Sun Cluster Data Service for SAP Guide for Solaris OS

Sun microsystems

Contents

Preface

Sun Cluster Data Service for SAP Guide for Solaris OS explains how to install and configure Sun™ Cluster HA for SAP.

Note – This Sun 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 Sun 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 Solaris™ Operating System (Solaris OS) and expertise with the volume-manager software that is used with Sun Cluster software.

Using UNIX Commands

This document contains information about commands that are specific to installing and configuring Sun 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 Solaris Operating System
- 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 `.login` file.
		Use `ls -a` to list all files.
		`machine_name% you have mail.`
AaBbCc123	What you type, contrasted with onscreen computer output	`machine_name%` **su**
		`Password:`
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is `rm` *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*.
		A *cache* is a copy that is stored locally.
		Do *not* 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 the C shell, Bourne shell, and Korn shell.

TABLE P–2 Shell Prompts

Shell	Prompt
C shell	`machine_name%`
C shell for superuser	`machine_name#`
Bourne shell and Korn shell	`$`
Bourne shell and Korn shell for superuser	`#`

Related Documentation

Information about related Sun Cluster topics is available in the documentation that is listed in the following table. All Sun Cluster documentation is available at `http://docs.sun.com`.

Topic	Documentation
Data service administration	*Sun Cluster Data Services Planning and Administration Guide for Solaris OS*
	Individual data service guides
Concepts	*Sun Cluster Concepts Guide for Solaris OS*
Overview	*Sun Cluster Overview for Solaris OS*
Software installation	*Sun Cluster Software Installation Guide for Solaris OS*
System administration	*Sun Cluster System Administration Guide for Solaris OS*
Hardware administration	*Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS*
	Individual hardware administration guides
Data service development	*Sun Cluster Data Services Developer's Guide for Solaris OS*
Error messages	*Sun Cluster Error Messages Guide for Solaris OS*
Command and function reference	*Sun Cluster Reference Manual for Solaris OS*

For a complete list of Sun Cluster documentation, see the release notes for your release of Sun 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 – Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun 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. Sun 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

The Sun web site provides information about the following additional resources:

- Documentation (http://www.sun.com/documentation/)
- Support (http://www.sun.com/support/)
- Training (http://www.sun.com/training/)

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. To share your comments, go to http://docs.sun.com and click Feedback.

Getting Help

If you have problems installing or using Sun 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 Solaris Operating System (for example, Solaris 10)
- The release number of Sun Cluster (for example, Sun Cluster 3.2)

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

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

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

Installing and Configuring Sun Cluster HA for SAP

This chapter explains how to install and configure Sun Cluster 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. Sun Cluster HA for SAP is supported in non-global zones.

This chapter contains the following sections.

Sun Cluster HA for SAP Overview

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

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

Sun Cluster 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 Sun Cluster configuration. You can configure Sun Cluster HA for SAP as a failover application or a scalable application.

TABLE 1   Protection of SAP Components

SAP Component	Protected by
SAP database	The data service for the database that you are using, for example: ■ For the MaxDB database, the data service is Sun Cluster HA for MaxDB. ■ For the Oracle database, the data service is Sun Cluster HA for Oracle.
SAP central instance	Sun Cluster HA for SAP The resource type is `SUNW.sap_ci` or `SUNW.sap_ci_v2`.
SAP application server	Sun Cluster HA for SAP The resource type is `SUNW.sap_as` or `SUNW.sap_as_v2`.
NFS file system	Sun Cluster HA for NFS

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

Installing and Configuring Sun Cluster HA for SAP

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

TABLE 2   Task Map: Installing and Configuring Sun Cluster HA for SAP

Task	Instructions
Plan the SAP installation	"Planning the Sun Cluster HA for SAP Installation and Configuration" on page 12
Upgrade Sun Cluster HA for SAP	"How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource" on page 17
Prepare the nodes and disks	"How to Prepare the Nodes" on page 18

TABLE 2 Task Map: Installing and Configuring Sun Cluster HA for SAP   *(Continued)*

