



Sun Cluster 3.1 Data Service for SAP

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817-1536-10
May 2003 Revision A

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



030312@5533



Contents

Preface 5

Installing and Configuring Sun Cluster HA for SAP 11

Sun Cluster HA for SAP Overview 12

Installing and Configuring Sun Cluster HA for SAP 13

Planning the Sun Cluster HA for SAP Installation and Configuration 14

 Configuration Restrictions 14

 Configuration Requirements 15

 Standard Data Service Configurations 16

 Configuration Considerations 17

 Configuration Planning Questions 18

 Packages and Support 18

Upgrading Sun Cluster HA for SAP 19

 ▼ How to Upgrade a Resource Type or Convert a Failover Application Resource
 to a Scalable Application Resource 20

Preparing the Nodes and Disks 20

 ▼ How to Prepare the Nodes 21

Installing and Configuring SAP and Database 21

 ▼ How to Install SAP and the Database 21

 ▼ How to Install an SAP Scalable Application Server 22

 ▼ How to Enable Failover SAP Instances to Run in a Cluster 25

Configuring Sun Cluster HA for DBMS 26

 Where to Go From Here 26

Verifying the SAP Installation 27

 ▼ How to Verify SAP and the Database Installation with Central Instance 27

 ▼ How to Verify an SAP Failover Application Server 28

▼ How to Verify an SAP Scalable Application Server	29
Installing the Sun Cluster HA for SAP Packages	29
▼ How to Install the Sun Cluster HA for SAP Packages	29
Registering and Configuring Sun Cluster HA for SAP	30
Sun Cluster HA for SAP Extension Properties	30
▼ How to Register and Configure Sun Cluster HA for SAP with Central Instance	37
▼ How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service	38
▼ How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service	39
Setting Up a Lock File	40
▼ How to Set Up a Lock File for Central Instance or the Failover Application Server	41
▼ How to Set Up a Lock File for Scalable Application Server	42
Verifying the Sun Cluster HA for SAP Installation and Configuration	43
▼ How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance	43
▼ How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Failover Data Service	43
▼ How to Verify Sun Cluster HA for SAP Installation and Configuration of as a Scalable Data Service	44
Understanding Sun Cluster HA for SAP Fault Monitor	45
Sun Cluster HA for SAP Fault Probes for Central Instance	45
Sun Cluster HA for SAP Fault Probes for Application Server	46
Index	49

Preface

The Sun Cluster 3.1 Data Service for SAP contains procedures to install and configure the Sun Cluster data services.

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 document assume knowledge of the Solaris™ operating environment and expertise with the volume manager software used with Sun Cluster.

UNIX Commands

This document contains information on commands specific to installing and configuring Sun Cluster data services. It might not contain information on basic UNIX® commands and procedures, such as shutting down the system, booting the system, and configuring devices. For that information, see one or more of the following:

- Online documentation for the Solaris software environment
- Solaris operating environment man pages
- Other software documentation that you received with your system

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Installation	<i>Sun Cluster 3.1 Software Installation Guide</i>	816-3388

Application	Title	Part Number
Data Services Installation and Configuration	<i>Sun Cluster 3.1 Data Service Planning and Administration Guide</i>	817-1526
	<i>Sun Cluster 3.1 Data Service for Apache</i>	817-1532
	<i>Sun Cluster 3.1 Data Service for BroadVision One-To-One Enterprise</i>	817-1542
	<i>Sun Cluster 3.1 Data Service for Domain Name Service (DNS)</i>	817-1533
	<i>Sun Cluster 3.1 Data Service for Netbackup</i>	817-1538
	<i>Sun Cluster 3.1 Data Service for Network File System (NFS)</i>	817-1534
	<i>Sun Cluster 3.1 Data Service for Oracle</i>	817-1527
	<i>Sun Cluster 3.1 Data Service for Oracle Parallel Server/Real Application Clusters</i>	817-1535
	<i>Sun Cluster 3.1 Data Service for SAP</i>	817-1536
	<i>Sun Cluster 3.1 Data Service for SAP liveCache</i>	817-1539
	<i>Sun Cluster 3.1 Data Service for Siebel</i>	817-1540
	<i>Sun Cluster 3.1 Data Service for Sun ONE Application Server</i>	817-1530
	<i>Sun Cluster 3.1 Data Service for Sun ONE Directory Server</i>	817-1529
	<i>Sun Cluster 3.1 Data Service for Sun ONE Message Queue</i>	817-1531
	<i>Sun Cluster 3.1 Data Service for Sun ONE Web Server</i>	817-1528
	<i>Sun Cluster 3.1 Data Service for Sybase ASE</i>	817-1537
	<i>Sun Cluster 3.1 Data Service for WebLogic Server</i>	817-1541
	<i>Sun Cluster 3.1 Data Service for DHCP</i>	817-1716
	<i>Sun Cluster 3.1 Data Service for Samba</i>	817-1715
	<i>Sun Cluster 3.1 Data Service for WebSphere MQ</i>	817-1714
<i>Sun Cluster 3.1 Data Service for WebSphere MQ Integrator</i>	817-1713	
Hardware	<i>Sun Cluster 3.x Hardware Administration Manual</i>	817-0168
	Sun Cluster 3.x Hardware Administration Collection at http://docs.sun.com/db/coll/1024.1	
API development	<i>Sun Cluster 3.1 Data Services Developer's Guide</i>	816-3385

