



Sun Cluster 2.2 7/00 Data Services Update: SAP With Oracle, SAP With Informix

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 U.S.A.
650-960-1300

Part No. 806-4886
August 2000, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 U.S.A. All rights reserved.

This product or document is 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, AnswerBook2, docs.sun.com, 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 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est 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 des systèmes 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, AnswerBook2, docs.sun.com, 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.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Adobe PostScript

Contents

Preface	v
Before You Read This Book	v
Typographic Conventions	vi
Shell Prompts	vi
Related Documentation	vii
Ordering Sun Documentation	vii
Accessing Sun Documentation Online	vii
1. Setting Up and Administering Sun Cluster HA for SAP	1
Sun Cluster HA for SAP Overview	2
Supported Configurations	3
Two-Node Cluster With One Logical Host	3
Two-Node Cluster With One Logical Host, Application Servers, and Separate NFS Cluster	5
Two-Node Cluster With Two Logical Hosts	6
Two-Node Cluster With Two Logical Hosts, Application Servers, and Separate NFS Cluster	7
Configuration Guidelines for Sun Cluster HA for SAP	9
Space Considerations	9
Sun Cluster Software Upgrade Considerations	10

Configuration Options for Application Servers and Test/Development Systems	10
Sun Cluster HA for NFS Considerations	13
SAP With Oracle	14
Overview of Procedures (SAP With Oracle)	14
Installation Worksheet for Sun Cluster HA for SAP (SAP With Oracle)	16
Installing and Configuring SAP and the Database (SAP With Oracle)	17
Configuring Sun Cluster HA for SAP (SAP With Oracle)	38
Configuration Parameters for Sun Cluster HA for SAP (SAP With Oracle)	40
Setting Data Service Dependencies for SAP (SAP With Oracle)	43
SAP With Informix	46
Installation and Configuration Overview (SAP With Informix)	46
Installation Worksheet for Sun Cluster HA for SAP (SAP With Informix)	48
Installation and Configuration Procedures (SAP With Informix)	49
Configuration Parameters for Sun Cluster HA for SAP (SAP With Informix)	67
Setting Data Service Dependencies for SAP (SAP With Informix)	69

Preface

Sun™ Cluster 2.2 supports SAP with Oracle and SAP with Informix. Either can be used with the Sun Cluster HA for SAP data service included with Sun Cluster 2.2.

This AnswerBook™ provides installation and configuration instructions for SAP with Oracle, SAP with Informix, and Sun Cluster HA for SAP. These instructions are designed for an experienced system administrator with knowledge of Sun Cluster and SAP.

For the most current information about supported versions of Solaris software, databases, and data services, see the latest version of the *Sun Cluster 2.2 Release Notes*.

Before You Read This Book

Before using the procedures described in this AnswerBook, you should already be familiar with cluster concepts such as logical hosts, multihost disks, switchovers and failovers, as described in the *Sun Cluster 2.2 Software Installation Guide*. The procedures in this book are intended for system administrators with cluster experience.

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> 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
Sun Cluster 2.2 software	<i>Sun Cluster 2.2 Software Installation Guide</i>	806-4648
	<i>Sun Cluster 2.2 System Administration Guide</i>	806-4649
	<i>Sun Cluster 2.2 API Developer's Guide</i>	806-4650
	<i>Sun Cluster 2.2 Cluster Volume Manager Guide</i>	806-4651
	<i>Sun Cluster 2.2 Error Messages Manual</i>	805-4242
	<i>Sun Cluster 2.2 Release Notes</i>	806-4652
Sun Cluster 2.2 hardware	<i>Sun Cluster 2.2 Hardware Site Preparation, Planning, and Installation Guide</i>	806-4748
	<i>Sun Cluster 2.2 Hardware Service Manual</i>	806-4749

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc. For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

<http://www1.fatbrain.com/documentation/sun>

Accessing Sun Documentation Online

The `docs.sun.com`SM 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>

Setting Up and Administering Sun Cluster HA for SAP

Sun Cluster HA for SAP is SAP components made highly available by running in a Sun Cluster environment. This chapter provides instructions for planning and configuring Sun Cluster HA for SAP on Sun Cluster servers.

Sun Cluster HA for SAP Overview

The Sun Cluster HA for SAP data service eliminates single points of failure in a SA system and also provides fault monitoring and failover mechanisms for the SAP application.

These basic services of the SAP system should be placed within the Sun Cluster framework:

- Database instance
- Central instance (consisting of message server, enqueue server, and dispatcher)
- NFS file service

In a Sun Cluster configuration, protection of SAP components is best provided as described in TABLE 1-1.

TABLE 1-1 Protection of SAP Components

SAP Component	Protected by...
SAP database instance	Sun Cluster HA for Oracle or Sun Cluster HA for Informix
SAP central instance	Sun Cluster HA for SAP
NFS file service	Sun Cluster HA for NFS
SAP application servers	SAP, through redundant configuration

The Sun Cluster HA for SAP data service can be installed during or after initial cluster installation. Before you register and start Sun Cluster HA for SAP, you must have a functioning cluster that already contains logical hosts and associated IP addresses and disk groups.

See Chapter 3, "Installing and Configuring Sun Cluster Software," in the *Sun Cluster 2.2 Software Installation Guide* for details about initial installation of clusters and data services. The Sun Cluster HA for SAP data service can be registered after the basic components of the Sun Cluster and SAP software have been installed.

Supported Configurations

See your Enterprise Services representative for the most current information about supported SAP versions. More information on each configuration type is provided in the following sections.

Two-Node Cluster With One Logical Host

The simplest SAP cluster configuration is a two-node cluster with one logical host, as illustrated in FIGURE 1-1. In this asymmetric configuration, the SAP central instance and database instance (collectively called the central system), are both placed on one node. NFS is also placed on the same node. This configuration is relatively easy to configure and administer. A drawback is that the backup node is underutilized. In case of failover, the central instance, database instance, and NFS service are switched to the backup node.

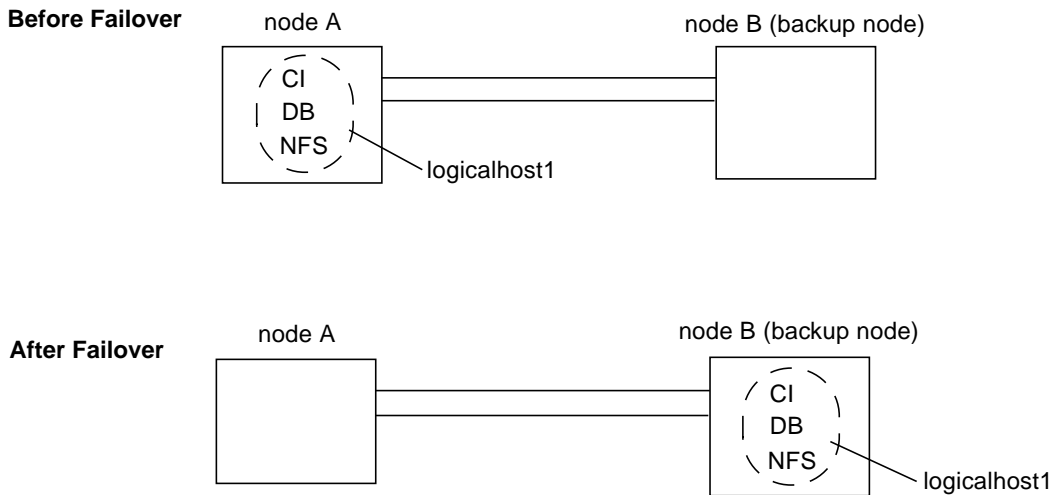


FIGURE 1-1 Asymmetric SAP Configuration

Two-Node Cluster With One Logical Host and Development or Test System

In this configuration, the central system (the central instance and database instance) is placed on one node and a development or test system is placed on a backup node. The development or test system remains running until a failover of the logical host

moves the central system to the backup node. This scenario is illustrated in FIGURE 1-2. In this configuration, you must customize the Sun Cluster HA for SAP `hasap_stop_all_instances` script such that the development or test system is shut down before the SAP central instance is switched over and brought up. See the `hasap_stop_all_instances(1M)` man page and “Configuration Options for Application Servers and Test/Development Systems” for more information.

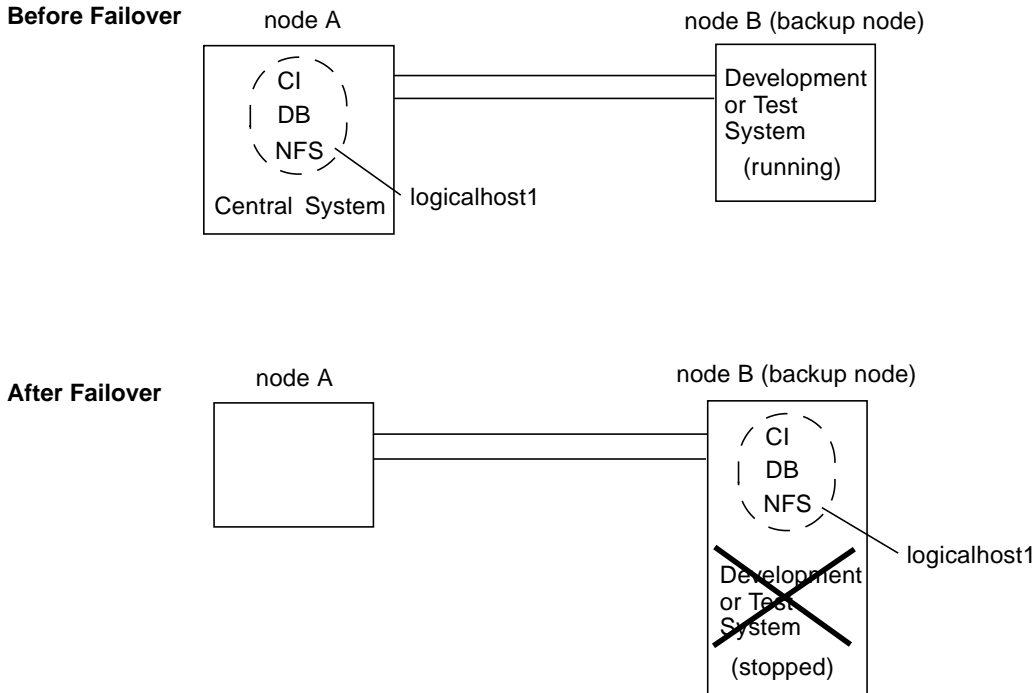


FIGURE 1-2 Asymmetric SAP Configuration With Development or Test System

Two-Node Cluster With One Logical Host, Application Servers, and Separate NFS Cluster

You can also place SAP application servers on one or both physical hosts. In this configuration, you must provide NFS services from a host outside the cluster. Set up the application servers to NFS-mount the file systems from the external NFS cluster, as illustrated in FIGURE 1-3. In case of failover, the logical host containing the central system (the central instance and database instance) switches to the backup node. The application servers do not migrate with the logical host, but are instead started or shut down depending on where the logical host is mastered. This prevents the application servers from competing for resources with the central instance and database.

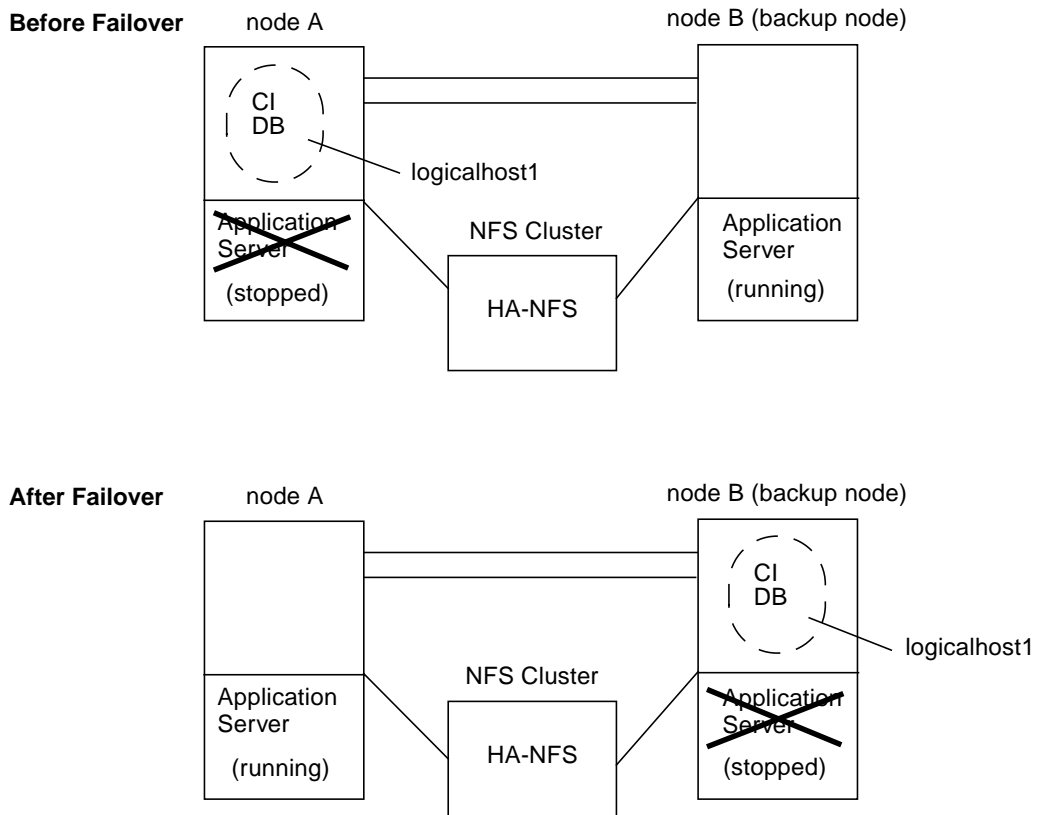


FIGURE 1-3 Asymmetric SAP Configuration With Application Servers and External HA-NFS

Two-Node Cluster With Two Logical Hosts

A two-node cluster with two logical hosts can be configured with the SAP central instance on one logical host and the SAP database instance on the other logical host, as illustrated in FIGURE 1-4. In this configuration, the nodes are load-balanced and both are utilized. In case of failover, the central instance or database instance is switched to the sibling node.

