAP Extension Properties for the Central Instance

SAP configuration extension properties are as follows:

SAPSID

SAP system ID or SAPSID.

Default: None

Tunable: When disabled

Ci_instance_id

Two-digit SAP system number.

Default: 00

Tunable: When disabled

Ci_services_string

String of central instance services.

Default: DVEBMGS

Tunable: When disabled

Starting SAP extension properties are as follows:

`Ci_start_retry_interval`

The interval in seconds to wait between attempting to connect to the database before starting the central instance.

Default: 30

Tunable: When disabled

`Ci_startup_script`

Name of the SAP startup script for this instance in your *SAPSIDadm* home directory.

Default: None

Tunable: When disabled

Stopping SAP extension properties are as follows:

`Stop_sap_pct`

Percentage of stop-timeout variables that are used to stop SAP processes. The SAP shutdown script is used to stop processes before calling Process Monitor Facility (PMF) to terminate and then kill the processes.

Default: 95

Tunable: When disabled

`Ci_shutdown_script`

Name of the SAP shutdown script for this instance in your *SAPSIDadm* home directory.

Default: None

Tunable: When disabled

Probe extension properties are as follows:

`Message_server_name`

The name of the SAP Message Server.

Default: *sapms SAPSID*

Tunable: When disabled

`Lgtst_ms_with_logicalhostname`

How to check the SAP Message Server with the SAP *lgtst* utility. The *lgtst* utility requires a hostname (IP address) as the location for the SAP Message Server. This hostname can be either a Oracle Solaris Cluster logical hostname or a local host (loopback) name. If you set this resource property to TRUE, use a logical hostname. Otherwise, use a localhost name.

Default: TRUE

Tunable: Any time

Check_ms_retry

Maximum number of times the SAP Message Server check fails before a total failure is reported and the Resource Group Manager (RGM) starts.

Default: 2

Tunable: When disabled

Probe_timeout

Timeout value in seconds for the probes.

Default: 120

Tunable: Any time

Monitor_retry_count

Number of PMF restarts that are allowed for the fault monitor.

Default: 4

Tunable: Any time

Monitor_retry_interval

Time interval in minutes for the fault monitor restarts.

Default: 2

Tunable: Any time

Development system extension properties are as follows:

Shutdown_dev

Whether the RGM should shut down the development system before starting up the central instance.

Default: FALSE

Tunable: When disabled

Dev_sapsid

SAP system name for the development system. If you set Shutdown_dev to TRUE, HA for SAP requires this property.

Default: None

Tunable: When disabled

Dev_shutdown_script

Script that is used to shut down the development system. If you set Shutdown_dev to TRUE, HA for SAP requires this property.

Default: None

Tunable: When disabled

Dev_stop_pct

Percentage of startup timeouts HA for SAP uses to shut down the development system before starting the central instance.

Default: 20

Tunable: When disabled

HA for SAP Extension Properties for the Application Servers

SAP configuration extension properties are as follows:

SAPSID

SAP system name or *SAPSID* for the application server.

Default: None

Tunable: When disabled

As_instance_id

Two-digit SAP system number for the application server.

Default: None

Tunable: When disabled

As_services_string

String of application server services.

Default: D

Tunable: When disabled

Starting SAP extension properties are as follows:

As_db_retry_interval

The interval in seconds to wait between attempting to connect to the database and starting the application server.

Default: 30

Tunable: When disabled

As_startup_script

Name of the SAP startup script for the application server.

Default: None

Tunable: When disabled

Stopping SAP extension properties are as follows:

`Stop_sap_pct`

Percentage of stop-timeout variables that are used to stop SAP processes. The SAP shutdown script is used to stop processes before calling Process Monitor Facility (PMF) to terminate and then kill the processes.

Default: 95

Tunable: When disabled

`As_shutdown_script`

Name of the SAP shutdown script for the application server.

Default: None

Tunable: When disabled

Probe extension properties are as follows:

`Probe_timeout`

Timeout value in seconds for the probes.

Default: 60

Tunable: Any time

`Monitor_retry_count`

Number of PMF restarts that the probe allows for the fault monitor.

Default: 4

Tunable: Any time

`Monitor_retry_interval`

Time interval in minutes for fault monitor restarts.

Default: 2

Tunable: Any time

▼ How to Register and Configure HA for SAP With Central Instance

Use this procedure to configure HA for SAP with central instance.