Application	Title	Part Number
Administration	<i>Sun Cluster 3.1 5/03 System Administration Guide</i>	816-3384
Cluster concepts	<i>Sun Cluster 3.1 Concepts Guide</i>	816-3383
Error Messages	<i>Sun Cluster 3.1 Error Messages Guide</i>	816-3382
Man Pages	<i>Sun Cluster 3.1 Reference Manual</i>	816-5251
Release notes	<i>Sun Cluster 3.1 Data Service Release Notes</i>	817-1790
	<i>Sun Cluster 3.1 Release Notes</i>	816-5317
	<i>Sun Cluster 3.1 Release Notes Supplement</i>	816-3381

Sun Documentation Online

The `docs.sun.comSM` Web site enables you to access Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at `http://docs.sun.com`.

Help

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

- Your name and E-mail address (if available)
- Your company name, address, and phone number
- The model and serial numbers of your systems
- The release number of the operating environment (for example, Solaris 8)
- The release number of Sun Cluster (for example, Sun Cluster 3.0)

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.

Command	Function
<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>scinstall -pv</code>	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 contains the procedures on how to install and configure Sun Cluster HA for SAP.

This chapter contains the following procedures.

- “How to Upgrade a Resource Type or Convert a Failover Application Resource to a Scalable Application Resource” on page 20
- “How to Prepare the Nodes” on page 21
- “How to Install SAP and the Database” on page 21
- “How to Install an SAP Scalable Application Server” on page 22
- “How to Verify an SAP Scalable Application Server” on page 29
- “How to Enable Failover SAP Instances to Run in a Cluster” on page 25
- “How to Verify SAP and the Database Installation with Central Instance” on page 27
- “How to Verify an SAP Failover Application Server” on page 28
- “How to Install the Sun Cluster HA for SAP Packages” on page 29
- “How to Register and Configure Sun Cluster HA for SAP with Central Instance” on page 37
- “How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service” on page 38
- “How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service” on page 39
- “How to Set Up a Lock File for Central Instance or the Failover Application Server” on page 41
- “How to Set Up a Lock File for Scalable Application Server” on page 42
- “How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance” on page 43
- “How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Failover Data Service” on page 43
- “How to Verify Sun Cluster HA for SAP Installation and Configuration of as a Scalable Data Service” on page 44

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 on failover and scalable services, see the *Sun Cluster 3.1 Concepts Guide*.

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 1-1 Protection of SAP Components

SAP Component	Protected by
SAP database	Sun Cluster HA for Oracle Use Oracle as your database.
SAP central instance	Sun Cluster HA for SAP The resource type is <code>SUNW.sap_ci</code> or <code>SUNW.sap_ci_v2</code> .
SAP application server	Sun Cluster HA for SAP The resource type is <code>SUNW.sap_as</code> or <code>SUNW.sap_as_v2</code> .
NFS file system	Sun Cluster HA for NFS

Use the `scinstall(1M)` command to install Sun Cluster HA for SAP. Sun Cluster HA for SAP requires a functioning cluster with the initial cluster framework already installed. See the *Sun Cluster 3.1 Software Installation Guide* 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 1-2 lists the tasks for installing and configuring Sun Cluster HA for SAP. Perform these tasks in the order that they are listed.