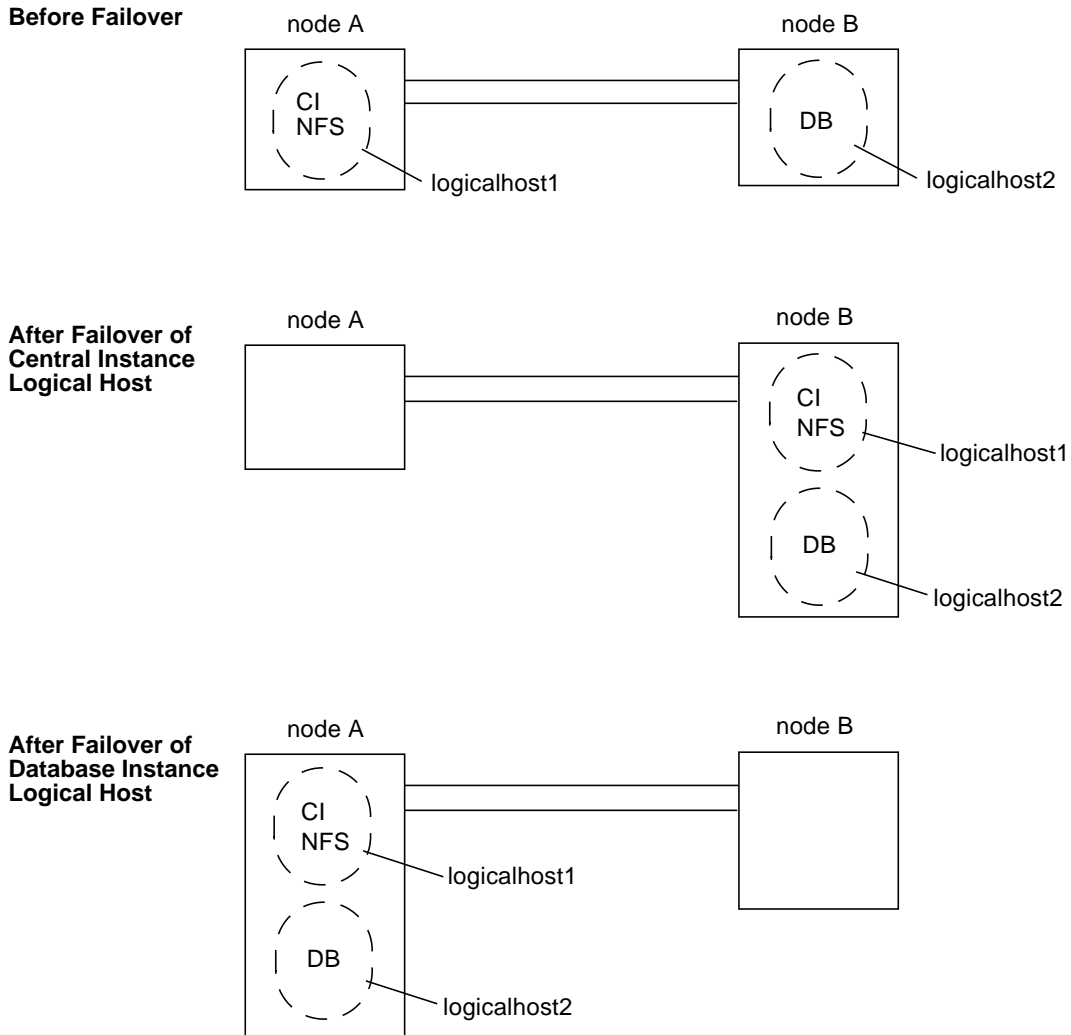


FIGURE 1-4 Symmetric SAP Configuration With Two Logical Hosts

Two-Node Cluster With Two Logical Hosts, Application Servers, and Separate NFS Cluster

A two-node cluster with two logical hosts can be configured with SAP application servers on one or both physical hosts. In this configuration, you must provide NFS services from a host outside the cluster. Set up the application servers to NFS-mount the file systems from the external NFS cluster, as illustrated in FIGURE 1-5. In this case, both nodes are utilized and load-balanced.

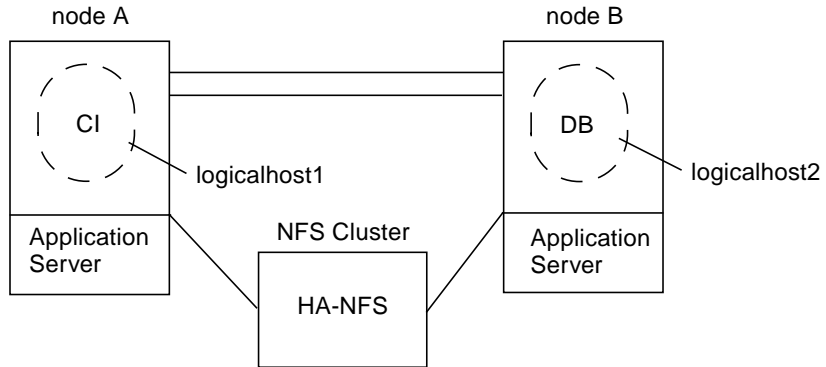
In case of failover, the logical hosts switch over to the sibling node. The application servers do not fail over.

If the central instance logical host fails over, the application server can be shut down through the `hasap_stop_all_instances` script.

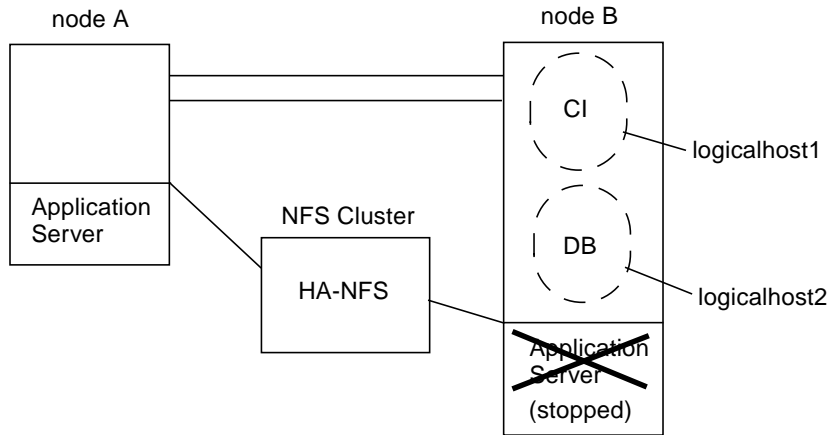
There are no customizable scripts to start and stop application servers in case of failover of the database logical host. If the database logical host fails over, the application servers cannot be shut down to release resources for the database logical host. Therefore, you must size your configuration to allow for the possible scenario in which the central instance, database instance, and application server are all running on the same node simultaneously.

In this configuration, NFS is protected by Sun Cluster HA for NFS. For more information, see “Sun Cluster HA for NFS Considerations”.

Before Failover



After Failover of Central Instance Logical Host



After Failover of Database Instance Logical Host

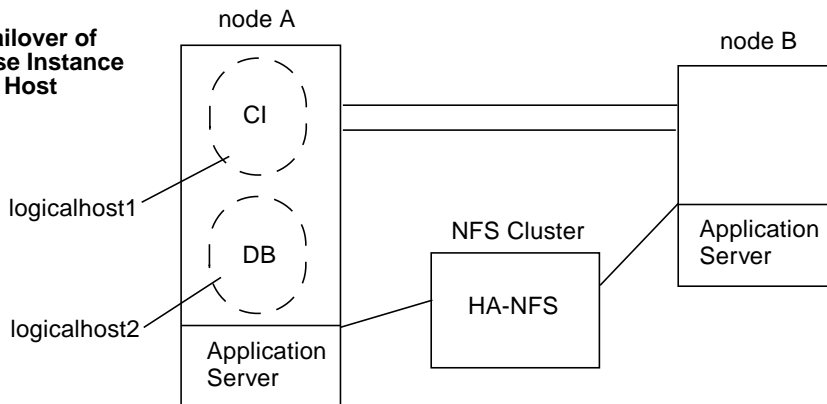


FIGURE 1-5 Symmetric SAP Configuration With Two Logical Hosts and Application Servers

Configuration Guidelines for Sun Cluster HA for SAP

Consider these general guidelines when designing a Sun Cluster HA for SAP configuration:

- Be generous in estimating the total possible load on standby servers in case of failover. Allocate ample resources for CPU, swap, shared memory, and I/O bandwidth on the standby server, because in case of failover, the central instance and database instance might co-exist on the standby.
- Use a logging file system:
 - If your volume manager is SSVN, use VxFS and Dirty Region Logging.
 - If your volume manager is Solstice DiskSuite, use either Solaris UFS logging or Solstice DiskSuite UFS logging.
- Configure separate disk groups for SAP software and the database. The `scinstall(1M)` command cannot configure more than one disk group per logical host. Therefore, do not set up logical hosts with `scinstall(1M)` during initial cluster installation. Instead, set up logical hosts with `scconf(1M)` after the cluster is up. See the `scconf(1M)` man page for details.
- Limit host names to eight characters or less, if possible. If your host names are longer than eight characters, modify the `/etc/hosts` file to alias the actual host names to shorter names.
- As per SAP guidelines, limit the central instance profile to Enqueue, Message, one Dialog and one Update work process. Do not permit SAP users to connect to SAP through the central instance. Instead, encourage all users to connect to an alternate application server. System administrators and Sun Cluster HA for SAP can connect to the central instance through the single Dialog work process.

Space Considerations

SAP and the database use a large amount of memory and swap space. Consult your SAP and database documentation for additional recommendations.

- On all potential masters of the central instance logical host, set aside space in `/var/opt/informix` or `/var/opt/oracle` for the database binaries. At least 280 Megabytes is required. See your SAP documentation for details.

Sun Cluster Software Upgrade Considerations

Note these SAP-related issues before performing an upgrade to Sun Cluster 2.2 from HA 1.3 or Sun Cluster 2.1.

- On each node, if you customized `hasap_start_all_instances` or `hasap_stop_all_instances` scripts in HA 1.3 or Sun Cluster 2.1, save them to a safe location before beginning the upgrade to Sun Cluster 2.2. Restore the scripts after completing the upgrade. Do this to prevent loss of your customizations when Sun Cluster 2.2 removes the old scripts.
- The configuration parameters implemented in Sun Cluster 2.2 are different from those implemented in HA 1.3 and Sun Cluster 2.1. Therefore, after upgrading to Sun Cluster 2.2, you will have to re-configure Sun Cluster HA for SAP by running the `hadsconfig(1M)` command.
- Before starting the upgrade, view the existing configuration and note the current configuration variables. For HA 1.3, use the `hainetconfig(1M)` command to view the configuration. For Sun Cluster 2.1, use the `hadsconfig(1M)` command to view the configuration. After upgrading to Sun Cluster 2.2, use the `hadsconfig(1M)` command to re-create the instance.
- In Sun Cluster 2.2, the `hareg -n` command shuts down the entire Sun Cluster HA for SAP data service, including all instances and fault monitors. In previous releases, the `hareg -n` command, when used with Sun Cluster HA for SAP, shut down only the fault monitors.

Additionally, before turning on the Sun Cluster HA for SAP data service with `hareg -y`, you must stop the SAP central instance. Otherwise, the Sun Cluster HA for SAP data service will not be able to start and monitor the instance properly.

Configuration Options for Application Servers and Test/Development Systems

Conventionally you stop and restart the application server instances manually after the central instance is restarted. Sun Cluster HA for SAP provides hooks that are called whenever the central instance logical host switches over or fails over. These hooks are provided by calling the `hasap_stop_all_instances` and `hasap_start_all_instances` scripts. The scripts must be idempotent.

If you configure application servers and want to control them automatically when the logical host switches over or fails over, you can create start and stop scripts according to your needs. Sun Cluster provides sample scripts that can be copied and customized:

```
/opt/SUNWcluster/ha/sap/hasap_stop_all_instances.sample and  
/opt/SUNWcluster/ha/sap/hasap_start_all_instances.sample.
```

Customization examples are included in these scripts. Copy the sample scripts, rename them by removing the “.sample” suffix, and modify them as appropriate.

After failovers, Sun Cluster HA for SAP will invoke the customized scripts to restart the application servers. The scripts control the application servers from the central instance, and are invoked by the full path name.

If you include a test or development system in your configuration, modify the `hasap_stop_all_instances` script to stop the test or development system in case of failover of the central instance logical host.

During a central instance logical host switchover or failover, the scripts are called in the following sequence:

1. Stopping the application server instances and test or development systems by calling `hasap_stop_all_instances`
2. Stopping the central instance
3. Switching over the logical host(s) and disk group(s)
4. Calling `hasap_stop_all_instances` again to make sure all application servers and test or development systems have stopped
5. Starting the central instance
6. Starting the application server instances by calling `hasap_start_all_instances` (see the `hasap_start_all_instances(1M)` and `hasap_stop_all_instances(1M)` man pages for more information)

Additionally, you must enable root access to the SAP administrative account (`<sapsid>adm`) on all SAP application servers and test or development systems from all logical hosts and all physical hosts in the cluster. For test or development systems, also grant root access to the database administrative account (`ora<sapsid>`). Accomplish this by creating `.rhosts` files for these users. For example:

```
...
phys-hahost1 root
phys-hahost2 root
phys-hahost3 root
hahost1      root
hahost2      root
hahost3      root
...
```

In configurations including several application servers or a test or development system, consider increasing the timeout value of the `STOP_NET` method for Sun Cluster HA for SAP. Increasing the `STOP_NET` timeout value is necessary only if the

hasap_stop_all_instances script takes longer to execute than 60 seconds, because 60 seconds is the default timeout value for the STOP_NET method. If the hasap_stop_all_instances script does not finish within the 60 second timeout, then increase the STOP_NET timeout value.

Check the timeout value of the STOP_NET method by using the following command:

```
# haresg -q sap -T STOP_NET
```

Note – The hasap_dbms command can be used only when Sun Cluster HA for SAP is registered but is in the off state. Run the command on only one node, while that node is a member of the cluster. See the hasap_dbms(1M) man page for more information.



Caution – If the hasap_dbms(1M) command returns an error stating that it cannot add rows to or update the CCD, it might be because another cluster utility is also trying to update the CCD. If this occurs, re-run hasap_dbms(1M) until it runs successfully. After the hasap_dbms(1M) command runs successfully, verify that all necessary rows are included in the resulting CCD by running the command haresg -q sap.

If the haresg(1M) command returns an error, then first restore the original method timeouts by running the command hasap_dbms -f. Second, restore the default dependencies by running the command hasap_dbms -r. After both commands complete successfully, retry the original hasap_dbms(1M) command to configure new dependencies and method timeouts. See the hasap_dbms(1M) man page for more information.

Increase the STOP_NET timeout value by using the following command:

```
# /opt/SUNWcluster/ha/sap/hasap_dbms -t STOP_NET=new_timeout_value
```

If you increase the STOP_NET method timeout value, you also must increase the timeouts that the Sun Cluster framework uses when remastering logical hosts during cluster reconfiguration. Use the sconfg(1M) command to increase logical host timeout values. Refer to Section 3.15, “Configuring Timeouts for Cluster Transition Steps,” in the *Sun Cluster 2.2 System Administration Guide* for details about how to increase the timeouts for the cluster framework. Make sure that the loghost_timeout value is at least double the new STOP_NET timeout value.

Sun Cluster HA for NFS Considerations

If you have application servers outside the cluster, you must configure Sun Cluster HA for NFS on the central instance logical host. Application servers outside the cluster must NFS-mount the SAP profile directories and executable directories from the SAP central instance. See Chapter 11, “Setting Up and Administering Sun Cluster HA for NFS,” in the *Sun Cluster 2.2 Software Installation Guide* for detailed procedures on setting up Sun Cluster HA for NFS, and note the following SAP-specific guidelines:

- Do not configure any node to be an NFS client of another node within the same cluster.
- If you will run application servers within the cluster, you must set up an external cluster running NFS. The application servers and central instance will mount files from this NFS cluster.
- There are start order dependencies among Sun Cluster HA for NFS, HA-DBMS, and Sun Cluster HA for SAP data services. You can use special scripts to manage these dependencies. See “Setting Data Service Dependencies for SAP (SAP With Oracle)”, or “Setting Data Service Dependencies for SAP (SAP With Informix)” for more information.
- Usually, you should share the following directories to all SAP instances:
 - `/usr/sap/trans`
 - `/sapmnt/<SAPSID>/exe`
 - `/sapmnt/<SAPSID>/global`
 - `/sapmnt/<SAPSID>/profile`

SAP With Oracle

Use the information in the following sections to install and configure SAP with Oracle. For information on installing and configuring SAP with Informix, see "SAP With Informix".

Overview of Procedures (SAP With Oracle)

TABLE 1-1 summarizes the tasks you must complete to configure SAP.