Task	Instructions
Install SAP and the database	"How to Install SAP and the Database" on page 19
Install the SAP application server	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 23
	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 20
Configure your highly available database	"Configuring Your Highly Available Database" on page 27
Verify the SAP installation	If you plan to install a failover SAP application server, go to the following sections: ■ "How to Verify SAP and the Database Installation With Central Instance" on page 28 ■ "How to Verify an SAP Failover Application Server" on page 29 If you plan to install a scalable SAP application server, go to "Verifying an SAP Scalable Application Server" on page 31.
Install Sun Cluster HA for SAP packages	"Installing the Sun Cluster HA for SAP Packages" on page 31
Register and configure Sun Cluster HA for SAP	If you plan to configure Sun Cluster HA for SAP as a failover data service, go to: ■ "How to Register and Configure Sun Cluster HA for SAP With Central Instance" on page 40 ■ "How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service" on page 41 If you plan to configure Sun Cluster HA for SAP as a scalable data service, go to: ■ "How to Register and Configure Sun Cluster HA for SAP With Central Instance" on page 40 ■ "How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service" on page 42
Set up a lock file	"Setting Up a Lock File" on page 43
Verify Sun Cluster HA for SAP installation and configuration	If you plan to configure Sun Cluster HA for SAP as a failover data service, go to: ■ "How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance" on page 45 ■ "How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Failover Data Service" on page 46 If you plan to configure Sun Cluster HA for SAP as a scalable data service, go to "How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Scalable Data Service" on page 47.

TABLE 2   Task Map: Installing and Configuring Sun Cluster HA for SAP   *(Continued)*

Task	Instructions
Understand Sun Cluster HA for SAP fault monitor	"Understanding Sun Cluster HA for SAP Fault Monitor" on page 47

Planning the Sun Cluster HA for SAP Installation and Configuration

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

Note – Sun Cluster 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 Sun Cluster HA for SAP. This section provides a list of software and hardware configuration restrictions that apply to Sun Cluster HA for SAP.

For restrictions that apply to all data services, see the release notes for your release of Sun 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 Sun Cluster HA for SAP. These requirements apply to Sun Cluster HA for SAP only. You must meet these

requirements before you proceed with your Sun Cluster 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 Sun Cluster Data Services," in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

- **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 Sun Cluster HA for NFS on the cluster that hosts the central instance if you plan to install some external SAP application servers.** See *Sun Cluster Data Service for NFS Guide for Solaris OS* for details about how to configure Sun Cluster 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, Sun Cluster 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.** Sun Cluster 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 Sun Cluster HA for SAP. Sun Cluster HA for SAP supports the standard configurations in this section. Sun Cluster HA for SAP might support additional configurations. However, you must contact your Enterprise Services representative for information about additional configurations.
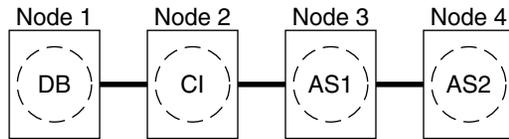
CLUSTER 1

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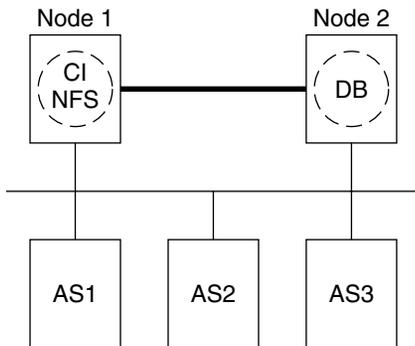
CLUSTER 1

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

Note – The configuration in Figure 2 was a common configuration under previous Sun Cluster releases. To use the Sun Cluster software to the fullest extent, configure SAP as shown in Figure 1 or Figure 3.
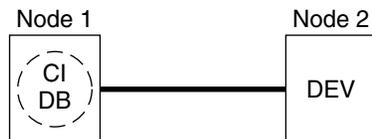
CLUSTER 1

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

Failover and Scalable Applications

- **Retrieve the latest patch for the** `sapstart` **executable**. This patch enables Sun Cluster 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 43.