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

Task	For Instructions, Go To
Plan the SAP installation	Chapter 1 of <i>Sun Cluster 3.1 Data Service for SAP</i> “Planning the Sun Cluster HA for SAP Installation and Configuration” on page 14
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 20
Prepare the nodes and disks	“How to Prepare the Nodes” on page 21
Install SAP, SAP failover application server, and the database	“How to Install SAP and the Database” on page 21 “How to Enable Failover SAP Instances to Run in a Cluster” on page 25
Configure the Sun Cluster HA for DBMS	“Configuring Sun Cluster HA for DBMS” on page 26
Verify SAP Installation	“How to Verify SAP and the Database Installation with Central Instance” on page 27 “How to Verify an SAP Failover Application Server” on page 28
or	
Install SAP, SAP scalable application server, and the database	“How to Install SAP and the Database” on page 21 “How to Install an SAP Scalable Application Server” on page 22 “Configuring Sun Cluster HA for DBMS” on page 26
Configure the Sun Cluster HA for DBMS	“How to Verify an SAP Scalable Application Server” on page 29
Verify SAP Installation	
Install Sun Cluster HA for SAP packages	“How to Install the Sun Cluster HA for SAP Packages” on page 29

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

Task	For Instructions, Go To
Register and configure Sun Cluster HA for SAP as a failover data service	“How to Register and Configure Sun Cluster HA for SAP with Central Instance” on page 37 “How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service” on page 38
or	
Register and configure Sun Cluster HA for SAP as a scalable data service	“How to Register and Configure Sun Cluster HA for SAP with Central Instance” on page 37 “How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service” on page 39
Set up a lock file	“Setting Up a Lock File” on page 40
Verify Sun Cluster HA for SAP installation and configuration	“How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance” on page 43 “How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Failover Data Service” on page 43 “How to Verify Sun Cluster HA for SAP Installation and Configuration of as a Scalable Data Service” on page 44
Understand Sun Cluster HA for SAP fault monitor	“Understanding Sun Cluster HA for SAP Fault Monitor” on page 45

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.

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 *Sun Cluster 3.1 Release Notes*.

- **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 requirements that apply to all data services, see “Configuring and Administering Sun Cluster Data Services” in *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

- **After you create all of 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 of 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 on how to set up the database and SAP file systems.
- **Create the required groups and users on all of 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 on 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 “Installing and Configuring Sun Cluster HA for NFS” in *Sun Cluster 3.1 Data Service for Network File System (NFS)* for details on 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 on additional configurations.



CLUSTER 1


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 Sun Cluster releases. To use the Sun Cluster software to the fullest extent, configure SAP as shown in Figure 1-1 or Figure 1-3.



CLUSTER 1

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 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 on the benefits of this patch in your cluster environment, see “Setting Up a Lock File” on page 40.
- **Read all of 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 `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.
- **Place the application servers into multiple resource groups if you want to use the RGOffload resource type to shut down one or more application servers when a higher priority resource is failing over** - This consideration provides flexibility and availability if you want to use the RGOffload resource type to offload one or more application servers for the database. The value you gain from this consideration supersedes the ease of use you gain from placing the application servers into one large group. See "Freeing Node Resources by Offloading Non-critical Resource Groups" in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for more information on using the RGOffload resource type.
- **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 the *Sun Cluster 3.1 Release Notes*. See "Configuration Considerations" on page 17 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 "Determining the Location of the Application Binaries" on page 3 of the *Sun Cluster 3.1 Data Service for SAP* for the advantages and disadvantages of placing the liveCache binaries on the local file system as opposed to the cluster file system.

Packages and Support

Table 1-3 and Table 1-4 lists the packages that Sun Cluster HA for SAP supports.

TABLE 1-3 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 1-4 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 1-3.
SUNW.sap_as	Same as Sun Cluster 3.0 7/01. See Table 1-3.
SUNW.sap_ci_v2	Added the <code>Network_resources_used</code> resource property to the Resource Type Registration (RTR) file. Retained support for failover central instance.
SUNW.sap_as_v2	Added the <code>Network_resources_used</code> 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 1-3 and Table 1-4 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 down time.
 - 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 20 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 20 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 down time.

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. **Which task are you 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.**

Where to Go From Here

Go to “How to Prepare the Nodes” on page 21.

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 of the nodes.**
2. **Configure the `/etc/nsswitch.conf` so that Sun Cluster HA for SAP starts and stops correctly in the event of a switchover or a failover.**

On each node 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. When you add one of the entries for `group` in the `/etc/nsswitch.conf` file, you ensure that the `su(1M)` command does not refer to the NIS/NIS+ name services if the network information name service is unavailable.

Where to Go From Here

Go to "How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service" on page 39.

Installing and Configuring SAP and Database

This section contains the procedures 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.**

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 of the SAP-required kernel parameter changes, edit the `/etc/system` file on all of 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 on kernel parameter changes.

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

See “How to Install an SAP Scalable Application Server” on page 22 for the procedure on how to install a scalable application server in a Sun Cluster environment.

Where to Go From Here

Go to “How to Enable Failover SAP Instances to Run in a Cluster” on page 25 or “How to Install an SAP Scalable Application Server” on page 22.