- 1 **Become superuser on one of the nodes in the cluster that hosts the central instance.**

2 Register the resource type for the central instance.

```
# clresourcetype register SUNW.sap_ci | SUNW.sap_ci_v2
```

3 Add the HASStoragePlus resource to the resource group for the central instance.

```
# clresourcetype register SUNW.HASStoragePlus
# clresource create -g sap-ci-resource-group \
-t SUNW.HASStoragePlus -p filesystemmountpoints=mountpoint,... ci-storage-resource
```

The resource is created in the enabled state.

For more details about how to set up an HASStoragePlus resource, see “Enabling Highly Available Local File Systems” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

4 Create SAP central instance resources in this failover resource group.

```
# clresource create -d \
-g sap-ci-failover-resource-group \
-t SUNW.sap_ci | SUNW.sap_ci_v2 \
-p SAPSID=SAPSID -p Ci_instance_id=ci-instance-id \
-p Ci_startup_script=ci-startup-script \
-p Ci_shutdown_script=ci-shutdown-script \
-p resource_dependencies=ci-storage-resource,db-resource sap-ci-resource
```

The resource is created in the enabled state.

See “HA for SAP Extension Properties” on page 37 for a list of extension properties.

Note – SAP application requires that the database should be online before the SAP application can start. Setting the resource dependency value to the database will ensure that it checks for the resource to be online before starting the SAP central instance. This dependency check is necessary only if the database is on the same cluster as that of the SAP central instance.

Next Steps Go to “How to Register and Configure HA for SAP as a Failover Data Service” on page 42 or “How to Register and Configure HA for SAP as a Scalable Data Service” on page 43.

▼ How to Register and Configure HA for SAP as a Failover Data Service

Use this procedure to configure HA for SAP as a failover data service.

1 Become superuser on one of the nodes in the cluster that hosts the application server.**2 Register the resource type for the failover application server.**

```
# clresourcetype register SUNW.sap_as

# clresourcetype register SUNW.sap_as_v2
```

3 Add the HASStoragePlus resource to the resource group for the failover application server.

```
# clresource register SUNW.HASStoragePlus
# clresource create -g sap-as-fo-resource-group \
-t SUNW.HASStoragePlus \
-p filesystemmountpoints=mountpoint,... sap-as-storage-resource
```

The resource is created in the enabled state.

For more details about how to set up an HASStoragePlus resource, see “[Enabling Highly Available Local File Systems](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

4 Create SAP application server resources in their failover resource group.

```
# clresource create \
-g sap-as-fo-resource-group \
-t SUNW.sap_as | SUNW.sap_as_v2 \
-p SAPSID=SAPSID -p As_instance_id=as-instance-id \
-p As_startup_script=as-startup-script \
-p As_shutdown_script=as-shutdown-script \
-p resource_dependencies=sap-as-storage-resource,db-resource sap-as-resource
```

The resource is created in the enabled state.

See “[HA for SAP Extension Properties](#)” on [page 37](#) for a list of extension properties.

Note – SAP application requires that the database should be online before the SAP application can start. Setting the resource dependency value to the database will ensure that it checks for the resource to be online before starting the SAP application server. This dependency check is necessary only if the database is on the same cluster as that of the SAP application server.

Next Steps Go to “[How to Verify HA for SAP Installation and Configuration and Central Instance](#)” on [page 47](#).

▼ How to Register and Configure HA for SAP as a Scalable Data Service

Use this procedure to configure HA for SAP as a scalable data service.

- 1 Become superuser on one of the nodes in the cluster that hosts the application server.**
- 2 Create a scalable resource group for the application server.**

```
# clresourcegroup create \
-p Maximum_primaries=value \
-p Desired_primaries=value sap-as-sa-appinstanceid-resource-group
```

Note – Solaris Cluster HA for SAP as a scalable data service does not use shared addresses because the SAP logon group performs the load balancing of the application server.

3 Register the resource type for the scalable application server.

```
# clresourcetype register SUNW.sap_as_v2
```

4 Add the HASStoragePlus resource to the failover application server resource group.

```
# clresourcetype register SUNW.HASStoragePlus
# clresource create \
-g sap-as-sa-appinstanceid-resource-group \
-t SUNW.HASStoragePlus \
-p filesystemmountpoints=mountpoint,... sap-as-storage-resource
```

The resource is created in the enabled state.