- **Read all the related SAP online service-system notes for the SAP software release and database that you are installing on your Sun 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** *SAPSID*adm **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 Sun Cluster HA for SAP. Insert the answers to these questions into the data service worksheets in Appendix D, "Data Service Configuration Worksheets and Examples," in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*. See "Configuration Considerations" on page 15 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 Sun Cluster Data Services" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* 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 3 and Table 4 list the packages that Sun Cluster HA for SAP supports.

TABLE 3   Sun Cluster HA for SAP Packages From Sun 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 Sun Cluster HA for SAP. The *_v2 resource types are a superset of the original RTs. Whenever possible, use the latest RTs provided.

TABLE 4   Sun Cluster HA for SAP Package From Sun Cluster 3.0 12/01

Resource Type	Description
SUNW.sap_ci	Same as Sun Cluster 3.0 7/01. See Table 3.
SUNW.sap_as	Same as Sun Cluster 3.0 7/01. See Table 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.

TABLE 4   Sun Cluster HA for SAP Package From Sun Cluster 3.0 12/01   *(Continued)*

Resource Type	Description
`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 Sun Cluster HA for SAP

As Table 3 and Table 4 illustrate, the Sun Cluster HA for SAP package from Sun 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 17 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 17 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 18.

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 Sun Cluster 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 Sun Cluster 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
```

Sun Cluster 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 Sun Cluster HA for SAP as a Scalable Data Service" on page 42.

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 Sun 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 20 for the procedure on how to install a scalable application server in a Sun Cluster environment.

Next Steps   Go to "How to Enable Failover SAP Instances to Run in a Cluster" on page 23 or "How to Install and Enable an SAP Scalable Application Server" on page 20.

▼ 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 Sun Cluster 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 Sun Cluster HA for SAP, first complete this procedure. Then complete the additional steps in "How to Configure an SAP J2EE Engine Cluster With Your Sun Cluster HA for SAP With an Application Server" on page 25.

To configure an SAP Web dispatcher with your Sun Cluster HA for SAP agent, first complete this procedure. Then complete the additional steps in "How to Configure an SAP Web Dispatcher With Your Sun Cluster HA for SAP" on page 27.

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 *SAPSID*adm 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 44.

- 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 `startup` **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 *sapsid*adm**.**

b.   Make a copy of the `startsap` **script and the** `stopsap` **script, and save these files in the** *SAPSID*adm **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** `startup` **log file and the** `stop` **log file will be node specific under the home directories of users** *sapsid*adm **and** ora*sapsid* **or** ora*dbsid*.

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** *SAPSID*adm **home directory as described in "How to Install and Configure MaxDB" in** *Sun Cluster Data Service for MaxDB Guide for Solaris OS***.**

   ___

   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** *SAPSID*adm **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 *SAPSID*adm 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 Sun Cluster HA for SAP, go to "How to Configure an SAP J2EE Engine Cluster With Your Sun Cluster HA for SAP With an Application Server" on page 25.

To configure an SAP Web dispatcher with your Sun Cluster HA for SAP agent, go to "How to Configure an SAP Web Dispatcher With Your Sun Cluster HA for SAP" on page 27.

If you are not using a J2EE engine cluster or an SAP Web dispatcher with your Sun Cluster HA for SAP, go to "Configuring Your Highly Available Database" on page 27.

▼ 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 Sun 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 Sun 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 Sun Cluster HA for SAP With Central Instance" on page 26.

To configure a J2EE engine cluster with your Sun Cluster HA for SAP with an SAP Application Server, go to "How to Configure an SAP J2EE Engine Cluster With Your Sun Cluster HA for SAP With an Application Server" on page 25.

To configure SAP Web dispatcher with your Sun Cluster HA for SAP agent, first complete this procedure. Then complete the additional steps in "How to Configure an SAP Web Dispatcher With Your Sun Cluster HA for SAP" on page 27.

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** *SAPSID*adm **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 *sapsid*adm**, and then perform the following tasks to change references of physical server name to logical hostname.**

- In the *SAPSID*adm 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 *SAPSID*adm 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 *sapsid*adm**, 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 *sapsid*adm**, 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 ora*sapsid* **or** ora*dbsid***, and then perform the following tasks.**

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

- In the ora*SAPSID* or ora*dbsid* home directory, modify all the file names that reference a physical server name.