TABLE 1-2 Sun Cluster HA for SAP Installation Overview (SAP With Oracle)

Task	Description	For Instructions, Go To...
Plan the SAP installation	- Read through all guidelines and procedures	"Sun Cluster HA for SAP Overview" and "Configuration Guidelines for Sun Cluster HA for SAP"
	- Complete the SAP installation worksheet	"Installation Worksheet for Sun Cluster HA for SAP (SAP With Oracle)"
Prepare the environment for SAP	<ul style="list-style-type: none">- Perform all pre-requisite installation tasks- Set up Solaris- Set up the volume manager- Create disk groups or disksets- Create volumes and file systems- Install Sun Cluster- Set up PNM- Set up logical hosts and mount points- Set up HA-NFS, if necessary- Adjust kernel parameters- Create swap space- Create user and group accounts	<p>"How to Prepare the Cluster Environment for SAP and the Database (SAP With Oracle)"</p> <p>See also: Chapter 3, "Installing and Configuring Sun Cluster Software," Appendix B, "Configuring Solstice DiskSuite," and Appendix C, "Configuring Sun StorEdge Volume Manager and Cluster Volume Manager," all in the <i>Sun Cluster 2.2 Software Installation Guide</i></p>
Install and configure SAP and the database	<ul style="list-style-type: none">- Install the SAP central instance and database instance- Load the database- Load all reports- Install the GUI	"How to Install SAP and the Database (SAP With Oracle)"

TABLE 1-2 Sun Cluster HA for SAP Installation Overview (SAP With Oracle) (Continued)

Task	Description	For Instructions, Go To...
Enable SAP to run in the cluster	<ul style="list-style-type: none"> - Set up the SAP central instance admin environment - Modify SAP profile files - Modify the database environment - Update <code>/etc/services</code> and create <code>/usr/sap/tmp</code> - Test the SAP installation 	"How to Enable SAP to Run in the Cluster Environment (SAP With Oracle)"
Configure the HA-DBMS	<ul style="list-style-type: none"> - Shut down SAP and the database - Adjust Oracle alert files and listener files - Register and activate the database - Set up the database instance - Start fault monitoring for the database - Test switchover of the database 	"How to Configure Sun Cluster HA for Oracle"
Configure Sun Cluster HA for SAP	<ul style="list-style-type: none"> - Install and register Sun Cluster HA for SAP 	"How to Configure Sun Cluster HA for SAP (SAP With Oracle)"
	<ul style="list-style-type: none"> - Configure Sun Cluster HA for SAP 	"How to Configure Sun Cluster HA for SAP (SAP With Oracle)", and "Configuration Parameters for Sun Cluster HA for SAP (SAP With Oracle)"
	<ul style="list-style-type: none"> - Set dependencies, if necessary 	"Setting Data Service Dependencies for SAP (SAP With Oracle)"
	<ul style="list-style-type: none"> - Test switchover of Sun Cluster HA for SAP 	"How to Configure Sun Cluster HA for SAP (SAP With Oracle)"
	<ul style="list-style-type: none"> - Customize and test start and stop scripts for the application servers and test/development systems 	"Configuration Options for Application Servers and Test/Development Systems"

Installation Worksheet for Sun Cluster HA for SAP (SAP With Oracle)

Complete the following worksheet before beginning the Sun Cluster HA for SAP installation.

TABLE 1-3 Sun Cluster HA for SAP Installation Worksheet (SAP With Oracle)

Name of Cluster	
Number of logical hosts	
Name and IP address of all physical hosts that are potential masters of the CI logical host	
Name and IP address of CI logical host	
SAP system ID (<SAPSID>)	
SAP system number	
Name and IP address of all physical hosts that are potential masters of the DB logical host	
Name and IP address of DB logical host (In asymmetric configurations, this is identical to the CI logical host)	
Name of NFS logical host (see Note, below)	
SAP license for each potential master of the CI logical host	

Note – If all application servers are external to the cluster, the name of the NFS logical host is the central instance logical host. If the application servers are inside the cluster, the NFS logical host is the logical host that provides NFS service from the external NFS cluster. See “Sun Cluster HA for NFS Considerations”.

Installing and Configuring SAP and the Database (SAP With Oracle)

This section describes how to install and configure SAP with Oracle. For instructions on installing SAP with Informix, see “SAP With Informix”.

▼ How to Prepare the Cluster Environment for SAP and the Database (SAP With Oracle)

Before installing SAP and Oracle, perform the following tasks.

1. On all nodes, install the Solaris operating environment and Solaris patches.

See Chapter 3, “Installing and Configuring Sun Cluster Software” in the *Sun Cluster 2.2 Software Installation Guide*.

2. On all nodes, install Volume Manager software and any required Volume Manager patches.

See Chapter 3, “Installing and Configuring Sun Cluster Software” in the *Sun Cluster 2.2 Software Installation Guide*.

3. On the node on which you will install SAP and Oracle, create Solstice DiskSuite disksets or SSVM disk groups.

Separate disk groups for the SAP central instance and database instance are recommended, for ease of administration.

4. On the node on which you will install SAP and Oracle, create volumes according to Sun Cluster guidelines:

- Mirror volumes across controllers
- With SSVM, use Dirty Region Logging for faster mirror resynchronization

- Use a logging file system for faster logical host failover

Use TABLE 1-4 as a worksheet to capture the name of the volume that corresponds to each file system used for the SAP central instance. Refer to the SAP installation guide for the file system sizes recommended for your particular configuration. These are database-independent file systems.

TABLE 1-4 Worksheet: File Systems and Volume Names for the SAP Central Instance (SAP With Oracle)

File System Name / Mount Point	Volume Name
/oracle/805_32 (for SAP 4.6B and SAP 4.6C only)	
/usr/sap/trans	
/sapmnt/<SAPSID>	
/usr/sap/<SAPSID>	

Use TABLE 1-5 as a worksheet to capture the name of the volume that corresponds to each file system used for the database instance. Refer to the SAP installation guide for the file system sizes recommended for your particular configuration. These are database-dependent file systems.

TABLE 1-5 Worksheet: File Systems and Volume Names for the SAP Database Instance (SAP With Oracle)

File System Name / Mount Point	Volume Name
/oracle/<SAPSID>	
/oracle/stage/stage_<version>	
/oracle/<SAPSID>/origlogA	
/oracle/<SAPSID>/origlogB	
/oracle/<SAPSID>/mirrlogA	
/oracle/<SAPSID>/mirrlogB	
/oracle/<SAPSID>/saparch	
/oracle/<SAPSID>/sapreorg	
/oracle/<SAPSID>/sapdata1	
/oracle/<SAPSID>/sapdata2	
/oracle/<SAPSID>/sapdata3	

TABLE 1-5 Worksheet: File Systems and Volume Names for the SAP Database Instance (SAP With Oracle) (Continued)

File System Name / Mount Point	Volume Name
/oracle/<SAPSID>/sapdata4	
/oracle/<SAPSID>/sapdata5	
/oracle/<SAPSID>/sapdata6	

5. On all nodes, install Sun Cluster, Sun Cluster HA for SAP, Sun Cluster HA for Oracle, and any required patches.

Use the procedures described in Chapter 3, “Installing and Configuring Sun Cluster Software” in the *Sun Cluster 2.2 Software Installation Guide*, but do not set up logical hosts with `scinstall(1M)` during this installation. Instead, set up logical hosts with `scconf(1M)` after the cluster is up. Set up two disksets per logical host.

6. On all nodes, configure PNM.

For detailed procedures, see Chapter 3, “Installing and Configuring Sun Cluster Software” in the *Sun Cluster 2.2 Software Installation Guide*, and Chapter 6, “Administering Network Interfaces” in the *Sun Cluster 2.2 System Administration Guide*.

7. Start the cluster.

Run the following command on one node.

```
# scadmin startcluster physicalhost clustername
```

Run the following command on all other nodes, sequentially.

```
# scadmin startnode
```

8. (SSVM only) Verify that all disk groups are deported.

9. (Solstice DiskSuite only) Release ownership of all disksets.

10. On the node on which you installed SAP, create logical hosts with `scconf(1M)`.

The number of logical host depends on your particular configuration. See Section 3.5, “Adding and Removing Logical Hosts,” in the *Sun Cluster 2.2 System Administration Guide*. You will need:

- Logical host name(s)
- Physical host names of potential masters of logical host(s)
- Names of the primary public network controllers for the potential masters of the logical host(s)

- Disk group name(s)

When you create logical hosts, disable the automatic failback mechanism by using the `-m` option to `scconf(1M)`.

11. (SSVM, two-node configurations only) Configure the shared CCD.

See Appendix C, “Configuring Sun StorEdge Volume Manager and Cluster Volume Manager” in the *Sun Cluster 2.2 Software Installation Guide*.

12. After creating the logical host(s), create the logical host administrative file system.

For detailed procedures, see Appendix B, “Configuring Solstice DiskSuite” or Appendix C, “Configuring Sun StorEdge Volume Manager and Cluster Volume Manager” in the *Sun Cluster 2.2 Software Installation Guide*.

13. Create mount points for the central instance and database instance volumes, and enter them into the respective `vfstab.logicalhost` files on all potential masters of each logical host.

The `vfstab.logicalhost` files are located in `/etc/opt/SUNWcluster/conf/hanfs`.

TABLE 1-6 lists the suggested file system mount points for the disk groups (SSVM) or disksets (Solstice DiskSuite) associated with the central instance and database instance. Note that separating the central instance and database instance file systems into separate disk groups or disksets (even if using a single logical host) may provide more configuration flexibility in the future.

TABLE 1-6 File Systems and Mount Points for the SAP Central Instance and Database Instance (SAP With Oracle)

Disk Group (SSVM)	Diskset (Solstice DiskSuite)	Volume Name	Mount Point
ci_dg	<i>CIloghost</i>	sap	<code>/usr/sap/<SAPSID></code>
ci_dg	<i>CIloghost</i>	saptrans	<code>/usr/sap/trans</code>
ci_dg	<i>CIloghost</i>	sapmnt	<code>/sapmnt/<SAPSID></code>
db_dg	<i>DBloghost</i>	oracle	<code>/oracle/<SAPSID></code>
db_dg	<i>DBloghost</i>	stage	<code>/oracle/stage/stage_<version></code>
db_dg	<i>DBloghost</i>	origlogA	<code>/oracle/<SAPSID>/origlogA</code>
db_dg	<i>DBloghost</i>	origlogB	<code>/oracle/<SAPSID>/origlogB</code>
db_dg	<i>DBloghost</i>	mirrlogA	<code>/oracle/<SAPSID>/mirrlogA</code>
db_dg	<i>DBloghost</i>	mirrlogB	<code>/oracle/<SAPSID>/mirrlogB</code>
db_dg	<i>DBloghost</i>	saparch	<code>/oracle/<SAPSID>/saparch</code>
db_dg	<i>DBloghost</i>	sapreorg	<code>/oracle/<SAPSID>/sapreorg</code>
db_dg	<i>DBloghost</i>	sapdata1	<code>/oracle/<SAPSID>/sapdata1</code>

TABLE 1-6 File Systems and Mount Points for the SAP Central Instance and Database Instance (SAP With Oracle) (Continued)

Disk Group (SSVM)	Diskset (Solstice DiskSuite)	Volume Name	Mount Point
db_dg	DBloghost	sapdata2	/oracle/<SAPSID>/sapdata2
db_dg	DBloghost	sapdata3	/oracle/<SAPSID>/sapdata3
db_dg	DBloghost	sapdata4	/oracle/<SAPSID>/sapdata4
db_dg	DBloghost	sapdata5	/oracle/<SAPSID>/sapdata5
db_dg	DBloghost	sapdata6	/oracle/<SAPSID>/sapdata6

- 14. If SAP application servers will be configured outside the cluster, then configure Sun Cluster HA for NFS and enter the appropriate shared file systems into the `dfstab.logicalhost` files on all potential masters of each logical host.**

These files are located in `/etc/opt/SUNWcluster/conf/hanfs`. See “Configuration Options for Application Servers and Test/Development Systems” and Chapter 11, “Setting Up and Administering Sun Cluster HA for NFS” in the *Sun Cluster 2.2 Software Installation Guide* for more information.

Share the following file systems to SAP application servers outside the cluster. These are general guidelines. See the SAP documentation for more information.

TABLE 1-7 File Systems to Share in HA-NFS to External SAP Application Servers (SAP With Oracle)

File Systems to Share to External Application Servers
<code>/usr/sap/trans</code>
<code>/sapmnt/<SAPSID>/exe</code>
<code>/sapmnt/<SAPSID>/profile</code>
<code>/sapmnt/<SAPSID>/global</code>

- 15. Test the functionality and mount points of the logical host(s) by switching them between all potential masters.**

This verifies that all mount points have been created correctly.

- 16. Adjust kernel parameters on all potential masters, as per the “R/3 Installation on UNIX: OS Dependencies” guidelines in the SAP documentation.**

In configurations where the central instance and database instance may coexist with each other or with other instances, be sure to size the kernel parameters accordingly.

17. **Create appropriately sized permanent swap areas on all potential master nodes.**
See the "Installation Requirements Checklist" in your SAP documentation for swap guidelines. Use the SAP-supplied memlimits utility to assist you in sizing the swap space. See the "R/3 Installation on UNIX" guidelines in the SAP documentation for more information on this utility.
18. **Stop the cluster and reboot all nodes after adjusting kernel parameters and swap space.**
19. **Create SAP and database user and group accounts on all potential masters of the logical hosts.**
Refer to the "R/3 Installation on UNIX: OS Dependencies" guidelines in the SAP documentation for details. User and group IDs must be identical on all nodes. Create the home directories for these users on the shared diskset. TABLE 1-8 shows suggested home directory paths for the user accounts.

TABLE 1-8 Home Directory Paths for SAP User Accounts (SAP With Oracle)

User	Home directory
<sapsid>adm	/usr/sap/<SAPSID>/home
ora<sapsid>	/oracle/<SAPSID>

Note – For SAP 4.0b, read OSS note 0100125 for special steps required when creating user home directories outside of the /home location.

Now proceed to "How to Install SAP and the Database (SAP With Oracle)".

▼ How to Install SAP and the Database (SAP With Oracle)

1. Verify that you have completed all tasks listed in “How to Prepare the Cluster Environment for SAP and the Database (SAP With Oracle)”.
2. Verify that all nodes are running in the cluster.
3. Switch over all logical hosts to the node from which you will install SAP and the database.

```
# scadmin switch clustername phys-hahost1 Cloghost DBloghost ...
```

4. Create the SAP installation directory and begin SAP installation.

Refer to the “R/3 Installation on UNIX” guidelines in the SAP documentation for details.

Note – Read all SAP OSS notes prior to beginning the SAP installation.

- a. **Install the central instance and database instance on the node currently mastering the central instance and database instance logical host.**

(For SAP 3.1x only) When installing SAP using R3INST, specify the physical host name of the current master of the database logical host when prompted for “Database Server.” After the installation is complete, you must manually adjust various files to refer to the logical host where the database resides.

(For SAP 4.0x only) When installing SAP using R3SETUP, select the CENTRDB.SH script to generate the installation command file.

- b. **Continue the SAP installation to install the central instance, to create and load the database, to load all reports, and to install the R/3 Frontend (GUI).**

▼ How to Enable SAP to Run in the Cluster Environment (SAP With Oracle)

1. Set up the SAP central instance administrative environment.

During SAP installation, SAP creates files and shell scripts on the server on which the SAP central instance is installed. These files and scripts use physical host names. Follow these steps to replace all occurrences of physical host names with logical host names.

Note – Make backup copies of all files before performing the following steps.

First, shut down the SAP central instance and database using the following command:

```
# su - <sapsid>adm
$ stopsap all
...
# su - ora<sapsid>
$ lsnrctl stop
```

Note – Become the <sapsid>adm user before editing these files.

a. Revise the `.cshrc` file in the <sapsid>adm home directory.

On the server on which the SAP central instance is installed, the `.cshrc` file contains aliases that use Sun Cluster physical host names. Replace the physical host names with the central instance logical host name.