▼ How to Install 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 and SAP 6.20 users to ensure that Sun Cluster HA for SAP can manage and bring online SAP 6.10 and SAP 6.20 services. SAP 6.10 and SAP 6.20 create one `startsap` script and one `stopsap` script. Other SAP versions create one of each of these scripts for each service you create. This difference accounts for the additional steps for SAP 6.10 and SAP 6.20 users.

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

- Cluster File Systems

- `/sapmnt/SID`

- `/usr/sap/SID` -> all subdirectories except the *app-instance* subdirectory

- `/usr/sap/SID/home` -> the SAPSIDadm home directory

- `/usr/sap/trans`

- Local File Systems

- `/usr/sap/local/SID/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 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 41.

- 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 on how to install the central instance and the database.

3. On all nodes that will 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/SID/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 and the two application servers will overwrite each other.

4. Set up a link to point to the local application server directory from a cluster file system, so the application server and the startup log file and the stop log file will be installed on the local file system.

Example:

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

5. Install the application server.

6. Are you using SAP 6.10 or SAP 6.20?

- If no, skip to Step 11
- If yes, proceed to Step 7.

7. Become user *sapsidadm*.

8. Make a copy of the *startsap* script and the *stopsap* script, and save these files in the *SAPSIDadm* home directory. The filenames that you choose specify this instance.

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

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

9. Make backup copies of the following files because you will modify them. In the SAP profile directory, modify all the filenames for this instance. The filenames that you choose must be specific to this instance, and they must follow the same naming convention you chose in Step 8.

```
# mv SAPSID_Service-StringSystem-Number_physical-hostname \  
SAPSID_Service-StringSystem_instance-number
```

```
# mv START_Service-StringSystem-Number_physical-hostname \  
START_Service-StringSystem_instance-number
```

10. Modify the contents of the files you created in Step 9 to replace any reference to the physical host with the instance number.



Caution – It is important that you make your updates consistent so that you can start and stop this application server instance from all the nodes that will run this scalable application server. For example, if you make these changes for SAP instance number 02, then use 02 where this instance number appears. If you do not use a consistent naming convention you will be unable start and stop this application server instance from all the nodes that will run this scalable application server.

11. 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 *sapsidadm* and *orasapsid*.

Example:

```
# vi startsap_D03
```

Before:

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

After:

```
LOGFILE=$R3S_LOGDIR/`basename $0`_`uname -n`.log
```

12. 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 scalable application server resource group *nodelist*.

13. Ensure that you can startup and stop the application server from each node, and verify that the log files are in the correct location.

14. Create the SAP logon group if you use a logon group.

Where to Go From Here

Go to “Configuring Sun Cluster HA for DBMS” on page 26.

▼ 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.

1. **Make backup copies of the files you will modify in Step 5 through Step 8.**
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. **Are you using SAP 6.10 or SAP 6.20?**
 - If no, skip to Step 6.
 - If yes, proceed to Step 5.
5. **Make a copy of the `startsap` script and the `stopsap` script, and save these files in the `SAPSIDadm` home directory. The filenames that you choose must specify this instance.**

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

6. **Become user `sapsidadm`, and then perform the following tasks.**
 - In the `SAPSIDadm` home directory, modify all of the file names that reference a physical server name.
 - In the `SAPSIDadm` home directory, modify all of the file contents—except log file contents—that reference a physical server name.
 - In the SAP profile directory, modify all of the file names that reference a physical server name.
7. **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`

8. Become user `orasapsid`, and then perform the following tasks.

- In the `orasapsid` home directory, modify all of the file names that reference a physical server name.
- In the `orasapsid` home directory, modify all of the file contents—except log file contents—that reference a physical server name.

9. Ensure that the `/usr/sap/tmp` directory owned by user `sapsidadm` and group `sapsys` exists on all nodes that can master the failover SAP instance.

Where to Go From Here

Go to “Configuring Sun Cluster HA for DBMS” on page 26.

Configuring Sun Cluster HA for DBMS

SAP supports various databases. See the appropriate chapter of this book for details on how to configure the resource type, resource group, and resource for your highly available database. For example, see “Installing and Configuring Sun Cluster HA for Oracle” in *Sun Cluster 3.1 Data Service for Oracle* for more information if you plan to use Oracle with SAP.

Additionally, see the appropriate chapter of this book and the appropriate chapter of your database installation book for details on other resource types to configure with your database. This book includes details on how to configure other resource types for Oracle databases. For instance, set up the `SUNW.HAStoragePlus` resource type if you use Oracle. See the procedure “Synchronizing the Startups Between Resource Groups and Disk Device Groups” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for more information.