For more details on how to set up an HASStoragePlus resource, see “[Enabling Highly Available Local File Systems](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

5 Create SAP application server resources in this scalable resource group.

```
# clresource create -d \
-g sap-as-sa-appinstanceid-resource-group \
-t SUNW.sap_as_v2 \
-p SAPSID=SAPSID \
-p As_instance_id=as-instance-id \
-p As_startup_script=as-startup-script \
-p As_shutdown_script=as-shutdown-script \
-p resource_dependencies=sap-as-storage-resource,db-resource sap-as-resource
```

See “[HA for SAP Extension Properties](#)” on page 37 for a list of extension properties.

Note – SAP application requires that the database should be online before it can start. Setting the resource dependency value to the database resource will ensure that it checks for the resource to be online before starting the SAP application server. This dependency check is necessary only if the database is on the same cluster as that of the SAP application server.

6 Enable the scalable resource group that now includes the SAP application server resource.

```
# clresourcegroup online -M sap-as-sa-appinstanceid-resource-group
```

Next Steps Go to “[How to Verify HA for SAP Installation and Configuration and Central Instance](#)” on page 47.

Setting Up a Lock File

Use the procedure in this section to perform the following tasks.

- Set up a lock file for the central instance or the failover application server.
- Set up a lock file for a scalable application server.

Set up a lock file to prevent multiple startups of the SAP instance when the instance is already active on one node. Multiple startups of the same instance cause each startup to fail. Furthermore, the failure prevents SAP shutdown scripts from performing a clean shutdown of the instances, which might cause data corruption.

If you set up a lock file, when you start the SAP instance the SAP software locks the file `startup_lockfile`. If you start the same instance outside the Oracle Solaris Cluster environment and then try to start SAP under the Oracle Solaris Cluster environment, the HA for SAP data service will attempt to start the same instance. However, because of the file-locking mechanism, this attempt will fail. The data service will log appropriate error messages in `/var/adm/messages`.

The only difference between the lock file for the central instance or the failover application server and the lock file for a scalable application server is that the lock file for the scalable application server resides on the local file system and the lock file for the central instance or the failover application server resides on a cluster file system.

▼ How to Set Up a Lock File for Central Instance or the Failover Application Server

Use this procedure to set up a lock file for the central instance or the failover application server.

- 1 **Install the latest patch for the `sapstart` executable, which enables HA for SAP users to configure a lock file.**
- 2 **Set up the lock file for the central instance or the failover application server on a *cluster* file system.**
- 3 **Edit the profile that `sapstart` uses to start the instance such that you add the new SAP parameter, `sapstart/lockfile`, for the central instance or failover application server. This profile is the one that is passed to `sapstart` as a parameter in the `startsap` script.**

For the central instance, type the following.

```
sapstart/lockfile =/usr/sap/SAPSID/ \
Service-StringSystem-Number/work/startup_lockfile
```

For the failover application server, type the following.

```
sapstart/lockfile =/usr/sap/SAPSID/ Dinstance-id/work/startup_lockfile
```

`sapstart/lockfile`

New parameter name.

`/usr/sap/SAPSID/Service-StringSystem-Number/work`

Work directory for the central instance.

`/usr/sap/SAPSID/Dinstance-id/work`

Work directory for failover application server.

`startup_lockfile`

Lock file name that HA for SAP uses.

SAP creates the lock file.

Note – You must locate the lock file path on a cluster file system. If you locate the lock file path locally on the nodes, a startup of the same instance from multiple nodes cannot be prevented.

▼ How to Set Up a Lock File for Scalable Application Server

Use this procedure to set up a lock file for a scalable application server.

- 1 Install the latest patch for the `sapstart` executable, which enables HA for SAP users to configure a lock file.**
- 2 Set up the application server lock file on the *local* file system.**
- 3 Edit the profile that `sapstart` uses to start the instance such that you add the new SAP parameter, `sapstart/lockfile`, for scalable application server. This profile is the one that is passed to `sapstart` as a parameter in the `startsap` script.**

`sapstart/lockfile =/usr/sap/local/SAPSID/Dinstance-id/work/startup_lockfile`

`sapstart/lockfile`

New parameter name

`/usr/sap/local/SAPSID/Dinstance-id/work`

Work directory for the scalable application server

`startup_lockfile`

Lock file name that HA for SAP uses

SAP creates the lock file.

Note – The lock file will reside on the local file system. The lock file does not prevent multiple startups from *other* nodes, but the lock file does prevent multiple startups on the *same* node.