(For SAP 3.1x only) The resulting `.cshrc` file should look similar to the following example, in which *Clloghost* is the logical host containing the central instance and *DBloghost* is the logical host containing the database. If the central instance and database are on the same logical host, then use that logical host name for the substitutions.

```
# aliases
alias startsap "$HOME/startsap_Clloghost_00"
alias stopsap "$HOME/stopsap_Clloghost_00"

# RDBMS environment
if (-e $HOME/.dbenv_DBloghost.csh) then
    source $HOME/.dbenv_DBloghost.csh
else if (-e $HOME/.dbenv.csh) then
    source $HOME/.dbenv.csh
endif
```

(For SAP 4.0x only) The resulting `.cshrc` file should look similar to the following example, in which *Clloghost* is the logical host containing the central instance and *DBloghost* is the logical host containing the database. If the central instance and database are on the same logical host, then use that logical host name for the substitutions:

```
if ( -e $HOME/.sapenv_Clloghost.csh ) then
    source $HOME/.sapenv_Clloghost.csh
else if ( -e $HOME/.sapenv.csh ) then
    source $HOME/.sapenv.csh
endif

# RDBMS environment
if ( -e $HOME/.dbenv_DBloghost.csh ) then
    source $HOME/.dbenv_DBloghost.csh
else if ( -e $HOME/.dbenv.csh ) then
    source $HOME/.dbenv.csh
endif
```

- b. (For SAP 4.0x only) Rename the file `.sapenv_physicalhost.csh` to `.sapenv_Clloghost.csh`, and edit it to replace occurrences of the physical host name with the logical host name.**

First rename the file, replacing the physical host name with the central instance logical host name.

```
$ mv .sapenv_physicalhost.csh .sapenv_Clloghost.csh
```

Then edit the aliases in the file. For example:

```
alias startsap "$HOME/startsap_Clloghost_00"  
alias stopsap "$HOME/stopsap_Clloghost_00"
```

- c. Rename the `.dbenv_physicalhost.csh` file.**

Rename the `.dbenv_physicalhost.csh` file to `.dbenv_DBloghost.csh`. If the central instance and database are on the same logical host, use that logical host name for the substitution.

```
$ mv .dbenv_physicalhost.csh .dbenv_DBloghost.csh
```

- d. (For SAP 4.0x only) Edit the `.dbenv_DBloghost.csh` file to set the `ORA_NLS` environment variable to point to the appropriate subdirectories of `/var/opt/oracle` for the database client configuration files. Also, set the `TNS_ADMIN` environment variable to point to the `/var/opt/oracle` directory.**

The `.dbenv_DBloghost.csh` file is located in the `<sapsid>adm` home directory.

```
#setenv ORA_NLS /oracle/<SAPSID>/ocommon/NLS_723/admin/data  
setenv ORA_NLS /var/opt/oracle/ocommon/NLS_723/admin/data  
  
#setenv ORA_NLS32 /oracle/<SAPSID>/ocommon/NLS_733/admin/data  
setenv ORA_NLS32 /var/opt/oracle/ocommon/NLS_733/admin/data  
  
#setenv ORA_NLS33 /oracle/<SAPSID>/ocommon/NLS_804/admin/data  
setenv ORA_NLS33 /var/opt/oracle/ocommon/NLS_804/admin/data  
  
...  
  
# setenv TNS_ADMIN @TNS_ADMIN@  
setenv TNS_ADMIN /var/opt/oracle  
...
```

- e. (For SAP 4.6B only) Edit the `.dbenv_DBloghost.csh` file to set the `ORA_NLS` environment variable to point to the appropriate subdirectories of `/var/opt/oracle` for the database client configuration files. Set the `TNS_ADMIN` environment variable to point to the `/var/opt/oracle` directory. Also, set the `LD_LIBRARY_PATH` to `/var/opt/oracle`.

The `.dbenv_DBloghost.csh` file is located in the `<sapsid>adm` home directory.

```
...
#setenv ORA_NLS      /oracle/D01/ocommon/NLS_723/admin/data
#setenv ORA_NLS32   /oracle/D01/ocommon/NLS_733/admin/data
#setenv ORA_NLS33   /oracle/D01/ocommon/nls/admin/data
setenv ORA_NLS      /var/opt/oracle/ocommon/NLS_723/admin/data
setenv ORA_NLS32    /var/opt/oracle/ocommon/NLS_733/admin/data
setenv ORA_NLS33    /var/opt/oracle/ocommon/nls/admin/data
...
# setenv TNS_ADMIN @TNS_ADMIN@
setenv TNS_ADMIN /var/opt/oracle
...
    default:
if ( ! $?LD_LIBRARY_PATH ) then
#setenv LD_LIBRARY_PATH /oracle/805_32/lib
    setenv LD_LIBRARY_PATH /var/opt/oracle/805_32/lib
else
#foreach d ( /oracle/805_32/lib )
    foreach d ( /var/opt/oracle/805_32/lib )
...

```

- f. (For SAP 4.6C only) Edit the `.dbenv_DBloghost.csh` file to set the `ORA_NLS` environment variable to point to the correct subdirectories of `/var/opt/oracle` for the database client configuration files. Set the `TNS_ADMIN` environment variable to point to the `/var/opt/oracle` directory. If the `TNS_ADMIN` variable doesn't exist, create it.

The `.dbenv_DBloghost.csh` file is located in the `<sapsid>adm` home directory.

```
...
#setenv ORA_NLS      /oracle/D01/ocommon/NLS_723/admin/data
#setenv ORA_NLS32   /oracle/D01/ocommon/NLS_733/admin/data
#setenv ORA_NLS33   /oracle/D01/ocommon/nls/admin/data
setenv ORA_NLS      /var/opt/oracle/ocommon/NLS_723/admin/data
setenv ORA_NLS32    /var/opt/oracle/ocommon/NLS_733/admin/data
setenv ORA_NLS33    /var/opt/oracle/ocommon/nls/admin/data
...
# setenv TNS_ADMIN @TNS_ADMIN@
setenv TNS_ADMIN /var/opt/oracle
...

```

g. Rename and revise the SAP instance `startsap` and `stopsap` shell scripts in the `<sapsid>adm` home directory.

On the server on which the SAP central instance is installed, the `<sapsid>adm` home directory contains shell scripts that include physical host names. Rename these shell scripts by replacing the physical host names with logical host names. In this example, `Clloghost` represents the logical host name of the central instance:

```
$ mv startsap_physicalhost_00 startsap_Clloghost_00
$ mv stopsap_physicalhost_00 stopsap_Clloghost_00
```

The `startsap_Clloghost_00` and `stopsap_Clloghost_00` shell scripts specify physical host names in their `START_PROFILE` parameters. Replace the physical host name with the central instance logical host name in the `START_PROFILE` parameters in both files.

```
...
START_PROFILE="START_DVEBMGS00_Clloghost"
...
```

h. Revise the SAP central instance profile files.

During SAP installation, SAP creates three profile files on the server on which the SAP central instance is installed. These files use physical host names. Use these steps to replace all occurrences of physical host names with logical host names. To revise these files, you must be user `<sapsid>adm`, and you must be in the profile directory.

- Rename the `START_DVEBMGS00_physicalhost` and `<SAPSID>_DVEBMGS00_physicalhost` profile files.

In the `/sapmnt/<SAPSID>/profile` directory, replace the physical host name with the logical host name. In this example, the `<SAPSID>` is `HA1`:

```
$ cdpro; pwd
/sapmnt/HA1/profile
$ mv START_DVEBMGS00_physicalhost START_DVEBMGS00_Clloghost
$ mv HA1_DVEBMGS00_physicalhost HA1_DVEBMGS00_Clloghost
```

- In the `START_DVEBMGS00_Clloghost` profile file, replace occurrences of the physical host name with the central instance logical host name for all `'pf='` arguments.

In this example, the <SAPSID> is HA1:

```
...
Execute_00 =local $(DIR_EXECUTABLE)/sapmscsa -n \
pf=$(DIR_PROFILE)/HA1_DVEBMGS00_Clloghost
Start_Program_01 =local $_MS pf=$(DIR_PROFILE)/HA1_DVEBMGS00_Clloghost
Start_Program_02 =local $_DW pf=$(DIR_PROFILE)/HA1_DVEBMGS00_Clloghost
Start_Program_03 =local $_CO -F pf=$(DIR_PROFILE)/HA1_DVEBMGS00_Clloghost
Start_Program_04 =local $_SE -F pf=$(DIR_PROFILE)/HA1_DVEBMGS00_Clloghost
...
```

- Edit the <SAPSID>_DVEBMGS00_Clloghost file to add a new entry for the SAPLOCALHOST parameter.

Add this entry only for the central instance profile. Set the SAPLOCALHOST parameter to be the central instance logical host name. This parameter allows external application servers to locate the central instance by using the logical host name.

```
...
SAPLOCALHOST          =Clloghost
...
```

- Edit the DEFAULT.PFL file to replace occurrences of the physical host name with the logical host name.

For each of the rdisp/ parameters, replace the physical host name with the central instance logical host name. For the SAPDBHOST parameter, enter the logical host name of the database. If the central instance and database are installed on the same logical host, enter the central instance logical host name. If the database is installed on a different logical host, use the database logical host name instead. In this example, Clloghost represents the logical host name of the central instance, DBloghost represents the logical host name of the database, and HA1 is the <SAPSID>:

```
...
SAPDBHOST              =DBloghost
rdisp/mshost           =Clloghost
rdisp/sna_gateway      =Clloghost
rdisp/vbname           =Clloghost_HA1_00
rdisp/enqname          =Clloghost_HA1_00
rdisp/btcname          =Clloghost_HA1_00
...
```

i. Revise the `TPPARAM` transport configuration file.

Change to the directory containing the transport configuration file.

```
# cd /usr/sap/trans/bin
```

Replace the database physical host name with the database logical host name. In this example, *DBloghost* represents the database logical host name and HA1 is the `<SAPSID>`. For example:

```
...  
HA1/dbhost = DBloghost  
...
```

j. (For SAP 4.0x only) In the `TPPARAM` file, also set `/var/opt/oracle` to be the location for the database client configuration files.

```
...  
HA1/dbconfpath = /var/opt/oracle  
...
```

2. Modify the environment for the SAP database user.

During SAP installation, SAP creates Oracle files that use Sun Cluster physical host names. Replace the physical host names with logical host names using the following steps.

Note – Become the `ora<sapsid>` user before editing these files.

a. **Revise the .cshrc file in the ora<sapsid> home directory.**

The .cshrc file on the server in which SAP was installed contains aliases that use Sun Cluster physical host names. Replace the physical host names with logical host names.

(For SAP 3.1x only) The resulting file should look similar to the following example, in which *Clloghost* represents the central instance logical host and *DBloghost* is the database logical host. If the central instance and database reside on the same logical host, use the central instance logical host name for each of the substitutions:

```
# aliases
alias  startsap  "$HOME/startsap_Clloghost_00"
alias  stopsap   "$HOME/stopsap_Clloghost_00"

# RDBMS environment
if ( -e $HOME/.dbenv_DBloghost.csh ) then
    source $HOME/.dbenv_DBloghost.csh
else if ( -e $HOME/.dbenv.csh ) then
    source $HOME/.dbenv.csh
endif
```

(For SAP 4.0x only) The resulting .cshrc file should look similar to the following example, in which *Clloghost* is the central instance logical host and *DBloghost* is the database logical host. If the central instance and database reside on the same logical host, use the central instance logical host name for each of the substitutions:

```
if ( -e $HOME/.sapenv_Clloghost.csh ) then
    source $HOME/.sapenv_Clloghost.csh
else if ( -e $HOME/.sapenv.csh ) then
    source $HOME/.sapenv.csh
endif

# RDBMS environment
if ( -e $HOME/.dbenv_DBloghost.csh ) then
    source $HOME/.dbenv_DBloghost.csh
else if ( -e $HOME/.dbenv.csh ) then
    source $HOME/.dbenv.csh
endif
```

- b. (For SAP 4.0x only) Rename the `.sapenv_physicalhost.csh` to `.sapenv_Clloghost.csh`.**

In this example, *Clloghost* represents the central instance logical host name.

```
$ mv .sapenv_physicalhost.csh .sapenv_Clloghost.csh
```

- c. Rename the `.dbenv_physicalhost.csh` file.**

Replace the physical host name with the database logical host name in the `.dbenv_physicalhost.csh` file name. If the central instance and database are on the same logical host, use the central instance logical host name for the substitution. In this example, *DBloghost* represents the database logical host:

```
$ mv .dbenv_physicalhost.csh .dbenv_DBloghost.csh
```

- d. (For SAP 4.0x only) Edit the `.dbenv_DBloghost.csh` file to set the `ORA_NLS` environment variable to point to the appropriate subdirectories of `/var/opt/oracle` for the database client configuration files. Also, set the `TNS_ADMIN` environment variable to point to the `/var/opt/oracle` directory.**

The `.dbenv_DBloghost.csh` file is located in the `ora<sapsid>` home directory.

```
#setenv ORA_NLS /oracle/<SAPSID>/ocommon/NLS_723/admin/data
setenv ORA_NLS /var/opt/oracle/ocommon/NLS_723/admin/data

#setenv ORA_NLS32 /oracle/<SAPSID>/ocommon/NLS_733/admin/data
setenv ORA_NLS32 /var/opt/oracle/ocommon/NLS_733/admin/data

#setenv ORA_NLS33 /oracle/<SAPSID>/ocommon/NLS_804/admin/data
setenv ORA_NLS33 /var/opt/oracle/ocommon/NLS_804/admin/data

...

# setenv TNS_ADMIN @TNS_ADMIN@
setenv TNS_ADMIN /var/opt/oracle

...
```


- e. **(For SAP 4.6B and SAP 4.6C only) Edit the `.dbenv_DBloghost.csh` file to set the `ORA_NLS` environment variable to point to the appropriate subdirectories of `/var/opt/oracle` for the database client configuration files. Set the `TNS_ADMIN` environment variable to point to the `/var/opt/oracle` directory.** The `.dbenv_DBloghost.csh` file is located in the `ora<sapsid>` home directory.

```
...
#setenv ORA_NLS           /oracle/D01/ocommon/NLS_723/admin/data
#setenv ORA_NLS32        /oracle/D01/ocommon/NLS_733/admin/data
#setenv ORA_NLS33        /oracle/D01/ocommon/nls/admin/data
setenv ORA_NLS           /var/opt/oracle/ocommon/NLS_723/admin/data
setenv ORA_NLS32         /var/opt/oracle/ocommon/NLS_733/admin/data
setenv ORA_NLS33         /var/opt/oracle/ocommon/nls/admin/data
...
# setenv TNS_ADMIN @TNS_ADMIN@
setenv TNS_ADMIN /var/opt/oracle
...
```

3. **Edit the Oracle listener configuration files to replace occurrences of the physical host name with the database logical host name.**

If the central instance and database instance are on the same logical host, use the central instance logical host name for the substitutions.

4. **Make the Oracle listener configuration files locally accessible on every potential master.**

Use the following steps to accomplish this.

- a. **Replace all occurrences of physical host names with the database logical host name in the `listener.ora` and `tnsnames.ora` files.**

(For SAP 3.1x only) The `listener.ora` file is located at `/etc/listener.ora`. The `tnsnames.ora` file is located at `/usr/sap/trans/tnsnames.ora`.

(For SAP 4.0x only) The `listener.ora` file is located at `/oracle/<SAPSID>/network/admin/listener.ora`. The `tnsnames.ora` file is located at `/oracle/<SAPSID>/network/admin/tnsnames.ora`.