- In the ora*SAPSID* or ora*dbsid* 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** *sapsid*adm **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 Sun Cluster HA for SAP with a Central Instance, go to "How to Configure an SAP J2EE Engine With Your Sun Cluster HA for SAP With Central Instance" on page 26.

To configure a J2EE engine cluster with your Sun Cluster HA for SAP with an SAP Application Server, go to "How to Configure an SAP J2EE Engine Cluster With Your Sun Cluster HA for SAP With an Application Server" on page 25.

To configure an SAP Web dispatcher with your Sun Cluster HA for SAP agent, go to "How to Configure an SAP Web Dispatcher With Your Sun Cluster HA for SAP" on page 27.

If you are not using a J2EE engine cluster or an SAP Web dispatcher with your Sun Cluster HA for SAP agent, go to "Configuring Your Highly Available Database" on page 27.

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

After you have completed "How to Install and Enable an SAP Scalable Application Server" on page 20, perform the following steps.

Note – The SAP J2EE engine is started by the SAP dispatcher, which is under the protection of the Sun 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 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

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

After you have completed "How to Enable Failover SAP Instances to Run in a Cluster" on page 23 or "How to Enable Failover SAP Instances to Run in a Cluster" on page 23, perform the following steps.

Note – The SAP J2EE engine is started by the SAP dispatcher which is under the protection of the Sun 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

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

After you have configured the SAP Web dispatcher with your Sun 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 Sun 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 27.

Configuring Your Highly Available Database

SAP supports various databases. See the appropriate Sun 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 *Sun Cluster Data Service for MaxDB Guide for Solaris OS*.
- If you plan to use Oracle with SAP, see *Sun Cluster Data Service for Oracle Guide for Solaris OS*.

Additionally, see the appropriate Sun 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 Sun 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 *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* and "Configuring the HAStoragePlus Resource Type to Work With Sun Cluster HA for MaxDB" in *Sun Cluster Data Service for MaxDB Guide for Solaris OS*.

- If you plan to use the Oracle database, see "Synchronizing the Startups Between Resource Groups and Device Groups" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Where to Go From Here

Go to "How to Verify SAP and the Database Installation With Central Instance" on page 28 or "Verifying an SAP Scalable Application Server" on page 31.

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** *SAPSID*adm **home directory as described in** "How to Install and Configure MaxDB" in *Sun Cluster Data Service for MaxDB Guide for Solaris OS*. **Verify that the** .XUSER.62 **file is correctly configured by using the** -d **option of the** R3trans **command. You must be an** *SAPSID*adm **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 29.

▼ 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 *SAPSID*adm home directory as described in "How to Install and Configure MaxDB" in *Sun Cluster Data Service for MaxDB Guide for Solaris OS*.

8 Verify that the .XUSER.62 file is correctly configured by using the -d option of the R3trans command.

Note – You must be an *SAPSID*adm 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 Sun Cluster HA for SAP Packages" on page 31.

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 20, 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 20.

Where to Go From Here

Go to "Installing the Sun Cluster HA for SAP Packages" on page 31.

Installing the Sun Cluster HA for SAP Packages

If you did not install the Sun Cluster HA for SAP packages during your initial Sun Cluster installation, perform this procedure to install the packages. To install the packages, use the Sun Java™ Enterprise System Installation Wizard.

Note – You need to install the Sun Cluster HA for SAP packages in the global cluster and not in the zone cluster.

▼ How to Install the Sun Cluster HA for SAP Packages

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

You can run the Sun Java Enterprise System Installation Wizard 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 Sun Java Availability Suite DVD-ROM.

If you intend to run the Sun Java Enterprise System Installation Wizard 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 Sun Java Availability Suite DVD-ROM 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 Sun Java Enterprise System 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 Sun Java Enterprise System Installation Wizard.

`# ./installer`

5 When you are prompted, accept the license agreement.

If any Sun Java Enterprise System components are installed, you are prompted to select whether to upgrade the components or install new software.

6 From the list of Sun 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 Sun Java Enterprise System 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 Sun Java Enterprise System Installation Wizard with the CLI, omit this step.

11 Exit the Sun Java Enterprise System Installation Wizard.

12 **Unload the Sun Java Availability Suite DVD-ROM 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

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 Sun 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 Sun Cluster HA for SAP

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

Sun Cluster HA for SAP Extension Properties