Verifying the HA for SAP Installation and Configuration

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

▼ How to Verify HA for SAP Installation and Configuration and Central Instance

Use this procedure to verify the HA for SAP installation and configuration and central instance.

- 1 **Log in to the node that hosts the resource group that contains the SAP central instance resource.**
- 2 **Start the SAP GUI to check that HA for SAP is functioning correctly.**
- 3 **As user *sapsidadm*, use the central instance `stopsap` script to shut down the SAP central instance.**
The Oracle Solaris Cluster software restarts the central instance.
- 4 **As user *root*, switch the SAP resource group to another cluster member.**
`# clresourcegroup switch -n node2sap-ci-resource-group`
- 5 **Verify that the SAP central instance starts on this node.**
- 6 **Repeat [Step 1](#) through [Step 5](#) until you have tested all the potential nodes on which the SAP central instance can run.**

Next Steps Go to “[How to Verify the Installation and Configuration of HA for SAP as a Failover Data Service](#)” on page 47 or “[How to Verify the Installation and Configuration of HA for SAP as a Scalable Data Service](#)” on page 48.

▼ How to Verify the Installation and Configuration of HA for SAP as a Failover Data Service

Use this procedure to verify the installation and configuration of HA for SAP as a failover data service.

- 1 Log in to the node that currently hosts the resource group that contains the SAP application server resource.
- 2 As user *sapsidadm*, start the SAP GUI to check that the application server is functioning correctly.
- 3 Use the application server `stopsap` script to shut down the SAP application server on the node that you identified in [Step 1](#).
The Oracle Solaris Cluster software restarts the application server.
- 4 As user *root*, switch the resource group that contains the SAP application server resource to another cluster member.

```
# clresourcegroup switch -n node2sap-as-resource-group
```
- 5 Verify that the SAP application server starts on the node that you identified in [Step 4](#).
- 6 Repeat [Step 1](#) through [Step 5](#) until you have tested all the potential nodes on which the SAP application server can run.

▼ How to Verify the Installation and Configuration of HA for SAP as a Scalable Data Service

Use this procedure to verify the installation and configuration of HA for SAP as a scalable data service.

- 1 Log on to one of the nodes that runs the application server.
- 2 Become user *sapsidadm*.
- 3 Start the SAP GUI to check that the application server is functioning correctly.
- 4 Use the application server `stopsap` script to shut down the SAP application server on the node that you identified in [Step 1](#).
The Oracle Solaris Cluster software restarts the application server.
- 5 Repeat [Step 1](#) through [Step 3](#) until you have tested all the potential nodes on which the SAP application server can run.

Understanding HA for SAP Fault Monitor

The HA for SAP fault monitor checks SAP process and database availability. SAP process availability impacts SAP resources' failure history. SAP resources' failure history in turn drives the fault monitor's actions, which include no action, restart, or failover.

In contrast to SAP process availability, SAP database unavailability has no impact on SAP resources' failure history. Any change in the SAP database availability does, however, trigger the SAP fault monitor to log any `syslog` messages to `/var/adm/messages` and to set the status accordingly for the SAP resources that use the database.

HA for SAP Fault Probes for Central Instance

For the central instance, the fault probe executes the following steps.

1. Retrieves the process IDs for the SAP Message Server and the dispatcher
2. Loops infinitely (sleeps for `Thorough_probe_interval`)
3. Checks the availability of the SAP resources
 - a. **Abnormal exit** – If the Process Monitor Facility (PMF) detects that the SAP process tree has failed, the fault monitor treats this problem as a complete failure. The fault monitor restarts or fails over the SAP resource to another node based on the resources' failure history.
 - b. **Availability check of the SAP resources through probe** – The probe uses the `ps(1)` command to check the SAP Message Server and main dispatcher processes. If any of the SAP Message Server or main dispatcher processes are missing from the system's active processes list, the fault monitor treats this problem as a complete failure.

If you configure the parameter `Check_ms_retry` to have a value greater than zero, the probe checks the SAP Message Server connection. If you have set the extension property `Lgtst_ms_with_logicalhostname` to its default value `TRUE`, the probe completes the SAP Message Server connection test with the utility `lgtst`. The probe uses the logical hostname interface that is specified in the SAP resource group to call the SAP-supplied utility `lgtst`. If you set the extension property `Lgtst_ms_with_logicalhostname` to a value other than `TRUE`, the probe calls `lgtst` with the node's local hostname (loopback interface).