Where to Go From Here

Go to “How to Verify SAP and the Database Installation with Central Instance” on page 27 or “How to Verify an SAP Scalable Application Server” on page 29.

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 SAP central instance. Perform the following steps on all of the potential nodes on which the central instance can run.

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

```
# scrgadm -a -g sap-ci-resource-group [-h nodelist]
```

Note – Use the `-h` option to the `scrgadm(1M)` 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 of the network resources that you use.**
3. **Add a network resource (logical hostname) to the failover resource group.**

```
# scrgadm -a -L -g sap-ci-resource-group  
-l ci-logical-hostname [-n netiflist]
```

4. **Enable the resource group.**

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

```
# scswitch -Z -g sap-ci-resource-group
```

5. **Log in to the cluster member that hosts the central instance resource group.**
6. **Ensure that the database is running.**
7. **Manually start the central instance.**

8. **Start the SAP GUI using the logical hostname, and verify that SAP initializes correctly.**

The default dispatcher port is 3200.

9. **Manually stop the central instance.**

10. Switch this resource group to another cluster member that can host the central instance.

```
# scswitch -z -h node -g sap-ci-resource-group
```

11. Repeat Step 5 through Step 9 until you verify startup and shutdown of the central instance on each cluster node that can host the central instance.

Where to Go From Here

Go to “How to Verify an SAP Failover Application Server” on page 28.

▼ 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 of 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.

```
# scrgadm -a -g sap-as-fo-resource-group
```

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

```
# scrgadm -a -g sap-as-fo-resource-group \  
[-h nodelist]
```

2. Verify that you added to your name service database all of the network resources that you use.

3. Add a network resource (logical hostname) to the failover resource group.

```
# scrgadm -a -L -g sap-as-fo-resource-group \  
-l as-fo-logical-hostname [-n netiflist]
```

4. Enable the resource group.

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

```
# scswitch -Z -g sap-as-fo-resource-group
```

5. Log in to the cluster member that hosts the application server resource group.

6. Manually start the application server.

7. Start the SAP GUI using the logical hostname, and verify that SAP initializes correctly.
8. Manually stop the application server.
9. Switch this resource group to another cluster member that can host the application server.

```
# scswitch -z -h node -g sap-as-fo-resource-group
```
10. Repeat Step 5 through Step 7 until you verify startup and shutdown of the application server on each cluster node that can host the application server.

Where to Go From Here

Go to “How to Install the Sun Cluster HA for SAP Packages” on page 29.

▼ How to Verify an SAP Scalable Application Server

If you installed scalable application server instances in “How to Install an SAP Scalable Application Server” on page 22, you verified the installation of an SAP scalable application server in Step 13 of “How to Install an SAP Scalable Application Server” on page 22.

Where to Go From Here

Go to “How to Install the Sun Cluster HA for SAP Packages” on page 29.

Installing the Sun Cluster HA for SAP Packages

This section contains the procedure you need to install the Sun Cluster HA for SAP packages.

▼ How to Install the Sun Cluster HA for SAP Packages

Use this procedure to install the Sun Cluster HA for SAP packages. You need the Sun Cluster Agents CD-ROM to perform this procedure. This procedure assumes that you did not install the data service packages during your initial Sun Cluster installation.

1. **Load the Sun Cluster Agents CD-ROM into the CD-ROM drive.**
2. **Run the `scinstall` utility with no options.**
This step starts the `scinstall` utility in interactive mode.
3. **Choose the Add Support for New Data Service to This Cluster Node menu option.**
The `scinstall` utility prompts you for additional information.
4. **Provide the path to the Sun Cluster Agents CD-ROM.**
The utility refers to the CD-ROM as the “data services cd.”
5. **Specify the data service to install.**
The `scinstall` utility lists the data service that you selected and asks you to confirm your choice.
6. **Exit the `scinstall` utility.**
7. **Unload the CD-ROM from the drive.**

Where to Go From Here

Go to “Registering and Configuring Sun Cluster HA for SAP” on page 30.

Registering and Configuring Sun Cluster HA for SAP

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

Sun Cluster HA for SAP Extension Properties

Use the extension properties in Table 1–5 and Table 1–6 to create your resources. Use the command line `scrgadm -x parameter=value` to configure extension properties when you create your resource. Use the procedure in “Administering Data Service Resources” in *Sun Cluster 3.1 Data Service 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 for details on all Sun Cluster properties.