- b. **Relocate the Oracle listener configuration files on the node where the database is installed.**

(For SAP 3.1x only) During installation, SAP places the `listener.ora` file in the local `/etc` directory of the node where the installation took place, and creates a soft link in `/usr/sap/trans`. Move the `listener.ora` file to

/var/opt/oracle. Reset soft links in /usr/sap/trans to point to the new location. Move the tnsnames.ora and sqlnet.ora files to the /var/opt/oracle directory.

```
$ su
# mv /etc/listener.ora /var/opt/oracle
# rm /usr/sap/trans/listener.ora
# ln -s /var/opt/oracle/listener.ora /usr/sap/trans
# mv /usr/sap/trans/tnsnames.ora /var/opt/oracle
# ln -s /var/opt/oracle/tnsnames.ora /usr/sap/trans
# mv /usr/sap/trans/sqlnet.ora /var/opt/oracle
# ln -s /var/opt/oracle/sqlnet.ora /usr/sap/trans
```

(For SAP 4.0x and SAP 4.6C only) SAP places the listener.ora file in the default directory under \$ORACLE_HOME/network/admin. Use the steps below to move the listener.ora file to /var/opt/oracle, and re-set soft links in the original directory to point to the new location. Move all other Oracle listener configuration files to the new location and re-set links to point to the new location.

```
$ su
# mv /oracle/<SAPSID>/network/admin/listener.ora /var/opt/oracle
# ln -s /var/opt/oracle/listener.ora /oracle/<SAPSID>/network/admin
# mv /oracle/<SAPSID>/network/admin/tnsnames.ora /var/opt/oracle
# ln -s /var/opt/oracle/tnsnames.ora /oracle/<SAPSID>/network/admin
# mv /oracle/<SAPSID>/network/admin/sqlnet.ora /var/opt/oracle
# ln -s /var/opt/oracle/sqlnet.ora /oracle/<SAPSID>/network/admin
# mv /oracle/<SAPSID>/network/admin/protocol.ora /var/opt/oracle
# ln -s /var/opt/oracle/protocol.ora /oracle/<SAPSID>/network/admin
```

Note – When the database is Oracle 8.1.6 (upgraded from oracle 8.0.6) the *.ora files have to be copied from /oracle/<SAPSID>/816_32/network/admin directory.

c. (For SAP 4.0x and SAP 4.6C only) Copy the Oracle client configuration files to the common /var/opt/oracle directory.

```
# cd /var/opt/oracle; mkdir rdbms ocommon lib
# cd /var/opt/oracle/rdbms; cp -R /oracle/<SAPSID>/rdbms/mesg .
# cd /oracle/<SAPSID>/ocommon
# tar -cvf - NLS* | (cd /var/opt/oracle/ocommon ; tar xvf -)
# cd /var/opt/oracle/lib; cp /oracle/<SAPSID>/lib/libclntsh.so.1.0 .
```

Note – When the database is Oracle 8.1.6 (upgraded from oracle 8.0.6) the ocommon directory mentioned below is in the directory /oracle/<SAPSID>/816_32.

d. (For SAP 4.6B only) Copy the Oracle client configuration files to the /var/opt/oracle directory.

```
# cd /var/opt/oracle; mkdir rdbms ocommon lib nls
# cd /var/opt/oracle/rdbms; cp -R /oracle/<SAPSID>/rdbms/mesg .
# cd /oracle/<SAPSID>/ocommon
# tar -cf - NLS* | (cd /var/opt/oracle/ocommon ; tar xf -)
# cd /var/opt/oracle/lib; cp /oracle/<SAPSID>/lib/libclntsh.so.1.0 .
# cp -R /oracle/805_32 /var/opt/oracle
# cd /oracle/<SAPSID>/ocommon
# tar -cf - nls | (cd /var/oracle/ocommon ; tar xf -)
```

e. Distribute the Oracle listener configuration files to all potential masters of the central instance and database instance.

Copy or transfer the Oracle configuration files from the node on which the database was initially installed into the local directory /var/opt/oracle on all potential central instance and database masters. In this example, *physicalhost2* represents the name of the backup physical host.

```
$ su
# tar cf - /var/opt/oracle | rsh physicalhost2 tar xf -
```

Note – As part of the maintenance of HA-DBMS, the configuration files must be synchronized on all potential master nodes, whenever modifications are made.

5. Update the /etc/services files on all potential masters to include the new SAP service entries.

The /etc/services files must be identical on all nodes.

6. Create the /usr/sap/tmp directory on all nodes.

The saposcol program will rely on this directory.

7. Test the SAP installation.

Test the SAP installation by manually shutting down SAP, manually switching the logical host between the potential master nodes, and then manually starting SAP on the backup node. This will verify that all kernel parameters, service port entries, file systems and mount points, and user/group permissions are properly set on all potential masters of the logical hosts.

a. Start the central instance and database.

```
# su - ora<sapsid>
$ lsnrctl start
...
# su - <sapsid>adm
$ startsap all
```

b. Run the GUI and verify that SAP comes up correctly.

In this example, the dispatcher port number is 3200.

```
# su - <sapsid>adm
$ setenv DISPLAY your_workstation:0
$ sapgui /H/Clloghost/S/3200
```

c. Verify that SAP can connect to the database.

```
# su - <sapsid>adm
$ R3trans -d
```

d. Run the `saplicense` utility to get a CUSTOMER KEY for the current node.

You will need a SAP license for all potential masters of the central instance logical host.

e. Stop SAP and the database.

```
# su - <sapsid>adm
$ stopsap all
...
# su - ora<sapsid>
$ lsnrctl stop
```

8. For each remaining node that is a potential master of the central instance logical host, switch the central instance logical host to that node and repeat the test sequence described in Step 1.

```
# scadmin switch clustername phys-hahost2 Clloghost
```

Next, proceed to “How to Configure Sun Cluster HA for Oracle”.

▼ How to Configure Sun Cluster HA for Oracle

1. Shut down SAP and the database.

```
# su - <sapsid>adm
$ stopsap all
...
# su - ora<sapsid>
$ lsnrctl stop
```

2. (For SAP 3.1x only) Adjust the Oracle alert file parameter in the `init<SAPSID>.ora` file.

By default, SAP uses the prefix “`?.?.?`” in the `init<SAPSID>.ora` file to denote the relative path from `$ORACLE_HOME`. The Sun Cluster fault monitors cannot parse the prefix, but instead require the full path name to the alert file. Therefore, you must edit the `/oracle/<SAPSID>/dbs/init<SAPSID>.ora` file and define the dump destination parameters as follows:

```
background_dump_dest = /oracle/<SAPSID>/saptrace/background
```

3. Register and activate the database.

Run the `hareg(1M)` command from only one node. For example, for Oracle:

```
# hareg -s -r oracle -h DBloghost
# hareg -y oracle
```

4. Set up the database instance.

See Chapter 5, “Setting Up and Administering Sun Cluster HA for Oracle” in the *Sun Cluster 2.2 Software Installation Guide* for more information.

For example, for Oracle:

```
# haoracle insert <SAPSID> DBloghost 60 10 120 300 \
user/password /oracle/<SAPSID>/dbs/init<SAPSID>.ora LISTENER
```

5. Start fault monitoring for the database instance.

For example:

```
# haoracle start <SAPSID>
```

6. Test switchover of the HA-DBMS.

For example:

```
# scadmin switch clustername phys-hahost2 DBloghost
```

Next, proceed to “Configuring Sun Cluster HA for SAP (SAP With Oracle)”.

Configuring Sun Cluster HA for SAP (SAP With Oracle)

This section describes how to register and configure Sun Cluster HA for SAP.

▼ How to Configure Sun Cluster HA for SAP (SAP With Oracle)

1. **If Sun Cluster HA for SAP has not yet been installed, install it now by running `scinstall(1M)` on all nodes and adding the Sun Cluster HA for SAP data service.**

See Section 3.2, “Installation Procedures” in the *Sun Cluster 2.2 Software Installation Guide* for details. If the cluster is already running, you must stop it before installing the data service.

2. **Register the Sun Cluster HA for SAP data service by running the `hareg(1M)` command.**

Run this command on only one node:

```
# hareg -s -r sap -h Cloghost
```

3. **Verify that all nodes are running in the cluster.**
4. **Create a new Sun Cluster HA for SAP instance using the `hadsconfig(1M)` command.**

The `hadsconfig(1M)` command is used to create, edit, and delete instances of the Sun Cluster HA for SAP data service. The configuration parameters are described in “Configuration Parameters for Sun Cluster HA for SAP (SAP With Oracle)”.

Run this command on only one node, while all nodes are running in the cluster:

```
# hadsconfig
```

5. If Sun Cluster HA for SAP is dependent upon other data services within the same logical host, set dependencies on those data services.

See “How to Set a Data Service Dependency for SAP (SAP With Oracle)”. If you do set dependencies, start all services on which SAP depends before proceeding.

6. Stop the central instance before starting SAP under the control of Sun Cluster HA for SAP.

```
# su - <sapsid>adm
$ stopsap r3
```



Caution – The SAP central instance must be stopped before Sun Cluster HA for SAP is turned on.

7. Turn on the Sun Cluster HA for SAP instance.

```
# hareg -y sap
```

8. Test switchover of Sun Cluster HA for SAP.

For example:

```
# scadmin switch clustername phys-hahost2 Cloghost
```

9. (Optional) If you have application servers or a test/development system, customize and test the `hasap_start_all_instances` and `hasap_stop_all_instances` scripts.

See “Configuration Options for Application Servers and Test/Development Systems” for details. Test switchover of Sun Cluster HA for SAP and verify start and stop of application servers. Verify that the test/development system stops when the central instance logical host is switched to the test/development system physical host.

```
# scadmin switch clustername phys-hahost1 Cloghost
```

Configuration Parameters for Sun Cluster HA for SAP (SAP With Oracle)

This section describes the information you supply to `hadsconfig(1M)` to create configuration files for the Sun Cluster HA for SAP data service. The `hadsconfig(1M)` command uses templates to create these configuration files. The templates contain some default, some hard coded, and some unspecified parameters. You must provide values for all parameters that are unspecified.

The fault probe parameters, in particular, can affect the performance of Sun Cluster HA for SAP. Tuning the probe interval value too low (increasing the frequency of fault probes) might encumber system performance, and also might result in false takeovers or attempted restarts when the system is simply slow.

Note – The Sun Cluster HA for SAP parameter `LOG_DB_WARNING` determines whether warning messages should be displayed if the Sun Cluster HA for SAP probe cannot connect to the database. When `LOG_DB_WARNING` is set to `-y` and the probe cannot connect to the database, a message is logged at the warning level in the `local0` facility. By default, the `syslogd(1M)` daemon does not display these messages to `/dev/console` or to `/var/adm/messages`. To see these warnings, you must modify the `/etc/syslog.conf` file to display messages of `local0.warning` priority. After modifying the file, you must restart `syslogd(1M)`. See the `syslog.conf(1M)` and `syslogd(1M)` man pages for more information.

Configure Sun Cluster HA for SAP by supplying the `hadsconfig(1M)` command with parameters listed in TABLE 1-9.

TABLE 1-9 Sun Cluster HA for SAP Configuration Parameters (SAP With Oracle)

Name of the Instance	Nametag used internally as an identifier for the instance. The log messages generated by Sun Cluster refer to this nametag. The <code>hadsconfig(1M)</code> command prefixes the package name to the value you supply here. You can use the <code>SAPSID</code> for this nametag. For example, if you specify <code>HA1</code> , <code>hadsconfig(1M)</code> produces <code>SUNWscsap_HA1</code> .
Logical Host	Name of the logical host that provides service for this instance of Sun Cluster HA for SAP. This name should be the logical host name for the central instance.
Time Between Probes	The interval, in seconds, of the fault probing cycle. The default value is 60 seconds.
SAP R/3 System ID	This is the SAP system name or <code><SAPSID></code> .
Central Instance ID	This is the SAP system number or Instance ID. For example, the CI is normally "00."

TABLE 1-9 Sun Cluster HA for SAP Configuration Parameters (SAP With Oracle)

SAP Admin Login Name	The name used by Sun Cluster HA for SAP to log in to the SAP central instance administrative account. This name must exist on all central instance and application server hosts. This is the <sapsid>adm. For example, "ha1adm."
Database Admin Login Name	This is the SAP database administrator's account. For SAP with Oracle, this is the ora<sapsid>. For example, oraha1.
Database Logical Host Name	Name of the logical host for the database used by SAP. This might be the same as the logical host name used for the central instance, depending on your configuration.
Log Database Warnings	Possible values are "y" or "n." If set to "y" and the Sun Cluster HA for SAP probe detects that it cannot connect to the database during a probe cycle, a warning message appears saying the database is unavailable. For example, this occurs if the database logical host is in maintenance mode or if the database is being relocated to another node in the cluster. If the parameter is set to "n," then no messages appear if the probe cannot connect to the database.
Central Instance Start Retry Count	This must be an integer greater than or equal to 1. This is the number of times Sun Cluster HA for SAP should attempt to start the central instance before giving up. This value is also the number of times the Sun Cluster HA for SAP fault monitor will probe in grace mode before entering normal probe mode. While in grace mode, the probe will not perform a restart or initiate a failover of the central instance if the probe detects that the central instance is not yet up. Instead, the fault monitor will report the status of all probes and will continue in grace mode until all probes pass, or until the retry count has been exhausted.
Central Instance Start Retry Interval	This is the number of seconds Sun Cluster HA for SAP should wait between each attempt to start the central instance. This value is also the number of seconds that the Sun Cluster HA for SAP fault monitor will sleep (between probe attempts) while in grace mode.
Time Allowed to Stop All Instances Before Central Instance Starts	This must be an integer greater than or equal to 0. This parameter dictates for how much time (in seconds) the <code>hasap_stop_all_instances</code> script should be run before starting the central instance. If set to 0, then <code>hasap_stop_all_instances</code> is run in the background while the central instance is being started. If set to a positive integer, then <code>hasap_stop_all_instances</code> is run for that amount of time in the foreground before the central instance is started.

TABLE 1-9 Sun Cluster HA for SAP Configuration Parameters (SAP With Oracle)

<p>Allow the Central Instance to Start if Foregrounded Stop All Instances Returns Error</p>	<p>This flag should be set to either “y” or “n”. This value determines whether the central instance should be started in the case where the <code>hasap_stop_all_instances</code> script returns a non-zero exit code or does not complete in the time specified by the “Time Allowed to Stop All Instances Before Central Instance Starts” parameter. If set to “n” <i>and</i> the value for “Time Allowed to Stop All Instances Before Central Instance Starts” is greater than 0, <i>and</i> if the <code>hasap_stop_all_instances</code> script does not complete in the time configured above or the <code>hasap_stop_all_instances</code> script returns a non-zero exit status, the central instance will not be started and the fault monitors will take action based on the other configuration parameters. If set to “y,” then the central instance will be started regardless of whether <code>hasap_stop_all_instances</code> returns an error code or finishes within the timeout specified above.</p>
<p>Number of Central Instance Restarts on Local Node:</p>	<p>This must be an integer greater than or equal to 0. This dictates how many times the SAP central instance will be restarted on the local node before giving up, after a failure has been detected. When this number of restarts has been exhausted, Sun Cluster HA for SAP either issues a failover request, if permitted by the “Allow Central Instance Failover” parameter, or does nothing to correct the failure detected by the fault monitor.</p>
<p>Number of Probe Successes to Reset the Restart Count</p>	<p>This parameter should be an integer that is greater than or equal to 0. If set to a positive integer, then after that many consecutive successful probes, the count of restarts done so far on the local node will be reset to 0. For example, if the value for “Number of Central Instance Restarts on Local Node” parameter is 1 and the value for “Number of Probe Successes to Reset the Restart Count” is 60, then after the first failure occurs, the probe will try to restart the central instance on the local node. If this restart succeeds, then after 60 successful probes, the restart count will be reset to 0, allowing the probe to do another restart if it detects another failure. If the parameter “Number of Probe Successes to Reset the Restart Count” is set to 0, then the restart count is never reset. This means that the number of restarts set in the parameter “Number of Central Instance Restarts on Local Node” is the absolute number of restarts that will be done on the local node before failing over.</p>
<p>Allow Central Instance Failover</p>	<p>Possible values are “y” or “n.” If set to “y” <i>and</i> Sun Cluster HA for SAP detects an error in the SAP instance it is monitoring <i>and</i> the “Number of central instance Restarts on Local Node” has been exhausted, then Sun Cluster HA for SAP issues a request to relocate the instance’s logical host to another cluster node. If this flag is set to “n,” then even if an error is detected and all of the local restarts have been exhausted, Sun Cluster HA for SAP will not cause a relocation of this instance's logical host. When this occurs, the central instance is left in the its failed state, and the probe exits.</p>

Setting Data Service Dependencies for SAP (SAP With Oracle)

Setting a dependency with `hasap_dbms` is only necessary to specify the order that data services are started and stopped within a single logical host. There is no mechanism for setting dependencies for data services configured on two different logical hosts.