If the `lgtst` utility call fails, the SAP Message Server connection is not functioning. In this situation, the fault monitor considers the problem to be a partial failure and does not trigger an SAP restart or a failover immediately. The fault monitor counts two partial failures as a complete failure if the following conditions occur.

- i. You configure the extension property `Check_ms_retry` to be 2.
- ii. The fault monitor accumulates two partial failures that happen within the retry interval that the resource property `Retry_interval` sets.

A complete failure triggers either a local restart or a failover, based on the resource's failure history.

- c. **Database connection status through probe** – The probe calls the SAP-supplied utility `R3trns` to check the status of the database connection. HA for SAP fault probes verify that SAP can connect to the database. HA for SAP depends, however, on the highly available database fault probes to determine database availability. If the database connection status check fails, the fault monitor logs the message, `Database might be down`, to `/var/adm/messages`. The fault monitor then sets the status of the SAP resource to `DEGRADED`. If the probe checks the status of the database again and the connection is reestablished, the fault monitor logs the message, `Database is up`, to `/var/adm/messages` and sets the status of the SAP resource to `OK`.
4. Evaluates the failure history
- Based on the failure history, the fault monitor completes one of the following actions.
- No action
 - Local restart
 - Failover

HA for SAP Fault Probes for Application Server

For the application server, the fault probe executes the following steps.

1. Retrieves the process ID for the main dispatcher
2. Loops infinitely (sleeps for `Thorough_probe_interval`)
3. Checks the availability of the SAP resources
 - a. **Abnormal exit** – If the Process Monitor Facility (PMF) detects that the SAP process tree has failed, the fault monitor treats this problem as a complete failure. The fault monitor restarts or fails over the SAP resource to another node, based on the resources' failure history.
 - b. **Availability check of the SAP resources through probe** – The probe uses the `ps(1)` command to check the SAP Message Server and main dispatcher processes. If the SAP main dispatcher process is missing from the system's active processes list, the fault monitor treats the problem as a complete failure.
 - c. **Database connection status through probe** – The probe calls the SAP-supplied utility `R3trns` to check the status of the database connection. HA for SAP fault probes verify that SAP can connect to the database. HA for SAP depends, however, on the highly available database fault probes to determine database availability. If the database connection status check fails, the fault monitor logs the message, `Database might be down`, to `/var/adm/messages` and sets the status of the SAP resource to `DEGRADED`. If the probe checks the status of the database again and the connection is reestablished, the fault monitor logs the message, `Database is up`, to `/var/adm/messages`. The fault monitor then sets the status of the SAP resource to `OK`.

4. Evaluates the failure history

Based on the failure history, the fault monitor completes one of the following actions.

- No action
- Local restart
- Failover

If the application server resource is a failover resource, the fault monitor fails over the application server.

If the application server resource is a scalable resource, after the number of local restarts are exhausted, RGM will start the application server on a different node if another node is available in the cluster.

Index

C

- `clnode` command, 9
- commands, node information, 8
- configuring
 - SAP database, 20–21
 - Solaris Cluster HA for SAP, 41–42, 42–43, 43–44

E

- extension properties
 - Solaris Cluster HA for SAP
 - `As_db_retry_interval` (application server), 40–41
 - `As_instance_id` (application server), 40–41
 - `As_services_string` (application server), 40–41
 - `As_shutdown_script` (application server), 40–41
 - `As_startup_script` (application server), 40–41
 - `Check_ms_retry` (central instance), 37–40
 - `Ci_instance_id` (central instance), 37–40
 - `Ci_services_string` (central instance), 37–40
 - `Ci_shutdown_script` (central instance), 37–40
 - `Ci_start_retry_interval` (central instance), 37–40
 - `Ci_startup_script` (central instance), 37–40
 - `Dev_shutdown_script` (central instance), 37–40
 - `Dev_stop_pct` (central instance), 37–40
 - `Lgtst_ms_with_logicalhostname` (central instance), 37–40
 - `Message_server_name` (central instance), 37–40

extension properties, Solaris Cluster HA for SAP

(Continued)