TABLE 1-5 Sun Cluster HA for SAP Extension Properties for the Central Instance

Property Category	Property Name	Description
SAP Configuration	SAPSID	SAP system ID or SID. 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	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 <i>SID</i> adm home directory. Default: None Tunable: When disabled

TABLE 1-5 Sun Cluster HA for SAP Extension Properties for the Central Instance
(Continued)

Property Category	Property Name	Description
Stopping SAP	Stop_sap_pct	<p>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.</p> <p>Default: 95</p> <p>Tunable: When disabled</p>
	Ci_shutdown_script	<p>Name of the SAP shutdown script for this instance in your <i>SIDadm</i> home directory.</p> <p>Default: None</p> <p>Tunable: When disabled</p>

TABLE 1-5 Sun Cluster HA for SAP Extension Properties for the Central Instance
(Continued)

Property Category	Property Name	Description
Probe	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 Default: 4 Tunable: Any time
	Monitor_retry_interval	Time interval in minutes for the fault monitor restarts. Default: 2 Tunable: Any time

TABLE 1-5 Sun Cluster HA for SAP Extension Properties for the Central Instance
(Continued)

Property Category	Property Name	Description
Development System	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

TABLE 1-6 Sun Cluster HA for SAP Extension Properties for the Application Servers

Property Category	Property Name	Description
SAP Configuration	SAPSID	SAP system name or <i>SAPSID</i> 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	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

TABLE 1-6 Sun Cluster HA for SAP Extension Properties for the Application Servers
(Continued)

Property Category	Property Name	Description
Stopping SAP	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	Probe_timeout	Time-out 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.

```
# scrgadm -a -t SUNW.sap_ci | SUNW.sap_ci_v2
```

3. Add the HAStoragePlus resource to the central instance resource group.

```
# scrgadm -a -t SUNW.HAStoragePlus
# scrgadm -a -j ci-storage-resource \
-g sap-ci-resource-group \
-t SUNW.HAStoragePlus -x filesystemmountpoints=mountpoint, ... |
```

For more details on how to set up an HAStoragePlus resource, see “Enabling Highly Available Local File Systems” in *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

4. Enable the central instance storage resource.

```
# scswitch -e -j ci-storage-resource
```

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

```
# scrgadm -a -j sap-ci-resource \
-g sap-ci-resource-group\
-t SUNW.sap_ci | SUNW.sap_ci_v2
-x SAPSID=SAPSID -x Ci_instance_id=ci-instance-id \
-x Ci_startup_script=ci-startup-script \
-x Ci_shutdown_script=ci-shutdown-script \
-y resource_dependencies=ci-storage-resource
```

See *Sun Cluster 3.1 Data Service Planning and Administration Guide* for a list of extension properties.

6. Enable the failover resource group that now includes the SAP central instance resource.

```
# scswitch -Z -g sap-ci-resource-group
```

If you configure the central instance resource to shut down a development system, you will receive the following console message.

```
ERROR : SAPSYSTEMNAME not set
Please check environment and restart
```

This message displays when the central instance starts on a node that does not have the development system installed and that is not meant to run the central instance. SAP renders this message, and you can safely ignore it.

Where to Go From Here

Go to “How to Register and Configure Sun Cluster HA for SAP as a Failover Data Service” on page 38 or “How to Register and Configure Sun Cluster HA for SAP as a Scalable Data Service” on page 39.

▼ 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.**

```
# scrgadm -a -t SUNW.sap_as | SUNW.sap_as_v2
```

3. **Add the HAStoragePlus resource to the failover application server resource group.**

```
# scrgadm -a -t SUNW.HAStoragePlus
# scrgadm -a -j sap-as-storage-resource -g sap-as-fo-resource-group \
-t SUNW.HAStoragePlus \
-x filesystemmountpoints=mountpoint, ...
```

For more details on how to set up an HAStoragePlus resource, see “Enabling Highly Available Local File Systems” in *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

4. **Enable the failover application server storage resource.**

```
# scswitch -e -j sap-as-storage-resource
```

5. **Create SAP application server resources in their failover resource group.**

```
# scrgadm -a -j sap-as-resource \
-g sap-as-fo-resource-group \
-t SUNW.sap_as | SUNW.sap_as_v2
-x SAPSID=SAPSID -x As_instance_id=as-instance-id \
-x As_startup_script=as-startup-script \
-x As_shutdown_script=as-shutdown-script \
-y resource_dependencies=sap-as-storage-resource
```

See “Sun Cluster HA for SAP Extension Properties” on page 30 for a list of extension properties.

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

```
# scswitch -z -g sap-as-fo-resource-group
```

Where to Go From Here

Go to “How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance” on page 43.