If Sun Cluster HA for Oracle or Sun Cluster HA for NFS are configured on the same logical host as Sun Cluster HA for SAP, then you should set a dependency for Sun Cluster HA for SAP on those data services. You can use the `hasap_dbms` command to create or remove such a dependency. These dependencies affect the order that the services are started and stopped. Sun Cluster HA for Oracle and Sun Cluster HA for NFS should always be started before Sun Cluster HA for SAP is started. Similarly, Sun Cluster HA for SAP should always be stopped before the other data services are stopped.



Caution – If Sun Cluster HA for Oracle or Sun Cluster HA for NFS is *not* configured on the same logical host as Sun Cluster HA for SAP, then do not use the `hasap_dbms` command.

▼ How to Set a Data Service Dependency for SAP (SAP With Oracle)

To set a data service dependency, issue *one* of the `hasap_dbms` commands described below.

Note – The `hasap_dbms` command can be used only when Sun Cluster HA for SAP is registered but is in the `off` state. Run the command on only one node, while that node is a member of the cluster. See the `hasap_dbms(1M)` man page for more information.



Caution – If the `hasap_dbms(1M)` command returns an error stating that it cannot add rows to or update the CCD, it might be because another cluster utility is also trying to update the CCD. If this occurs, re-run `hasap_dbms(1M)` until it runs successfully. After the `hasap_dbms(1M)` command runs successfully, verify that all necessary rows are included in the resulting CCD by running the command `hareg -q sap`.

If the `hareg(1M)` command returns an error, then first restore the original method timeouts by running the command `hasap_dbms -f`. Second, restore the default dependencies by running the command `hasap_dbms -r`. After both commands

complete successfully, retry the original `hasap_dbms (1M)` command to configure new dependencies and method timeouts. See the `hasap_dbms (1M)` man page for more information.

1. Set the data service dependency using *one* of the following commands.

If you are using only Sun Cluster HA for NFS and Sun Cluster HA for SAP on the same logical host, use the following command:

```
# /opt/SUNWcluster/ha/sap/hasap_dbms -d nfs
```

If you are using only Sun Cluster HA for Oracle and Sun Cluster HA for SAP on the same logical host, use the following command:

```
# /opt/SUNWcluster/ha/sap/hasap_dbms -d oracle
```

If you are using Sun Cluster HA for Oracle, Sun Cluster HA for NFS, and Sun Cluster HA for SAP on the same logical host, use the following command:

```
# /opt/SUNWcluster/ha/sap/hasap_dbms -d oracle,nfs
```

2. Check the dependencies set for Sun Cluster HA for SAP using the following command:

```
# hares -q sap -D
```

▼ How to Remove a Data Service Dependency for SAP (SAP With Oracle)

The dependencies set for Sun Cluster HA for SAP can be removed by running the `hasap_dbms -r` command. Issuing this command causes all of the dependencies set for Sun Cluster HA for SAP to be removed.

Note – The `hasap_dbms` command can be used only when Sun Cluster HA for SAP is registered but is in the `off` state. Run the command on only one node, while that node is a member of the cluster. See the `hasap_dbms(1M)` man page for more information.



Caution – If the `hasap_dbms(1M)` command returns an error stating that it cannot add rows to or update the CCD, it might be because another cluster utility is also trying to update the CCD. If this occurs, re-run `hasap_dbms(1M)` until it runs successfully. After the `hasap_dbms(1M)` command runs successfully, verify that all necessary rows are included in the resulting CCD by running the command `hareg -q sap`.

If the `hareg(1M)` command returns an error, then first restore the original method timeouts by running the command `hasap_dbms -f`. Second, restore the default dependencies by running the command `hasap_dbms -r`. After both commands complete successfully, retry the original `hasap_dbms(1M)` command to configure new dependencies and method timeouts. See the `hasap_dbms(1M)` man page for more information.

1. Remove all of the dependencies set for Sun Cluster HA for SAP, using the following command:

```
# /opt/SUNWcluster/ha/sap/hasap_dbms -r
```

2. Check the dependencies set for Sun Cluster HA for SAP, using the following command:

```
# hareg -q sap -D
```

SAP With Informix

Use the information in the following sections to install and configure SAP with Informix. For information about installing and configuring SAP with Oracle, see “SAP With Oracle”.

Installation and Configuration Overview (SAP With Informix)

TABLE 1-10 summarizes the tasks you must complete to install and configure SAP and Sun Cluster HA for SAP.

TABLE 1-10 Installation Overview for Sun Cluster HA for SAP (SAP With Informix)

Task	See ...
Prepare the cluster environment for SAP and Informix: <ul style="list-style-type: none">- Install Solaris- Install and configure the volume manager- Create disksets or disk groups- Create volumes and file systems- Install Sun Cluster and the data services- Set up public network monitoring (PNM)- Set up logical hosts and mount points- Configure the shared CCD (SSVM, 2-node only)- Set up HA-NFS, if necessary- Adjust kernel parameters- Create links for Informix- Create and modify the administrative file systems- Configure Sun Cluster HA for NFS- Configure swap space and paging space- Create user and group accounts	“How to Prepare the Cluster Environment for SAP and the Database (SAP With Informix)”
Install SAP and Informix: <ul style="list-style-type: none">- Install SAP, Informix- Install other components as necessary, such as application servers- Install the SAP GUI- Shut down SAP and Informix	“How to Install SAP and the Database (SAP With Informix)”

TABLE 1-10 Installation Overview for Sun Cluster HA for SAP (SAP With Informix)

Task	See ...
<p>Enable SAP and Informix to run in the cluster environment:</p> <ul style="list-style-type: none"> - Modify the Informix configuration files - Set up the SAP central instance environment - Modify the SAP database user environment - Update <code>/etc/services</code> and create <code>/usr/sap/tmp</code> - Test the SAP installation 	<p>“How to Enable SAP and the Database to Run in the Cluster Environment (SAP With Informix)”</p>
<p>Configure Sun Cluster HA for Informix:</p> <ul style="list-style-type: none"> - Shut down SAP and Informix - Register and activate HA-Informix - Bring Informix under the control of HA-Informix - Start HA-Informix - Test switchover of the database 	<p>“How to Configure Sun Cluster HA for Informix”</p>
<p>Configure Sun Cluster HA for SAP:</p> <ul style="list-style-type: none"> - Register Sun Cluster HA for SAP - Configure Sun Cluster HA for SAP with <code>hadscfg(1M)</code> - Turn on Sun Cluster HA for SAP - Test switchover of SAP - Customize start and stop scripts for application servers - Set data service dependencies for SAP 	<p>“How to Configure Sun Cluster HA for SAP (SAP With Informix)” and “How to Set a Data Service Dependency for SAP (SAP With Informix)”</p>

Installation Worksheet for Sun Cluster HA for SAP (SAP With Informix)

Complete the following worksheet before beginning the installation procedures.

TABLE 1-11 Sun Cluster HA for SAP Installation Worksheet (SAP With Informix)

Name of the cluster	
Number of logical hosts	
Name and IP address of all physical hosts that are potential masters of the CI logical host	
Name and IP address of CI logical host	
SAP system ID (<SAPSID>)	
SAP system number	
Name and IP address of all physical hosts that are potential masters of the DB logical host	
Name and IP address of DB logical host (In asymmetric configurations, this is identical to the CI logical host.)	
Name of NFS logical host (see Note below)	
SAP license for each potential master of the CI logical host	

Note – If all application servers are external to cluster, the name of the NFS logical host is the same as the central instance logical host. If the application servers are inside the cluster, name of the NFS logical host is the logical host that provides NFS service from the external NFS cluster. See “Sun Cluster HA for NFS Considerations”.

Installation and Configuration Procedures (SAP With Informix)

Perform the procedures in the order indicated in TABLE 1-10.

▼ How to Prepare the Cluster Environment for SAP and the Database (SAP With Informix)

Before installing SAP and Informix, perform the following tasks.

1. On all nodes, install the Solaris operating environment and Solaris patches.

See Chapter 3, “Installing and Configuring Sun Cluster Software” in the *Sun Cluster 2.2 Software Installation Guide*.

2. On all nodes, install Volume Manager software and any required Volume Manager patches.

See Chapter 3, “Installing and Configuring Sun Cluster Software” in the *Sun Cluster 2.2 Software Installation Guide*.

3. On the node on which you will install SAP and Informix, create Solstice DiskSuite disksets or SSVM disk groups.

Separate disk groups for the SAP central instance and database instance are recommended, for ease of administration.

4. On the node on which you will install SAP and Informix, create volumes according to Sun Cluster guidelines:

- Mirror volumes across controllers
- With SSVM, use Dirty Region Logging for faster mirror resynchronization
- Use a logging file system for faster logical host failover

Use the following table as a worksheet to capture the name of the volume that corresponds to each file system used for the SAP central instance and database instance. Refer to the SAP installation guide for the file system sizes recommended for your particular configuration. The central instance file systems are database-independent; the database instance file systems are database-dependent. Use raw partitions for the database instances.

TABLE 1-12 Worksheet: File Systems and Volume Names for the SAP Instances (SAP With Informix)

File System Name	Volume Name
/usr/sap/trans	
/sapmnt/<SAPSID>	
/usr/sap/<SAPSID>	
/informix/<SAPSID>	

5. On all nodes, install Sun Cluster, Sun Cluster HA for SAP, Sun Cluster HA for Informix, and any required patches.

Use the procedures described in Chapter 3, “Installing and Configuring Sun Cluster Software” in the *Sun Cluster 2.2 Software Installation Guide*, but do not set up logical hosts with `scinstall(1M)` during this installation (you will set up logical hosts with `sconf(1M)` in Step 10).

6. On all nodes, configure PNM.

For detailed procedures, see Chapter 3, “Installing and Configuring Sun Cluster Software” in the *Sun Cluster 2.2 Software Installation Guide*, and Chapter 6, “Administering Network Interfaces” in the *Sun Cluster 2.2 System Administration Guide*.

7. Start the cluster.

Run the following command on one node.

```
# scadmin startcluster physicalhost clustername
```

Run the following command on all other nodes, sequentially.

```
# scadmin startnode
```

8. (SSVM only) Verify that all disk groups are deported.

9. (Solstice DiskSuite only) Release ownership of all disksets.

10. On the node on which you installed SAP, create logical hosts with `sconf(1M)`.

The number of logical hosts depends on your particular configuration. You should set up two disk groups: one for SAP and one for Informix. You can place both disk groups on the same logical host, or configure one disk group per logical host (in a

configuration with two logical hosts). See Chapter 3, “Installing and Configuring Sun Cluster Software” in the *Sun Cluster 2.2 Software Installation Guide* for more information.

You will need:

- Logical host name(s)
- Physical host names of potential masters of logical host(s)
- Names of the primary public network controllers for the potential masters of the logical host(s)
- Disk group name(s)

When you create logical hosts, disable the automatic failback mechanism by using the `-m` option to `scconf(1M)`.

11. (SSVM, two-node configurations only) Configure the shared CCD.

See Appendix C, “Configuring Sun StorEdge Volume Manager and Cluster Volume Manager” in the *Sun Cluster 2.2 Software Installation Guide*.

12. Create mount points for the central instance and database instance volumes, and update the `vfstab.logicalhost` files on all potential masters of each logical host.

The `vfstab.logicalhost` files are located in `/etc/opt/SUNWcluster/conf/hanfs`.

The following table lists the suggested file system mount points for the disk groups (SSVM) or disksets (Solstice DiskSuite) associated with the central instance and database instance. Note that separating the central instance and database instance file systems into different disk groups or disksets (even when you use a single logical host) may provide more configuration flexibility in the future

TABLE 1-13 Disk Groups/Disksets and Mount Points for the SAP Central Instance and Database Instance (SAP With Informix)

Disk Group or Diskset for the ...	Mount Point
Central instance	<code>/usr/sap/<SAPSID></code>
Central instance	<code>/usr/sap/trans</code>
Central instance	<code>/sapmnt/<SAPSID></code>
Database instance	<code>/informix/<SAPSID></code>

13. On all nodes, create directories for Informix.

```
# mkdir /informix
# mkdir -p /var/opt/informix
# cd /var/opt/
# chown informix:informix informix
```

14. On the node on which you installed SAP and Informix, create Informix data directories and set up soft links.

See your SAP installation documentation for more information. For example:

```
# mkdir /informix/<SAPSID>/sapdata
# mkdir /informix/<SAPSID>/sapdata/physdev<n>
...
# ln -s /dev/vx/rdisk/dbdg/vol01 /informix/<SAPSID>/sapdata/physdev1/data1
# ln -s /dev/vx/rdisk/dbdg/vol02 /informix/<SAPSID>/sapdata/physdev1/data2
# ln -s /dev/vx/rdisk/dbdg/vol03 /informix/<SAPSID>/sapdata/physdev1/data3
# ln -s /dev/vx/rdisk/dbdg/vol04 /informix/<SAPSID>/sapdata/physdev1/data4
# ln -s /dev/vx/rdisk/dbdg/vol05 /informix/<SAPSID>/sapdata/physdev2/data5
# ln -s /dev/vx/rdisk/dbdg/vol06 /informix/<SAPSID>/sapdata/physdev2/data6
# ln -s /dev/vx/rdisk/dbdg/vol07 /informix/<SAPSID>/sapdata/physdev2/data7
# ln -s /dev/vx/rdisk/dbdg/vol08 /informix/<SAPSID>/sapdata/physdev2/data8
# ln -s /dev/vx/rdisk/dbdg/vol09 /informix/<SAPSID>/sapdata/physdev3/data9
# ln -s /dev/vx/rdisk/dbdg/vol10 /informix/<SAPSID>/sapdata/physdev3/data10
# ln -s /dev/vx/rdisk/dbdg/vol11 /informix/<SAPSID>/sapdata/physdev3/data11
# ln -s /dev/vx/rdisk/dbdg/vol12 /informix/<SAPSID>/sapdata/physdev3/data12
```

15. On all nodes, create links from /var/opt/informix to the appropriate directory on the shared disk.

For example:

```
# ln -s /informix/<SAPSID>/sapdata /var/opt/informix/sapdata
# ln -s /informix/<SAPSID>/sapreorg /var/opt/informix/sapreorg
```

16. On all nodes, create logical host administrative file systems, using `scconf(1M)`.

For detailed procedures, see Appendix B, “Configuring Solstice DiskSuite” and Appendix C, “Configuring Sun StorEdge Volume Manager and Cluster Volume Manager” in the *Sun Cluster 2.2 Software Installation Guide*.

- 17. If SAP application servers will be configured outside the cluster, configure Sun Cluster HA for NFS and enter the appropriate shared file systems into the `dfstab.logicalhost` files on all potential masters of each logical host.**

These files are located in `/etc/opt/SUNWcluster/conf/hanfs`. See “Configuration Options for Application Servers and Test/Development Systems” for more information.

Share the following file systems to SAP application servers outside the cluster. These are general guidelines. See your SAP documentation for more information.