Use the extension properties in "Sun Cluster HA for SAP Extension Properties for the Central Instance" on page 35 and "Sun Cluster HA for SAP Extension Properties for the Application Servers" on page 38 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 *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* 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 B, "Standard Properties," in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for details about all Sun Cluster properties.

Sun Cluster 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 *SAPSID*adm 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 *SAPSID*adm 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 Sun 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, Sun Cluster 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, Sun Cluster HA for SAP requires this property.

Default: None

Tunable: When disabled

Dev_stop_pct

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

Default: 20

Tunable: When disabled

Sun Cluster 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 Sun Cluster HA for SAP With Central Instance

Use this procedure to configure Sun Cluster 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 `HAStoragePlus` **resource to the resource group for the central instance.**

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

The resource is created in the enabled state.

For more details about how to set up an `HAStoragePlus` resource, see "Enabling Highly Available Local File Systems" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

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 "Sun Cluster HA for SAP Extension Properties" on page 35 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 Sun Cluster HA for SAP as a Failover Data Service" on page 41 or "How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service" on page 42.

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

Use this procedure to configure Sun Cluster 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** HAStoragePlus **resource to the resource group for the failover application server.**

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

The resource is created in the enabled state.

For more details about how to set up an HAStoragePlus resource, see "Enabling Highly Available Local File Systems" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

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 "Sun Cluster HA for SAP Extension Properties" on page 35 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 Sun Cluster HA for SAP Installation and Configuration and Central Instance" on page 45.

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

Use this procedure to configure Sun Cluster 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 – Sun 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 HAStoragePlus **resource to the failover application server resource group.**

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

The resource is created in the enabled state.

For more details on how to set up an HAStoragePlus resource, see "Enabling Highly Available Local File Systems" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

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 "Sun Cluster HA for SAP Extension Properties" on page 35 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 Sun Cluster HA for SAP Installation and Configuration and Central Instance" on page 45.

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 Sun Cluster environment and then try to start SAP under the Sun Cluster environment, the Sun Cluster 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 Sun Cluster 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 Sun Cluster 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 Sun Cluster 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***`/D`***instance-id***`/work/startup_lockfile`**

`sapstart/lockfile`
New parameter name

`/usr/sap/local/`*SAPSID*`/D`*instance-id*`/work`
Work directory for the scalable application server

`startup_lockfile`
Lock file name that Sun Cluster 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 Sun Cluster 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 Sun Cluster HA for SAP Installation and Configuration and Central Instance

Use this procedure to verify the Sun Cluster 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 Sun Cluster HA for SAP is functioning correctly.**

3 **As user** *sapsid*adm**, use the central instance** `stopsap` **script to shut down the SAP central instance.**
The Sun 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 Sun Cluster HA for SAP as a Failover Data Service" on page 46 or "How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Scalable Data Service" on page 47.

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

Use this procedure to verify the installation and configuration of Sun Cluster 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** *sapsid*adm**, 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 Sun 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 Sun Cluster HA for SAP as a Scalable Data Service

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

1   **Log on to one of the nodes that runs the application server.**

2   **Become user** *sapsid*adm**.**

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 Sun 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 Sun Cluster HA for SAP Fault Monitor

The Sun Cluster 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.

Sun Cluster 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 R3trans to check the status of the database connection. Sun Cluster HA for SAP fault probes verify that SAP can connect to the database. Sun Cluster 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

Sun Cluster 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 `R3trans` to check the status of the database connection. Sun Cluster HA for SAP fault probes verify that SAP can connect to the database. Sun Cluster 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