▼ 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.**

```
# scrgadm -a -g sap-as-sa-appinstanceid-resource-group \  
-y Maximum primaries=value \  
-y Desired primaries=value
```

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.

Note – If you are using the SUNW.RGoffload resource type to offload an application server within this scalable application server resource group, then set `Desired primaries=0`. See “Freeing Node Resources by Offloading Non-critical Resource Groups” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for more information about using the SUNW.RGoffload resource type.

3. **Register the resource type for the scalable application server.**

```
# scrgadm -a -t SUNW.sap_as_v2
```

4. **Add the HAStoragePlus resource to the failover application server resource group.**

```
# scrgadm -a -t SUNW.HAStoragePlus  
# scrgadm -a -j sap-as-storage-resource -g \  
-g sap-as-sa-appinstanceid-resource-group \  
-t SUNW.HAStoragePlus \  
-x filesystemmountpoints=mountpoint, ... \  

```

For more details on how to set up an HAStoragePlus resource, see “Enabling Highly Available Local File Systems” in *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

5. Enable the failover application server storage resource.

```
# scswitch -e -j sap-as-storage-resource
```

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

```
# scrgadm -a -j sap-as-resource \  
-g sap-as-sa-appinstanceid-resource-group \  
-t SUNW.sap_as_v2 \  
-x SAPSID=SAPSID \  
-x As_instance_id=as-instance-id \  
-x As_startup_script=as-startup-script \  
-x As_shutdown_script=as-shutdown-script \  
-y resource_dependencies=sap-as-storage-resource
```

See “Sun Cluster HA for SAP Extension Properties” on page 30 for a list of extension properties.

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

- If you *do not* use the `RGoffload` resource type with this application server, use the following command.

```
# scswitch -Z -g sap-as-sa-appinstanceid-resource-group
```

- If you use the `RGoffload` resource type with this application server, use the following command.

```
# scswitch -z -h node1, node2 -g sap-as-sa-appinstanceid-resource-group
```

Note – If you use the `SUNW.RGoffload` resource type with this application server, you must specify which node you want to bring the resource online by using the `-z` option instead of the `-j` option.

Where to Go From Here

Go to “How to Verify Sun Cluster HA for SAP Installation and Configuration and Central Instance” on page 43.

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 crash each other. Furthermore, the crash 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 up the same instance outside of the Sun Cluster environment and then try to bring up SAP under the Sun Cluster environment, the Sun Cluster HA for SAP data service will attempt to start up 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 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 central instance lock file or the failover application server lock file 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 central instance or failover application server. This profile is the one that is passed to `sapstart` as a parameter in the `startsap` script.**

For central instance, enter the following.

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

For failover application server, enter the following.

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

```
sapstart/lockfile
```

New parameter name.

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

Work directory for the central instance.

```
/usr/sap/SID/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/SID/Dinstance-id/work/startup_lockfile
```

```
sapstart/lockfile
```

New parameter name.

```
/usr/sap/local/SID/Dinstance-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 *sapsidadm*, 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.**

```
# scswitch -z -h node2 -g sap-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 of the potential nodes on which the SAP central instance can run.**

Where to Go From Here

Go to “How to Verify the Installation and Configuration of Sun Cluster HA for SAP as a Failover Data Service” on page 43 or “How to Verify Sun Cluster HA for SAP Installation and Configuration of as a Scalable Data Service” on page 44.

▼ 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 *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 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.

```
# scswitch -z -h node2 -g sap-as-resource-group
```
5. Verify that the SAP application server starts on the node you identified in Step 4.
6. Repeat Step 1 through Step 5 until you have tested all of the potential nodes on which the SAP application server can run.

▼ How to Verify Sun Cluster HA for SAP Installation and Configuration of 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 *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 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 of 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 availability uses has no impact on SAP resources' failure history. 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 resource that uses 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. Evaluate 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 bring up the application server on a different node if there is another node available in the cluster.

Index

A

adding
 See modifying

C

changing
 See modifying
commands
 prtconf -v, 10
 prtdiag -v, 10
 psrinfo -v, 10
 showrev -p, 10
configuring
 SAP database, 21
 Sun Cluster HA for SAP, 37, 38, 39

D

deleting
 See modifying
disabling
 See stopping
displaying
 See listing

E

editing
 See modifying
enabling
 See starting
executing
 See starting
extension properties
 Sun Cluster HA for SAP
 As_db_retry_interval (Application Servers), 37
 As_instance_id (Application Servers), 37
 As_services_string (Application Servers), 37
 As_shutdown_script (Application Servers), 37
 As_startup_script (Application Servers), 37
 Check_ms_retry (Central Instance), 35
 Ci_instance_id (Central Instance), 35
 Ci_services_string (Central Instance), 35