- `Monitor_retry_count` (application server), 40–41
- `Monitor_retry_count` (central instance), 37–40
- `Monitor_retry_interval` (central instance), 37–40
- `Probe_timeout` (application server), 40–41
- `Probe_timeout` (central instance), 37–40
- `SAPSID` (central instance), 37–40
- `shutdown_dev` (central instance), 37–40
- `Stop_sap_pct` (application server), 40–41
- `Stop_sap_pct` (central instance), 37–40
- Sun Cluster HA for SAP
 - `Dev_sapsid` (central instance), 37–40
 - `Monitor_retry_interval` (application server), 40–41
 - `SAPSID` (application server), 40–41

F

- fault monitor, Solaris Cluster HA for SAP, 49–51

G

- global zone, 33

H

- HA for SAP
 - installing, 33–34
 - software packages, installing, 33–34
- help, 8–9

I

- installing
 - HA for SAP, 33–34
 - SAP database, 20–21
 - SAP scalable application server, 21–24
 - Solaris Cluster HA for SAP
 - database, 29

L

- local zones, *See* non-global zones
- lock file
 - for central instance
 - Solaris Cluster HA for SAP, 45–46, 46–47
 - Solaris Cluster HA for SAP, 45–47

M

- messages file, 9

N

- non-global zones, 11, 33

O

- overview, Solaris Cluster HA for SAP, 11–12

P

- packages, 33–34
- prtconf -v command, 9

- prtdiag -v command, 9
- psrinfo -v command, 9

R

- registering
 - HA for SAP as a failover data service, 42–43
 - HA for SAP as a scalable data service, 43–44
 - HA for SAP with central instance, 41–42
- restrictions, zones, 33

S

SAP

- See also* Solaris Cluster HA for SAP
- database
 - installing and configuring, 20–21
- installing
 - enabling failover SAP instances, 24–27
 - preparing the nodes, 20
- scalable application server
 - installing, 21–24
- scalable installation
 - verifying, 32
- verifying installation
 - database with central instance, 30–31
 - failover application server, 31–32
- show- rev subcommand, 9
- showrev -p command, 9
- software packages, 33–34
- Solaris Cluster HA for SAP
 - See also* SAP
 - application server extension properties
 - As_db_retry_interval, 40–41
 - As_instance_id, 40–41
 - As_services_string, 40–41
 - As_shutdown_script, 40–41
 - As_startup_script, 40–41
 - Monitor_retry_count, 40–41
 - Monitor_retry_interval, 40–41
 - Probe_timeout, 40–41
 - SAPSID, 40–41
 - Stop_sap_pct, 40–41

Solaris Cluster HA for SAP (*Continued*)

- central instance extension properties
 - Check_ms_retry, 37–40
 - Ci_instance_id, 37–40
 - Ci_services_string, 37–40
 - Ci_shutdown_script, 37–40
 - Ci_start_retry_interval, 37–40
 - Ci_startup_script, 37–40
 - Dev_sapsid, 37–40
 - Dev_shutdown_script, 37–40
 - Dev_stop_pct, 37–40
 - Lgtst_ms_with_logicalhostname, 37–40
 - Message_server_name, 37–40
 - Monitor_retry_count, 37–40
 - Monitor_retry_interval, 37–40
 - Probe_timeout, 37–40
 - SAPSID, 37–40
 - shutdown_dev, 37–40
 - Stop_sap_pct, 37–40
- configuration
 - considerations, 16–17
 - database, 29
 - planning, 14–18
 - planning questions, 17–18
 - requirements, 14–15
 - restrictions, 14
 - standard, 15–16
- fault monitor, 49–51
 - central instance, 49–50
 - fault probes, 50–51
- overview, 11–12
- registering and configuring
 - as a failover data service, 42–43
 - as a scalable data service, 43–44
 - with central instance, 41–42
- setting up a lock file, 45–47
 - central instance, 45–46, 46–47
- supported packages, 18
- task map, 12–14
- upgrading, 19–20
- verifying installation
 - central instance, 47
 - failover installation, 47–48
 - scalable installation, 48

Sun Cluster HA for SAP

- installing
- planning, 14–18

T

- task map, Solaris Cluster HA for SAP, 12–14
- technical support, 8–9

U

- upgrading, Solaris Cluster HA for SAP, 19–20

V

- /var/adm/messages file, 9
- verifying
 - SAP installation
 - database installation with central instance, 30–31
 - failover application server, 31–32
 - SAP scalable installation, 32
 - Solaris Cluster HA for SAP, 47–48
 - central instance installation, 47
 - failover installation, 47–48
 - scalable installation, 48

Z

- zones, 11, 33