TABLE 1-14 File Systems to Share to External Application Servers (SAP With Informix)

File Systems to Share to External Application Servers
<code>/usr/sap/trans</code>
<code>/sapmnt/<SAPSID>/exe</code>
<code>/sapmnt/<SAPSID>/profile</code>
<code>/sapmnt/<SAPSID>/global</code>

- 18. Test the functionality and mount points of the logical host(s) by switching them between all potential masters.**

This verifies that all mount points have been created correctly.

- 19. Adjust kernel parameters in the `/etc/system` files on all potential masters of the logical hosts.**

Follow the “R/3 Installation on UNIX: OS Dependencies” guidelines in the SAP documentation.

In configurations where the central instance and database instance may coexist with each other or with other instances, be sure to size the kernel parameters accordingly.

- 20. Create permanent swap areas on all potential masters of the logical hosts.**

See the “Installation Requirements Checklist” in your SAP documentation for swap size guidelines.

- 21. On all nodes, check the paging space size.**

Use the SAP-supplied `memlimits` utility to assist you in checking the address space. See the “R/3 Installation on UNIX” guidelines in the SAP documentation for more information on this utility. As a general rule, swap should be at least three times the memory on a given node. See your SAP installation documentation for details.

- 22. Stop the cluster and reboot all nodes.**

23. On all nodes, verify system resources.

See your SAP installation documentation for details.

```
# ulimit -a
...
time(seconds) unlimited
file(blocks) unlimited
data(kbytes) 2097148
stack(kbytes) 8192
coredump(blocks) unlimited
nofiles(descriptors) 64
memory(kbytes) unlimited
```

24. Create SAP and Informix groups, users, passwords, and home directories on all potential masters of the logical hosts.

Create user home directories.

```
# mkdir /export/home/<sapsid>adm
# mkdir /export/home/sapr3
# mkdir /export/home/informix
```

Add the following users and groups. Refer to the “R/3 Installation on UNIX: OS Dependencies” guidelines in the SAP documentation for details. User and group IDs must be identical on all nodes.

```
# groupadd -g 10000 sapsys
# groupadd -g 10002 informix
# groupadd -g 10004 super_archive
# groupadd -g 10006 super_
# groupadd -g 10008 bargroup (for SAP4.5B only)
# useradd -g sapsys -G super_archive,super_,root,informix,bargroup -s \
/usr/bin/csh -d /export/home/<sapsid>adm -u 2001 <sapsid>adm
# useradd -g sapsys -G super_archive,super_,root,informix -s /usr/bin/csh -d \
/export/home/sapr3 -u 2002 sapr3
# useradd -g informix -G super_archive,super_,root,sapsys -s /usr/bin/csh -d \
/export/home/informix -u 2004 informix
```

Create passwords for the users.

```
# passwd sapr3
# passwd informix
# passwd <sapsid>adm
```

This completes preparation of the cluster environment for SAP and Informix. Now proceed to “How to Install SAP and the Database (SAP With Informix)”.

▼ How to Install SAP and the Database (SAP With Informix)

1. Verify that you have completed all tasks described in “How to Prepare the Cluster Environment for SAP and the Database (SAP With Informix)”.
2. Verify that all nodes are running in the cluster.
3. Switch over all logical hosts to the node from which you will install SAP and the database.

```
# scadmin switch clustername phys-hahost1 Clogicalhost DBlogicalhost ...
```

4. Create the SAP installation directory and install SAP, the database, other components such as application servers, and the SAP front-end GUI.

Use your SAP documentation to perform the installation and refer to the “R/3 Installation on UNIX” guidelines in the SAP documentation for details.

This completes the installation of SAP and Informix. Next, proceed to “How to Enable SAP and the Database to Run in the Cluster Environment (SAP With Informix)”.

▼ How to Enable SAP and the Database to Run in the Cluster Environment (SAP With Informix)

1. Shut down the SAP central instance and database.

```
# su - <sapsid>adm  
$ stopsap
```

2. As root, copy the Informix files from the shared disk to all nodes.
 - a. On the node on which you installed SAP and Informix, create or edit the `/.rhosts` file to permit access from all nodes.
 - b. Change directories to the Informix directory on the shared disk.

```
# cd /informix/<SAPSID>
```

- c. Use `tar(1M)` to package the Informix directories and copy them to the local Informix directory on the node on which you installed SAP and Informix (in this example, `phys-hahost1`).

The directories and files present in the directory depend on the version of SAP. Include all files and directories except the data directories (`sapdata` and `sapreorg`). For example:

```
# tar cf - aadir bin console.phys-hahost1.<SAPSID>.log dbssodir \  
forms gls incl help installconn installserver ism IVODBC.LIC lib locale \  
messages release snmp | ( cd /var/opt/informix ; tar xf -)
```

3. On all nodes, modify the Informix configuration files.

Log in as user `informix` to perform the following tasks.

Note – Make backup copies of all files before performing the following steps.

- a. Rename the `sqlhosts.tli` file to `sqlhosts`, for Informix use.

```
# mv /var/opt/informix/etc/sqlhosts.tli /var/opt/informix/etc/sqlhosts
```

- b. In the `sqlhosts` file, replace all occurrences of the physical host name with the database instance logical host name.

For example:

```
CIlogicalhost<sapsid>shm onipcshm DBlogicalhost sapinf<SAPSID>  
CIlogicalhost<sapsid>tcp ontlitcp DBlogicalhost sapinf<SAPSID>
```

- c. Modify the `/export/home/informix/.rhosts` file to allow user `informix` to access the database from all nodes.

Create entries similar to the following, with one entry for each host.

```
phys-hahost1 informix  
phys-hahost2 informix  
CIlogicalhost informix  
DBlogicalhost informix
```


d. Rename the Informix onconfig file to replace the physical host name with the database instance logical host name.

Rename `/var/opt/informix/etc/onconfig.physicalhost.<sapsid>` to `/var/opt/informix/etc/onconfig.Cllogicalhost.<sapsid>`.

e. Modify the onconfig file for Informix.

Modify the file `/var/opt/informix/onconfig.Cllogicalhost.<sapsid>` to direct all Informix paths to `/var/opt/informix` rather than to the shared diskset, for the following parameters:

```
ROOTPATH
MIRRORPATH
MSGPATH
CONSOLE
ALARMPROGRAM
DRLOSTFOUND
SYSALARMPROGRAM
```

The resulting entry should look similar to the following:

```
# original entry
# ROOTPATH /informix/<SAPSID>/sapdata/physdev1/data1

# new entry
ROOTPATH /var/opt/informix/sapdata/physdev1/data1
```

Additionally, replace the physical host name with the logical host name in the database server fields. For example:

```
DBSERVERNAME Cllogicalhost<sapsid>tcp
DBSERVERALIASES Cllogicalhost<sapsid>shm
```

f. Create the `/var/opt/informix/inftab` file.

The file format is `$ONCONFIG:$INFORMIXDIR`. For example:

```
onconfig.Cllogicalhost.<sapsid>:/var/opt/informix
```

4. Copy the Informix directories to the local Informix directory on all nodes other than the node on which SAP and Informix is installed (in this example, `phys-hahost1`).

```
# rsh phys-hahost1 tar cfB - /var/opt/informix | tar xfB -
```

5. On all nodes, set up the administrative environment for the SAP database user (user informix).

a. On all nodes, rename the `.dbenv_physicalhost.csh` file to `.dbenv.csh`.

```
$ mv .dbenv_physicalhost.csh .dbenv.csh
```

b. On all nodes, edit the `.dbenv.csh` files as follows.

Modify the file so that `$INFORMIXDIR` points to `/var/opt/informix` and change the `ONCONFIG` value to `onconfig.Cllogicalhost.<sapsid>`.

Also, modify the file to specify use of TCP for `$INFORMIXSERVER` and `ping(1M)` to check the status of the database logical host. This is necessary to enable dynamic reset of the `$INFORMIXSERVER` parameter in case of switchover or failover.

Note – In asymmetric configurations, the use of TCP and loopback might reduce performance. If so, you can set `$INFORMIXSERVER` to use shared memory instead.

The resulting file should look similar to the following sample, in which the modified fields are in bold type:

```
...
setenv INFORMIXDIR /var/opt/informix
setenv ONCONFIG onconfig.Cllogicalhost.<sapsid>
...
  case Sun*:
    setenv INFORMIXSHMBASE 0x01000000
    setenv LC_CTYPE iso_8859_1
    setenv INFORMIXSQLHOSTS $INFORMIXDIR/etc/sqlhosts
# use TCP for connection prototype always because connection
# cannot be reset dynamically between shared memory and TCP in
# the Sun Cluster environment.
    setenv INFORMIXSERVER `grep 'Cllogicalhost<sapsid>.*ontlitcp'
$INFORMIXSQLHOSTS | awk '{print $1}'`

    /usr/sbin/ping DBlogicalhost >& /dev/null
    if ( $status != 0 ) then
        echo dbserver DBlogicalhost is not alive.
    endif
```

- c. **On all nodes, rename the `.sapenv_physicalhost.csh` file to `.sapenv.csh`, and edit it to replace occurrences of the physical host name with the logical host name.**

First rename the file.

```
$ mv .sapenv_physicalhost.csh .sapenv.csh
```

Then edit the `startsap` and `stopsap` aliases in the `.sapenv.csh` file to specify the central instance logical host in the `set hostname=` field.

```
...  
set hostname='Clogicalhost'  
...
```

6. Modify the SAP configuration files.

Perform the tasks in these substeps on all nodes except the application server. Log in as user `<sapsid>adm` to perform the following tasks.

- a. **Rename and revise the SAP instance `startsap` and `stopsap` shell scripts in the `<sapsid>adm` home directory.**

On the server on which the SAP central instance is installed, the `<sapsid>adm` home directory contains shell scripts that include physical host names. Rename these shell scripts by replacing the physical host names with logical host names. In this example, `Clogicalhost` represents the logical host name of the central instance:

```
$ mv startsap_physicalhost_00 startsap_Clogicalhost_00  
$ mv stopsap_physicalhost_00 stopsap_Clogicalhost_00
```

The `startsap_Clogicalhost_00` and `stopsap_Clogicalhost_00` shell scripts specify physical host names in their `START_PROFILE` parameters. Replace the physical host name with the central instance logical host name in the `START_PROFILE` parameters in both files.

```
...  
START_PROFILE="START_DVEBMGS00_Clogicalhost"  
...
```

b. Revise the SAP central instance profile files.

Replace all occurrences of physical host names with logical host names, in the three profile files created by SAP during installation. You must be user `<sapsid>adm`, and you must be in the `profile` directory.

Rename the `START_DVEBMGS00_physicalhost` and `<SAPSID>_DVEBMGS00_physicalhost` profile files.

```
$ cd /sapmnt/<SAPSID>/profile
$ mv START_DVEBMGS00_physicalhost START_DVEBMGS00_Cllogicalhost
$ mv <SAPSID>_DVEBMGS00_physicalhost <SID>_DVEBMGS00_Cllogicalhost
```

In the `START_DVEBMGS00_Cllogicalhost` profile file, replace occurrences of the physical host name with the central instance logical host name for all `pf=` arguments.

```
Execute_00 =local $(DIR_EXECUTABLE)/sapmscsa -n \  
pf=$(DIR_PROFILE)/<SAPSID>_DVEBMGS00_Cllogicalhost  
Start_Program_01 =local $(MS) pf=$(DIR_PROFILE)/<SAPSID>_DVEBMGS00_Cllogicalhost  
Start_Program_02 =local $(DW) pf=$(DIR_PROFILE)/<SAPSID>_DVEBMGS00_Cllogicalhost  
Start_Program_03 =local $(CO) -F pf=$(DIR_PROFILE)/  
<SAPSID>_DVEBMGS00_Cllogicalhost  
Start_Program_04 =local $(SE) -F pf=$(DIR_PROFILE)/  
<SAPSID>_DVEBMGS00_Cllogicalhost  
...
```

Edit the `<SAPSID>_DVEBMGS00_Cllogicalhost` file to add a new entry for the `SAPLOCALHOST` parameter.

Add this entry only for the central instance profile. Set the `SAPLOCALHOST` parameter to be the central instance logical host name. This parameter allows external application servers to locate the central instance by using the logical host name.

```
...  
SAPLOCALHOST          =Cllogicalhost  
...
```

c. Edit the `DEFAULT.PFL` file to replace occurrences of the physical host name with the logical host name.

For each of the `rdisp` parameters, replace the physical host name with the central instance logical host name. For the `SAPDBHOST` parameter, enter the logical host name of the database. If the central instance and database are installed on the same logical host, enter the central instance logical host name. If the database is

installed on a different logical host, use the database logical host name instead. In this example, *Cllogicalhost* represents the logical host name of the central instance, and *DBlogicalhost* represents the logical host name of the database:

```
$ vi /sapmnt/<SAPSID>/profile/DEFAULT.PFL
...
SAPDBHOST                =DBlogicalhost
rdisp/mshost              =Cllogicalhost
rdisp/sna_gateway         =Cllogicalhost
rdisp/vbname              =Cllogicalhost_<SAPSID>_00
rdisp/enqname             =Cllogicalhost_<SAPSID>_00
rdisp/btcname            =Cllogicalhost_<SAPSID>_00
...
```

d. Rename the `.dbenv_physicalhost.csh` file to `.dbenv.csh`.

```
$ mv .dbenv_physicalhost.csh .dbenv.csh
```

e. Rename the `.sapenv_physicalhost.csh` file to `.sapenv.csh`.

```
$ mv .sapenv_physicalhost.csh .sapenv.csh
```

f. Edit the `startsap` and `stopsap` aliases in the `.sapenv.csh` file to specify the central instance logical host in the `'set hostname='` field.

```
...
set hostname='Cllogicalhost'
...
```

g. Modify the `.dbenv.csh` file to specify use of TCP for `$INFORMIXSERVER` and to use `ping(1M)` to check the status of the database logical host.

This is necessary to enable dynamic reset of the `$INFORMIXSERVER` parameter in case of switchover or failover.

Note – In asymmetric configurations, the use of TCP and loopback might reduce performance. If so, you can set `$INFORMIXSERVER` to use shared memory instead.

Modify the file to point `$INFORMIXDIR` to the shared disk, and to modify `$INFORMIXSERVER` to use TCP and `ping(1M)`. The resulting file should look similar to the following sample, in which the modified fields are in bold type:

```
...
setenv INFORMIXDIR /var/opt/informix
setenv ONCONFIG onconfig.Cllogicalhost.<sapsid>
...
  case Sun*:
    setenv INFORMIXSHMBASE 0x01000000
    setenv LC_CTYPE iso_8859_1
    setenv INFORMIXSQLHOSTS $INFORMIXDIR/etc/sqlhosts
# use TCP for connection prototype always because connection
# cannot be reset dynamically between shared memory and TCP in
# the Sun Cluster environment.
    setenv INFORMIXSERVER `grep 'Cllogicalhost<sapsid>.*ontlitcp'
$INFORMIXSQLHOSTS | awk '{print $1}'`

    /usr/sbin/ping DBlogicalhost >& /dev/null
    if ( $status != 0 ) then
        echo dbserver DBlogicalhost is not alive.
    endif
```

h. Modify the `/export/home/<sapsid>adm/.rhosts` file to allow user `<sapsid>adm` to access the database from all nodes.

Create entries similar to the following, with one entry for each physical and logical host in the cluster.

```
phys-hahost1 <sapsid>adm
phys-hahost2 <sapsid>adm
Cllogicalhost <sapsid>adm
DBlogicalhost <sapsid>adm
```

7. Create the `/usr/sap/tmp` directory on all nodes.

The `saposcol` program will rely on this directory.

8. Copy the SAP-specific `/etc/services` entries from the node on which SAP and Informix are installed to the `/etc/services` files on all other nodes.