Sun Cluster HA for SAP (Continued)

- Ci_shutdown_script (Central Instance), 35
- Ci_start_retry_interval (Central Instance), 35
- Ci_startup_script (Central Instance), 35
- Dev_sapsid (Central Instance), 35
- Dev_shutdown_script (Central Instance), 35
- Dev_stop_pct (Central Instance), 35
- Lgtst_ms_with_logicalhostname (Central Instance), 35
- Message_server_name (Central Instance), 35
- Monitor_retry_count (Application Servers), 37
- Monitor_retry_count (Central Instance), 35
- Monitor_retry_interval (Application Servers), 37
- Monitor_retry_interval (Central Instance), 35
- Probe_timeout (Application Servers), 37
- Probe_timeout (Central Instance), 35
- SAPSID (Application Servers), 37
- SAPSID (Central Instance), 35
- shutdown_dev (Central Instance), 35
- Stop_sap_pct (Application Servers), 37
- Stop_sap_pct (Central Instance), 35

F

- fault monitor, Sun Cluster HA for SAP, 45

I

- installing
 - SAP database, 21
 - SAP scalable application server, 22
 - Sun Cluster HA for SAP, 29
 - DBMS, 26

L

- lock file
 - for central instance
 - Sun Cluster HA for SAP, 41
 - for scalable application server
 - Sun Cluster HA for SAP, 42
 - Sun Cluster HA for SAP, 40

O

- overview, Sun Cluster HA for SAP, 12

P

- printing
 - See listing*
- prtconf -v command, 10
- psrinfo -v command, 10
- pv command, 10

R

- registering
 - Sun Cluster HA for Oracle as a failover data service, 38
 - Sun Cluster HA for Oracle as a scalable data service, 39
 - Sun Cluster HA for Oracle with central instance, 37
- running
 - See enabling*

S

- SAP
 - See also Sun Cluster HA for SAP*
 - database
 - installing and configuring, 21
 - installation
 - enabling failover SAP instances, 25
 - preparing the nodes, 21
 - verifying, 27
 - verifying failover application server, 28

SAP (Continued)

- scalable application server
 - installing, 22
- scalable installation
 - verifying, 29

Sun Cluster HA for SAP

See also SAP

Application servers extension properties

- As_db_retry_interval, 37
- As_instance_id, 37
- As_services_string, 37
- As_shutdown_script, 37
- As_startup_script, 37
- Monitor_retry_count, 37
- Monitor_retry_interval, 37
- Probe_timeout, 37
- SAPSID, 37
- Stop_sap_pct, 37

Central instance extension properties

- Check_ms_retry, 35
- Ci_instance_id, 35
- Ci_services_string, 35
- Ci_shutdown_script, 35
- Ci_start_retry_interval, 35
- Ci_startup_script, 35
- Dev_sapsid, 35
- Dev_shutdown_script, 35
- Dev_stop_pct, 35
- Lgtst_ms_with_logicalhostname, 35
- Message_server_name, 35
- Monitor_retry_count, 35
- Monitor_retry_interval, 35
- Probe_timeout, 35
- SAPSID, 35
- shutdown_dev, 35
- Stop_sap_pct, 35

configuration

- considerations, 17
- DBMS, 26
- planning, 14
- planning questions, 18
- requirements, 15
- restrictions, 14
- standard, 16

fault monitor, 45

- central instance, 45
- fault probes, 46

Sun Cluster HA for SAP (Continued)

- installation
 - planning, 14
 - Sun Cluster HA for SAP packages, 29
 - verifying, 43
 - verifying central instance, 43
 - verifying failover installation, 43
 - verifying scalable installation, 44
- overview, 12
- registering and configuring
 - as a failover data service, 38
 - as a scalable data service, 39
 - with central instance, 37
- setting up a lock file, 40
 - central instance, 41
 - scalable application server, 42
- supported packages, 18
- task map, 13
- upgrading, 19

T

- task map, Sun Cluster HA for SAP, 13
- turning off
 - See* disabling
- turning on
 - See* enabling

U

- upgrading, Sun Cluster HA for SAP, 19
- `/usr/cluster/bin/scinstall -pv`, 10

V

- verifying
 - SAP installation
 - Database installation with central instance, 27
 - Failover application server, 28
 - SAP scalable installation, 29
 - Sun Cluster HA for SAP, 43
 - central instance installation, 43
 - failover installation, 43

Sun Cluster HA for SAP (Continued)
scalable installation, 44