Copy these entries from the `/etc/services` files:

```
sapms<SID>      3601/tcp
sapdp00         3200/tcp
sapdp00s       4700/tcp
sapgw00        3300/tcp
sapgw00s       4800/tcp
```

9. Test the SAP installation.

Test the SAP installation by manually shutting down SAP, manually switching the logical host between the potential master nodes, and then manually starting SAP on the backup node. This will verify that all kernel parameters, service port entries, file systems and mount points, and user/group permissions are properly set on all potential masters of the logical hosts.

- a. As user `<sapsid>adm`, start the central instance and database.

```
# startsap
```

- b. Run the GUI and verify that SAP comes up correctly.

```
# su - <sapsid>adm
$ setenv DISPLAY workstation:0
$ sapwin phys-hahost1 instancenumbr
```

- c. Verify that SAP can connect to the database.

```
# su - <sapsid>adm
$ R3trans -d
```

- d. Run the `saplicense` utility to get a `CUSTOMER KEY` for the current node.

You need a SAP license for all potential masters of the central instance logical host.

- e. Stop SAP and the database.

```
# su - <sapsid>adm
$ stopsap
```

- f. On all nodes (except the application servers), set up links for the Informix library files.

You must be root to perform these commands.

```
# unlink iosm07a.so
# unlink ipldd07a.so
# unlink ismdd07b.so
# ln -s /var/opt/informix/lib/iosm07a.so /usr/lib/iosm07a
# ln -s /var/opt/informix/lib/ipldd07a.so /usr/lib/ipldd07a.so
# ln -s /var/opt/informix/lib/ismdd07b.so /usr/lib/ismdd07b.so
```

10. For each remaining node that is a potential master of the central instance logical host, switch the central instance logical host to that node and repeat the test sequence described in Step 9.

```
# scadmin switch clustername phys-hahost2 Cllogicalhost
```

Next, proceed to “How to Configure Sun Cluster HA for Informix”.

▼ How to Configure Sun Cluster HA for Informix

1. On all nodes, bring up the Informix database and make sure it's running.

```
# oninit
...
# dbaccess
```

2. From only one node, as root, register Sun Cluster HA for Informix.

```
# hareg -s -r informix [-h DBlogicalhost]
```

3. From only one node, activate Sun Cluster HA for Informix.

```
# hareg -y informix
```


4. From only one node, bring Informix under the control of Sun Cluster HA for Informix.

See the `hainformix(1M)` man page for details.

```
# hainformix insert onconfig.CIlogicalhost.<sapsid> DBlogicalhost \  
60 10 120 300 sysmaster CIlogicalhost<sapsid>tcp
```

5. From only one node, bring Sun Cluster HA for Informix into service.

```
# hainformix start onconfig.CIlogicalhost.<sapsid>
```

6. Verify that the database is working properly under the control of Sun Cluster HA for Informix.

Perform a switchover of the database and make sure the `oninit` processes are stopped on the old master and restarted on the new master. The database should be accessible from all potential masters.

Next, proceed to “How to Configure Sun Cluster HA for SAP (SAP With Informix)”.

▼ How to Configure Sun Cluster HA for SAP (SAP With Informix)

1. Register the Sun Cluster HA for SAP data service by running the `hareg(1M)` command.

If you configure one logical host for the central instance and one for the database instance, you must register and unregister data services in the following order to preserve the order in which Informix and SAP are started and stopped:

```
hareg -n informix (do this if you had informix registered before sap)
```

```
hareg -u informix (do this if you had informix registered before sap)
```

```
hareg -s -r sap -h CIlogicalhost
```

```
hareg -s -r informix -h DBlogicalhost
```

Activate the data service and bring it under the control of Sun Cluster using the procedures “How to Enable SAP and the Database to Run in the Cluster Environment (SAP With Informix)” and “How to Configure Sun Cluster HA for Informix”.

Run the command on only one node. For example:

```
# hareg -s -r sap -h CIlogicalhost
```

2. Verify that all nodes are running in the cluster.

3. **Create a new Sun Cluster HA for SAP instance using the `hadsconfig(1M)` command.**

The `hadsconfig(1M)` command is used to create, edit, and delete instances of the Sun Cluster HA for SAP data service. The configuration parameters are described in “Configuration Parameters for Sun Cluster HA for SAP (SAP With Informix)”.

Run this command on only one node, while all nodes are running in the cluster:

```
# hadsconfig
```

4. **Stop the central instance before starting SAP under the control of Sun Cluster HA for SAP.**

```
# su - <sapsid>adm
$ stopsap r3
```



Caution – The SAP central instance must be stopped before Sun Cluster HA for SAP is turned on.

5. **Turn on the Sun Cluster HA for SAP instance.**

```
# hareg -y sap
```

6. **Test switchover of Sun Cluster HA for SAP.**

For example:

```
# scadmin switch clustername phys-hahost2 Clogicalhost
```

7. **(Optional) If you have application servers or a test/development system, customize and test the `hasap_start_all_instances` and `hasap_stop_all_instances` scripts.**

See “Configuration Options for Application Servers and Test/Development Systems” for details. Test switchover of Sun Cluster HA for SAP, and verify start and stop of application servers. Verify that the test/development system stops when the central instance logical host is switched to the test/development system physical host.

```
# scadmin switch clustername phys-hahost1 Clogicalhost
```

Next, proceed to “Setting Data Service Dependencies for SAP (SAP With Oracle)”, if you want to specify the start and stop order of data services within a logical host.

Configuration Parameters for Sun Cluster HA for SAP (SAP With Informix)

This section describes the information you supply to `hadsconfig(1M)` to create configuration files for the Sun Cluster HA for SAP data service. The `hadsconfig(1M)` command uses templates to create these configuration files. The templates contain some default, some hard coded, and some unspecified parameters. You must provide values for all parameters that are unspecified.

The fault probe parameters, in particular, can affect the performance of Sun Cluster HA for SAP. Tuning the probe interval value too low (increasing the frequency of fault probes) might encumber system performance, and also might result in false takeovers or attempted restarts when the system is simply slow.

Configure Sun Cluster HA for SAP by supplying the `hadsconfig(1M)` command with parameters listed in the following table.

TABLE 1-15 Sun Cluster HA for SAP Configuration Parameters (SAP With Informix)

Name of the Instance	Nametag used internally as an identifier for the instance. The log messages generated by Sun Cluster refer to this nametag. The <code>hadsconfig(1M)</code> command prefixes the package name to the value you supply here. You can use the <code><SAPSID></code> for this nametag. For example, if you specify <code>HA1</code> , <code>hadsconfig(1M)</code> produces <code>SUNWscsap_HA1</code> .
Logical Host	Name of the logical host that provides service for this instance of Sun Cluster HA for SAP. This name should be the logical host name for the central instance.
Time Between Probes	The interval, in seconds, of the fault probing cycle. The default value is 60 seconds.
SAP SID	This is the SAP system name or <code><SAPSID></code> .
Central Instance ID	This is the SAP system number or central instance ID. The default value is 00.
SAP Admin Login Name	The name used by Sun Cluster HA for SAP to log in to the SAP central instance administrative account. This name must exist on all central instance and application server hosts. This is the <code><sapsid>adm</code> . For example, <code>ha1adm</code> .
Database Admin Login Name	This is the SAP database administrator’s account. For SAP with Informix, this is <code>informix</code> .
Database Logical Host Name	Name of the logical host for the database used by SAP. This might be the same as the logical host name used for the central instance, depending on your configuration.

TABLE 1-15 Sun Cluster HA for SAP Configuration Parameters (SAP With Informix) (Continued)

Log Database Warnings	Possible values are y or n. If set to y and the Sun Cluster HA for SAP probe detects that it cannot connect to the database during a probe cycle, a warning message appears saying the database is unavailable. For example, this occurs if the database logical host is in maintenance mode or if the database is being relocated to another node in the cluster. If the parameter is set to n, then no messages appear if the probe cannot connect to the database.
Central Instance Start Retry Count	This must be an integer greater than or equal to 1. The default value is 10. This is the number of times Sun Cluster HA for SAP should attempt to start the central instance before giving up. This value is also the number of times the Sun Cluster HA for SAP fault monitor will probe in grace mode before entering normal probe mode. While in grace mode, the probe will not perform a restart or initiate a failover of the central instance if the probe detects that the central instance is not yet up. Instead, the fault monitor will report the status of all probes and will continue in grace mode until all probes pass, or until the retry count has been exhausted.
Central Instance Start Retry Interval	This is the number of seconds Sun Cluster HA for SAP should wait between each attempt to start the central instance. This value is also the number of seconds that the Sun Cluster HA for SAP fault monitor will sleep (between probe attempts) while in grace mode. The default value is 30.
Time Allowed to Stop All Instances Before Central Instance Starts	This must be an integer greater than or equal to 0. The default value is 60. This parameter dictates for how much time (in seconds) the <code>hasap_stop_all_instances</code> script should be run before starting the central instance. If set to 0, then <code>hasap_stop_all_instances</code> is run in the background while the central instance is being started. If set to a positive integer, then <code>hasap_stop_all_instances</code> is run for that amount of time in the foreground before the central instance is started.
Allow the Central Instance to Start if Foregrounded Stop All Instances Returns Error	This flag should be set to either y or n. The default value is n. This value determines whether the central instance should be started in the case where the <code>hasap_stop_all_instances</code> script returns a non-zero exit code or does not complete in the time specified by the "Time Allowed to Stop All Instances Before Central Instance Starts" parameter. If set to n and the value for "Time Allowed to Stop All Instances Before Central Instance Starts" is greater than 0, and if the <code>hasap_stop_all_instances</code> script does not complete in the time configured above or the <code>hasap_stop_all_instances</code> script returns a non-zero exit status, the central instance will not be started and the fault monitors will take action based on the other configuration parameters. If set to y, then the central instance will be started regardless of whether <code>hasap_stop_all_instances</code> returns an error code or finishes within the timeout specified above.
Number of Central Instance Restarts on Local Node	This must be an integer greater than or equal to 0. The default value is 1. This dictates how many times the SAP central instance will be restarted on the local node before giving up, after a failure has been detected. When this number of restarts has been exhausted, Sun Cluster HA for SAP either issues a failover request, if permitted by the "Allow Central Instance Failover" parameter, or does nothing to correct the failure detected by the fault monitor.

TABLE 1-15 Sun Cluster HA for SAP Configuration Parameters (SAP With Informix) (Continued)

Number of Probe Successes to Reset the Restart Count	This parameter should be an integer that is greater than or equal to 0. The default value is 60. If set to a positive integer, then after that many consecutive successful probes, the count of restarts done so far on the local node will be reset to 0. For example, if the value for “Number of Central Instance Restarts on Local Node” parameter is 1 and the value for “Number of Probe Successes to Reset the Restart Count” is 60, then after the first failure occurs, the probe will try to restart the central instance on the local node. If this restart succeeds, then after 60 successful probes, the restart count will be reset to 0, allowing the probe to do another restart if it detects another failure. If the parameter “Number of Probe Successes to Reset the Restart Count” is set to 0, then the restart count is never reset. This means that the number of restarts set in the parameter “Number of Central Instance Restarts on Local Node” is the absolute number of restarts that will be done on the local node before failing over.
Allow Central Instance Failover	Possible values are <i>y</i> or <i>n</i> . The default value is <i>y</i> . If set to <i>y</i> and Sun Cluster HA for SAP detects an error in the SAP instance it is monitoring and the “Number of Central Instance Restarts on Local Node” has been exhausted, then Sun Cluster HA for SAP issues a request to relocate the instance’s logical host to another cluster node. If this flag is set to <i>n</i> , then even if an error is detected and all of the local restarts have been exhausted, Sun Cluster HA for SAP will not cause a relocation of this instance’s logical host. When this occurs, the central instance is left in the failed state, and the probe exits.

Setting Data Service Dependencies for SAP (SAP With Informix)

Setting a dependency with `hasap_dbms` is only necessary to specify the order that data services are started and stopped within a single logical host. There is no mechanism for setting dependencies for data services configured on two different logical hosts.

If Sun Cluster HA for Informix or Sun Cluster HA for NFS are configured on the same logical host as Sun Cluster HA for SAP, then you should set a dependency for Sun Cluster HA for SAP on those data services. You can use the `hasap_dbms` command to create or remove such a dependency. These dependencies affect the order that the services are started and stopped. Sun Cluster HA for Informix and Sun Cluster HA for NFS should always be started before Sun Cluster HA for SAP is started. Similarly, Sun Cluster HA for SAP should always be stopped before the other data services are stopped.



Caution – If Sun Cluster HA for Informix or Sun Cluster HA for NFS is not configured on the same logical host as Sun Cluster HA for SAP, then do not use the `hasap_dbms` command.

▼ How to Set a Data Service Dependency for SAP (SAP With Informix)

To set a data service dependency, issue one of the `hasap_dbms` commands described below.

Note – The `hasap_dbms` command can be used only when Sun Cluster HA for SAP is registered but is in the off state. Run the command on only one node, while that node is a member of the cluster. See the `hasap_dbms(1M)` man page for more information.



Caution – If the `hasap_dbms(1M)` command returns an error stating that it cannot add rows to or update the CCD, it might be because another cluster utility is also trying to update the CCD. If this occurs, re-run `hasap_dbms(1M)` until it runs successfully. After the `hasap_dbms(1M)` command runs successfully, verify that all necessary rows are included in the resulting CCD by running the command `hareg -q sap`.

If the `hareg(1M)` command returns an error, then first restore the original method timeouts by running the command `hasap_dbms -f`. Second, restore the default dependencies by running the command `hasap_dbms -r`. After both commands complete successfully, retry the original `hasap_dbms(1M)` command to configure new dependencies and method timeouts. See the `hasap_dbms(1M)` man page for more information.

1. Set the data service dependency using one of the following commands.

If you are using only Sun Cluster HA for NFS and Sun Cluster HA for SAP on the same logical host, use the following command:

```
# /opt/SUNWcluster/ha/sap/hasap_dbms -d nfs
```

If you are using only Sun Cluster HA for Informix and Sun Cluster HA for SAP on the same logical host, use the following command:

```
# /opt/SUNWcluster/ha/sap/hasap_dbms -d informix
```

If you are using Sun Cluster HA for Informix, Sun Cluster HA for NFS, and Sun Cluster HA for SAP on the same logical host, use the following command:

```
# /opt/SUNWcluster/ha/sap/hasap_dbms -d informix,nfs
```

2. Check the dependencies set for Sun Cluster HA for SAP using the following command:

```
# hareg -q sap -D
```

▼ How to Remove a Data Service Dependency for SAP (SAP With Informix)

The dependencies set for Sun Cluster HA for SAP can be removed by running the `hasap_dbms -r` command. Issuing this command causes all of the dependencies set for Sun Cluster HA for SAP to be removed.

Note – The `hasap_dbms` command can be used only when Sun Cluster HA for SAP is registered but is in the off state. Run the command on only one node, while that node is a member of the cluster. See the `hasap_dbms (1M)` man page for more information.



Caution – If the `hasap_dbms (1M)` command returns an error stating that it cannot add rows to or update the CCD, it might be because another cluster utility is also trying to update the CCD. If this occurs, re-run `hasap_dbms (1M)` until it runs successfully. After the `hasap_dbms (1M)` command runs successfully, verify that all necessary rows are included in the resulting CCD by running the command `hareg -q sap`.

If the `hareg (1M)` command returns an error, then first restore the original method timeouts by running the command `hasap_dbms -f`. Second, restore the default dependencies by running the command `hasap_dbms -r`. After both commands complete successfully, retry the original `hasap_dbms (1M)` command to configure new dependencies and method timeouts. See the `hasap_dbms (1M)` man page for more information.

1. Remove all of the dependencies set for Sun Cluster HA for SAP, using the following command:

```
# /opt/SUNWcluster/ha/sap/hasap_dbms -r
```

2. Check the dependencies set for Sun Cluster HA for SAP, using the following command:

```
# hareg -q sap -D
```

